AGT

Algebraic & Geometric Topology

Volume 26 (2026)

Homoclinic leaves, Hausdorff limits and homeomorphisms

Ian Biringer and Cyril Lecuire

msp

Algebraic & Geometric Topology **26**:5 (2026) 1801–1866
DOI: 10.2140/agt.2026.26.1801
Published: May 29, 2026

Homoclinic leaves, Hausdorff limits and homeomorphisms

IAN BIRINGER AND CYRIL LECUIRE

We show that except for one exceptional case, a lamination on the boundary of a handlebody H is commensurable to a Hausdorff limit of meridians if and only if it is commensurable to a lamination with a "homoclinic leaf". This is an "if and only if" version of a theorem called Casson's criterion. Applications of our techniques include a characterisation of when a nonminimal lamination is a Hausdorff limit of meridians, in terms of properties of its minimal components, and a related characterisation of which reducible self-homeomorphisms of ∂H have powers that extend to subcompression bodies of H.

1 Introduction

Let H be a 3-dimensional handlebody[1] with genus $g \geq 2$ and let $S := \partial H$, which we usually equip with a reference hyperbolic metric. A simple closed curve m on S is called a *meridian* if it bounds an embedded disk in H but not in S. These curves play a fundamental role when studying the topology and geometry of handlebodies and more general 3-manifolds with compressible boundary. Our work here is centred around geodesic laminations $\lambda \subset S$ that are *Hausdorff limits of meridians*, i.e., λ where there is a sequence of geodesic meridians (m_i) on S such that $m_i \to \lambda$ in the Hausdorff topology on the set of closed subsets of S. See Section 2.8 for definitions.

Limits of meridians are important in the geometry and topology of hyperbolic 3-manifolds [2; 13; 42], convergence of sequences of Kleinian groups [12; 35; 38] and the action of $\mathrm{Out}(F_n)$ on the character variety $\chi(F_n, \mathrm{PSL}_2(\mathbb{C}))$; see [31; 36]. In many of the above references, the focus is on a closely related set called the "Masur domain", defined in [46], but to work with that one usually has to understand Hausdorff limits of meridians too.

We have three main goals. First, we investigate the relationship between Hausdorff limits of meridians and "homoclinic leaves", defined by Casson in unpublished notes and later exploited by Otal [56] in his thesis. Using this analysis, we study Hausdorff limits of meridians that are *nonminimal* laminations, in essence reducing their classification to the minimal case. Finally, using similar techniques we characterise when a reducible homeomorphism $f : S \to S$ extends (partially, up to a power) into H, in terms of extension properties of its Nielsen–Thurston components.

Another motivation for writing this paper is that it provides all the technical topological machinery necessary for our forthcoming follow-up paper, called *Iterations in Schottky space*. In that paper, we

[1] The body of this paper is written in greater generality, with (H, S) replaced by a compact, orientable, 3-manifold M with hyperbolisable interior, together with an essential connected subsurface $S \subset \partial M$ such that the multicurve ∂S is incompressible in M. However, everything we do is just as interesting in the handlebody case.

MSC2020: 57K30, 57K32.

© 2026 The Authors, under license to MSP (Mathematical Sciences Publishers). Distributed under the Creative Commons Attribution License 4.0 (CC BY). Open Access made possible by subscribing institutions via Subscribe to Open.

fix a point $[X] \in \mathcal{T}(S)$ and describe the geometric limits of sequences of convex cocompact hyperbolic structures on H whose conformal boundaries are the iterates $f^n([X]) \in \mathcal{T}(S)$. This is a compressible version of Jeff Brock's thesis [8]. Brock's result was a first step toward work of Minsky [52] and Brock, Canary and Minsky [9], in which they developed bilipschitz models for hyperbolic structures on $S \times \mathbb{R}$ in terms of their end invariants. Our work on iterations in Schottky space will likewise be a step toward a bilipschitz model theory for hyperbolic structures on a handlebody.

1.1 Homoclinic leaves

Let $\widetilde{H}$ be the universal cover of H, which is homeomorphic to a thickened infinite tree. We equip H with any Riemannian metric, and equip $\widetilde{H}$ with the lift of this metric. A geodesic $\ell \subset S$ is called *homoclinic*[2] if it has a (possibly periodic) parametrisation $h : \mathbb{R} \to \lambda \subset S = \partial H$ that lifts to a path $\tilde{h} : \mathbb{R} \to \partial \widetilde{H}$ where there are sequences $s_i, t_i \in \mathbb{R}$ such that

$$|s_i - t_i| \to \infty \quad \text{and} \quad \sup_i d_{\widetilde{H}}(\tilde{\ell}(s_i), \tilde{\ell}(t_i)) < \infty.$$

Intuitively, ℓ is homoclinic if it travels very inefficiently in H, even though on the boundary of H it is a geodesic. As an example, one can check that a simple closed geodesic ℓ on S is homoclinic if and only if ℓ is a meridian. Similarly, a biinfinite geodesic that spirals onto a meridian will also be homoclinic.

The following was first written by Otal [56], building on work in an unpublished manuscript of Casson.

Theorem 1.1 ("Casson's criterion") *Let λ be a geodesic lamination on S. If λ is a Hausdorff limit of meridians, then λ contains a homoclinic leaf.*

The converse of Theorem 1.1 is not true. Of course, any Hausdorff limit of meridians is connected, and there are disconnected laminations on S that have homoclinic leaves, e.g., the union of two disjoint simple closed curves, one of which is a meridian. There are also connected laminations containing a meridian as a leaf (say) that are not even Hausdorff limits of simple closed curves. And there are also some more interesting examples that are related more intimately to the structure of H; see the beginning of Section 7.

Nevertheless, in certain particular cases homoclinic leaves have been used to construct limits of sequences of meridians (or annuli) in H; see [41; 43; 56] and particularly [38; 42]. And in Otal's thesis, Casson's criterion was even stated (incorrectly) in passing as an "if and only if". Our goal is to state precisely a result that is closest to a converse of Theorem 1.1. This will tie together, extend and explain the partial results in the papers above.

We say that μ_1, μ_2 are *strongly commensurable* if they contain a common sublamination ν such that for both i, the difference $\mu_i \setminus \nu$ is the union of finitely many isolated leaves, none of which are simple closed curves. Additionally, let's say that a lamination λ on S is *exceptional* if S has genus 2, there is a

[2]This is not quite what Casson and Otal call "homoclinic" (in French, "homoclinique"), but rather what Otal calls "faiblement homoclinique", or "weakly homoclinic". However, the definition we give has been adopted in almost all subsequent papers. See Section 5.1 for an explanation of the difference between the two definitions.

separating meridian m on S that does not intersect λ transversely, and λ intersects transversely the two nonseparating meridians disjoint from m.

Theorem 1.2 (a bidirectional Casson's criterion; see Theorem 7.2) *If λ is a geodesic lamination on S that is not exceptional, then λ is strongly commensurable to a Hausdorff limit of meridians if and only if it is strongly commensurable to a lamination with a homoclinic leaf.*

This is the first part of Theorem 7.2. In our view, it is the strongest converse for Casson's criterion that is likely to be true for arbitrary laminations.

The main tool in the proof of Theorem 1.2 is a complete characterisation of the minimal laminations onto which the two ends of a homoclinic simple geodesic on S can accumulate.

Theorem 1.3 (limits of homoclinic geodesics; see Corollary 6.2) *Suppose that h is a homoclinic simple biinfinite geodesic on S and that the two ends of h limit onto minimal laminations $\lambda_-, \lambda_+ \subset S$. Then either*

(1) *the two ends of h are asymptotic on S,*

(2) *one of λ_-, λ_+ is an intrinsic limit of meridians, or*

(3) *λ_-, λ_+ are contained in incompressible subsurfaces $S_-, S_+ \subset S$ that bound an essential interval bundle $B \subset H$ through which λ_- and λ_+ are homotopic.*

Here, a minimal lamination $\lambda \subset S$ is *an intrinsic limit of meridians* if it is strongly commensurable to the Hausdorff limit of a sequence of meridians that are contained in the smallest essential subsurface $S(\lambda) \subset S$ containing λ; see Proposition 5.12 for a number of equivalent definitions. We refer the reader to Theorem 6.1 and Corollary 6.2 for more precise and more general versions of the above that apply both to homoclinic biinfinite geodesics, and also to pairs of "mutually homoclinic" geodesic rays on S, as well as many examples.

Interval bundles as in (3) are essential to the study of meridians on handlebodies, and it is no surprise that they appear in Theorem 1.3. For example, subsurfaces bounding such interval bundles are the "incompressible holes" studied by Masur and Schleimer [49], and interval bundles appear frequently in Hamenstädt's work on the disk set; see, e.g., [25; 26; 27]. We note that the interval bundles B appearing in Theorem 1.3 may be twisted interval bundles over nonorientable surfaces, in which case $\lambda_- = \lambda_+$ and $S_- = S_+$. See Proposition 4.5 for background on interval bundles, and Section 6 for examples of (3).

1.2 Hausdorff limits via their minimal sublaminations

The previous two theorems suggest that if a lamination λ that is a Hausdorff limits of meridians, one might expect to see minimal sublaminations of λ that are intrinsic limits of meridians, or pairs of components that are homotopic through essential interval bundles in H. In fact, we show the following.

Theorem 1.4 (the second part of Theorem 7.2) *Suppose that $\lambda \subset S$ is a nonexceptional geodesic lamination that is a finite union of minimal components. Then λ is strongly commensurable to a Hausdorff limit of meridians if and only if either*

(1) *λ is disjoint from a meridian on S*,

(2) *some component of λ is an intrinsic limit of meridians*, *or*

(3) *there are components $\lambda_\pm \subset \lambda$ that fill incompressible subsurfaces $S_\pm \subset S$, such that $S_\pm$ bound an essential interval bundle $B \subset H$, the laminations $\lambda_\pm$ are essentially homotopic through B, and there is a compression arc α for B that is disjoint from λ.*

In (3), a *compression arc* for B is an arc from ∂S_- to ∂S_+ that is homotopic in H, keeping its endpoints in $\partial S_\pm$, to a fibre of the interval bundle B. See Section 2.7 and Figure 2 for more explanation.

Note that Theorem 1.4 does not say anything interesting about which minimal filling laminations λ on S are strongly commensurable to Hausdorff limits of meridians. Indeed, for minimal filling laminations, it is not clear that there should be an easy way to "identify" Hausdorff limits of meridians. The point of Theorem 1.4, though, is that it reduces the characterisation of Hausdorff limits of meridians to the minimal filling case. We note that one should also be able to replace the part of the proof of Theorem 1.4 that references homoclinic geodesics with arguments inspired by Masur and Schleimer's paper [49].

1.3 Extension of reducible maps to compression bodies

Lackenby [39] studied a generalisation of Dehn filling in which one starts with a compact 3-manifold M with (say, connected) genus-g boundary Σ, and glues a genus-g handlebody H via a homeomorphism $\phi : \Sigma \to S := \partial H$. Among other results, he showed that when an arbitrary ϕ is precomposed by a high power of a sufficiently "generic" homeomorphism of S, the gluing is hyperbolic:

Theorem 1.5 (Lackenby [39]) *Suppose that M is hyperbolisable and acylindrical with incompressible boundary*, *and $f : S \to S$ is a homeomorphism such that no nonzero power of f extends to a nontrivial subcompression body of H. Then for infinitely many n, the gluing $M \cup_{\phi\circ f^n} H$ is hyperbolic.*[3]

See also the work of Namazi and Souto [55] and Brock, Minsky, Namazi and Souto [10] for more general hyperbolisation theorems inspired by Lackenby's result.

Above, a *subcompression body* of H is a 3-dimensional submanifold $C \subset H$ with $S \subset \partial C$ that is obtained by choosing a finite collection Γ of disjoint meridians on S, taking a regular neighbourhood of S and a collection of discs in H with boundary Γ, and adding in any complementary components that are topological 3-balls. We say C is obtained by *compressing* Γ. We usually consider subcompression bodies only up to isotopy, and we allow the case that $\Gamma = \varnothing$, in which case we recover the *trivial subcompression body*, which is just a regular neighbourhood of S. See Section 2.4 for details.

In light of Lackenby's theorem, it is natural to investigate which homeomorphisms $f : S \to S$ have powers that extend to a nontrivial subcompression body of H. Biringer, Johnson and Minsky [2] showed that for pseudo-Anosov f, this condition is equivalent to the attracting lamination λ_+ being a (projective) limit of meridians. In fact, Lackenby also proved a version of his theorem for pseudo-Anosovs with the latter condition on λ_+, not knowing the equivalence of the two properties at that time.

[3] Lackenby's paper was written pregeometrisation, so he actually proves that the hypotheses of the hyperbolisation theorem are satisfied, not that the gluing is hyperbolic.

Here, we show that extension of powers of a homeomorphism $f : S \to S$ to subcompression bodies can be detected by looking at extension of powers of its components in the Nielsen–Thurston decomposition. More precisely, recall that f is *pure* if there are disjoint, nonisotopic,[4] essential subsurfaces $S_i \subset S$, such that $f = \mathrm{id}$ on $S_{\mathrm{id}} := S \setminus \bigcup_i S_i$, and where for each i, if we set $f_i := f|_{S_i}$, then either

(1) S_i is an annulus and f_i is a nonzero power of a Dehn twist, or

(2) f_i is a pseudo-Anosov map on S_i.

It follows from the Nielsen–Thurston classification [21] that every homeomorphism of S has a power that is isotopic to a pure homeomorphism.

Theorem 1.6 (partial extension of reducible maps; see Theorem 9.2) *Let $f : S \to S$ be a pure homeomorphism. Then f has a power that extends to a nontrivial subcompression body of H if and only if either*

(1) *there is a meridian in S_{id},*

(2) *for some i, the map $f_i : S_i \to S_i$ has a power that extends to a nontrivial subcompression body of H that is obtained by compressing a set of meridians in S_i, or*

(3) *there are (possibly equal) indices i, j such that S_i, S_j bound an essential interval bundle B in H, such that some power of $f|_{S_i \cup S_j}$ extends to B, and there is a compression arc α for B whose interior lies in S_{id}.*

Moreover, Theorem 8.1 says that when f_i is a pseudo-Anosov map on S_i, then (2) above is equivalent to the condition that the attracting lamination of f_i is a (projective) limit of meridians in S_i. This is a relative version of Biringer, Johnson and Minsky's article [2], that generalises their theorem from pseudo-Anosov maps on S to partial pseudo-Anosovs.

In (3), note that if (for simplicity) B is a trivial interval bundle, then $f|_{S_i \cup S_j}$ extends to B exactly when f_i, f_j become isotopic maps when S_i, S_j are identified through B. More generally, a power of $f|_{S_i \cup S_j}$ extends to B when f_j is obtained from f_i by multiplying by a periodic map that commutes with f_i.

1.4 Other results of interest

There are two other theorems in this paper that we should mention in the introduction.

In Section 3 we study the *disk set* $\mathcal{D}(S, M)$ of all isotopy classes of meridians in an essential subsurface $S \subset \partial M$ with ∂S incompressible, where here M is a compact, irreducible 3-manifold with boundary. We show in Proposition 3.1 that either $\mathcal{D}(S, M)$ is *small*, meaning that it is either empty, has a single element, or has a single nonseparating element and infinitely many separating elements that one can explicitly describe, or $\mathcal{D}(S, M)$ is *large*, meaning that it has infinite diameter in the curve complex $\mathcal{C}(S)$. This result will probably not surprise any experts, but we have never seen it in the literature.

In Section 4 we show how essential interval bundles in a compact 3-manifold with boundary M can be seen in the limit sets in $\partial\mathbb{H}^3$ associated to hyperbolic metrics on $\mathrm{int}(M)$. This picture was originally

[4] This is only important for annuli.

known to Thurston [59], and was studied previously under more restrictive assumptions by Walsh [60] and Lecuire [40]. We need a more general theorem in the proof of Theorem 6.1: in particular, we need a version that allows accidental parabolics. This is Theorem 4.1. Our proof is also more direct and more elementary than those of [40; 60]. See Section 4 for more context and details.

1.5 Outline of the paper

Section 2 contains all the necessary background for the rest of the paper. We discuss the curve complex, the disc set, compression bodies, interval bundles, the Jaco–Shalen and Johannson characteristic submanifold theory, compression arcs, and geodesic laminations. Sections 3 and 4 are described in the previous subsection. Section 5 contains a discussion of homoclinic geodesics, intrinsic limits of meridians, and some of their basic properties. In Section 6 we explain how to build essential discs and annulus from a homotopic leaf. Section 7 is devoted to Theorem 7.2, which characterises Hausdorff limits of meridians and their structure; see Theorem 1.2 above. Section 8 contains our extension of [2] to partial pseudo-Anosovs, and Section 9 contains the proof of Theorem 9.2, which generalises Theorem 1.6 above.

2 Preliminaries

2.1 Subsurfaces with geodesic boundary

Suppose S is a finite-type hyperbolic surface with geodesic boundary. A *connected subsurface with geodesic boundary* in S is by definition either

(1) a simple closed geodesic X on S, which is the degenerate case, or

(2) an immersed surface $X \to S$ such that the restriction to $\mathrm{int}(X)$ and to each component of ∂X is an embedding, and where each component of ∂X maps to a simple closed geodesic on S.

In (2), the point is that our surface is basically an embedding, except that we allow two boundary components of X to map to the same geodesic in S. We will usually suppress the immersion and write $X \subset S$, abusing notation. We consider $X, Y \subset S$ to be *equal* if they are either the same simple closed geodesic, or if they are both immersions as in (2) and the interiors of their domains have the same images. We say X, Y are *essentially disjoint* if either

- X, Y are disjoint simple closed geodesics,
- X is a simple closed geodesic, Y is not, and X is disjoint from $\mathrm{int}(Y)$, or vice versa with X, Y exchanged, or
- X, Y have nonempty disjoint interiors.

More generally, we define a (possibly disconnected) *subsurface with geodesic boundary* in S to be a finite union of essentially disjoint connected subsurfaces with geodesic boundary.

Any connected essential subsurface $T \subset S$ that is not an annulus homotopic into a cusp of S determines a unique connected subsurface with geodesic boundary X such that the images of $\pi_1 T$ and $\pi_1 X$ in $\pi_1 S$

are conjugate. Here, we say that X is obtained by *tightening* T. More generally, we can tighten a disconnected T to a disconnected X by tightening all its components.

Tightening is performed as follows. If T is an annulus, then we let X be the unique simple closed geodesic homotopic to the core curve of T. Otherwise, we obtain X by homotoping T so that every component of ∂T is either geodesic or bounds a cusp in $S \setminus T$, and then adding in any components of $S \setminus T$ that are cusp neighbourhoods. Alternatively, let $\widetilde{T}$ be a component of the preimage of T in the universal cover $\widetilde{S}$, which is isometric to a convex subset of $\mathbb{H}^2$, let $\Lambda_T \subset \partial \mathbb{H}^3$ be the set of limit points of $\widetilde{T}$ on $\partial_\infty \mathbb{H}^2$, and let $\widetilde{X}$ be the convex hull of Λ_T within $\widetilde{S}$. Then $\widetilde{X}$ projects to an X as desired.

Conversely, suppose X is a subsurface with geodesic boundary in S. Then there is a compact essential subsurface $T \hookrightarrow S$, unique up to isotopy and called a *resolution* of X, that tightens to X. When X is a simple closed geodesic, we take T to be a regular neighbourhood of X. Otherwise, construct T by deleting half-open collar neighbourhoods of all boundary components of X, and deleting open neighbourhoods of all cusps of T.

Note that subsurfaces with geodesic boundary X, Y are essentially disjoint if and only if they admit disjoint resolutions.

2.2 The curve complex

Let S be a compact orientable surface, possibly with boundary, and assume that S is not an annulus.

Definition 2.1 The *curve complex* of S, written $\mathcal{C}(S)$, is the graph whose vertices are homotopy classes of nonperipheral, essential simple closed curves on S and whose edges connect homotopy classes that intersect minimally.

When S is a 4-holed sphere, minimally intersecting simple closed curves intersect twice, while on a punctured torus they intersect once. Otherwise, edges in $\mathcal{C}(S)$ connect homotopy classes that admit disjoint representatives.

Masur and Minsky [47] have shown that the curve complex is Gromov hyperbolic, when considered with the path metric in which all edges have unit length. Klarreich [37] (see also [24]) showed that the Gromov boundary $\partial_\infty \mathcal{C}(S)$ is homeomorphic to the space of *ending laminations* of S: i.e., filling, measurable geodesic laminations on S with the topology of Hausdorff superconvergence.

2.3 The disc set

Suppose $S \subset \partial M$ is an essential subsurface of the boundary of a compact, irreducible 3-manifold M, and that ∂S is incompressible in M. An essential simple closed curve γ on M is called a *meridian* if it bounds an embedded disc in M. By the loop theorem, γ is a meridian if and only if it is homotopically trivial in M.

Definition 2.2 The *disc set* of S in M, written $\mathcal{D}(S, M)$, is the (full) subgraph of $\mathcal{C}(S)$ whose vertices are the meridians of S in M.

When convenient, we will sometimes regard $\mathcal{D}(S, M)$ as a subset of the space of projective measured laminations $\mathcal{PML}(S)$, instead of as a graph.

The following is an extension of a theorem of Masur and Minsky [48, Theorem 1.1], which they prove in the case that S is an entire component of ∂M.

Theorem 2.3 (Masur–Minsky) *The subset $\mathcal{D}(S, M)$ of $\mathcal{C}(S)$ is quasiconvex.*

To prove Theorem 2.3 as stated above, one follows the outline of [48]: given $a, b \in \mathcal{D}(S, M)$, the goal is to construct a *well-nested curve replacement sequence* from $a = a_1, \dots, a_n = b$ consisting of meridians, which must be a quasigeodesic by their Theorem 1.2. The sequence (a_i) is created by successive surgeries along innermost discs, and the only difference here is that one needs to ensure that none of the surgeries create peripheral curves. However, the surgeries create meridians and S has incompressible boundary.

2.4 Compression bodies

We refer the reader to Section 2 of [3] for a more detailed discussion of compression bodies, and state here only a few definitions that will be used later on.

A *compression body* is a compact, orientable, irreducible 3-manifold C with a π_1-surjective boundary component $\partial_+ C$, called the *exterior boundary* of C. The complement $\partial C \setminus \partial_+ C$ is called the *interior boundary*, and is written $\partial_- C$. Note that the interior boundary is incompressible. For if an essential simple closed curve on $\partial_- C$ bounds a disk $D \subset C$, then $C \setminus D$ has either one or two components, and in both cases, Van Kampen's theorem implies that $\partial_+ C$, which is disjoint from D, cannot π_1-surject.

Suppose M is a compact irreducible 3-manifold with boundary, let Σ be a component of ∂M. A *subcompression body of* (M, Σ) is a compression body that is embedded as a submanifold $C \subset M$ with exterior boundary Σ. Up to isotopy, any such subcompression body can be constructed as follows. Choose a set Γ of disjoint, pairwise nonhomotopic simple closed curves on Σ that are all meridians in M. Let $C' \subset M$ be the union of Σ with a set of disjoint disks in M whose boundaries are the components of Γ, and define $C \subset M$ to be the union of a regular neighbourhood of $C' \subset M$ together with any components of the complement of this neighbourhood that are topological 3-balls. Here, we say that $C \subset M$ is obtained *by compressing* Γ. Note that the irreducibility of M implies that no component of ∂C is a 2-sphere, and hence that C is irreducible, and therefore a compression body. See [3, Section 2] for details about this construction. More generally, if $S \subset \Sigma$ is an essential subsurface, a *subcompression body of* (M, S) is a subcompression body of (M, Σ) obtained by compressing a multicurve $\Gamma \subset S$.

Two subcompression bodies of a handlebody H are illustrated in Figure 1. On the left, we compress a separating meridian $\Gamma = \{m_1\}$ and obtain a subcompression body of H that has two "interior" boundary components contained in $\operatorname{int}(H)$; these are the tori drawn in grey. On the right, we compress a nonseparating meridian $\Gamma = \{m_2\}$ and obtain a subcompression body with a single torus interior boundary component. Note that compressing $\Gamma = \{m_1, m_2\}$ gives the same subcompression body as compressing m_2, because we fill in complementary components that are balls.

When the compressing set Γ is empty, we obtain the *trivial subcompression body* of (M, S), which is just a regular neighbourhood of $\Sigma \subset \partial M$. (In other words, the trivial compression body $\Sigma \times [0, 1]$ can be considered as a subcompression body of (M, S) for any subsurface $S \subset \Sigma$.) At the other extreme, we

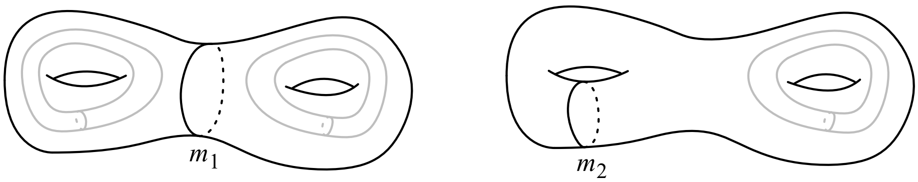

Figure 1: Compression bodies inside a genus-2 handlebody.

can compress a maximal $\Gamma \subset S$, which gives the "characteristic compression body" of (M, S), defined via the following fact.

Fact 2.4 (the characteristic compression body) *Suppose M is an irreducible compact 3-manifold, that Σ is a component of ∂M and that $S \subset \Sigma$ is an essential subsurface such that the multicurve ∂S is incompressible in M. Then there is a unique (up to isotopy) subcompression body*

$$C := C(M, S) \subset M$$

of (M, S), called the ***characteristic compression body*** *of (M, S), such that a curve γ in S is a meridian in C if and only if it is a meridian in M.*

Moreover, C can be constructed by compressing any maximal set of disjoint, pairwise nonhomotopic meridians in S.

This is a version of a construction of Bonahon [5], except that he only defines the characteristic compression body when S is an entire boundary component of M. In that case, the interior boundary components of M are incompressible in M, so Bonahon's construction can be used to reduce problems about 3-manifolds to problems about compression bodies and about 3-manifolds with incompressible boundary.

The reader can also compare Fact 2.4 to Lemma 2.1 in [3], which is the special case of the fact where M is a compression body and S is its exterior boundary, so that $C = M$ is obtained by compressing any maximal set of disjoint, nonhomotopic meridians in M.

Proof Let Γ be a maximal set of disjoint, pairwise nonhomotopic M-meridians on S, and define C by compressing Γ. We have to check that any curve in S that is an M-meridian is also a C-meridian. Suppose not, and take an M-meridian $m \subset S$ that is not a C-meridian, and that intersects Γ minimally. Since Γ is maximal, m intersects some component $\gamma \subset \Gamma$. Then there is an arc $\alpha \subset \gamma$ with endpoints on m and interior disjoint from m, that is homotopic rel endpoints in M to the arcs $\beta', \beta'' \subset m$ with the same endpoints. (Here α is an "outermost" arc of intersection on a disk bounded by γ, where the intersection is with the disk bounded by m; see, e.g., Lemma 2.8 in [3].) Since m is in minimal position with respect to Γ, the curves $m' = \alpha \cup \beta'$ and $m'' = \alpha \cup \beta''$ are both essential, and are M-meridians in S that intersects Γ fewer times than m. So by minimality of m, both m', m'' are C-meridians, implying that α is homotopic rel endpoints to β' and β'' in C. This implies m is a C-meridian, contrary to assumption.

For uniqueness, suppose we have two subcompression bodies C_1, C_2 of (M, S) in which all curves in S that are meridians in M are also meridians in C_1, C_2. Since C_1, C_2 are subcompression bodies of (M, S), the kernels of the maps

$$\pi_1 \Sigma \to \pi_1 C_i$$

induced by inclusion are both normally generated by the set of all elements of $\pi_1 \Sigma$ that represent simple closed curves in S that are meridians in M. Hence, the disk sets $\mathcal{D}(\Sigma, C_i)$ are the same for $i = 1, 2$. It follows that C_1, C_2 are isotopic in M, say by Corollary 2.2 of [3]. □

2.5 Interval bundles

In this paper, an *interval bundle* always means a fibre bundle $B \to Y$, where Y is a compact surface with boundary, and where all fibres are closed intervals I. Regarding the fibres as "vertical", we call the associated ∂I-bundle over Y the *horizontal boundary* of B, written $\partial_H B$. An interval bundle that is isomorphic to $Y \times [-1, 1]$ is called *trivial*, and we often call nontrivial interval bundles *twisted*.

All 3-manifolds in this paper are assumed to be orientable, but even when the total space B of an interval bundle is orientable, the base surface Y may not be. Indeed, let Y be a compact nonorientable surface and let $\pi : \widehat{Y} \to Y$ be its orientation cover. Then the mapping cylinder

$$B := \widehat{Y} \times [0, 1]/\sim, \quad (x, 1) \sim (x', 1) \iff \pi(x) = \pi(x'),$$

is orientable, and is a twisted interval bundle over Y, where the fibre over $y \in Y$ is obtained by gluing together the two intervals $\{x\} \times [0, 1]$ and $\{x'\} \times [0, 1]$ along $(x, 1)$ and $(x', 1)$, where $\pi^{-1}(y) = \{x, x'\}$. The horizontal boundary $\partial_H B$ here is $\widehat{Y} \times \{0\}$, which is homeomorphic to the orientable surface $\widehat{Y}$. Note that B is double covered by the trivial interval bundle $\widehat{Y} \times [-1, 1]$.

Fact 2.5 *Suppose that $B \to Y$ is an interval bundle and B is orientable. If Y is orientable, then B is a trivial interval bundle. If Y is nonorientable, then B is isomorphic to the mapping cylinder of the orientation cover of Y.*

Proof If Y and B are orientable, so is the line bundle, so the bundle is trivial. If Y is nonorientable, the horizontal boundary $\partial_H B \subset \partial B$ is an orientable surface that double covers Y, and from there it's easy to construct the desired isomorphism to the mapping cylinder of the projection $\partial_H B \to Y$. □

An interval bundle $B \to Y$ comes with a *canonical involution* σ, which is well defined up to isotopy, and which is defined as follows. If $B \cong Y \times [0, 1]$ is a trivial interval bundle, we define

$$\sigma : Y \times [-1, 1] \to Y \times [-1, 1], \quad \sigma(y, t) = (y, -t).$$

And if B is the twisted interval bundle $B \cong \widehat{Y} \times [0, 1]/\sim$ above, we define

$$\sigma : \widehat{Y} \times [0, 1]/\sim \to \widehat{Y} \times [0, 1]/\sim, \quad \sigma(\hat{y}, t) = (\iota(\hat{y}), t),$$

where ι is the nontrivial deck transformation of the orientation cover. Note that σ is always an orientation reversing involution of B, so in particular, when we give the surface $\partial_H B$ its boundary orientation, the restriction $\sigma|_{\partial_H B}$ is also orientation reversing.

We also recall the following well-known fact.

Fact 2.6 *Suppose $B \to Y$ is an interval bundle (as always, over a compact surface with boundary) and B is orientable. Then B is homeomorphic to a handlebody.*

It's a nice topology exercise to visualise the homeomorphism. Regard Y as the union of a polygon and a collection of bands (long, skinny rectangles), each of which is glued along its short sides to two sides of the polygon. Thickening, the picture becomes a ball with 1-handles attached, so since B is orientable, it is a handlebody.

Note that if $S = S_{g,b}$ has genus g and b boundary components, then the handlebody $S \times [-1,1]$ has genus $2g + b - 1$, since that is the rank of the free group $\pi_1(S \times [-1,1]) \cong \pi_1 S$.

Finally, suppose $\pi : B \to Y$ is an interval bundle and $f : \partial_H B \to \partial_H B$ is a homeomorphism. We say that f *extends to* B if there is a homeomorphism $F : B \to B$ such that $F|_{\partial_H B} = f$. We leave the following to the reader.

Fact 2.7 *The following are equivalent:*

(1) *f extends to B.*

(2) *$f \circ \sigma$ is isotopic to f on $\partial_H B$.*

(3) *After isotoping f, there is a homeomorphism $\bar{f} : Y \to Y$ such that $\pi \circ f = \bar{f} \circ \pi$.*

(4) *There is a homeomorphism from B to either*

$$Y \times [-1,1] \quad \text{or} \quad \widehat{Y} \times [0,1]/\sim,$$

taking horizontal boundary to horizontal boundary, such that $f = F|_{\partial_H B}$, and where either

$$F : Y \times [-1,1] \to Y \times [-1,1], \quad F(y,t) = (\bar{f}(y), t),$$

for some homeomorphism $\bar{f} : Y \to Y$, or

$$F : \widehat{Y} \times [0,1]/\sim \to \widehat{Y} \times [0,1]/\sim, \quad F(y,t) = (\bar{f}(y), t),$$

for some homeomorphism $\bar{f} : \widehat{Y} \to \widehat{Y}$ commuting with the deck group of $\widehat{Y} \to Y$, and hence covering a homeomorphism of Y.

2.6 The characteristic submanifold of a pair

Suppose that M is a compact, orientable 3-manifold and that $S \subset \partial M$ is an incompressible subsurface. In the late 1970s, Jaco and Shalen [30] and Johannson [32] described a "characteristic" submanifold of (M, S) that contains the images of all nondegenerate maps from interval bundles and Seifert fibred spaces.

Theorem 2.8 (Jaco–Shalen [29, page 138]) *There is a perfectly embedded Seifert pair $(X, \Sigma) \subset (M, S)$, unique up to isotopy and called the* **characteristic submanifold** *of (M, S), such that any nondegenerate map $(B, F) \to (M, S)$ from a Seifert pair (B, F) is homotopic as a map of pairs into (X, Σ).*

A *Seifert pair* is 3-manifold pair that is a finite disjoint union of interval bundle pairs $(B, \partial_H B)$ and S^1-bundle pairs. Here, an S^1*-bundle pair* (B, F) is a 3-manifold B fibred by circles, where $F \subset \partial B$ is a compact subsurface saturated by fibres. A Seifert pair $(X, \Sigma) \subset (M, S)$ is *well embedded* if $X \cap \partial M = \Sigma \subset S$ and the frontier of X in M is a π_1-injective surface, and is *perfectly embedded* if it is well embedded, no component of the frontier of X in M is homotopic into S, and no component of X is homotopic into another component.

When (B, F) is a connected Seifert pair, a map $f : (B, F) \to (M, S)$ is *essential* if it is not homotopic as a map of pairs into S. Notice that this only depends on the image of f and not on f itself. One says f is *nondegenerate* if it is essential, its π_1-image is nontrivial, its π_1-image is noncyclic when $F = \varnothing$, and no fibre of B is nullhomotopic in (M, S). For disconnected (B, F), one says f is nondegenerate if its restriction to every component is nondegenerate.

The following is very well known.

Fact 2.9 *If $\mathrm{int}(M)$ is hyperbolisable and (B, F) is an S^1-bundle pair that is perfectly embedded in (M, S), then either*

(1) *(B, F) is a "fibred solid torus", i.e., B is an S^1-bundle over a disk, and $F \subset \partial B \cong T^2$ is a collection of fibred parallel annuli, or*

(2) *(B, F) is a "thickened torus", i.e., B is an S^1-bundle over an annulus, so is homeomorphic to $T^2 \times [0, 1]$, and each component of F is either a torus or a fibred annulus.*

So in particular, the components of the characteristic submanifold of (M, S) are either interval bundles, solid tori, or thickened tori.

Proof Suppose that (B, F) is a perfectly embedded S^1-bundle pair in M. Then $B \to Y$ is an S^1-bundle, where Y is a compact 2-orbifold, and the cyclic subgroup $Z \subset \pi_1 B$ corresponding to a regular fibre is normal in $\pi_1 B$. In a hyperbolic 3-manifold, any subgroup of π_1 that has a cyclic normal subgroup is elementary, say by a fixed point analysis on $\partial_\infty \mathbb{H}^3$. So, $\pi_1 B$ is either cyclic or isomorphic to $\mathbb{Z}^2$. It follows that Y is a disc, in which case B is a fibred solid torus, or Y is an annulus, in which case B is a thickened torus. □

In this paper we will mostly be interested in interval bundles. For brevity, we'll use the following terminology, which differs slightly from the terminology above used by Jaco and Shalen.

Definition 2.10 An *essential interval bundle* in (M, S) is an essential, well-embedded interval bundle pair $(B, \partial_H B) \hookrightarrow (M, S)$.

Note that the horizontal boundary of any essential interval bundle is an incompressible subsurface of S.

The definition above differs from a well-embedded interval bundle pair in that we are excluding boundary-parallel interval bundles over annuli, and differs from a perfectly embedded interval bundle pair in that we are allowing components of the frontier of an interval bundle over a surface that is not an annulus to be boundary parallel. For instance, if Y is a surface with boundary and $Y' \subset Y$ is obtained by deleting collar neighbourhoods of the boundary components, and we set $M = Y \times [-1,1]$, which is a handlebody, then $(Y' \times [-1,1], Y' \times \{-1,1\})$ is an essential interval bundle in $(M, \partial M)$, but is not perfectly embedded. However, note that any essential interval bundle $(B, \partial_H B) \hookrightarrow (M, S)$ is perfectly embedded in $(M, \partial_H B)$.

2.7 Compression arcs

Suppose $(B, \partial_H B) \subset (M, S)$ is an essential interval bundle. An arc $\alpha \subset S$ with endpoints on $\partial(\partial_H B)$ and interior disjoint from $\partial_H B$ is called a *compression arc* if it is homotopic in M to a fibre of B, while keeping its endpoints on $\partial(\partial_H B)$. See Figure 2. To link this definition with more classical ones, it is easy to see that there is a compression arc for B if and only if $\overline{\mathrm{Fr}(B)}$ is boundary compressible; see [29, pages 36–37] for a definition.

Write our interval bundle as $\pi : B \to Y$. Let α be a compression arc for B. After isotoping the bundle map π, we can assume that α is homotopic rel endpoints to a fibre $\pi^{-1}(y)$, where $y \in Y$. Suppose c is an oriented, two-sided, essential, simple closed loop Y based at y, and suppose that either c is nonperipheral in Y, or that Y is an annulus or Möbius band. Write $\pi^{-1}(c) = c_- \cup c_+$, where $c_\pm$ are disjoint simple closed oriented loops in X based at $y_\pm$, and where the orientations of $c_\pm$ project to that of c.

Claim 2.11 *The concatenation $m(c) := c_- \cdot \alpha \cdot c_+^{-1} \cdot \alpha^{-1}$ is homotopic to a meridian on S.*

So, a compression arc α allows one to make compressible curves on S from essential curves on Y. See Figure 2.

Proof Since α is homotopic rel endpoints to the fibre $\pi^{-1}(y)$, the curve $m(c)$ is homotopic in M to a curve in B that projects under π to $c \cdot c^{-1}$, and hence $m(c)$ is nullhomotopic in M. Checking orientations, one can see that $m(c)$ is homotopic to a simple closed curve on S. So, we only have to prove that $m(c)$ is homotopically essential on S.

Suppose that c_-, c_+ are freely homotopic on $\partial_H B$ as oriented curves. (This happens exactly when the curve $c \subset Y$ bounds a Möbius band in Y.) Then $m(c)$ is homotopic to the commutator of two essential simple closed curves on S that intersect once, and hence is essential since S is not a torus.

We can now assume that $c_\pm$ are not freely homotopic in $\partial_H B$ as oriented curves. If $m(c)$ is inessential, then $c_\pm$ *are* freely homotopic on S, so $c_\pm$ are homotopic in $\partial_H B$ to boundary components $c'_\pm \subset \partial_H B$ that bound an annulus in $S \setminus \partial_H B$. In this case $c_\pm$ are peripheral, so we may assume that Y is either an annulus or a Möbius band. If Y is a Möbius band, we are in the situation of the previous paragraph and are done. So, Y is an annulus, and $\partial_H B$ is a pair of disjoint annuli on S, where $c'_\pm$ lie in different components of $\partial_H B$. Since $c'_\pm$ bound an annulus in $S \setminus \partial_H B$, the interval bundle B is inessential, contrary to our assumption. □

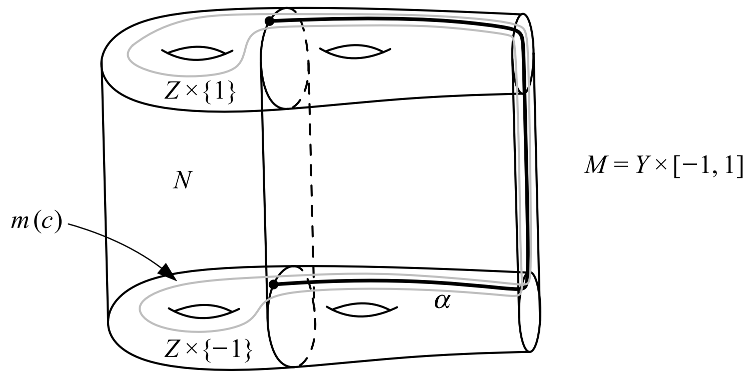

Figure 2: $Z \subset Y$ is a compact surface, the interval bundle $B = Z \times [-1,1]$ embeds in the manifold $M = Y \times [-1,1]$, and α above is a compression arc. Also pictured in light grey is a meridian as described in Claim 2.11.

In fact, more is true.

Fact 2.12 (arcs that produce meridians) *Suppose $(B, \partial_H B) \subset (M, S)$ is an essential interval bundle and let $\alpha \subset S$ be an arc with endpoints on $\partial_H B$ and interior disjoint from $\partial_H B$. Let $X \subset S$ be a regular neighbourhood of $\alpha \cup \partial_H B$ within S. Then there is a meridian in X if and only if we have either*

(1) *the endpoints of α lie on the same component c of $\partial(\partial_H B)$, and there is an arc $\beta \subset c$ such that $\alpha \cup \beta$ is a meridian, or*

(2) *α is a compression arc.*

Note that in the second case the endpoints of α lie on distinct components of $\partial(\partial_H B)$, so in particular the two cases are mutually exclusive.

The reason we say X "contains a meridian" instead of "is compressible" is that X may not be an essential subsurface of S, and we want to emphasise that the essential curve in X that is compressible in M is actually essential in S. For example, let Y be a compact surface with boundary, $Y' \subset Y$ be obtained by deleting a collar neighbourhood of ∂Y, set $B = Y' \times [-1,1]$ and $M = Y \times [-1,1]$, and let α be a spanning arc of B in ∂M.

Proof The "if" direction is immediate: in case (1) we are essentially given a meridian in X, and in case (2) we can appeal to Claim 2.11.

We now work on the "only if" direction. Write our regular neighbourhood of $\partial_H B \cup \alpha$ as $X = \partial_H B \cup R$ where R is a rectangle with two opposite "short" sides on the boundary of $\partial_H B$. Let $D \subset M$ be an essential disc whose boundary is contained in X, and where D intersects the frontier $\mathrm{Fr}(B) \subset M$ in a minimal number of components. Let $a \subset D \cap \mathrm{Fr}(B)$ be an arc that is "outermost" in D, i.e., there is some arc $a' \subset \partial D$ with the same endpoints as a such that a, a' bound an open disk in D that does not intersect $\mathrm{Fr}(B)$.

We claim that $a' \subset R$. If not, then $a' \subset \partial_H B$, and bounds a disk in B with the arc $a \subset \partial B$. Writing the interval bundle as $\pi : B \to Y$, the projection $\pi(a \cup a')$ in Y is then also nullhomotopic, so $\pi(a')$ is homotopic rel endpoints into ∂Y. Lifting this homotopy through the covering map $\partial_H B \to Y$ we get that

a' is inessential in $\partial_H B$, i.e., is homotopic in $\partial_H B$ rel endpoints into $\partial(\partial_H B)$. Lifting this homotopy through the covering map $\partial_H Y \to Y$ $\pi(a) \subset \partial Y$, it follows that $\pi(a')$ is an inessential arc in Y. We can then decrease the number of components of $D \cap \mathrm{Fr}(B)$, contradicting that this number is minimal.

So, $a' \subset R$. Again by minimality of the intersection, the endpoints of a' lie on opposite short sides of R, so α is homotopic to a' through arcs in R with endpoints on $\mathrm{Fr}(B)$. Since a' is homotopic rel endpoints to $a \subset \mathrm{Fr}(B)$, it follows that α is homotopic rel endpoints into $\mathrm{Fr}(B)$. If the two endpoints of α lie on the same component of $\partial(\partial_H B)$, we are in case (1), and otherwise we are in case (2). □

2.8 Laminations

We assume the reader is familiar with geodesic and measured laminations on finite-type hyperbolic surfaces. See, e.g., [16; 33].

Suppose λ is a connected geodesic lamination on a finite-type hyperbolic surface S with geodesic boundary. We say that λ *fills* an essential subsurface $T \subset S$ if $\lambda \subset T$ and λ intersects every essential, nonperipheral simple closed curve in T.

Fact 2.13 *For every connected λ, there is a unique subsurface with geodesic boundary (as in Section 2.1) that is filled by λ, which we denote by $S(\lambda)$. It is the minimal subsurface with geodesic boundary in S that contains λ.*

Here, $S(\lambda)$ can be constructed by taking a component $\tilde{\lambda} \subset \widetilde{S} \subset \mathbb{H}^2$ of the preimage of λ, letting $C \subset \mathbb{H}^2$ be the convex hull of the set of endpoints of leaves of $\tilde{\lambda}$ in $\partial\mathbb{H}^2$, and projecting C into S.

Suppose that M is a compact, orientable irreducible 3-manifold let $S \subset \partial M$ be an essential subsurface. The *limit set* of (S, M) is the closure

$$\Lambda(S, M) = \overline{\{\text{meridians } \gamma \subset S\}} \subset \mathcal{PML}(S),$$

where $\mathcal{PML}(S)$ is the space of projective measured laminations on S. The limits set was first studied by Masur [46] in the case that M is a handlebody, with S its entire boundary. In this case, Kerckhoff [34] later proved that the limit set has measure zero in $\mathcal{PML}(S)$, although a mistake in his argument was later found and fixed by Gadre [23].

In some ways, $\Lambda(S, M)$ acts as a dynamical limit set. For instance, let $\mathrm{Map}(S)$ be the mapping class group of S, and let $\mathrm{Map}(S, M) \subset \mathrm{Map}(S)$ be the subgroup consisting of mapping classes represented by restrictions of homeomorphisms of M. Then we have:

Fact 2.14 (1) *If $\Lambda(S, M)$ is nonempty, it is the smallest nonempty closed subset of $\mathcal{PML}(S)$ that is invariant under* $\mathrm{Map}(S, H)$.

(2) *If* $\mathrm{Map}(S, M)$ *contains a pseudo-Anosov map on S, then $\Lambda(S, M)$ is the closure of the set of the attracting and repelling laminations of pseudo-Anosov elements of* $\mathrm{Map}(S, M)$.

Note that $\mathrm{Map}(S, M)$ contains a pseudo-Anosov map on S if and only if the disk set $\mathcal{D}(S, M)$ has infinite diameter in the curve complex $\mathcal{C}(S)$, where the latter condition was discussed earlier in Proposition 3.1. See also [2; 41].

Proof For the first part just note that Dehn twist T_m around meridians $m \subset S$ are in $\mathrm{Map}(S, M)$, so if $A \subset \mathcal{PML}(S)$ is nonempty and invariant, $\lambda \in A$ and m is a meridian, then $m = \lim_i T_m^i(\lambda)$ is also in A, implying $\Lambda(S, M) \subset A$.

For the second part, take a pseudo-Anosov $f \in \mathrm{Map}(S, M)$ with attracting lamination λ_+, say. If m is a meridian in S, then $T_m^i \circ f \circ T_m^{-i}$ are pseudo-Anosov maps on S and their attracting laminations converge to m, and then the argument finishes as before. □

2.9 Laminations on interval bundles

Suppose that Y is a compact hyperbolisable surface with boundary, and that $B \to Y$ is an interval bundle over Y. Endow Y and the horizontal boundary $\partial_H B$ with arbitrary hyperbolic metrics such that the boundary components are all geodesic.

Suppose we have two geodesic laminations $\lambda_\pm$ on $\partial_H B$.

Definition 2.15 We say that $\lambda_\pm$ are *essentially homotopic through* B if there is a lamination λ and a homotopy $h_t : \lambda \to B$, $t \in [-1, 1]$, such that $h_{\pm 1}$ is a homeomorphism onto $\lambda_\pm$, and where (h_t) is not homotopic into $\partial_H B$.

When B is a trivial interval bundle, $\lambda_\pm$ are essentially homotopic through B if and only if we can write $B \cong Y \times [0, 1]$ in such a way that $\lambda_\pm = \lambda \times \{\pm 1\}$ for some geodesic lamination on Y. This is an easy consequence of the fact that on a surface, homotopic laminations are isotopic. In general:

Fact 2.16 *Suppose that $\lambda_\pm$ are disjoint or equal geodesic laminations on $\partial_H B$. Then these are equivalent*:

(1) *$\lambda_\pm$ are essentially homotopic through B.*

(2) *$\lambda_\pm$ is isotopic on $\partial_H B$ to $\sigma(\lambda_\mp)$, where σ is the canonical involution of B discussed in Section 2.5.*

Moreover, (1) *and* (2) *imply*

(3) *there is a geodesic lamination $\bar{\lambda}$ on Y such that $\lambda_- \cup \lambda_+$ is isotopic on $\partial_H B$ to $(\pi|_{\partial_H B})^{-1}(\bar{\lambda})$.*

Here, (3) does not always imply (1) and (2), since it could be that $\bar{\lambda}$ has two components, $(\pi|_{\partial_H B})^{-1}(\bar{\lambda})$ has four, and these components are incorrectly partitioned into the two laminations $\lambda_\pm$. However, that's the only problem, so, for instance, if $\lambda_\pm$ are minimal then (1)–(3) are equivalent.

While we have phrased things more generally in the section, we can always assume in proofs that our hyperbolic metrics have been chosen so that the covering map $\pi|_{\partial_H B} : \partial_H B \to Y$ is locally isometric. Here, we're using the fact that given two hyperbolic metrics with geodesic boundary on a compact surface, a geodesic lamination with respect to one metric is isotopic to a unique geodesic lamination with respect to the other hyperbolic metric. In this case, we can remove the word "isotopic" from (2) and (3).

Proof The fact is trivial when B is a trivial interval bundle. When B is nontrivial, lift the homotopy to the trivial interval bundle $B' \to B$ that double covers B, giving homotopic laminations $\lambda'_\pm \subset \partial_H B'$. The statement (1) $\iff$ (2) follows since the canonical involution on B' covers that of B. For (2) $\implies$ (3), note that since $\lambda_\pm$ are disjoint or equal and differ by σ, their projections $\pi(\lambda_\pm) \subset Y$ are the same, and are a geodesic lamination $\bar{\lambda}$ on Y. □

3 Large and small disk sets and compression bodies

Suppose that $S \subset \partial M$ is an essential subsurface of the boundary of a compact, irreducible 3-manifold M, and that ∂S is incompressible in M. The following is probably known to some experts, but we don't think it appears anywhere in the literature, so we give a complete proof.

Proposition 3.1 (diameters of disk sets) *With M, S as above, either*

(1) *$\mathcal{D}(S, M)$ has infinite diameter in $\mathcal{C}(S)$,*

(2) *S has one nonseparating meridian δ, and every other meridian is a band sum of δ,*

(3) *S has a single meridian, which is separating, or*

(4) *$\mathcal{D}(S, M) = \varnothing$.*

In case (1) we will say that $\mathcal{D}(S, M)$ is *large*, and in cases (2)–(4), we will say that $\mathcal{D}(S, M)$ is *small*. Similarly, if $C(S, M)$ is the characteristic compression body defined in Fact 2.4, then $C(S, M)$ is said to be large or small depending on whether $\mathcal{D}(S, M)$ is large or small. See also the discussion of small compression bodies in Section 3 of [3].

Here, recall that a *band sum* of a meridian δ is the boundary of a regular neighbourhood of $\delta \cup \beta$, where β is a simple closed curve on S that intersects δ once. Any such band sum must be a meridian: for instance, as an element of $\pi_1 M$ it is a commutator with a trivial element. Also, (3) includes the case when M is a solid torus and $S = \partial M$, in which case there is only one (nonseparating) meridian. When M is not a solid torus, though, every nonseparating curve has infinitely many band sums. Similarly, a *band sum* of two disjoint meridians δ and γ over an arc k joining δ to γ is the boundary of a regular neighbourhood of $\delta \cup k \cup \gamma$.

Before beginning the proof, we first establish the following:

Claim 3.2 *Suppose S is not a torus, $\gamma \subset S$ is a meridian on S and δ is a meridian that lies in a component $T \subset S \setminus \gamma$. If γ is not a band sum of δ, there is a pseudo-Anosov $f : T \to T$ that extends to a homeomorphism of M.*

The condition that γ is not a band sum is necessary. For if M is a handlebody with $S = \partial M$, and γ is a separating meridian that bounds a compressible punctured torus $T \subset S$, then T has only a single meridian δ. This δ is nonseparating and γ is a band sum of δ. Any map $T \to T$ that extends to a homeomorphism of M must then fix δ, so cannot be pseudo-Anosov.

Similarly, if S is a torus and γ is a meridian, the complement of γ is an annulus, which does not admit any pseudo-Anosov map.

Proof of Claim 3.2 Suppose first that γ is not separating. Any simple closed curve that intersects γ once can be used to create a band sum. Now S is not a torus, and cannot be a punctured torus either, since then its boundary would be compressible. So, there are a pair α, β of band sums of γ that fill $S \setminus \gamma$. By a theorem of Thurston [22, III.3 in 13], the composition of twists $T_\alpha \circ T_\beta^{-1}$ is pseudo-Anosov. Each twist extends to M, because twist about meridians can be extended to twists along the disks they bound.

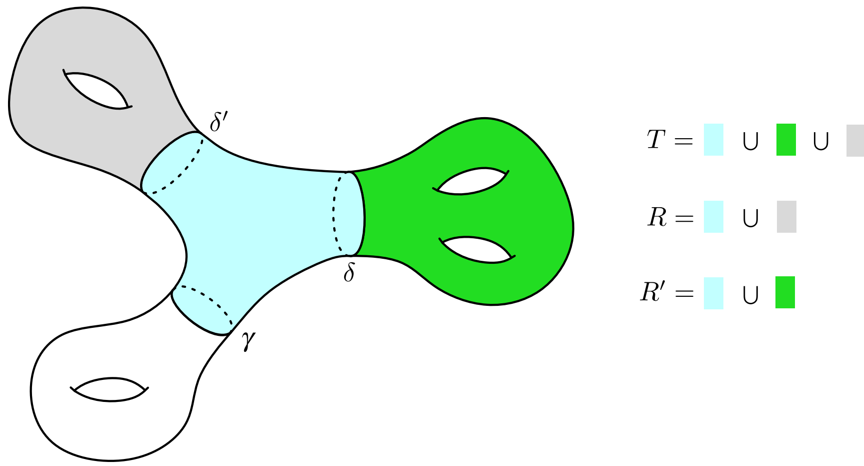

Figure 3: The surfaces R and R' fill T.

Now suppose γ separates S, and suppose that R is the component of $T \setminus (\gamma \cup \delta)$ adjacent to γ and δ. Any curve in R that bounds a pair of pants with γ and δ is also a meridian. Such curves are constructed as the boundary of a neighbourhood of the union of γ, δ and any arc in R joining the two. Therefore, there is a pair α, β of such curves that fills R. As before, $f = T_\alpha \circ T_\beta^{-1}$ is a pseudo-Anosov on R that extends to M.

However, there was nothing special about δ in the above construction. So if there is some (nonperipheral) meridian $\delta' \subset T$ with $\delta \neq \delta'$, there is also a pseudo-Anosov f' on the corresponding surface R', such that f' extends to M. Since R and R' fill T, [19, Theorem 6.1] says that for large i the composition $f^i (f')^i$ is a pseudo-Anosov on T. See Figure 3.

The only case left to consider is when δ is the only (nonperipheral) meridian in T. Since new meridians can be created by taking a band sum of δ and γ over any arc joining them, the only possibility here is that T is a punctured torus, in which case this construction always just produces γ again. But then γ is a band sum of δ. □

Proof of Proposition 3.1 When S is a torus, distinct curves have nonzero algebraic intersection number, so either there are no meridians or there is a single meridian. So, we assume $S \neq T^2$ below.

It follows easily from Claim 3.2 that if S contains two disjoint meridians, neither of which is a band sum of the other, then $\mathcal{D}(S, M)$ has infinite diameter in the curve complex. We first claim that if there are two disjoint meridians in S, neither of which is a band sum of the other, then $\mathcal{D}(S, M)$ has infinite diameter in the curve complex. To see this, suppose γ_1, γ_2 are such meridians. Claim 3.2 gives two pseudo-Anosov maps f_1, f_2, each defined on the component of $S \setminus \gamma_i$ that contains γ_j, where $i \neq j$. Since the component of $S \setminus \gamma_1$ containing γ_2 and the component of $S \setminus \gamma_2$ containing γ_1 together fill S, for large k the composition $f_1^k f_2^k$ is a pseudo-Anosov map on the entire surface S, by [19, Theorem 6.1]. Any such composition extends to M, so maps meridians to meridians. As pseudo-Anosovs act with unbounded orbits on the curve complex [47], this implies that the set of meridians has infinite diameter in $\mathcal{C}(S)$.

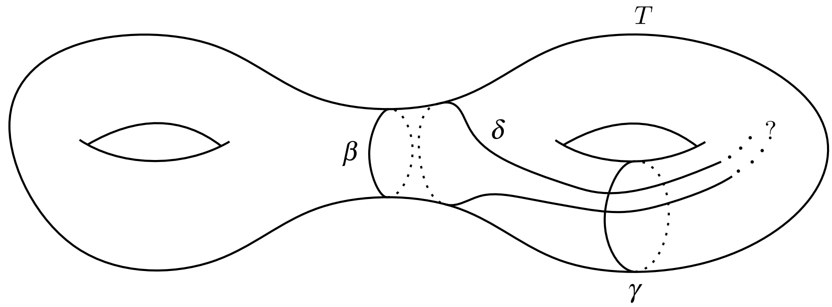

Figure 4: A surgery of a curve δ along a nonseparating γ cannot produce a curve β that is a band sum of γ.

Starting now with the proof of the proposition, suppose there are *no nonseparating meridians* in S. If γ, δ are distinct (separating) meridians, then an innermost disk surgery produces another separating meridian γ_2 disjoint from γ_1; see [3, Lemma 2.8]. By the previous paragraph, $\mathcal{D}(S, M)$ has infinite diameter in the curve complex. So, the only other options are if $\mathcal{D}(S, M) = \varnothing$, or if the only meridian is a single separating curve.

Suppose now that there is a nonseparating meridian γ in S. By Claim 3.2, unless the disc set has infinite diameter in the curve complex, any meridian disjoint from γ must be a band sum of γ. So, either we are in case (2) of the proposition, or there is some meridian δ that intersects γ. Any innermost disk surgery of δ along γ must produce a band sum β of γ. However, this β must then bound a punctured torus T containing γ, and δ is then forced to lie inside T, which gives a contradiction; see Figure 4. □

4 Windows from limit sets

Let $N = \Gamma \backslash \mathbb{H}^3$ be an orientable, geometrically finite hyperbolic 3-manifold, let $\Lambda \subset \partial \mathbb{H}^3$ be the limit set of Γ, and let

$$\mathrm{CC}(N) := \Gamma \backslash \mathrm{CH}(\Lambda) \subset N$$

be the convex core of N. Equip $\partial \mathrm{CC}(N)$ with its intrinsic length metric, which is hyperbolic; see, for instance, [58, Proposition 8.5.1].

Let $S_\pm$ be (possibly degenerate) incompressible subsurfaces with geodesic boundary in $\partial \mathrm{CC}(N)$ that are either equal or are essentially disjoint, as in Section 2.1. Let

$$\widetilde{S}_\pm \subset \partial \mathrm{CH}(\Lambda) \subset \mathbb{H}^3$$

be lifts of $S_\pm$, where if $S_- = S_+$, we require that $\widetilde{S}_- \neq \widetilde{S}_+$. Let $\Gamma_\pm \subset \Gamma$ be the stabilisers of $\widetilde{S}_\pm$, let $\Lambda_\pm \subset \partial \mathbb{H}^3$ be their limit sets and $\Delta = \Gamma_+ \cap \Gamma_-$.

The lift $\widetilde{S}_\pm$ is isometric to a convex subset of $\mathbb{H}^2$. Let $\partial_\infty S_\pm \subset \partial \mathbb{H}^2$ be the boundary of $\widetilde{S}_\pm$. By [53, Theorem 5.6], say, the inclusion $\widetilde{S}_\pm \hookrightarrow \mathbb{H}^3$ extends continuously to a $\Gamma_\pm$-equivariant quotient map

$$\iota_\pm : \partial_\infty \widetilde{S}_\pm \to \Lambda_\pm \subset \partial \mathbb{H}^3.$$

Theorem 4.1 (windows from limit sets) *We have $\Lambda_- \cap \Lambda_+ = \Lambda_\Delta$. Next, suppose Δ is nonempty and is not a cyclic group acting parabolically on either $\widetilde{S}_-$ or $\widetilde{S}_+$, and let $\widetilde{C}_\pm \subset \widetilde{S}_\pm$ be the convex hulls of the subsets $\iota_\pm^{-1}(\Lambda_\Delta) \subset \partial_\infty \widetilde{S}_\pm$. Then $\widetilde{C}_\pm$ are Δ-invariant, the quotients $C_\pm := \Delta \backslash \widetilde{C}_\pm$ are (possibly degenerate) subsurfaces with geodesic boundary in $S_\pm$, and there is an essential homotopy from C_- to C_+ in $\mathrm{CC}(N)$ that is the projection of a homotopy from $\widetilde{C}_-$ to $\widetilde{C}_+$.*

Above, $C_\pm$ are (possibly degenerate) subsurfaces with geodesic boundary in $S_\pm$, as defined in Section 2.1, but it follows from the above and Theorem 2.8 that there are "resolutions" (see Section 2.1) $C'_\pm \subset S_\pm$ such that $C'_\pm$ bound an interval bundle in $\mathrm{CC}(M)$. So informally, the theorem says that the intersection $\Lambda_- \cap \Lambda_+$ is exactly the limit set of the fundamental group of some essential interval bundle in $(\mathrm{CC}(M), S'_- \cup S'_+)$. The term "window" comes from Thurston [59] and refers to interval bundles; for example, one can "see through" a trivial interval bundle from one horizontal boundary component to the other.

The assumption that Δ is not cyclic and acting parabolically on either $\widetilde{S}_\pm$ is just for convenience in the statement of the theorem. (Just to be clear, note that an element $\gamma \in \Delta$ can act parabolically as an isometry of $\mathbb{H}^3$, but hyperbolically on the convex subsets $\widetilde{S}_\pm \subset \mathbb{H}^2$.) If Δ is cyclic and acts parabolically on $\widetilde{S}_+$ the subset $\widetilde{C}_+$ in the statement of the theorem will be empty. However, using the same proof one can construct a homotopy from a simple closed curve on S_+ bounding a cusp of S_+ to some simple closed curve on S_-.

As mentioned in the introduction, a version of Theorem 4.1 was known to Thurston; see his discussion of the only windows break theorem in [59]. Precise statements for geometrically finite N without accidental parabolics were worked out in Lecuire's thesis [40] and by Walsh [60]; note that Walsh uses the conformal boundary instead of the convex core boundary, but the two points of view are equivalent. However, for our applications in this paper, we need to allow accidental parabolics in $\widetilde{S}_\pm$, which are not allowed in those theorems. Also, our proof is more direct and natural[5] than those in [40; 60], despite the extra complication coming from parabolics.

Finally, the assumption that N is geometrically finite is not really essential for the theorem statement. With a bit more work dealing with degenerate ends, one can prove the theorem for all finitely generated Γ. Essentially, the point is to use Canary's covering theorem [14] to show that degenerate NP-ends in the covers $N_\pm := \Gamma_\pm \backslash \mathbb{H}^3$ have neighbourhoods that embed in N, and then to use this to prove that geodesic rays in $\mathbb{H}^3$ that converge to points in $\Lambda_- \cap \Lambda_+$ cannot exit degenerate ends in $N_\pm$. After showing this, the proof of Claim 4.3 extends to the general case. However, we don't have an application for that theorem in mind, so we'll spare the reader the details.

Proof of Theorem 4.1

We first focus on proving that $\Lambda_- \cap \Lambda_+ = \Lambda_\Delta$. For each $\xi \in \partial\mathbb{H}^3$, let $\Gamma_\pm(\xi) \subset \Gamma_\pm$ be the stabiliser of ξ.

[5] In both [40; 60], the authors focus on proving that the boundary components of $\widetilde{C}_\pm$ project to simple closed curves in $S_\pm$, but that isn't sufficient to say that $\widetilde{C}_\pm$ projects to a subsurface with geodesic boundary in $S_\pm$, which is what they then claim. E.g., in [60] it is stated that under a covering map, the boundary of a subset goes to the boundary of the image, but this isn't true.

Claim 4.2 *Let $\xi \in \partial\mathbb{H}^3$ and suppose that $\Gamma_-(\xi)$ and $\Gamma_+(\xi)$ are both nontrivial. Then they are equal.*

Proof By the tameness theorem [1; 11], we can identify $\mathrm{CC}(N)$ topologically with a subset of a 3-compact manifold with boundary M, where

$$\text{(1)} \qquad \mathrm{CC}(N) \supset \mathrm{int}(M), \quad \mathrm{CC}(N) \cap \partial M = \partial\mathrm{CC}(N),$$

and where $\partial\mathrm{CC}(N)$ is a collection of essential subsurfaces of $\partial\bar{N}$. Let $\partial_{\chi=0}M$ be the union of all torus boundary components of M, and let (X,Σ) be the characteristic submanifold of the pair $(M, S_- \cup S_+ \cup \partial_{\chi=0}M)$, as in Section 2.6.

Since $\Gamma_\pm$ are both contained in a discrete group Γ, both $\Gamma_\pm(\xi)$ are contained in the stabiliser $\Gamma(\xi)$, which is either infinite cyclic, or rank-2 parabolic.

Suppose first that $\Gamma(\xi)$ is rank-2 parabolic. The groups $\Gamma_\pm(\xi)$ are both cyclic, since $S_\pm$ are incompressible hyperbolic surfaces, so their fundamental groups do not contain $\mathbb{Z}^2$ subgroups. So, we can write $\Gamma_\pm(\xi) = \langle\gamma_\pm\rangle$ for closed curves $\gamma_\pm$ on $S_\pm$. Both $\gamma_\pm$ are homotopic into some fixed component $T \subset \partial_{\chi=0}M$, the component whose fundamental group can be conjugated to stabilise ξ. So, there is a component $(X_0,\Sigma_0) \subset (X,\Sigma)$ of the characteristic submanifold such that Σ_0 intersects T and both $\gamma_\pm$ are homotopic on $S_\pm$ into Σ_0. Since $M \not\cong T^2\times[0,x]$, the component (X_0,Σ_0) is either an interval bundle over an annulus (so, a fibred solid torus), or an S^1-bundle pair, so by Fact 2.9, X_0 is either a fibred solid torus or a thickened torus. In either case, Σ_0 intersects each of $S_\pm$ in a fibred annulus, and these annuli are disjoint, so they are parallel on a torus boundary component of X_0, implying that $\gamma_\pm$ are homotopic in M, and hence $\Gamma_\pm(\xi)$ are conjugate in Γ. But since $\Gamma_\pm(\xi)$ have the same fixed point at infinity, the conjugating element must fix ξ, and therefore commute with the two groups, implying $\Gamma_-(\xi) = \Gamma_+(\xi)$.

Now assume $\Gamma(\xi)$ is cyclic. Pick a basepoint $p \in S_-$, say, and let $\gamma_- \subset S_-$ be a loop based at p representing a generator of $\Gamma_-(\xi)$. Represent a generator of $\Gamma_+(\xi)$ as $\alpha\cdot\gamma_+\cdot\alpha^{-1}$, where α is an arc from $p \in S_-$ to a point in S_+, and γ_+ is a loop in S_+. Since $\Gamma_\pm$ stabilise distinct components $\widetilde{S}_\pm$, the arc α is not homotopic into $S_- \cup S_+$. So, α is a spanning arc of an essential map from an annulus, where the boundary components of the annulus map to powers of $\gamma_\pm$. It follows that the loops $\gamma_\pm$ are homotopic on $S_\pm$ into Σ_0 for some component $(X_0,\Sigma_0) \subset (X,\Sigma)$.

If X_0 is an I-bundle with horizontal boundary Σ_0, then as $\gamma_\pm$ are not proper powers in $\pi_1 S_\pm$, they are both primitive in $\pi_1 X_0$, and hence $\gamma_\pm$ (rather than their powers) are homotopic in $X_0 \subset M$. Similarly, if (X_0,Σ_0) is a fibred solid torus, Σ_0 is a collection of parallel annuli on ∂X_0, so since $\gamma_\pm$ are primitive in $\pi_1 S_\pm$, they are homotopic on $S_\pm$ to *simple* closed curves in these annuli, and hence are homotopic to each other in X_0.

It follows that there are generators for $\Gamma_\pm(\xi)$ that are conjugate in Γ, but since these generators both fix ξ, they are equal. □

Claim 4.3 *For all $\xi \in \Lambda_- \cap \Lambda_+$, we have $\Gamma_-(\xi) = \Gamma_+(\xi)$. Moreover,*

$$\Lambda_\Delta = \Lambda_- \cap \Lambda_+.$$

Proof Let $N_\pm \subset \mathbb{H}^3$ be the 1-neighbourhood of the convex hull of $\Lambda_\pm$, and for small $\epsilon > 0$, let $T_\pm(\epsilon) \subset \mathbb{H}^3$ be the set of all points that are translated less than ϵ by some parabolic element of $\Gamma_\pm$. If ϵ is at most the Margulis constant ϵ_0, then $T_\pm(\epsilon)$ is a disjoint union of horoballs in $\mathbb{H}^3$.

The sets $N_\pm$ and $T_\pm(\epsilon)$ are $\Gamma_\pm$ invariant. Since $\Gamma_\pm$ is a finitely generated subgroup of Γ, which is geometrically finite, $\Gamma_\pm$ is geometrically finite as well by [14]. So, the action of $\Gamma_\pm$ on $N_\pm \setminus T_\pm(\epsilon)$ is cocompact, see, e.g., Theorem 3.7 in [50], implying that either the function

$$D_+ : \mathbb{H}^3 \to \mathbb{R}_{>0}, \quad D_+(x) = \min\{d(x, \gamma(x)) \mid \gamma \in \Gamma_+ \text{ loxodromic}\},$$

is bounded above on $N_+ \setminus T_+(\epsilon)$ by some $B(\epsilon) > 0$, or Γ_+ is elementary parabolic. A similar statement holds for $-$ instead of $+$. With ϵ_0 the Margulis constant, the Margulis lemma then implies that *if $\epsilon > 0$ is sufficiently small with respect to $B(\epsilon_0)$, and Γ_+ is not elementary parabolic, then*

$$T_-(\epsilon) \cap N_+ \subset T_+(\epsilon_0), \tag{2}$$

and similarly with $-$, $+$ exchanged. Indeed, if not then we have (say) a point $p \in \mathbb{H}^3$ that is translated by less than ϵ by some parabolic $\gamma_- \in \Gamma_-$ and by at most B by some loxodromic $\gamma_+ \in \Gamma_+$. If ϵ is small with respect to B, then both γ_- and $[\gamma_+, \gamma_-]$ translates p by at most ϵ_0, so they generate an elementary discrete group by the Margulis lemma applied to Γ, implying that γ_+ fixes the fixed point of γ_-, which contradicts that they generate a discrete group.

Fix $\xi \in \Lambda_+ \cap \Lambda_-$. We claim that $\Gamma_-(\xi) = \Gamma_+(\xi)$. By Claim 4.2 it suffices to show that whenever $\Gamma_-(\xi)$ is nontrivial, say, so is $\Gamma_+(\xi)$.

First, assume that $\Gamma_-(\xi)$ is elementary parabolic. We claim that $\Gamma_+(\xi)$ is elementary parabolic as well. Assume not, and let α be a geodesic ray in $\mathbb{H}^3$ converging to ξ. Then $\alpha(t)$ lies in $T_-(\epsilon) \cap N_+$ for large t, and therefore in $T_+(\epsilon_0)$ for large t by (2), which implies ξ is a parabolic fixed point of Γ_+ as well, a contradiction.

Next, suppose that $\Gamma_-(\xi)$ is elementary loxodromic. If ξ is a parabolic fixed point of Γ_+, we are done, so let's assume this isn't the case. Let α be the axis of $\Gamma_-(\xi)$, parametrised so $\alpha(t) \to \xi$ as $t \to \infty$. Since ξ is not a Γ_+ parabolic fixed point, there are $t_i \to \infty$ such that $\alpha(t_i) \notin T_+(\epsilon)$ for all i. Since the action of Γ_+ on $N_+ \setminus T_+(\epsilon)$ is cocompact, if $p \in \mathbb{H}^3$ is a fixed basepoint, there are elements $\gamma_i^+ \in \Gamma_+$ such that $\sup_i d(\gamma_i^+(p), \alpha(t_i)) < \infty$. Since the action of $\Gamma_-(\xi)$ on α is cocompact, there are then elements $\gamma_i^- \in \Gamma_-(\xi)$ with

$$\sup_i d(\gamma_i^+(p), \gamma_i^-(p)) < \infty.$$

By discreteness of Γ, after passing to a subsequence we can assume $\gamma_i^+ = \gamma_i^- \circ g$ for some fixed $g \in \Gamma$. Hence, for all i we have

$$\gamma_i^+ \circ (\gamma_1^+)^{-1} = \gamma_i^- \circ (\gamma_1^-)^{-1} \in \Gamma_+ \cap (\Gamma_-(\xi)) \subset \Gamma_+(\xi),$$

so we are done.

Finally, we want to show that $\Lambda_- \cap \Lambda_+ = \Lambda_\Delta$. The inclusion $\Lambda_\Delta \subset \Lambda_- \cap \Lambda_+$ is clear. So, take $\xi \in \Lambda_- \cap \Lambda_+$. We can assume that $\Gamma_\pm(\xi) = 1$, since otherwise we're in the cases handled above. Let α be a geodesic ray in $\mathbb{H}^3$ converging to ξ. As in the previous case, since ξ is not a parabolic fixed point of Γ_+, there are $t_i \to \infty$ such that $\alpha(t_i) \notin T_+(\epsilon_0)$ for all i. Discarding finitely many i, we have $\alpha(t_i) \in N_+$, so it follows from (2) that $\alpha(t_i) \notin T_-(\epsilon)$. Fixing a base point $p \in \mathbb{H}^3$, as Γ_- acts cocompactly on $N_- \setminus T_-(\epsilon)$ and Γ_+ acts cocompactly on $N_+ \setminus T_+(\epsilon_0)$, there are elements $\gamma_i^\pm \in \Gamma_\pm$ such that

$$\sup_i d(\gamma_i^\pm(p), \alpha(t_i)) < \infty.$$

So passing to a subsequence, $\gamma_i^+ = \gamma_i^- \circ g$ for some fixed $g \in \Gamma$, and then

$$\gamma_i^+ \circ (\gamma_1^+)^{-1} = \gamma_i^- \circ (\gamma_1^-)^{-1} \in \Gamma_+ \cap \Gamma_- = \Delta$$

for all i. But applying this sequence to p and letting $i \to \infty$ gives a sequence of points in the orbit $\Delta(p)$ that converge to ξ, so $\xi \in \Lambda_\Delta$. □

Now assume that $\Delta \neq 1$. We want to construct the interval bundle W mentioned in the statement of the theorem. After an isotopy on $\partial \mathrm{CC}(N)$, let's assume that $S_\pm$ is a subsurface of $\partial \mathrm{CC}(N)$ with geodesic boundary. Consequently, we allow degenerate subsurfaces, where $S_\pm$ is a simple closed geodesic, as well as subsurfaces where only the interior is embedded and two boundary components can coincide. As $\partial \mathrm{CC}(N)$ may have cusps, we also must allow $S_\pm$ to be noncompact with finite volume, rather than compact.

Recall that $\widetilde{S}_\pm$ is isometric to a convex subset of $\mathbb{H}^2$, and that if $\partial_\infty S_\pm \subset \partial \mathbb{H}^2$ is the boundary of $\widetilde{S}_\pm$ the inclusion $\widetilde{S}_\pm \hookrightarrow \mathbb{H}^3$ extends continuously to a $\Gamma_\pm$-equivariant quotient map

$$\iota_\pm : \partial_\infty \widetilde{S}_\pm \to \Lambda_\pm \subset \partial \mathbb{H}^3.$$

Moreover, if $\xi, \xi' \in \partial_\infty \widetilde{S}_+$, say, we have $\iota_+(\xi) = \iota_+(\xi')$ if and only if there is an element $\gamma \in \Gamma_+$ that acts hyperbolically on $\widetilde{S}_+ \cup \partial_\infty \widetilde{S}_+$ with fixed points $\xi, \xi' \in \partial_\infty \widetilde{S}_+$, but acts parabolically on $\mathbb{H}^3$. By discreteness of the action $\Gamma_+ \curvearrowright \widetilde{S}_+$, each $\xi \in \partial_\infty \widetilde{S}_+$ has the same image under ι_+ as *at most one* other ξ'. Similar statements holds with $-$ instead of $+$. All this is a consequence (for instance) of Bowditch's theory of the boundary of a relatively hyperbolic group [6]: since the action $\Gamma_\pm \curvearrowright \mathbb{H}^3$ is geometrically finite, $\Lambda_\pm$ is a model for the Bowditch boundary of the group $\Gamma_\pm$ relatively to its maximal parabolic subgroups, so the statement above follows from Theorem 1.3 of [44], say.[6]

Let $\iota_\pm^{-1}(\Lambda_\Delta) \subset \partial_\infty \widetilde{S}_\pm$. Since $\Delta \neq 1$ and is not cyclic parabolic, $\iota_\pm^{-1}(\Lambda_\Delta)$ has at least two points, so it has a well-defined convex hull $\widetilde{C}_\pm \subset S_\pm$.

Claim 4.4 (convex hulls) *One of the following holds.*

(1) Δ *is cyclic and acts hyperbolically on* $\widetilde{S}_+$. *The convex hull* $\widetilde{C}_+$ *is its geodesic axis, which is precisely invariant under* $\Delta \subset \Gamma$, *so that the quotient* $C_+ := \Delta \backslash \widetilde{C}_+$ *embeds as a simple closed geodesic in* S_+.

(2) $\widetilde{C}_+$ *is a subsurface of* $\widetilde{S}_+$ *with geodesic boundary*, $\mathrm{int}(\widetilde{C}_+)$ *is precisely invariant under* $\Delta \subset \Gamma$, *and the quotient* $C_+ := \Delta \backslash \widetilde{C}_+$ *is a generalised subsurface of* S_+ *with compact geodesic boundary.*

A similar statement holds with $-$ *instead of* $+$.

[6] See also Theorem 5.6 of [53], which says that there is a continuous equivariant extension $\iota_\pm$ of the inclusion $\widetilde{S}_\pm \hookrightarrow \mathbb{H}^3$ as above. This theorem is stated in a much more general setting, though, and our statement is a trivial case.

Proof Let's work with $+$ for concreteness. If $g \in \Delta$, then $g(\Lambda_\Delta) = \Lambda_\Delta$, so g leaves $\iota_+^{-1}(\Lambda_\Delta)$ invariant by equivariance of ι_+. Hence g leaves $\widetilde{C}_+$ invariant.

Let's suppose first that $\widetilde{C}_+$ has nonempty interior, since that is the more interesting case. We'll address the case that $\widetilde{C}_+$ is a biinfinite geodesic at the end of the proof. Let $g \in \Gamma_+ \setminus \Delta$. We want to show

$$g(\mathrm{int}(\widetilde{C}_+)) \cap \mathrm{int}(\widetilde{C}_+) = \varnothing.$$

Assume this is not the case. By Claim 4.3, the fixed points of g in $\partial_\infty S_+$ lie outside $\iota_+^{-1}(\Lambda(\Delta))$. So, we cannot have $g(\widetilde{C}_+) \subset \widetilde{C}_+$, as then we'd have $g^n(\widetilde{C}_+) \subset \widetilde{C}_+$ for all n, contradicting that points of $\widetilde{S}_+$ converge to the fixed points of g under iteration. Considering backwards iterates, we also cannot have $\widetilde{C}_+ \subset g(\widetilde{C}_+)$. Therefore, $\partial\widetilde{C}_+$ and $\partial g(\widetilde{C}_+)$ intersect transversely.

Since $\widetilde{C}_+$ has nonempty interior, Δ is nonelementary, and therefore the fixed points of loxodromic isometries of Δ are dense in Λ_Δ. Loxodromic fixed points of Δ are in particular *not* parabolic fixed points in $\Gamma_\pm$, so any biinfinite geodesic in $\widetilde{C}_+$ is a limit of biinfinite geodesics in $\widetilde{C}_+$ whose endpoints are *not* fixed points of parabolic isometries of $\Gamma_\pm$. By the previous paragraph, there are then biinfinite geodesics α_+, β_+ in $\widetilde{C}_+$ such that $g(\alpha_+)$ and β_+ intersect transversely, and where the endpoints of α_+, β_+ project under ι_+ to points $\xi_\alpha, \xi'_\alpha, \xi_\beta, \xi'_\beta \in \Lambda_\Delta$ that are not parabolic fixed points in $\Gamma_\pm$.

Let α_- be the geodesic in $\widetilde{S}_-$ whose endpoints in $\partial_\infty\widetilde{S}_-$ map to the points ξ_α, ξ'_α under ι_-. Define β_- similarly. Then

$$\alpha := \alpha_+ \cup \{\xi_\alpha, \xi'_\alpha\} \cup \alpha_-, \quad \beta := \beta_+ \cup \{\xi_\beta, \xi'_\beta\} \cup \beta_-$$

are two *simple* closed curves on the closure $\mathrm{cl}(\partial \mathrm{CH}(\Lambda_\Gamma)) \subset \mathbb{H}^3 \cup \partial\mathbb{H}^3$, which is homeomorphic to a sphere. For instance, the arcs $\alpha_\pm$ are disjoint and $\xi_\alpha \neq \xi'_\alpha$, since the endpoints of α_+ are not parabolic fixed points.

Now consider how the two simple closed curves $g(\alpha)$, β intersect. The arcs β_- and $g(\alpha_+)$ are disjoint since $\widetilde{S}_- \neq \widetilde{S}_+$. The arcs $g(\alpha_-)$, β_- are disjoint since $g \notin \Gamma_-$ and hence $g(\alpha_-)$ lies on a different translate of $\widetilde{S}_-$ than $\beta_- \subset \widetilde{S}_-$. Moreover, since $g(\alpha_+)$, β_+ intersect transversely in $\widetilde{S}_+$, the endpoints of $g(\alpha_+)$ and β_+ are distinct in $\partial_\infty\widetilde{S}_+$, and since none of them are parabolic fixed points, the points $g(\xi_\alpha)$, $g(\xi'_\alpha)$, ξ_β, ξ'_β are all distinct. But by assumption, $g(\alpha_+)$ intersects β_+ transversely in a single point! This shows that $g(\alpha)$ and β intersect exactly once, transversely, which is a contradiction.

By precise invariance of the action on the interior, the quotient $\mathrm{int}(C_+) = \Delta \setminus \mathrm{int}(\widetilde{C}_+)$ embeds in the finite volume surface S_+, so C_+ has finite volume itself. So if ∂C_+ is noncompact, it must have two noncompact boundary components that are asymptotic. Lifting, we get two boundary components β_1, β_2 of $\widetilde{C}_+$ that are asymptotic. Since $\widetilde{C}_+$ is convex, it is contained in the subset of $\mathbb{H}^2$ bounded by β_1, β_2, and hence the common endpoint of β_1, β_2 is an isolated point of Λ_Δ, which is a contradiction since Δ is not elementary.

The case when $\widetilde{C}_+$ is a biinfinite geodesic is similar. Here, Δ must be cyclic, acting on $\widetilde{S}_+$ with axis $\widetilde{C}_+$, and acting either parabolically or loxodromically on $\mathbb{H}^3$. In the parabolic case, $\widetilde{C}_+$ compactifies to a simple closed curve on the sphere $\mathrm{cl}(\mathrm{CH}(\Lambda_\Gamma)) \subset \mathbb{H}^3 \cup \partial\mathbb{H}^3$, so no translate $g(\widetilde{C}_+)$, $g \in \Gamma_+$, can intersect $\widetilde{C}_+$ transversely, since if it did we'd get two simple closed curves on the sphere that intersect once. In the loxodromic case, we get a similar contradiction by looking at the simple closed curve $\mathrm{cl}(\widetilde{C}_+ \cup \widetilde{C}_-) \subset \mathrm{cl}(\partial \widetilde{M})$ and its g-image. So, $\widetilde{C}_+$ is precisely invariant under $\Delta \subset \Gamma$. The quotient $C_+ := \Delta \backslash \widetilde{C}_+$ is obviously compact, and is therefore a simple closed geodesic in S_+. □

We claim that C_- and C_+ are homeomorphic. If $C_\pm$ are isotopic in $\partial\mathrm{CC}(M)$ this is clear, and otherwise we argue as follows. The subgroups $\pi_1 C_\pm$ are both represented by Δ, so are conjugate in $\pi_1 M$. The fact that every curve in C_- is homotopic to a curve in C_+ (and vice versa) implies that $C_\pm$ are isotopic to subsurfaces $C'_\pm \subset \Sigma$ in the boundary Σ of a component (X, Σ) of the characteristic submanifold[7] of $(\mathrm{CC}(M), S_- \cup S_+)$, see Section 2.6, and that even within X every closed curve in C'_- is homotopic to a closed curve in C'_+, and vice versa. When X is a solid torus or thickened torus, $C'_\pm$ are annuli, while if X is an interval bundle, $C'_\pm$ bound a vertical interval bundle in X, and are homeomorphic.

So, let $f : C_- \to C_+$ be a homeomorphism, lift f to a Δ-equivariant homeomorphism $\tilde{f} : \widetilde{C}_- \to \widetilde{C}_+$ and let

$$F : \widetilde{C}_- \times [0,1] \to \mathrm{CH}(\Lambda),$$

where $F(x, \cdot)$ parametrises the geodesic from x to $f(x)$. Then F is Δ-equivariant, and projects to an essential homotopy from C_- to C_+, as desired.

4.1 An annulus theorem for laminations

Suppose M is a compact, orientable, hyperbolisable 3-manifold with nonempty boundary and let $S = \partial_{\chi<0} M$ be the union of all nontorus boundary components of M. When $\alpha, \beta \subset S$ are disjoint simple closed curves that are essential and homotopic in M, but not homotopic in S, the annulus theorem says that there is an essential embedded annulus $A \subset M$ with $\partial A = \alpha \cup \beta$; see Scott's paper [57].

More generally, equip S with an arbitrary hyperbolic metric. An *essential homotopy* between two geodesic laminations $\lambda_\pm$ on S is a map

$$H : (\lambda \times [-1,1], \lambda \times \{-1,1\}) \to (M, S),$$

where λ is a lamination, such that H maps $\lambda \times \{\pm 1\}$ homeomorphically onto $\lambda_\pm$, and where H is not homotopic rel $\lambda \times \{-1,1\}$ into ∂M.

Here is an "annulus theorem" for minimal laminations.

Proposition 4.5 (an annulus theorem for laminations) *Let λ_-, λ_+ be two minimal geodesic laminations on S that are either disjoint or equal, and assume that $S(\lambda_\pm)$ are incompressible in M. If $\lambda_\pm$ are essentially homotopic in (M, S), there is an essential interval bundle $(B, \partial_H B) \subset (M, S)$ such that $\lambda_\pm$ fill $\partial_H B$, and where $\lambda_\pm$ are essentially homotopic through B, as in Section 2.9.*

Here, $S(\lambda_\pm)$ are the subsurfaces with geodesic boundary filled by $\lambda_\pm$, as in Section 2.8. The assumption that they are incompressible generalises the assumption that α, β are homotopically essential in M in the annulus theorem.

[7] Really, we need to be using resolutions of our subsurfaces with geodesic boundary, as discussed in Section 2.1.

Proof Identify $M \setminus \partial_{\chi=0} M$ with the convex core of a geometrically finite hyperbolic 3-manifold. Set $S_\pm := S(\lambda_\pm)$. Lift the essential homotopy from λ_- to λ_+ to a homotopy from lifts $\tilde{\lambda}_- \subset \widetilde{S}_-$ to $\tilde{\lambda}_+ \subset \widetilde{S}_+$ in $\mathbb{H}^3$. Under the homotopy, which has bounded tracks, corresponding leaves of $\tilde{\lambda}_\pm$ have the same endpoints in $\partial \mathbb{H}^3$. The endpoints of $\tilde{\lambda}_\pm$ are dense in $\partial_\infty \widetilde{S}_\pm$, so this means that the subsurfaces $C_\pm \subset S_\pm$ constructed in Theorem 4.1 are just $C_\pm = S_\pm$. Passing to disjoint or equal resolutions $S'_\pm$ of $S_\pm$ and applying Theorem 2.8 gives an interval bundle B where $\lambda_\pm$ fill $\partial_H B = S'_- \cup S'_+$.

We claim that $\lambda_\pm$ are essentially homotopic through B. By Fact 2.16, it suffices to show that if σ is the canonical involution of B, as described in Proposition 4.5, then $\sigma(\lambda_\pm)$ is isotopic to $\lambda_\mp$ on $S'_\mp$. Using the notation of Theorem 4.1, σ lifts to a Δ-equivariant involution $\tilde{\sigma}$ of $\widetilde{B}$ that exchanges $\widetilde{S}'_-$ and $\widetilde{S}'_+$, where here $\Delta = \Gamma_- \cap \Gamma_+$. By equivariance, $\tilde{\sigma}$ extends continuously to the identity on Λ_Δ, so $\tilde{\sigma}(\tilde{\lambda}_-)$ is a lamination on $\widetilde{S}_+$ with all the same endpoints at infinity as $\tilde{\lambda}_+$, and hence equals $\tilde{\lambda}_+$. □

5 Laminations on the boundary

Suppose M is a compact, orientable 3-manifold with hyperbolisable interior and nonempty boundary ∂M. Equip M with an arbitrary Riemannian metric and lift it to a Riemannian metric on the universal cover $\widetilde{M}$. As in the introduction, a biinfinite path or ray h on $\partial \widetilde{M}$ is called *homoclinic* if there are points s^i, t^i with $|s^i - t^i| \to \infty$ such that

$$\sup_i d_{\widetilde{M}}(h(s^i), h(t^i)) < \infty.$$

Two rays h_+, h_- on $\partial \widetilde{M}$ are called *mutually homoclinic* if there are parameters $s^i_\pm \to \infty$ such that

$$\sup_i d_{\widetilde{M}}(h_+(s^i_+), h_-(s^i_-)) < \infty.$$

Here, a *ray* is a continuous map from an interval $[a, \infty)$, and a *biinfinite path* is a continuous map from $\mathbb{R}$. We will also call rays and paths on ∂M (mutually) homoclinic if they have lifts that are (mutually) homoclinic paths on $\partial \widetilde{M}$. We refer the reader to Section 5.1 for some comments on alternate definitions of homoclinic that exist in the literature.

Note that if we divide a biinfinite homoclinic path into two rays, then either one of the two rays is itself homoclinic, or the two rays are mutually homoclinic. Also, these definitions are metric independent: since M is compact, any two Riemannian metrics on M lift to quasi-isometric metrics on $\widetilde{M}$, and a path is homoclinic or mutually homoclinic with respect to one metric if and only if it is with respect to the other metric.

Here are some examples.

Example 5.1 (1) Suppose that D is a properly embedded disc in M, and $h : \mathbb{R} \to \partial M$ is a path that covers $\partial D \subset \partial M$. Then h is homoclinic: indeed, D lifts homeomorphically to $\widetilde{M}$, so h lifts to a path in $\widetilde{M}$ with compact image.

(2) Suppose that $\phi : (S^1 \times [0,1], S^1 \times \{0,1\}) \to (M, \partial M)$ is an essential embedded annulus. Then rays covering the two boundary components of the annulus are mutually homoclinic: indeed, ϕ lifts to

$$\tilde{\phi} : \mathbb{R} \times [0,1] \to \widetilde{M},$$

and we have $\sup_{t \in \mathbb{R}} d(\tilde{\phi}(t,0), \tilde{\phi}(t,1)) < \infty$, so restricting to $t \in [0,\infty)$ we get two mutually homoclinic rays in $\widetilde{M}$.

It will be convenient below to work with a particular choice of metric on M.

Definition 5.2 (an explicit metric on M) Let $\partial_{\chi<0} M$ be the union of all components of ∂M that have negative Euler characteristic, i.e., are not tori. Thurston's Haken hyperbolisation theorem, see [33], implies that there is a hyperbolic 3-manifold $N = \mathbb{H}^3/\Gamma$ homeomorphic to the interior of M, where every component of $\partial_{\chi<0} M$ corresponds to a convex cocompact end of N. A torus $T \subset \partial M$, on the other hand, determines a cusp of N. So, in other words, N is "minimally parabolic": the only parabolics come from torus boundary components of M. For each T, pick an open neighbourhood $N_T \subset N$ of the associated cusp that is the quotient of a horoball in $\mathbb{H}^3$ by a $\mathbb{Z}^2$-action. Then

$$M \cong \mathrm{CC}(N) \setminus \bigcup_{\text{tori } T \subset \partial M} N_T, \tag{3}$$

and we will identify M with the right-hand side everywhere below. Then

- $\widetilde{M} \subset \mathbb{H}^3$ is obtained from the convex hull $\mathrm{CH}(\Gamma) \subset \mathbb{H}^3$ of the limit set of Γ by deleting an equivariant collection of horoballs, and
- the path metric induced on $\partial_{\chi<0} M$ is hyperbolic [58, Proposition 8.5.1], and the path metric induced on every torus $T \subset \partial M$ is Euclidean.

We now specialise to the case of paths that are *geodesics* on ∂M. Recall from Example 5.1(1) above that one can make homoclinic paths by running around the boundaries of disks in ∂M. The following shows that discs are essential in such constructions.

Fact 5.3 *Suppose that $S \subset \partial M$ is an essential subsurface. Then the inclusion of any lift $\widetilde{S} \subset \partial \widetilde{M}$ is a quasi-isometric embedding into $\widetilde{M}$. Moreover if S is incompressible then any pair of mutually homoclinic infinite rays on S are asymptotic and no biinfinite geodesic γ in S is homoclinic.*

Proof Think of M as embedded in a complete hyperbolic 3-manifold N as in (3), write $N = \Gamma \backslash \mathbb{H}^3$, and let $\widetilde{M} \subset \mathbb{H}^3$ be the preimage of M, so that $\widetilde{M}$ is obtained from the convex hull $\mathrm{CH}(\Gamma)$ be deleting an equivariant collection of horoballs. Fix a subgroup $\Delta < \Gamma$ that represents the conjugacy class associated to the image of the fundamental group of $S \subset M$. To show that

$$\widetilde{S} \hookrightarrow \partial \widetilde{M}$$

is a quasi-isometric embedding, it suffices to show that Δ is undistorted in Γ. But since M is geometrically finite and Δ is finitely generated, it follows from a result of Thurston (see Proposition 7.1 in [54]) that the group Δ is geometrically finite, and geometrically finite subgroups of (say, geometrically finite) hyperbolic 3-manifold groups are undistorted; see Corollary 1.6 in [28].

For the "moreover" statement, assume S is incompressible, so that $\widetilde{S}$ is simply connected, and consider a pair of infinite rays

$$h^{\pm} : \mathbb{R}^+ \to \widetilde{S}$$

that are geodesic for the induced hyperbolic metric and $t_n^{\pm} \to +\infty$ such that $d_{\widetilde{M}}(h^+(t_n^+), h^-(t_n^-))$ is bounded. Since $\widetilde{S} \subset \partial \widetilde{M}$ is a quasi-isometric embedding, $d_{\widetilde{S}}(h^+(t_n^+), h^-(t_n^-))$ is also bounded. Since $\widetilde{S}$ is simply connected and hyperbolic, this is possible only if h^+ and h^- are asymptotic on S. Taking $h^+ = h^-$ we get that a geodesic ray on $\widetilde{S}$ cannot be homoclinic. Taking $h^+ \neq h^-$, we get that any pair of mutually homoclinic infinite rays on S are asymptotic. In particular two disjoint geodesic rays in a homoclinic geodesic should be asymptotic. This is impossible for a geodesic in a simply connected hyperbolic surface. □

Example 5.1(2) above shows how embedded annuli in M can be used to create mutually homoclinic rays. In analogy to Fact 5.3, one can show that annuli are essential in such a construction. For instance, suppose M is acylindrical. Then work of Thurston, see [33] and more generally [42], says that we can choose the hyperbolic manifold N so that $\partial \mathrm{CC}(N) \cong \partial_{\chi<0} M$ is totally geodesic. Hence, the preimage of $\partial_{\chi<0} M$ in $\widetilde{M} \subset \mathbb{H}^3$ is a collection of hyperbolic planes. Any geodesic ray on $\partial_{\chi<0} M$ then lifts to a geodesic in $\mathbb{H}^3$, and two geodesic rays on $\partial_{\chi<0} M$ are mutually homoclinic if and only if their geodesic lifts are asymptotic in $\mathbb{H}^3$, which implies that they were asymptotic on $\partial_{\chi<0} M$.

5.1 Alternate definitions of homoclinic

Above, we defined a path

$$h : I \to \partial \widetilde{M}$$

to be homoclinic if there is are $s^i, t^i \in I$ with $|s^i - t^i| \to \infty$ such that

$$\sup_i d_{\widetilde{M}}(h(s^i), h(t^i)) < \infty.$$

Some other papers use slight variants of this definition. For example, the definition of (*faiblement*) *homoclinique* in Otal's thesis [56] is almost the same as what is written above, except that distances are computed *in the intrinsic metric on* $\partial \widetilde{M}$ instead of in $\widetilde{M}$. This is equivalent to our definition, though: the nonobvious direction follows from Fact 5.3, which says that boundary components of $\widetilde{M}$ quasi-isometrically embed in $\widetilde{M}$. And in the definition of *homoclinique* in Lecuire's earlier work [42], distances are computed not in $\widetilde{M}$, but within $\mathbb{H}^3$, with respect to a given identification of $\widetilde{M}$ with the convex core of some minimally parabolic hyperbolic 3-manifold, as discussed in Definition 5.2. When M has tori in its boundary, the inclusion $\widetilde{M} \hookrightarrow \mathbb{H}^3$ is not a quasi-isometric embedding, but the following lemma says that $d_{\mathbb{H}^3}$ is bounded if and only if $d_{\widetilde{M}}$ is bounded, so Lecuire's earlier definition is equivalent to ours.

Lemma 5.4 *Whenever* $x, y \in \widetilde{M}$, *we have*

$$d_{\mathbb{H}^3}(x,y) \le d_{\widetilde{M}}(x,y) \le e^{d_{\mathbb{H}^3}(x,y)/2} d_{\mathbb{H}^3}(x,y).$$

Proof Set $N := \Gamma \backslash \mathbb{H}^3$, so that $\widetilde{M}$ is obtained from the convex hull $\mathrm{CH} := \mathrm{CH}(\Lambda(\Gamma))$ of the limit set of Γ by deleting horoball neighbourhoods around all rank-two cusps. Take a $\mathbb{H}^3$-geodesic γ from x to y. Then γ lies inside CH, and it can only penetrate the deleted horoball neighbourhoods to a depth of $d(x,y)/2$. Now, whenever $B \supset B'$ are horoballs in $\mathbb{H}^3$ such that $d_{\mathbb{H}^3}(\partial B, B') \le d(x,y)/2$, the closest point projection

$$\pi : B \setminus B' \to \partial B$$

is well defined and $e^{d(x,y)/2}$-lipschitz. (Indeed, it suffices to take B as the height 1 horoball in the upper half space model and B' as the height $e^{d(x,y)/2}$ horoball, and then the claim is obvious.) So, the parts of γ above that penetrate the deleted horoballs can be projected back into $\partial \widetilde{M}$, and if we do this the resulting path has length at most $e^{d(x,y)/2} d(x,y)$. □

We should mention the version of homoclinic defined in Casson's original unpublished notes. There, M is a handlebody, and if we regard $\partial \widetilde{M} \hookrightarrow \mathbb{H}^3$ as above, then a simple geodesic $h : I \to \partial \widetilde{M}$ is called *homoclinic* if when we subdivide h into two rays $h_\pm$, these rays limit onto subsets $A_\pm \subset \mathbb{H}^3 \cap \partial \mathbb{H}^3$ such that $A_+ \cap A_- \neq \varnothing$. This definition is stronger than all the ones mentioned above: if $A_+ \cap A_-$ contains a point on $\partial \widetilde{M}$, rather than at infinity, then the definition of homoclinic above is obviously satisfied. Otherwise, $h_\pm$ have to have a common accumulation point in $\partial \mathbb{H}^3$, which corresponds to an end ξ of $\widetilde{M}$, and one can use the treelike structure of the universal cover $\widetilde{M}$ of the handlebody M to say that $h_\pm$ have to both intersect a sequence of meridians (m_i) on $\widetilde{M}$ that cut off smaller and smaller neighbourhoods of ξ. The times $t_p^i m$ when $h_\pm$ intersects m_i then work in the definition of homoclinic above. In fact, Casson's definition is strictly stronger. For instance, if $h_\pm$ both spiral around disjoint meridians $\gamma_\pm \subset \partial \widetilde{M}$, then h is homoclinic by our definition but not by Casson's. However, Theorem 1.1 still fails using Casson's original definition, due to the examples in Figure 13 on page 1849.

5.2 Waves, tight position, and intrinsic limits

As in the previous section, let M be a compact, orientable hyperbolisable 3-manifold with nonempty boundary ∂M, which we think of as the convex core of a hyperbolic 3-manifold with horoball neighbourhoods of its rank-2 cusps deleted.

Definition 5.5 (waves and tight position) Suppose that m is a meridian multicurve on ∂M, and let $\gamma \subset \partial M$ be a simple closed geodesic, a simple geodesic ray, or a simple biinfinite geodesic. An *m-wave* is a segment $\beta \subset \gamma$ that has endpoints on m, and is homotopic rel endpoints in M to an arc of m. If γ has no m-waves, and every infinite length segment of γ intersects m, then we say that γ is in *tight position* with respect to m.

Waves and tight position were discussed previously in [38; 42], for instance. Note that in our definition, an m-wave β can intersect m in its interior. More generally, an m-wave of a lamination is an m-wave of one of its leaves, and a lamination is in *tight position* with respect to m if all of its leaves are. Note that from this perspective, if a geodesic γ is in tight position with respect to some multicurve m (regarded as a lamination), then it is in tight position with respect to some component of m.

As an example, a meridian γ can never be in tight position with respect to another meridian m: taking discs with boundaries γ and m that are transverse and intersect minimally, any arc of intersection of these disks terminates in a pair of intersection points of γ and m that bound a m-wave of γ.

More generally, we have the following fact.

Fact 5.6 (tight position $\Longrightarrow \mathbb{H}^3$ quasigeodesic) *Let γ be a simple geodesic ray or biinfinite geodesic on ∂M. If γ is in tight position with respect to some meridian m then any lift $\tilde{\gamma} \subset \partial \widetilde{M}$ of γ is a quasigeodesic in $\mathbb{H}^3$. In particular, $\tilde{\gamma}$ is an $\widetilde{M}$-quasigeodesic, and is not homoclinic.*

Proof Intersecting with m breaks γ into a union of finite arcs. By simplicity of γ, these arcs fall into only finitely many homotopy classes rel m, and there is a universal upper bound $L = L(\gamma, m)$ on their lengths. Let D be a disc with boundary m and let $\widetilde{D}$ be the entire preimage of D in $\widetilde{M}$. Tightness means that the path $\tilde{\gamma}$ intersects infinitely many components of $\widetilde{D}$, and intersects no single component more than once.

In the notation of Definition 5.2, we have that $\widetilde{M} \subset \mathbb{H}^3$ is obtained from $\mathrm{CH}(\Gamma)$ by deleting an equivariant collection of horoballs. Each component of $\widetilde{D}$ separates $\mathrm{CH}(\Gamma)$, so if γ' is a segment of γ, any geodesic in $\mathbb{H}^3$ joining the endpoints of γ' must intersect each of the discs that γ' intersects. Hence, if $\epsilon > 0$ is the minimum distance between any two components of $\widetilde{D}$, then $\tilde{\gamma}$ is a $(L/\epsilon, L)$-quasigeodesic. □

Given a lamination, we now describe how to create a system of meridians with respect to which the lamination is in tight position.

Definition 5.7 (surgery) Suppose that λ is a geodesic lamination on ∂M and $m = \bigsqcup_{i=1}^n m_i$ is a geodesic meridian multicurve on ∂M, and β is an m-wave in λ whose interior is disjoint from m. Then the pair of points $\partial\beta$ separates some component m_i of m into two arcs m_i^1, m_i^2, both of which are homotopic to β rel endpoints in M. We perform a λ*-surgery* on m by replacing m_i^1 (say) with β, thus constructing a new multicurve $m' := (\beta \cup m_i^2) \sqcup \bigsqcup_{j \neq i} m_j$.

This notion of surgery appears in many other references, e.g., [3; 18; 38; 42]. We summarise its elementary properties here:

Fact 5.8 *Suppose that λ is a geodesic lamination.*

(1) *If m is a meridian and λ has an m-wave, it also has an m-wave whose interior is disjoint from m, so a λ-surgery can be performed.*

(2) *Any curve m' obtained by λ-surgery on a meridian m as above is a meridian.*

(3) *If m is a* **cut system** *for M, i.e., a multicurve of meridians bounding discs that cut M into balls and 3-manifolds with incompressible boundary, then some λ-surgery on a component of m is another cut system.*

Proof For (1), suppose that λ has an m-wave β. Let $\tilde{m}$ be the entire preimage of m in the cover $\partial\widetilde{M}$, and lift β to an arc $\tilde{\beta}$ starting and ending on some fixed component $m_0 \subset \tilde{m}$. Since each component of $\tilde{m}$ separates $\partial\widetilde{M}$, there is some "outermost" subarc $\tilde{\beta}'$ that has endpoints on the same component of $\tilde{m}$, and that has interior disjoint from $\tilde{m}$. This $\tilde{\beta}'$ projects to an m-wave of λ whose interior is disjoint from m.

For (2), note that if $m' := \beta \cup m_2$ is obtained by λ-surgery on m, as above, then m, m' are homotopic in M, and hence m' is nullhomotopic. Also, if m' is inessential in ∂M, then β is homotopic on ∂M to m_2, implying that λ and m were not in minimal position, a contradiction since they are both geodesic. Hence m' is a meridian.

For (3), consider an m-wave in λ whose interior is disjoint from m and say that $\partial\beta \subset m_1$. Then $\partial\beta$ separates m_1 into two arcs m_1^1, m_1^2. It is not difficult to see that either $(\beta \cup m_1^1) \sqcup \bigsqcup_{j \neq 1} m_j$ or $(\beta \cup m_1^2) \sqcup \bigsqcup_{j \neq 1} m_j$ is a cut system. □

The following lemma is a modification of a result of Kleineidam and Souto [38, Lemmas 7 and 8] that is essential for everything below.

Lemma 5.9 (no waves, or a sequence of meridians) *Suppose λ is a geodesic lamination on $S = \partial M$ and m is a meridian multicurve. Then either*

(1) *there exists a finite sequence of λ-surgeries on m that terminates in some meridian multicurve m' where λ has no m'-waves*,

(2) *$S(\lambda)$ contains a sequence of meridians (γ_i) such that $i(\lambda, \gamma_i) \to 0$, with respect to every transverse measure on λ.*

Here, (2) makes sense even when λ admits no transverse measure of full support. Note that if λ is a minimal lamination and $\partial S(\lambda)$ is incompressible, then (2) implies that λ is an intrinsic limit of meridians.

Proof of Lemma 5.9 The two cases depend on whether λ contains infinitely many homotopy classes of m-waves, or not. Here, our homotopies are through arcs on S, keeping their endpoints on m.

If there are only finitely many classes of m-waves in λ, then a finite sequence of λ-surgeries converts m into a multicurve m' such that λ has no m'-waves, as each surgery decreases the number of waves by at least one. If there are infinitely many homotopy classes of m-waves in λ, then we can choose a sequence of parameterised m-waves $\alpha_i : [0,1] \to \mathbb{R}$ such that

(1) the two sequences of endpoints $(\alpha_i(0))$ and $(\alpha_i(1))$ both converge, and if either sequence converges into a simple closed curve $\gamma \subset \lambda$, then it approaches γ from only one side,

(2) no α_i and α_j are homotopic keeping their endpoints on m, for $i \neq j$.

To construct the desired sequence of meridians, let β_i^0 be the shortest geodesic on S from $\alpha_i(0)$ to $\alpha_{i+1}(0)$, and define β_i^1 similarly. Since $(\alpha_i(0)) \subset m$ and $(\alpha_i(1)) \subset m$ converge, we have $\beta_i^0, \beta_i^1 \subset m$ for i large enough. For large i, the union $\beta_i^0 \cup \alpha_i \cup \beta_i^1 \cup \alpha_{i+1}$ is an essential closed curve in $S(\lambda)$ that is nullhomotopic in M. It may not be simple, since β_i^0 and β_i^1 may overlap, but it has at most one self intersection. So by the loop theorem [29], one of the three simple closed curves obtained by surgery on it is a meridian γ_i.

Now, the fact that the endpoints can approach a simple closed curve in λ only from one side implies that for large i, the curves γ_i do not intersect any simple closed curve contained in λ. Since γ_i only intersects λ along the arcs β_i^0 and β_i^1, whose hyperbolic lengths converge to zero, it follows that $i(\gamma_i,\lambda)\to 0$ for any transverse measure on λ. □

Here is an important application of Lemma 5.9.

Lemma 5.10 (quasigeodesic or a sequence of meridians) *Suppose $\lambda\subset\partial_{\chi<0}M$ is a minimal geodesic lamination and that $\partial S(\lambda)$ is incompressible in M. Let $h\subset S(\lambda)$ be a simple geodesic ray or biinfinite geodesic that is disjoint from λ or contained in λ. Then either*

1. *any lift $\tilde h\subset\partial\widetilde M$ of h is a quasigeodesic in $\widetilde M$, or*
2. *$S(\lambda)$ contains a sequence of meridians (γ_i) such that $i(\lambda,\gamma_i)\to 0$, with respect to every transverse measure on λ.*

In particular, if h is homoclinic, then λ satisfies (2).

Proof Assume that (2) does not hold. If h is a geodesic ray, it is asymptotic to a geodesic ray $l^+\subset\lambda$ and any lift of h is a quasigeodesic if and only if any lift of l^+ is quasigeodesic. Let $\mu=\lambda\cup h$ if h is biinfinite and $\mu=\lambda$ otherwise. Given a cut system m for M, Lemma 5.9 and Fact 5.8(3) say that we can perform μ-surgeries until we obtain a new cut system m such that μ has no m-waves. If m intersects μ, then μ is in tight position with respect to m, so (1) follows from Fact 5.6. Therefore, we can assume m does not intersect μ λ. Up to isotopy, we can also assume that $S(\lambda)$ does not intersect m. Since $\partial S(\lambda)$ is assumed to be a collection of incompressible curves, it follows that $S(\lambda)$ is itself incompressible, so (1) follows from Fact 5.3. □

We now come to the central definition of the section.

Definition 5.11 A minimal geodesic lamination $\lambda\subset\partial_{\chi<0}M$ is an *intrinsic limit of meridians* if there is a transverse measure[8] on λ and a sequence of meridians (γ_i) contained in $S(\lambda)$ such that $\gamma_i\to\lambda$ in $\mathcal{PML}(S(\lambda))$.

Using Lemma 5.10, we can prove the following proposition, which gives several equivalent characterisations of intrinsic limits.

Proposition 5.12 (intrinsic limits) *Suppose $\lambda\subset S=\partial M$ is a minimal geodesic lamination and $\partial S(\lambda)$ is incompressible. The following are equivalent*:

(1) *λ is an intrinsic limit of meridians.*

(2) *Given (some/any) transverse measure on λ, there is a sequence of meridians (γ_i) in $S(\lambda)$ such that $i(\gamma_i,\lambda)\to 0$.*

[8] It is currently unknown whether the particular transverse measure matters: we might suspect that a measured lamination is a projective limit of meridians if and only if the same is true for any other measured lamination with the same support, but there could also very well be a counterexample.

(3) *There is a homoclinic geodesic in* $S(\lambda)$ *that is either a leaf of* λ, *or is disjoint from* λ.

(4) *Given any transverse measure on* λ, *there is a sequence of essential (possibly nonsimple) closed curves* (γ_i) *in* $S(\lambda)$ *such that each* γ_i *is nullhomotopic in* M, *and* $i(\gamma_i, \lambda) \to 0$.

Note that when we say $\partial S(\lambda)$ is incompressible, we mean that no closed curve that is a boundary component of $S(\lambda)$ is nullhomotopic in M. This condition is mainly here to make statements and proofs easier. For instance, without this assumption our proof of (4) $\implies$ (2) may produce peripheral meridians, but peripheral meridians can't be used in (2) $\implies$ (1).

Proof **(2) $\implies$ (1)** Fix some transverse measure on λ. By (2),

$$i(\gamma_i, \lambda)/\text{length}(\gamma_i) \to 0,$$

so after passing to a subsequence we can assume that (γ_i) converges to a measured lamination μ in $S(\lambda)$ that does not intersect λ transversely. As λ fills $S(\lambda)$, μ is supported on λ.

(1) $\implies$ (3) After passing to a subsequence, we can assume that (γ_i) converges in the Hausdorff topology to some lamination, which must then be an extension of λ by finitely many leaves. The statement (3) follows from an unpublished criterion of Casson, see Lecuire [42, Théorème B.1] for a proof, that states that any Hausdorff limit of meridians has a homoclinic leaf.

(3) $\implies$ (2) This is an immediate corollary of Lemma 5.10.

(4) $\iff$ (2) The direction $\Leftarrow$ is immediate, so suppose (γ_i) is a sequence of essential closed curves in $S(\lambda)$ that are nullhomotopic in H and $i(\lambda, \gamma_i) \to 0$. By Stallings' version of the loop theorem, for each i there is a meridian γ_i' that is obtained from γ_i by surgery at the self intersection points. Such surgeries can only decrease the intersection number with λ, so (2) follows. □

We will also need the following criterion in the next section.

Lemma 5.13 (intrinsic limits of annuli) *Suppose* $\lambda \subset \partial_{\chi<0} M$ *is a minimal lamination such that* $S(\lambda)$ *is compressible but* $\partial S(\lambda)$ *is incompressible, and that there is a sequence* (A_i) *of essential embedded annuli in* $(M, S(\lambda))$ *with* $i(\partial A_i, \lambda) \to 0$. *Then* λ *is an intrinsic limit of meridians.*

Proof Pick a meridian $m \subset S(\lambda)$. For each i, let $T_i : M \to M$ be the Dehn twist along the annulus A_i. Then for any sequence $n_i \in \mathbb{Z}$, the curves $T_i^{n_i}(m)$ are meridians, and if n_i grows sufficiently fast, then

$$i(T_i^{n_i}(m), \lambda)/\text{length}(T_i^{n_i}(m)) \to 0.$$

Hence, after passing to a subsequence $T_i^{n_i}(m)$ converges to a lamination λ' supported in $S(\lambda)$ with zero intersection number with λ, implying λ' and λ have the same support, so λ is an intrinsic limit of meridians. □

6 Limits of homoclinic rays

In this section we characterise the laminations onto which pairs of disjoint mutually homoclinic rays can accumulate.

Theorem 6.1 (mutually homoclinic rays) *Let M be a compact orientable hyperbolisable 3-manifold and equip $\partial_{\chi<0}M$ with an arbitrary hyperbolic metric. Let $h_\pm$ be two disjoint, mutually homoclinic simple geodesic rays on $\partial_{\chi<0}M$ that accumulate onto (possibly equal) minimal laminations $\lambda_\pm$, and where the multicurve $\partial S(\lambda_\pm)$ is incompressible in M. Then one of the following holds:*

(1) *one of λ_+ or λ_- is an intrinsic limit of meridians,*

(2) *h_+ and h_- are asymptotic on $\partial_{\chi<0}M$, and either*

 (a) *any two mutually homoclinic lifts $\tilde{h}_\pm$ to $\partial\widetilde{M}$ are asymptotic on $\partial\widetilde{M}$, or*

 (b) *$\lambda := \lambda_- = \lambda_+$ is a simple closed curve that is homotopic in M to a nontrivial power γ^n, $n > 1$ of some closed curve γ in M,*

(3) *$h_\pm$ are not asymptotic on $\partial_{\chi<0}M$, and there is an essential (possibly nontrivial) interval bundle $B \subset M$ such that $\lambda_\pm$ each fill a component of $\partial_H B$, and $\lambda_\pm$ are essentially homotopic through B, as in Section 2.9.*

The proof of Theorem 6.1 is given in Section 6.1. One can construct examples of mutually homoclinic rays in (1)–(3), as follows.

For (1), pick two meridians λ_-, λ_+ on M and let $h_\pm$ spiral onto $\lambda_\pm$. One can also produce similar examples by letting $\lambda_\pm$ be arbitrary laminations in disjoint subsurfaces $S(\lambda_\pm)$ that are spheres with at least 4 boundary components, all of which are compressible in M, and letting $h_\pm$ accumulate onto $\lambda_\pm$. In case (1), we expect it is possible that $S(\lambda_-)$ is incompressible, say, while λ_+ is an intrinsic limit of meridians. For instance, suppose C is a compression body with connected, genus-at-least-two interior boundary $\partial_- C$, and exterior boundary $\partial_+ C$. Let $f : C \to C$ be a homeomorphism such that $f|_{\partial_+ C}$ and $f|_{\partial_- C}$ are both pseudo-Anosov, with attracting laminations λ_+ and λ_-, respectively. We expect that there are rays $\ell_\pm \subset \lambda_\pm$ that are mutually homoclinic. But λ_+ is an intrinsic limit of meridians, while $S(\lambda_-)$ is incompressible.

For (2)(a), take M to be a handlebody, let λ be any simple closed curve on ∂M that is essential in M but has no nontrivial roots in $\pi_1 M$, and let $h_\pm$ spiral around λ in the same direction. For (2)(b), take $h_\pm$ to be the two ends of the homoclinic geodesic h on the left in Figure 5. In the picture, we have drawn a solid torus that is a boundary-connect-summand of M, which (say) is a handlebody. The rays $h_\pm$ both spiral onto a simple closed curve λ, the $(2,1)$-curve on the solid torus. This λ is homotopic to the square of the core curve of the solid torus. Although $h_\pm$ are asymptotic on ∂M, any lift $\tilde{h}$ in $\partial\widetilde{M}$ will have ends that are mutually homoclinic, but nonasymptotic. On the right, we have drawn the preimage $\tilde{\lambda}$ of λ, and two lifts $\tilde{h}_1, \tilde{h}_2$ of h. Note that, since $h_\pm$ are asymptotic on ∂M, one end of $\tilde{h}_1$ is asymptotic to an end of $\tilde{h}_2$.

When M is a compression body, Casson and Gordon [17, Theorem 4.1] proved that any simple closed curve $\lambda \subset \partial M$ that has a nontrivial root in $\pi_1 M$ lies on the boundary of a solid torus that is a boundary

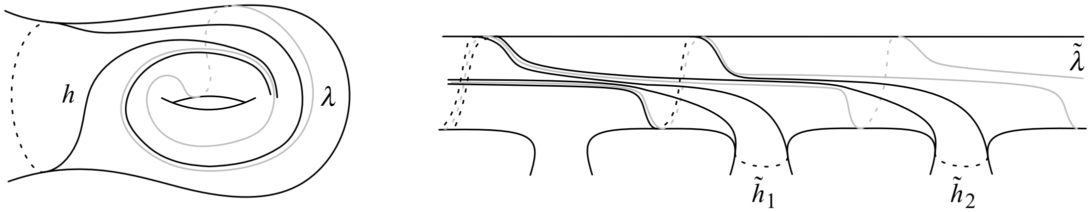

Figure 5: An example of (2)(b) in Theorem 6.1.

connect summand of M, exactly as in Figure 5. When M has incompressible boundary, such λ come from components of the characteristic submanifold of M, see Section 2.6, that are either solid tori or twisted interval bundles over nonorientable surfaces.

Examples of (3) are shown in Figure 6, with $h_\pm$ being the two ends of a homoclinic geodesic h. On the left, λ_-, λ_+ are simple closed curves that bound an embedded annulus A in M and B is a regular neighbourhood of A. The rays $h_\pm$ are mutually homoclinic since the annulus A lifts to an embedded infinite strip $\mathbb{R} \times [-1, 1] \subset \widetilde{H}$ and the $h_\pm$ are asymptotic to $\mathbb{R}_+ \times \{-1\}$ and $\mathbb{R}_+ \times \{1\}$, respectively. On the right, we write $M = Y \times [-1, 1]$ where Y is a genus-two surface with one boundary component. The laminations $\lambda_\pm$ are minimal (in the picture they are drawn as "train tracks") and fill $Z \times \{\pm 1\}$, where $Z \subset Y$ is a torus with two boundary components. Here, $B = Z \times [-1, 1]$. One can also construct similar examples of (3) where the interval bundle B is twisted.

The assumption that $\partial S(\lambda_\pm)$ is incompressible is necessary in Theorem 6.1. For instance, suppose M is a compression body with exterior boundary a genus-3 surface S, where the only meridian on S is a single separating curve γ. Let T be the component of $S \setminus \gamma$ that is a punctured genus-2 surface. Then there are distinct minimal geodesic laminations $\lambda, \lambda' \subset T$, each of which fills T, that are properly homotopic in M: just take distinct laminations that are identified when we cap off the puncture of T to get a closed genus-2 surface. Corresponding ends of corresponding leaves of λ, λ' are mutually homoclinic rays that accumulate onto λ, λ', respectively, but none of (2)–(3) hold. One could write down a version of Theorem 6.1 that omits the assumption that $\partial S(\lambda_\pm)$ is incompressible, but the conclusion would be relative to capping off $S(\lambda_\pm)$, and the statement would be more complicated.

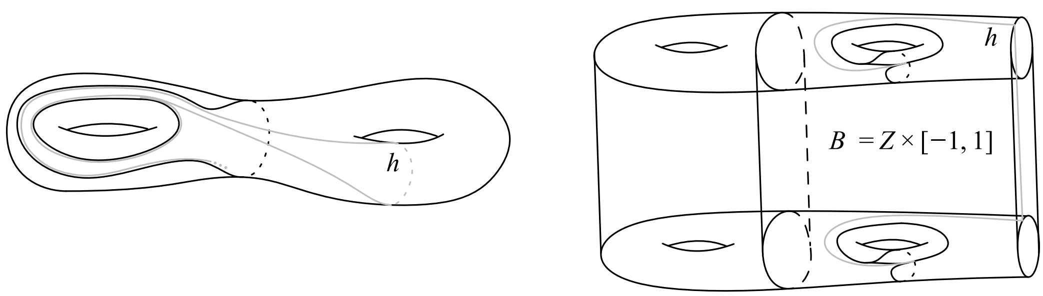

Figure 6: Examples of homoclinic geodesics (in grey) satisfying (3) in Theorem 6.1.

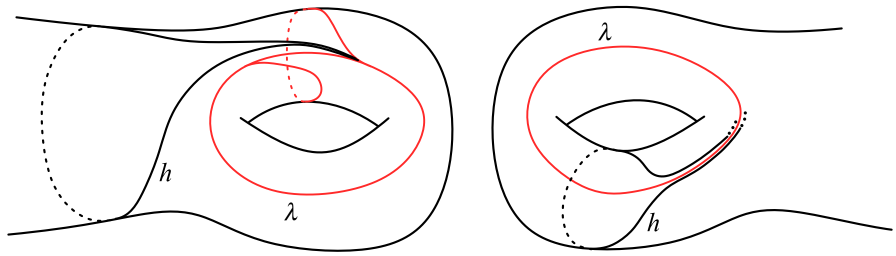

Figure 7: Homoclinic geodesics h as in cases (i) and (ii) in Corollary 6.2, respectively.

Here is a slightly more refined version of Theorem 6.1 that applies to homoclinic biinfinite geodesics on $\partial_{\chi<0}M$.

Corollary 6.2 (homoclinic biinfinite geodesics) *Suppose that M is as in Theorem 6.1, that h is a homoclinic biinfinite simple geodesic on some component $S \subset \partial_{\chi<0}M$, that $h_\pm$ are the two ends of h, that $h_\pm$ limit onto $\lambda_\pm$, and that $\partial S(\lambda_\pm)$ is incompressible in M.*

Suppose case (2) of Theorem 6.1 holds, in which case the two limiting laminations are equal, and we can set $\lambda := \lambda_\pm$. If this λ is not an intrinsic limit of meridians then either

(i) *after reparametrising h, we have that for all large s, the points $h(-s)$ and $h(s)$ are joined by a geodesic segment c with $h \cap \mathrm{int}(c) = \varnothing$, such that c, $h|_{(-\infty,-s)}$ and $h|_{[s,\infty)}$ bound an embedded geodesic triangle $\Delta \subset S$ with one ideal vertex, and $c \cup h([-s,s])$ is a meridian in M, or*

(ii) *λ is a simple closed curve on S, the two ends of h spiral around λ in the same direction, and any neighbourhood of the union $h \cup \lambda \subset S$ contains a meridian.*

See Figure 7. From this dichotomy, one can alternatively deduce that either

(i′) *any neighbourhood of $h \cup \lambda$ contains a meridian disjoint from λ, or*

(ii′) *λ is a simple closed curve and $h_\pm$ spiral around λ in the same direction but from opposite sides.*

Finally, if case (3) of Theorem 6.1 holds, we can choose the interval bundle B such that h contains a subarc $\alpha \subset h$ that is a compression arc for B.

Proof Let $\tilde{h}$ be a homoclinic lift of h on $\partial\widetilde{M}$. By Lemma 5.10, either one of $\lambda_\pm$ is an intrinsic limit of meridians in M, in which case we're in case (1) and are done, or both ends of $\tilde{h}$ are quasigeodesic in $\widetilde{M}$. Since $\tilde{h}$ is homoclinic, it follows that its two ends are mutually homoclinic, so we're in the setting of Theorem 6.1 and one of (2)–(3) holds.

Assume we're in case (2) of Theorem 6.1, and set $\lambda := \lambda_\pm$. Assume first that $\lambda := \lambda_\pm$ is a simple closed curve. Since the two ends of h are asymptotic, they spiral around λ in the same direction. Let U be a neighbourhood of $h \cup \lambda$ on $\partial_{\chi<0}M$. Then h is a homoclinic geodesic contained in U, so Fact 5.3 implies that U is compressible as desired in (ii).

Now suppose that λ is *not* a simple closed curve, in which case we're in case (2)(a) of Theorem 6.1. We show (i) holds. Let's start by constructing the desired geodesic triangle. Parametrise h, pick a universal covering map $\mathbb{H}^2 \to S$, lift h to a parametrised geodesic $\hat{h}$ in $\mathbb{H}^2$, and let

$$\xi = \lim_{t\to+\infty} \hat{h}(t) \in \partial_\infty \mathbb{H}^2.$$

Note that since λ is not a simple closed curve, ξ is not a fixed point of any deck transformation of $\mathbb{H}^2 \to S$. The two ends of h are asymptotic on S, so there is then a unique deck transformation $\gamma : \mathbb{H}^2 \to \mathbb{H}^2$ such that

$$\xi = \lim_{t\to-\infty} \gamma \circ \hat{h}(t).$$

It follows that if we use a particular arc-length parametrisation of h, we may assume that for each $t \in \mathbb{R}$, the points $\hat{h}(t), \gamma \circ \hat{h}(-t)$ lie on a common horocycle tangent to ξ. Fix a parameter $t = s$ large enough such that the geodesic segment $\hat{c}$ joining $\hat{h}(s)$ and $\gamma \circ \hat{h}(-s)$ is shorter than the injectivity radius of S, and therefore projects to a simple geodesic segment c in S.

Let $\hat{\Delta} \subset \mathbb{H}^2$ be the triangle bounded by $\hat{c}$ and the two rays $\hat{h}([s,\infty))$ and $\gamma \circ \hat{h}((-\infty, s])$. Let $g : \mathbb{H}^2 \to \mathbb{H}^2$ be a deck transformation. We claim that $g \circ \hat{h}(\mathbb{R}) \cap \mathrm{int}(\hat{\Delta}) = \varnothing$. If not, then since $\hat{\Delta}$ has geodesic sides, two of which are disjoint from $g \circ \hat{h}$, it follows that one of the two endpoints of $g \circ \hat{h}$ is ξ. If it's the positive endpoint, then g fixes ξ, and the axis of g projects to a (simple) closed curve on S, around which the two ends of h spiral, contradicting that λ isn't a simple closed curve. If the negative endpoint of $g \circ \hat{h}$ is ξ, then $g \circ \gamma^{-1}$ fixes ξ and we get a similar contradiction.

Next, we claim that we have $g(\mathrm{int}(\hat{\Delta})) \cap \mathrm{int}(\hat{\Delta}) = \varnothing$ as long as $g \neq \mathrm{id}$. Suppose that for some $g \neq \mathrm{id}$ the intersection is nonempty. Then $g(\xi) \neq \xi$, since otherwise we have a contradiction as in the previous paragraph. The previous paragraph implies that the sides of the triangles $g(\hat{\Delta}), \hat{\Delta}$ that are lifts of rays of h do not intersect the interior of the other triangle. So, the only way the interiors of $g(\hat{\Delta}), \hat{\Delta}$ can intersect is if $\hat{c}$ and $g(\hat{c})$ intersect. However, this does not happen since we chose s large enough so that $\hat{c}$ projects to a simple geodesic segment in S.

The previous two paragraphs imply that $\hat{\Delta}$ projects to an embedded geodesic triangle Δ in S whose interior is disjoint from h, as desired in (i). By construction, c and $h([-s,s])$ are simple geodesic segments, and since $g \circ \hat{h}(\mathbb{R}) \cap \mathrm{int}(\hat{\Delta}) = \varnothing$ for any $g \neq \mathrm{id}$, they are disjoint. It follows that $c \cup h([-s,s])$ is an essential simple closed curve on S.

We need to show that $c \cup h([-s,s])$ is nullhomotopic in M. For this, remember that since λ is not a simple closed curve, it cannot be that two distinct lifts of $h([s,\infty))$ to $\mathbb{H}^2$ are asymptotic, for if so they would differ by a deck transformation fixing the endpoint $\xi \in \partial\mathbb{H}^2$. Therefore, no two distinct lifts of $h([s,\infty))$ to $\partial\widetilde{M}$ can be asymptotic, and similarly for $h((-\infty,s])$. But the two ends of any lift of $\tilde{h} \subset \partial\widetilde{M}$ of h are mutually homoclinic, and hence asymptotic by the assertion in case (2) of Theorem 6.1. So, the projection $\hat{\Delta} \to \Delta$ factors through a geodesic triangle $\tilde{\Delta} \subset \partial\widetilde{M}$ bounded by $\tilde{h}([s,\infty))$, $\tilde{h}((-\infty,-s])$ and a geodesic segment $\tilde{c}$. The curve $c \cup h([-s,s])$ is the projection of the closed curve $\tilde{c} \cup \tilde{h}([-s,s]) \subset \widetilde{M}$, and so is nullhomotopic in M.

For the (i′) versus (ii′) dichotomy, note that if we're in case (i) then by taking s large, we can ensure that $c \cup h([-s,s])$ is inside a given neighbourhood of $h \cup \lambda$. Moreover, as h is disjoint from the interior of the geodesic triangle Δ mentioned in (i), so is λ, and hence after increasing s slightly we can assume $c \cup h([-s,s])$ is disjoint from λ. If we're in case (ii), and $h_\pm$ spiral onto λ *from the same side*, then λ is a peripheral curve in the specified regular neighbourhood of $h \cup \lambda$, so the meridian given in that neighbourhood can be taken disjoint from λ. Otherwise, we're in case (ii′).

Now assume we are in case (3). Let $S_\pm$ be the components of $\partial_H B$ containing $\lambda_\pm$. We may assume that h is in minimal position with respect to $\partial S_\pm$. Since h is simple and the ends of h limit onto minimal laminations that fill $S_\pm$, we have that h intersects $\partial S_- \cup \partial S_+$ at most twice. Furthermore, in the case that $S_- = S_+$, the homoclinic geodesic h cannot be contained entirely in the incompressible surface $S_\pm$, by Fact 5.3. So, h is the concatenation of two rays in $S_\pm$ and an arc α such that $\mathrm{int}(\alpha)$ lies outside $S_\pm$.

Let $X \subset S$ be the union of $S_\pm$ and a regular neighbourhood of α. Since h is homoclinic, there is a meridian on X by Fact 5.3. If the two endpoints of α lie on different boundary components of $\partial_H B$, then α is a compression arc for B by Fact 2.12. So, we may assume that the two endpoints of α lie on the same boundary component c of $\partial_H B$. Fact 2.12 then says that α is homotopic rel endpoints in M to an arc of c. So, if we make a new path $h' \subset \partial_H B$ from h by replacing α with that arc of c, then h' is still homoclinic, so it cannot be boundedly homotopic to a geodesic in $\partial_H B$ by Fact 5.3, which implies that its ends $h_\pm$ are asymptotic, a contradiction to the assumption in (3). □

If h is a parametrised biinfinite geodesic, let's denote by $h_\pm$ the associated *positive and negative rays*, namely $h_+ := h|_{[0,\infty)}$ and $h_- := h|_{(-\infty,0]}$.

Definition 6.3 (mutually bihomoclinic) We say that two biinfinite geodesics h, h' on some component $S \subset \partial_{\chi<0} M$ are *mutually bihomoclinic* if they have distinct lifts $\tilde{h}$, $\tilde{h}'$ on $\partial \widetilde{M}$ such that the associated rays $\tilde{h}_+$, $\tilde{h}'_+$ are mutually homoclinic, as are the rays $\tilde{h}_-$, $\tilde{h}'_-$.

Above, we allow $h_0 = h_1$, but we require the two lifts to be distinct. Here is a variant of Corollary 6.2 for pairs of mutually bihomoclinic rays that we will use in a sequel to this paper. For simplicity, we'll state only the analogue of (i′) versus (ii′) in Corollary 6.2.

Corollary 6.4 *Suppose that M is as in Theorem 6.1, that h, h' are simple biinfinite geodesics on a component $S \subset \partial_{\chi<0} M$ that are either disjoint or equal, and that h, h' are mutually bihomoclinic.*

Suppose case (2) of Theorem 6.1 holds both for the positive rays h_+, h'_+, and for the negative rays h_-, h'_-. Then either

- (i′) *in any neighbourhood of $h \cup h' \cup \lambda_- \cup \lambda_+$, there is a meridian m disjoint from λ_- and λ_+, or*
- (ii′) *for either $+$ or $-$, say $+$, we have that λ_+ is a simple closed curve and h_+, h'_+ spiral around λ_+ in the same direction but from different sides.*

The proof is similar to the proof of Corollary 6.4.

6.1 Proof of Theorem 6.1

The proof proceeds in a few cases. As in Definition 5.2, we identify M with the convex core $\mathrm{CC}(N)$ of a geometrically finite hyperbolic 3-manifold $N = \mathbb{H}^3/\Gamma$, and we identify the universal cover $\widetilde{M}$ with the preimage of $\mathrm{CC}(N)$ in $\mathbb{H}^3$, which is the convex hull of the limit set of Γ. Note that the closure of $\widetilde{M}$ in $\mathbb{H}^3 \cup \partial\mathbb{H}^3$ is a ball.

There are four cases to consider:

(A) Both $\lambda_\pm$ are simple closed curves. We show that either (2) or (3) holds.

(B) Both $\lambda_\pm$ are distinct, in which case the surfaces $S(\lambda_\pm)$ are disjoint, but one of these surfaces is compressible, say $S(\lambda_+)$. We show (1).

(C) At least one of $\lambda_\pm$ is not a simple closed curve, and both $S(\lambda_\pm)$ are incompressible. We show (2)(a) or (3) holds.

(D) $\lambda_- = \lambda_+$, which is not a simple closed curve, and $S(\lambda_\pm)$ is compressible. We show either (1) or (2)(a) holds.

Cases (A) and (B) above are the easiest. Our proof in case (C) involves a hyperbolic geometric interpretation of the characteristic submanifold of a pair, as discussed in Section 3 of [40] and by Walsh [60]; our argument is a bit more complicated than theirs, since we have to deal with accidental parabolics. In case (D), our argument adapts and fills some gaps in a surgery argument of Kleineidam and Souto [38] and Lecuire [42, Appendix C].

Proof of (A) Assume that both $\lambda_\pm$ are simple closed curves. Since $S(\lambda_\pm)$ is assumed to be incompressible, both $\lambda_\pm$ are incompressible in M. If $\lambda_- \neq \lambda_+$, they have mutually homoclinic lifts and hence are homotopic. Then we are in case (3) by the annulus theorem. So we may assume the two curves are the same, and write $\lambda := \lambda_\pm$.

We claim that $h_\pm$ spiral around λ in the same direction, so that they are asymptotic on ∂M. Suppose not, and pick mutually homoclinic lifts $\tilde{h}_\pm$ in $\widetilde{M}$. Then $\tilde{h}_-$ and $\tilde{h}_+$ are asymptotic to lifts $\tilde{\lambda}$ and $\alpha(\tilde{\lambda})$ of λ, where $\alpha \in \Gamma$ is a deck transformation. Any lift of λ is a quasigeodesic in $\widetilde{M}$, and hence in $\mathbb{H}^3$, so $\tilde{h}_\pm$ are quasigeodesic rays, and therefore have well-defined endpoints in $\partial\mathbb{H}^3$, which must be the same since the two rays are mutually homoclinic. Since $h_\pm$ spiral around λ in opposite directions, this means that $\alpha \in \Gamma$ takes one endpoint of $\tilde{\lambda}$ in $\partial\mathbb{H}^3$ to the other endpoint of $\tilde{\lambda}$. Since $\tilde{\lambda}$ is stabilised by a loxodromic isometry in Γ, and Γ is torsion-free and discrete, this is impossible.

Suppose we are not in case (2)(a), so there are mutually homoclinic lifts $\tilde{h}_\pm$ that are not asymptotic on $\partial\widetilde{M}$. As in the previous paragraph, we may assume that $\tilde{h}_-$ and $\tilde{h}_+$ are asymptotic to lifts $\tilde{\lambda}$ and $\alpha(\tilde{\lambda})$ for some deck transformation $\alpha \in \Gamma$. Since $\tilde{h}_\pm$ are not asymptotic, $\tilde{\lambda} \neq \alpha(\tilde{\lambda})$. As before, α fixes the common endpoint of $\tilde{h}_\pm$ in $\partial\mathbb{H}^3$, which is a fixed point of the cyclic group $\langle\beta\rangle \subset \Gamma$ of loxodromic isometries fixing $\tilde{\lambda}$. As Γ is discrete and torsion-free, and $\alpha \notin \langle\beta\rangle$, we have that α is a root of β or β^{-1} in Γ, and (2)(b) follows. □

Proof of (B) Suppose that $\lambda_\pm$ are distinct, in which case the surfaces $S(\lambda_\pm)$ are disjoint, but that one of these surfaces is compressible, say $S(\lambda_+)$. We claim that λ_+ is an intrinsic limit of meridians, in which case (1) holds and we are done. If not, take a meridian $m \subset S(\lambda_+)$ and apply Lemma 5.9. We obtain a new meridian $m' \subset S(\lambda_+)$ such that λ_+ has no m-waves. Since λ_+ fills $S(\lambda_+)$ and the boundary components $\partial S(\lambda_\pm)$ are incompressible, it follows that λ_+ is in tight position with respect to m. So after possibly restricting the domains, h_+ is in tight position with respect to m', while h_- never intersects m'. This contradicts the fact that $h_\pm$ are mutually homoclinic, since if $\tilde{h}_\pm$ are homoclinic lifts in $\widetilde{M}$, for large t the point $\tilde{h}_+(t)$ is separated from the image of $\tilde{h}_-$ by arbitrarily many lifts of m'. □

Proof of (C) Assume that at least one of $\lambda_\pm$ is not a simple closed curve, and that $S_\pm := S(\lambda_\pm)$ are incompressible. Note that $S_\pm$ are equal or have disjoint interiors. We want to prove that we're in case (2) or (3). Lift $h_\pm$ to mutually homoclinic rays $\tilde{h}_\pm \subset \partial\widetilde{M}$. Fact 5.3 implies that each inclusion $\widetilde{S}_\pm \hookrightarrow \widetilde{M}$ is a quasi-isometric embedding, so if $\widetilde{S}_- = \widetilde{S}_+$, then the two mutually homoclinic rays $\tilde{h}_\pm$ are actually asymptotic on $\partial\widetilde{M}$. If this is true for all lifts $\tilde{h}_\pm$, we are in case (2)(a) and are done. So, we can assume below that $\widetilde{S}_- \neq \widetilde{S}_+$. Note that it may still be that $\lambda_- = \lambda_+$ and $S_- = S_+$.

Let $\Gamma_\pm \subset \Gamma$ be the stabiliser of $\widetilde{S}_\pm$ and let $\Lambda_\pm \subset \partial\mathbb{H}^3$ be the limit set of $\Gamma_\pm$. Since $\Gamma_\pm$ acts cocompactly on $\widetilde{S}_\pm$, the inclusion $\widetilde{S}_\pm \hookrightarrow \mathbb{H}^3$ extends continuously to a map $\partial_\infty \widetilde{S}_\pm \to \Lambda_\pm \subset \partial\mathbb{H}^3$, by the main result of [53]. In particular, $\tilde{h}_\pm$ have well-defined endpoints in $\partial\mathbb{H}^3$, and since they're mutually homoclinic, they have the same endpoint $\xi \in \Lambda_- \cap \Lambda_+ \subset \partial\mathbb{H}^3$.

We now apply Theorem 4.1. Since $\xi \in \Lambda_- \cap \Lambda_+$, using the notation of Theorem 4.1, the rays $\tilde{h}_\pm$ are either eventually contained in the convex hulls $\widetilde{C}_\pm \subset \widetilde{S}_\pm$, or are asymptotic onto their boundaries. But $\widetilde{C}_\pm$ project to (possibly degenerate) generalised subsurfaces $C_\pm$ with geodesic boundary in $S_\pm$, and the rays $h_\pm$ limit onto filling laminations in $S_\pm$, so it follows that actually $C_\pm = S_\pm$, and that there is a homotopy from S_- to S_+ in M that is the projection of a homotopy from $\widetilde{S}_-$ to $\widetilde{S}_+$. Since one of $\lambda_\pm$ is not a simple closed curve, this means they are *both* not simple closed curves and the (a priori degenerate) subsurfaces with geodesic boundary $S_\pm$ are not simple closed geodesics.

Let $S'_\pm \subset \operatorname{int}(S_\pm)$ be obtained by deleting small collar neighbourhoods of $\partial S_\pm$, so that $S'_\pm$ are both actually embedded, still contain $\lambda_\pm$, and are either disjoint or equal. Since $S'_\pm$ are incompressible and homotopic in M, Theorem 2.8 implies that they bound an essential interval bundle $B \subset M$. Moreover, the fact that the homotopy from S_- to S_+ is the projection of a homotopy from $\widetilde{S}_-$ to $\widetilde{S}_+$ means that we can assume that there is a component $\widetilde{B} \subset \widetilde{M}$ of the preimage of B that intersects $\partial\widetilde{M}$ in $\widetilde{S}'_\pm$. Note that $\widetilde{B}$ is invariant under $\Delta = \Gamma_- \cap \Gamma_+$, since any element of Δ preserves $\widetilde{S}'_\pm$, and hence $\widetilde{B}$.

We claim that $\lambda_\pm$ are essentially homotopic through B. By Fact 2.16, it suffices to show that if σ is the canonical involution of B, as described in Proposition 4.5, then $\sigma(\lambda_\pm)$ is isotopic to $\lambda_\mp$ on $S'_\mp$. Well, σ lifts to a Δ-equivariant involution $\tilde{\sigma}$ of $\widetilde{B}$ that exchanges $\widetilde{S}'_-$ and $\widetilde{S}'_+$, where here $\Delta = \Gamma_- \cap \Gamma_+$. By equivariance, $\tilde{\sigma}$ extends continuously to the identity on Λ_Δ, so in particular its extension fixes ξ, and hence $\tilde{\sigma}(h_\pm)$ is properly homotopic to $h_\mp$ on $S_\mp$, which implies $\tilde{\sigma}(\lambda_\pm)$ is isotopic to $\lambda_\mp$ as desired.

If $h_\pm$ are not asymptotic on ∂M, then we are in case (3) and are done. So, assume $h_\pm$ are asymptotic. Then there is some $\gamma \in \Gamma$ such that $\gamma(\tilde{h}_-) \subset \widetilde{S}'_+$ and is asymptotic to $\tilde{h}_+$. This γ fixes the endpoint $\xi \in \partial\mathbb{H}^3$ of $\tilde{h}_\pm$. Moreover, $\gamma(\widetilde{B})$ is a component of the preimage of B that contains $\widetilde{S}'_+$, and therefore equals $\widetilde{B}$. So, γ exchanges $\widetilde{S}'_\pm$, and therefore $\gamma^2 \in \Delta$. But then $\tilde{h}_\pm$ are asymptotic to the axes of $\gamma^2 \curvearrowright \widetilde{S}_\pm$, implying that $h_\pm$ accumulate onto simple closed curves in ∂M, contradicting our assumption in (C). □

Proof of (D) Assume that $\lambda_- = \lambda_+$, write $\lambda = \lambda_\pm$ for brevity, assume that λ is not a simple closed curve, and that $S(\lambda)$ is compressible. We want to prove that either λ is an intrinsic limit of meridians, or $h_\pm$ are asymptotic, as are any pair of mutually homoclinic lifts $\tilde{h}_\pm$.

If λ is an intrinsic limit of meridians, we are done, so since $S(\lambda)$ is compressible with incompressible boundary, by Lemma 5.9 we can choose a meridian $m \subset S(\lambda)$ with respect to which λ is in tight position. Let $\widetilde{m}$ be its full preimage in $\partial\widetilde{M}$, and let $\tilde{h}_\pm$ be *any* pair of mutually homoclinic lifts in $\partial\widetilde{M}$. Truncating if necessary, we can assume that $h_\pm$ are in tight position with respect to m, and hence the lifts $\tilde{h}_\pm$ are quasigeodesic rays in $\mathbb{H}^3$, by Fact 5.6. Since they are mutually homoclinic, $\tilde{h}_-$ and $\tilde{h}_+$ converge to the same point $\xi \in \partial_\infty\mathbb{H}^3$, and tightness further implies that after restricting to appropriate subrays, $\tilde{h}_-$ and $\tilde{h}_+$ intersect exactly the same components of $\widetilde{m}$, in the same order. Reparametrising, we have

$$\tilde{h}_\pm : [0,\infty) \to \partial\widetilde{M}, \qquad \tilde{h}_+(i), \tilde{h}_-(i) \in \widetilde{m}_i \quad \text{for all } i \in \mathbb{N},$$

where each $\widetilde{m}_i$ is a component of $\widetilde{m}$, and where $\tilde{h}_\pm(t) \notin \widetilde{m}$ when $t \notin \mathbb{N}$. Let

$$d_i := d_{\widetilde{m}}(\tilde{h}_+(i), \tilde{h}_-(i))$$

be the distance along $\widetilde{m}$ between $\tilde{h}_+(i)$ and $\tilde{h}_-(i)$. □

Claim 6.5 *There is some uniform* $\epsilon > 0$, *independent of the particular chosen lifts* $\tilde{h}_\pm$, *such that either*

(1) $\tilde{h}_\pm$ *are asymptotic on* $\partial\widetilde{M}$, *and hence* $h_\pm$ *are asymptotic on* ∂M, *or*

(2) $\liminf_i d_i \geq \epsilon$.

Proof Let's assume that $\tilde{h}_+$ and $\tilde{h}_-$ are not asymptotic on $\partial\widetilde{M}$, and write $d = \liminf_i d_i$. Fix some transverse measure on λ. If d is small, we will construct meridians $\gamma \subset S(\lambda)$ with very small intersection number with λ. Since λ is not an intrinsic limit of meridians, there is some fixed lower bound for such intersection numbers, which will give a contradiction for small d.

Suppose d is small and pick $0 \ll i < j$ such that

$$d_i, d_j < 2d,$$

let b_i be the (unique) shortest path on $\widetilde{m}$ from $\tilde{h}_-(i)$ to $\tilde{h}_+(i)$, and define b_j similarly. Let $\tilde{\gamma}_{ij}$ be the loop on $\partial\widetilde{M}$ obtained by concatenating the four segments $\tilde{h}_+([i,j])$, $\tilde{h}_-([i,j])$, b_i and b_j in the obvious way.

We first claim that after fixing i, it is possible to choose j such that $\tilde{\gamma}_{ij}$ is homotopically essential on $\partial\widetilde{M}$. Assume not, let $\widetilde{S} \subset \partial\widetilde{M}$ be the component containing $\tilde{h}_\pm$, fix a universal covering map

$$\mathbb{H}^2 \to \widetilde{S},$$

and lift the rays $\tilde{h}_\pm|_{[i,\infty)}$ to rays

$$\mathfrak{h}_\pm : [i,\infty) \to \mathbb{H}^2$$

in such a way that b_i lifts to a segment connecting $\mathfrak{h}_-(i)$ to $\mathfrak{h}_+(i)$. Now, there are infinitely many $j > i$ with $d_j < 2d$. For each such j, we know that $\tilde{\gamma}_{ij}$ is homotopically inessential on $\partial\widetilde{M}$, so the points $\mathfrak{h}_-(j)$ to $\mathfrak{h}_+(j)$ are at most $2d$ away from each other in $\mathbb{H}^2$. This gives a sequences of points exiting the rays $\mathfrak{h}_\pm$

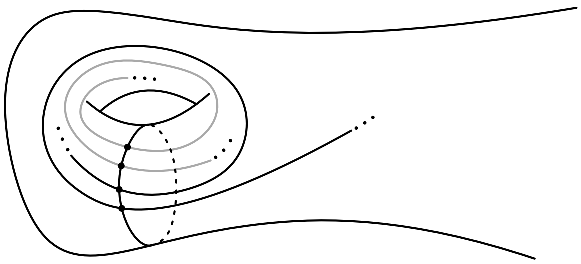

Figure 8: Two rays spiralling onto a simple closed curve (which is not allowed below), where the points in Claim 6.6 are linked.

that are always at most $2d$ apart, so $\mathfrak{h}_-$ is asymptotic to $\mathfrak{h}_+$. Hence, $\tilde{h}_-$ is asymptotic to $\tilde{h}_+$, contrary to our assumption.

We now fix large i, j such that $d_i, d_j < 2d$ and $\tilde{\gamma} := \tilde{\gamma}_{ij}$ is homotopically essential on $\partial\widetilde{M}$. Then $\tilde{\gamma}$ projects to a homotopially essential loop $\gamma \subset \partial M$ that is homotopically trivial in M ($\tilde{\gamma}$ is homotopically trivial in the simply connected space $\widetilde{M}$). Note that if i, j are chosen large enough and d is small, then $\gamma \subset S(\lambda)$. Furthermore, since the segments b_i, b_j are the only parts of $\tilde{\gamma}$ that intersect λ, and these segments have hyperbolic length less than $2d$, the intersection number $i(\gamma, \lambda)$ is small when d is small. (Recall that λ is a minimal lamination that is not a simple closed curve, so no leaves have positive weight, and hence hyperbolic length can be compared to intersection number.) But λ is not an intrinsic limit of meridians, so Proposition 5.12(4) says that there is some positive lower bound for the intersection numbers of λ with essential curves that are nullhomotopic in M. Hence, we get a contradiction if d is small. □

Suppose we have two pairs $\{a, b\}$ and $\{c, d\}$ of points in m, all four of which are distinct. We say the two pairs are *unlinked* in m if in the induced cyclic ordering on $\{a, b, c, d\} \subset m$, a is adjacent to b and c is adjacent to d, and we say that the two pairs are *linked* otherwise.

Claim 6.6 *If $i, j \in \mathbb{N}$, $i < j$, then the pairs $\{h_+(i), h_-(i)\}$ and $\{h_+(j), h_-(j)\}$ are unlinked in m.*

For an example where the pairs are *linked*, see Figure 8. The proof below works in general whenever $h_\pm$ are simple geodesic rays on ∂M in tight position with respect to m, where neither h_+ nor h_- spirals onto a simple closed curve.

Proof The essential observation used in the proof is that the closure

$$\mathrm{cl}(\partial\widetilde{M}) \subset \mathbb{H}^3 \cup \partial\mathbb{H}^3$$

is homeomorphic to a sphere: indeed, the closure of $\widetilde{M}$ in $\mathbb{H}^3 \cup \partial\mathbb{H}^3$ is a ball, since $\widetilde{M} \subset \mathbb{H}^3$ is convex with nonempty interior, and the closure of the boundary is the boundary of the closure. We obtain the unlinking property above by exploiting separation properties of arcs and curves on $\mathrm{cl}(\partial\widetilde{M})$.

Since $h_\pm$ are in tight position with respect to m, both lifts $\tilde{h}_\pm$ cross $\widetilde{m}_i$ exactly once. Since $\tilde{h}_\pm$ limit to the same point in $\partial\mathbb{H}^3$, they must then cross $\widetilde{m}_i$ in the same direction. In other words, the tangent vectors $h_+(i)'$, $h_-(i)'$ point to the same side of m. The same statement holds for j. This allows us to break into the following two cases:

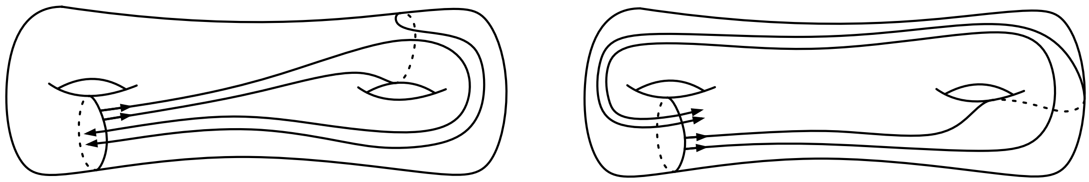

Figure 9: The cases (a) and (b) in the proof of Claim 6.6.

(a) the arcs $h_\pm|_{[i,j]}$ start and end on the same side of m, i.e., the vectors $h_\pm(i)'$ point to the opposite side of m as the vectors $h_\pm(j)'$, or

(b) the arcs $h_\pm|_{[i,j]}$ start and end on different sides of m, i.e., all four velocity vectors $h_\pm(i)', h_\pm(j)'$ point to the same side of m;

see Figure 9.

First, assume we're in case (a). Let

$$\alpha_\pm := \tilde{h}_\pm|_{[i,j]},$$

which we regard as oriented arcs in $\partial\widetilde{M}$ starting on $\widetilde{m}_i$ and ending on $\widetilde{m}_j$. Let $\gamma : \widetilde{M} \to \widetilde{M}$ be the deck transformation taking $\widetilde{m}_j$ to $\widetilde{m}_i$. Then the arcs

$$\beta_\pm := \gamma \circ \tilde{h}_\pm|_{[i,j]}$$

start on $\gamma(\widetilde{m}_i)$ and end on $\gamma(\widetilde{m}_j) = \widetilde{m}_i$, and since we're in case (a) they end on the *same side* of $\widetilde{m}_i$ as the arcs $\alpha_\pm$ start. Note that $\gamma(\widetilde{m}_i)$ is not $\widetilde{m}_i$ or $\widetilde{m}_j$. Indeed, if $\gamma(\widetilde{m}_i) = \widetilde{m}_i$ then we'd have $\gamma = \mathrm{id}$, contradicting that $\gamma(\widetilde{m}_j) = \widetilde{m}_i$. And if $\gamma(\widetilde{m}_i) = \widetilde{m}_j$, then γ^2 would leave $\widetilde{m}_i$ invariant, implying that $\gamma^2 = \mathrm{id}$, which is impossible since $\pi_1 M$ has no torsion.

We claim *the interiors of the arcs $\beta_\pm$ do not intersect $\widetilde{m}_i$ or $\widetilde{m}_j$, and the arcs $\alpha_\pm$ do not intersect $\gamma(\widetilde{m}_i)$.* Indeed, the interiors of $\beta_\pm$ don't intersect $\widetilde{m}_i$ because the arcs $\beta_\pm$ end on $\widetilde{m}_i$ and intersect each component of $\widetilde{m}$ at most once, by tight position. The interiors of $\beta_\pm$ don't intersect $\widetilde{m}_j$ because any arc from $\widetilde{m}_i$ to $\widetilde{m}_j$ intersect at least $j-i+1$ components of $\widetilde{m}$ (counting $\widetilde{m}_i$ and $\widetilde{m}_j$), while any proper subarc of $\beta_\pm$ intersects at most $j-i$ components of $\widetilde{m}$. Here, for the $j-i+1$ bound we are using tight position of $h_\pm$, the definitions of $\widetilde{m}_i, \widetilde{m}_j$, and the fact that each component of $\widetilde{m}$ separates $\partial\widetilde{M}$. The fact that the arcs $\alpha_\pm$ don't intersect $\gamma(\widetilde{m}_i)$ is similar: any arc from m_i to $\gamma(\widetilde{m}_i)$ must pass through at least $j-i+1$ components of $\widetilde{m}$, while any proper subarc of $\alpha_\pm$ intersects at most $j-i$ components, and $\alpha_\pm$ do not end on $\gamma(\widetilde{m}_i) \neq \widetilde{m}_j$.

Let $A \subset \mathrm{cl}(\partial\widetilde{M}) \cong S^2$ be the annulus that is the closure of the component of $\mathrm{cl}(\partial\widetilde{M}) \setminus (\widetilde{m}_i \cup \widetilde{m}_j)$ that contains the side of $\widetilde{m}_i$ on which the arcs $\alpha_\pm$ start and the arcs $\beta_\pm$ end. Then $\alpha_\pm$ are two disjoint arcs in A that join $\widetilde{m}_i$ to $\widetilde{m}_j$, and therefore $\alpha_\pm$ separate A into two rectangles. The component $\gamma(\widetilde{m}_i)$ on which the arcs $\beta_\pm$ start is contained in the interior of one of these two rectangles. Therefore, the two arcs $\beta_\pm$ must lie in the same component of $A \setminus (\alpha_+ \cup \alpha_-)$. Looking at endpoints, this means the pairs $\{\tilde{h}_+(i), \tilde{h}_-(i)\}$ and $\{\gamma \circ \tilde{h}_+(j), \gamma \circ \tilde{h}_-(j)\}$ are unlinked in $\widetilde{m}_i$, and the claim follows.

Now assume that we're in case (b). The curve $\widetilde{m}_i$ separates $\partial \widetilde{M}$, and we let $X \subset \partial \widetilde{M}$ be the closure of the component of $\partial \widetilde{M} \setminus \widetilde{m}_i$ into which the velocity vectors $\tilde{h}'_\pm(i)$ and $(\gamma \circ \tilde{h}_\pm)'(j)$ all point. The closure

$$\mathrm{cl}(X) \subset \mathbb{H}^3 \cup \partial \mathbb{H}^3$$

is homeomorphic to a disk, since $\mathrm{cl}(\partial \widetilde{M})$ is a sphere. As before, we let $\gamma : \widetilde{M} \to \widetilde{M}$ be the deck transformation taking $\widetilde{m}_j$ to $\widetilde{m}_i$. Then the rays

$$\alpha_\pm := \tilde{h}_\pm([i,\infty)), \quad \beta_\pm := \gamma \circ \tilde{h}_\pm([j,\infty))$$

are all contained in X. Note that $\alpha_\pm$ both limit to a point $\xi \in \partial \mathbb{H}^3$, while $\beta_\pm$ limit to $\gamma(\xi) \in \partial \mathbb{H}^3$.

The union $\alpha_- \cup \alpha_+$ compactifies to an arc in $\mathrm{cl}(X)$, since the two rays limit to the same point in $\mathbb{H}^3$. The same is true for $\beta_- \cup \beta_+$. *Hoping for a contradiction, suppose that the points in the statement of the claim are linked.* Then the pairs of endpoints of $\alpha_- \cup \alpha_+$ and $\beta_- \cup \beta_+$ are also linked on $\widetilde{m}_i = \partial \mathrm{cl}(X)$. We now have two arcs on the disk $\mathrm{cl}(X)$ with linked endpoints on $\partial \mathrm{cl}(X)$, so the two arcs must intersect. As $\alpha_\pm$, $\beta_\pm$ are all disjoint, the only intersection can be on $\partial \mathbb{H}^3$, so their endpoints at infinity must all agree, i.e., $\gamma(\xi) = \xi$.

Since $\gamma(\xi) = \xi$, all the rays $\gamma^k \circ \tilde{h}_+$ limit to ξ, where $k \in \mathbb{Z}$. Hence, all these (quasigeodesic) rays are pairwise mutually homoclinic, and for each pair k, l, the rays $\gamma^k \circ \tilde{h}_+$ and $\gamma^l \circ \tilde{h}_+$ eventually intersect the same components of $\widetilde{m}$, in the same order, although their initial behaviour may be different. In analogy with the setup of Claim 6.5, let $d_{k,l}$ be the liminf of the distances from $\gamma^k \circ \tilde{h}_+$ to $\gamma^l \circ \tilde{h}_+$ along the components of $\widetilde{m}$ that they both intersect.

We claim that there are k, l such that $d_{k,l} < \epsilon$, where ϵ is the constant from Claim 6.5. Indeed, for $N > \mathrm{length}(m)/\epsilon$, it is impossible to pack N points at least ϵ apart in any component of $\widetilde{m}$. So if we let k range over a set $F \subset \mathbb{Z}$ of size N, whenever a component of $\widetilde{m}$ intersects all $\gamma^k \circ \tilde{h}_+$, $k \in F$, two such intersections must be within ϵ of each other. There are infinitely many such components of $\widetilde{m}$ and F is finite, so we can pick $k, l \in S$ such that $\gamma^k \circ \tilde{h}_+$ and $\gamma^l \circ \tilde{h}_+$ are within ϵ on infinitely many such components.

Finally, $\gamma^k \circ \tilde{h}_+$ and $\gamma^l \circ \tilde{h}_+$ are mutually homoclinic lifts of $\tilde{h}_+$, and $d_{k,l} < \epsilon$, so the exact same argument as in Claim 6.5 shows that $\gamma^k \circ \tilde{h}_+$ and $\gamma^l \circ \tilde{h}_+$ are asymptotic on $\partial \widetilde{M}$. It follows that h_+ spirals onto a (simple) closed curve in ∂M in the homotopy class of (a primitive root of) γ^{l-k}. (Indeed, γ^{l-k} lifts to a deck transformation of the universal cover $\mathbb{H}^2 \to \partial M$, and the axis of this deck transformation is asymptotic to suitably chosen lifts of both $\gamma^k \circ \tilde{h}_+$ and $\gamma^l \circ \tilde{h}_+$.) This is a contradiction, though, since h_+ limits onto λ, which is not a simple closed curve. □

Assume now that our mutually homoclinic rays $\tilde{h}_\pm$ are not asymptotic on $\partial \widetilde{M}$, as otherwise we're in case (2) of the theorem and are done. By Claim 6.5, there is some $\epsilon > 0$ such that $d_{\widetilde{m}_i}(\tilde{h}_+(i), \tilde{h}_-(i)) \geq \epsilon$ for all i. We will show that λ is an intrinsic limit of annuli, in the sense of Lemma 5.13, which says that then λ is an intrinsic limit of meridians.

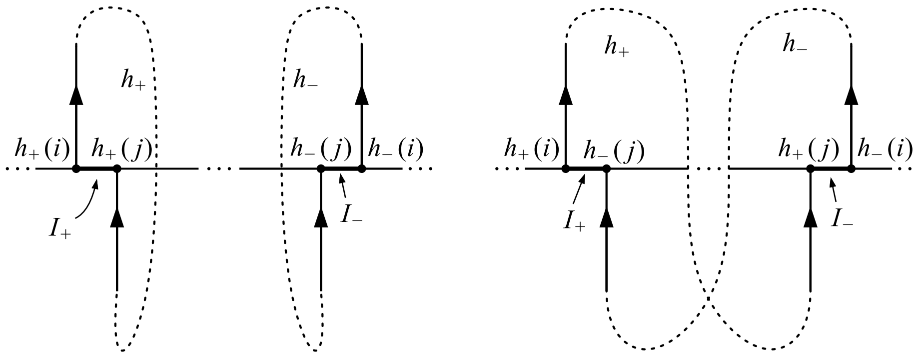

Figure 10: Above, the horizontal curve is always m.

The proof is an adaptation and correction of a surgery argument of Lecuire [42, Affirmation C.3]. As there are two gaps[9] in Lecuire's earlier argument, we give the proof in full detail below, without many citations of [42].

Claim 6.7 *Given $0 < \delta < \epsilon$, there are choices of $i < j$ such that either*:

(I) *The points $h_+(i)$ and $h_+(j)$ bound a segment $I_+ \subset m$ of length less than δ, and similarly with $-$ instead of $+$. The four velocity vectors $h'_+(i)$, $h'_+(j)$, $h'_-(i)$, $h'_-(j)$ all point to the same side of m, and the four segments $h_+([i,j])$, $h_-([i,j])$, I_+ and I_- have disjoint interiors. Therefore, the curves $\gamma_\pm := h_\pm([i,j]) \cup I_\pm \subset \partial M$ are simple and disjoint.*

(II) *The points $h_+(i)$ and $h_-(j)$ bound a segment $I_+ \subset m$ of length less than δ, and similarly the points $h_-(i)$ and $h_+(j)$ bound a segment $I_- \subset m$ of length less than δ. The four velocity vectors $h'_+(i)$, $h'_+(j)$, $h'_-(i)$, $h'_-(j)$ all point to the same side of m, and the four segments $h_+([i,j])$, $h_-([i,j])$, I_- and I_+ have disjoint interiors. Therefore, the curve $\gamma \subset \partial M$ obtained by concatenating all four segments is simple.*

See Figure 10 for a very useful picture. Note that in the picture, the velocity vectors of all paths intersecting m point to the same side of m, i.e., "up", and all 4-tuples of points are unlinked as in Claim 6.6.

Proof Start by fixing a circular order on m. Define "the right" to be the direction in m that is increasing with respect to the circular order, and define "the left" similarly. Since λ is minimal and not a simple closed curve, it has infinitely many leaves ℓ that are not boundary leaves. Fix some such ℓ, making sure that $h_\pm \not\subset \ell$ if the given rays happen to lie inside the lamination λ. The ray h_+ accumulates onto both sides of ℓ, so if we fix $p \in \ell \cap m$, the set $h_+(\mathbb{N})$ accumulates onto p from both sides, and similarly with $-$ instead of $+$. Fix an interval $J \subset m$ of length δ centred at p, and write $J = J_l \cup J_r$ as the union of the closed subintervals to the "left" and to the "right" of p. Note that $p \notin h_\pm(\mathbb{N})$, so each intersection of $h_\pm$ with J lies in exactly one of J_l or J_r.

[9]The first gap is that the sentence "*Quitte à extraire, la suite $(gh^{-1})^{2n} g(\tilde{l}^1)$ converge vers une géodésique $\tilde{\gamma} \subset p^{-1}(\alpha_1)$ dont la projection $l \subset \partial M$ est une courbe fermé.*" at the end of the proof of Affirmation C.3 isn't adequately justified; this is fixed in Claim 6.5. The second is that the assumption $d(l^2_+(y_i), l^2_+(y_j)) < \epsilon'$ in the statement of Affirmation C.3 is never actually verified, and does not come trivially from a compactness argument. This is fixed in Claim 6.7.

Let's call an index i *left-closest* if either $h_+(i)$ or $h_-(i)$ lies in J_l and is closer to p than any previous $h_\pm(k)$, $k < i$, that lies in J_l. *Right-closest* is defined similarly using J_r, and we call an index i *closest* if it is either left or right closest. Note that since $\delta < \epsilon$ we can never have both $h_+(i)$, $h_-(i)$ in J simultaneously, so no i is both right-closest and left-closest at the same time. Since there are infinitely many i of both types, at some point there will be a transition where some i_l is left-closest, some $i_r > i_l$ is right-closest, and there are no closest indices in between.

Let i_c be the smallest closest index that is bigger than i_r. (Here, c stands for "centre", since the corresponding point on J will lie between the points we get from the indices i_l and i_r.) We now have *three* points in J, so two of the corresponding velocity vectors point to the same side of m. Let $i, j \in \{i_l, i_r, i_c\}$ be the two corresponding indices, and for concreteness, *let's assume for the moment that $h_+(i)$ and $h_+(j)$ are the corresponding points in J*, deferring a discussion of the other cases to the end of the proof. Note that since the rays $h_\pm$ are mutually homoclinic and are in tight position with respect to m, the velocity vectors $h'_-(i)$ and $h'_-(j)$ point to the same sides of m as $h'_+(i)$ and $h'_+(j)$, respectively, and so all four vectors point to the same side. That is,

(a) $h'_+(i)$, $h'_+(j)$, $h'_-(i)$, $h'_-(j)$ all point to the same side of m,

(b) the segment $I_+ \subset J$ bounded by $h_+(i)$ and $h_+(j)$ contains no element $h_+(k)$ or $h_-(k)$ where k is between i and j.

Let $I_- \subset m \setminus J$ be the segment that is bounded by the points $h_-(i)$ and $h_-(j)$. *Suppose for a moment that we knew that I_- had length less than δ.* Then for each k, it is impossible that *both* $h_+(k)$ or $h_-(k)$ lie in I_-, as we're assuming that corresponding intersections of $h_\pm$ with m stay at least $\epsilon > \delta$ apart. In particular, if k is between i, j and we apply the unlinking condition of Claim 6.6 twice, once to i, k and once to j, k, we get from this and (b) above that *neither* element $h_+(k)$ or $h_-(k)$ is contained in I_-. So, the four segments $h_+([i, j])$, $h_-([i, j])$, I_+ and I_- have disjoint interiors, and we're in the situation of case (I) in the claim, as desired.

As constructed above, however, there is unfortunately no reason to believe that the interval I_- has length less than δ. To rectify this, recall that λ actually has infinitely many nonboundary leaves ℓ^n. For each such ℓ^n and $p^n \in \ell^n \cap m$, we can repeat the above construction using constants $\delta^n \to 0$, producing points (say) $h_+(i^n)$, $h_+(j^n)$ that lie within the length δ^n-interval $J^n \ni p^n$ and that satisfy properties (a) and (b) above. Choose the sequence $p^n \in \ell^n \cap m$ so that it is monotonic in the circular order induced on m, and let $\delta^n \to 0$ fast enough so that the associated intervals I^n_+ are all disjoint, so that in the circular order on m we have

$$h_+(i^1) < h_+(j^1) < h_+(i^2) < h_+(j^2) < \cdots < h_+(i^n),$$

and where each I^n_+ is the interval $[h_+(i^n), h_+(j^n)]$, rather than the complementary arc on m that has endpoints $h_+(i^n)$, $h_+(j^n)$. Then Claim 6.6 implies that

$$h_-(i^1) > h_-(j^1) > h_-(i^2) > h_-(j^2) > \cdots > h_-(i^n).$$

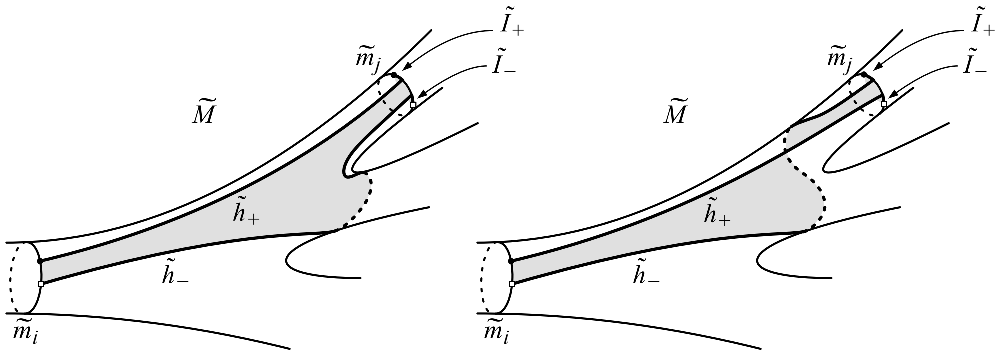

Figure 11: On the left, the two paths drawn in heavy ink project to the two simple closed curves $\gamma_\pm$ in Claim 6.7(I), shown on the left in Figure 10. The shaded region is a rectangle embedded in $\widetilde{M}$ that projects to an embedded annulus $A \hookrightarrow M$ with boundary $\gamma_- \cup \gamma_+$. On the right, the union of the two paths projects to the simple closed curve γ on M of Claim 6.7(II), and the shaded region projects to a Möbius band $B \hookrightarrow M$ with boundary γ.

Discarding finitely many n, we can assume all the point $h_+(i^n)$, $h_+(j^n)$ lie in an interval $U \subset m$ of length less than δ. Since the points $h_-(i^n)$, $h_-(j^n)$ are at least $\epsilon > \delta$ away from the corresponding $+$ points, they all lie in $m \setminus U$. Then since the interval I_-^n is defined to be disjoint from $I_+^n \subset U$, we must have $I_-^n = [h_-(j^n), h_-(i^n)]$, rather than the other interval with those endpoints. It follows that at least for large n, all the intervals I_-^n are disjoint. Since m is compact, we can then pick some n where I_-^n has length less than δ, as desired. Therefore, we are in case (I) in the statement of the claim, and are done.

In the argument above we simplified the notation by assuming that we have points $h_+(i^n), h_+(j^n) \in J^n$ satisfying conditions (a) and (b), which put us in case (I) at the end. Up to exchanging $+, -$, the only other relevant case is when our chosen points are $h_-(i^n), h_+(j^n) \in J^n$. After passing to a subsequence in n, if we are not in the case already addressed, then we may assume that our chosen points are $h_-(i^n), h_+(j^n) \in J^n$ for all n. And after exchanging $+$ with $-$ and passing to a further subsequence, we may assume

$$h_-(i^1) < h_+(j^1) < h_-(i^2) < h_+(j^2) < \cdots < h_-(i^1)$$

in the circular order on m, and that the interval $I_+^n = [h_-(i^n), h_+(j^n)]$. Everything from then on works exactly as above: if we set I_-^n to be the interval bounded by $h_+(i^n)$, $h_-(j^n)$ that is disjoint from I_+^n, then for some n we have that the length of I_-^n is less than δ, and it is easy to verify that we are in case (II) of the claim. □

We now finish the proof of Theorem 6.1. Suppose we are in case (I) of Claim 6.7. Then the two simple closed curves $\gamma_\pm$ drawn on the left in Figure 10 are the projections to M of the paths in $\widetilde{M}$ obtained by concatenating the arcs $\tilde{h}_\pm|_{[i,j]}$ with lifts $\widetilde{I}_\pm \subset \widetilde{m}_j$ of the intervals $I_\pm \subset m$; see Figure 11. We can homotope one path to the other in $\widetilde{M}$ while preserving the fact that the endpoints are points on $\widetilde{m}_i, \widetilde{m}_j$ that differ by the unique deck transformation taking $\widetilde{m}_i$ to $\widetilde{m}_j$. So projecting down, the simple closed curves $\gamma_\pm$ are freely homotopic in M, and hence bound an annulus $A \hookrightarrow M$. See the left part of Figure 11.

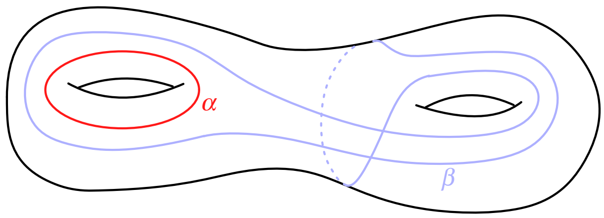

Figure 12: The two curves α and β are homotopic through the handlebody pictured, and therefore bound an essential singular annulus. The only annuli one can produce from surgery are inessential, but one can surger to obtain the "obvious" disc in the picture that separates the handlebody into two solid tori.

There is a uniform lower bound (depending on λ, m) for the angle at any intersection point of any leaf of λ with m, and the points $h_\pm(i)$ are at least ϵ away from each other in m. This implies that there is a uniform lower bound for the Hausdorff distance between $h_\pm|_{[i,j]}$ on ∂M. As long as the bound δ on the lengths of $I_\pm$ is small enough, the geodesics in the homotopy classes of $\gamma_\pm$ stay very close to $h_\pm|_{[i,j]}$, and are therefore distinct. So, the curves $\gamma_\pm$ are not homotopic in ∂M, and hence bound an *essential* annulus $A \hookrightarrow M$.

Choosing i, j to be large and δ to be very small, ∂A is contained in $S(\lambda)$ and its intersection number with λ is small. Hence, λ is an intrinsic limit of annuli, in the sense of Lemma 5.13, so we're done.

Case (II) is similar. Here, the single simple closed curve γ described in Claim 6.7(II) bounds a Möbius band $B \hookrightarrow M$; see the right side of Figure 11. Since ∂M is orientable, B is not boundary parallel, and hence by JSJ theory the boundary of a regular neighbourhood of B is an essential annulus $A \hookrightarrow M$ whose boundary consists of two disjoint curves that are both homotopic to γ on ∂M. As in case (I), we can make the intersection number of ∂A with λ arbitrarily small, so λ is an intrinsic limit of annuli, and we are done.

Remark 6.8 The proof (D) above is quite delicate. Most of this delicacy comes from Claims 6.5 and 6.7, which are needed to ensure that the annuli approximating λ that are produced immediately afterward are *embedded*. But while we are able to prove these claims using arguments involving the planarity of the closure of $\partial\widetilde{M}$ in $\mathbb{H}^3 \cup \partial_\infty\mathbb{H}^3$, one would not have to worry about these annuli being embedded if there was a strong "annulus theorem" guaranteeing that any essential singular annulus in an irreducible 3-manifold M can be surgered to give an essential embedded annulus. If this were true, Claims 6.5 and 6.7 could be replaced by a one paragraph compactness argument. Here, a *singular annulus* is a map $f:(A,\partial A)\to(M,\partial M)$ where $A = S^1\times[0,1]$. We say f is *essential* if it is not homotopic rel ∂A into ∂M.

Such an annulus theorem follows from the JSJ decomposition when M has incompressible boundary. When M has compressible boundary, there is a similar theorem as long as the original singular annulus has a spanning arc that is not homotopic rel ∂ into ∂M; see Cannon and Feustel's article [15]. However, our proof above does not provide such annuli, and indeed such annuli do not exist in compression bodies (the M of most interest to us), since *any* proper arc in a compressionbody M is homotopic rel ∂ into ∂M.

In fact, in a general M, one *cannot* always surger essential singular annuli to produce embedded essential annuli. For instance, the two curves in Figure 12 bound an essential singular annulus that cannot be surgered to give an embedded essential annulus. However, in that example, one *can* surger to get a meridian in the handlebody, so maybe an essential singular annulus can always be surgered to give either a meridian or an essential embedded annulus? This also turns out not to be true. Suppose P is a pair of pants and let $M = P \times [0,1]$, which is homeomorphic to a genus-two handlebody. If $\gamma \to P$ is an immersed figure-8 whose image forms a spine of P, then the singular annulus $\gamma \times [0,1] \to P \times [0,1]$ is essential, but the three embedded annuli that one can obtain from it by surgery are all inessential, and no meridian can be created by surgery either. However, we expect this is the *only* counterexample. The first author of this paper has spent considerable time trying to prove this with a tower argument, but pushing down the tower is very subtle, since if the obvious constructions fail, one has to characterise the figure-8 example.

7 Hausdorff limits of meridians

Let M be an orientable hyperbolisable compact 3-manifold, and equip $\partial_{\chi<0} M$ with a hyperbolic metric. Lecuire [42, Theorem B.1] showed that every lamination λ on $\partial_{\chi<0} M$ that is a Hausdorff limit of meridians contains a homoclinic leaf that is a homoclinic geodesic. This is a more general version of Casson's criterion for handlebodies, which was stated in the introduction.

The converse is not true: certainly in order to be a Hausdorff limit of meridians μ needs to be connected. And there are even connected laminations that contain homoclinic leaves but are not commensurable to Hausdorff limits of meridians. One way to do this is to just take a lamination that contains a meridian, but is not a limit even of simple closed curves, as on the left in Figure 13. There are also more subtle examples in genus 2, as pictured on the right in Figure 13: the reason they are not limits of meridians is as follows.

Lemma 7.1 *Suppose that S is a closed, genus-two surface, and λ is a geodesic lamination on S such that there is a separating meridian μ that does not intersect λ transversely, and λ intersects transversely the two nonseparating meridians disjoint from μ. Then λ is not a Hausdorff limit of meridians.*

Proof Let $T_\pm \subset S \setminus \mu$ be the two components of $S \setminus \mu$. Hoping for a contradiction, take a sequence of meridians (m_n) that Hausdorff converges to λ. We can assume after passing to a subsequence that m_n has a μ-wave in T_+ (say) for all n. Since T_+ is a compressible punctured torus, there is a *unique* homotopy class rel μ of μ-wave in T_+, so λ contains a leaf ℓ that either intersects T_+ in an arc in this homotopy class, or is contained in T_+ and is obtained by spinning an arc in this homotopy class around μ. But then ℓ

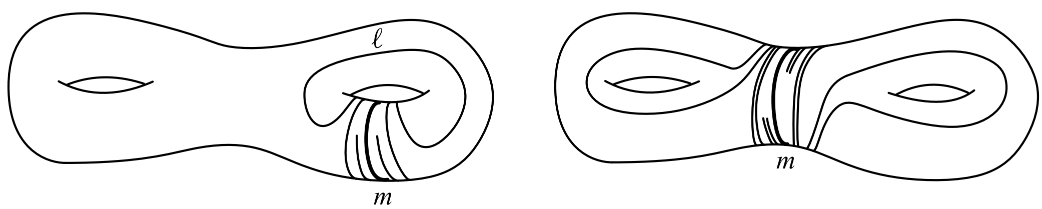

Figure 13: Two laminations on the boundary of a handlebody that have a meridian m as a leaf, but are not Hausdorff limits of meridians.

intersects nontrivially every nonperipheral minimal lamination in T_+ other than the unique nonseparating meridian μ_+ of T_+, so λ is disjoint from μ_+, contrary to assumption. □

In the examples in Figure 13, the problem is the spiralling leaves. So, maybe being a Hausdorff limit of meridians is the same as containing a homoclinic leaf if we ignore spiralling leaves? We say that two laminations μ_1, μ_2 are *commensurable* if they contain a common sublamination ν such that for both i, the difference $\mu_i \setminus \nu$ is the union of finitely many leaves. μ_1 and μ_2 are *strongly commensurable* if they contain a common ν such that for both i, the difference $\mu_i \setminus \nu$ is the union of finitely many leaves, none of which are simple closed curves.

Theorem 7.2 (Hausdorff limits of meridians) *Suppose that $S \subset \partial_{\chi<0} M$ is a connected subsurface with geodesic boundary, such that ∂S is incompressible, and that the disc set $\mathcal{D}(S, M)$ is "large", i.e., it has infinite diameter in the curve graph $C(S)$. Let λ be a geodesic lamination in* int(S) *that is a finite union of minimal laminations, and assume that the following does* **not** *hold*:

(⋆) *S is a closed, genus-two surface, there exists a separating meridian μ that does not intersect λ transversely, and λ intersects transversely the two nonseparating meridians disjoint from μ.*

Then λ is strongly commensurable to a Hausdorff limit of meridians in S if and only if λ is strongly commensurable to a lamination containing a homoclinic leaf, and this happens if and only if one of the following holds:

(1) *λ is disjoint from a meridian on S*,

(2) *some component of λ is an intrinsic limit of meridians*, *or*

(3) *there is an essential* (*possibly nontrivial*) *interval bundle $B \subset M$ over a compact surface Y that is not an annulus or Möbius band*, *and there are components $\lambda_\pm \subset \lambda$ that each fill a component of $\partial_h B$* (*possibly the same component*, *if $\partial_h B$ is connected*), *such that $\lambda_\pm$ are essentially homotopic through B*, *as in Section 2.9*, *and there is a compression arc α for B that is disjoint from λ.*

The case (⋆) above really is exceptional. Here, one should imagine a picture like the example on the right in Figure 13, but with the spiralling leaves replaced with minimal laminations in the two punctured tori. At least when $\mu \subset \lambda$ we have that λ contains a homoclinic leaf, but one can see that λ is *not* strongly commensurable to a Hausdorff limit of meridians, by using Lemma 7.1.

Recall from Proposition 3.1 that when $\mathcal{D}(S, M)$ does not have infinite diameter in $C(S)$, it is either empty, consists of a single separating meridian, or consists of a single nonseparating meridian m and all separating curves that are band sums of m. In these cases, it is obvious what the Hausdorff limits of meridians are. For instance, in the last case a finite union λ of minimal laminations in S is strongly commensurable to a Hausdorff limit of meridians if and only if either $\lambda = m$ or $\lambda \subset S \setminus m$. For the "if" direction, note that if $\lambda \subset S \setminus m$ then it can be approximated by an arc with endpoints on opposite sides of m, and doing a band sum with m gives a curve that approximates λ. For the "only if" direction, just note that all meridians are either equal to m or are contained in $S \setminus m$.

7.1 The proof of Theorem 7.2

Most of the proof of Theorem 7.2 is contained in the following results. Assume that $S \subset \partial_{\chi<0} M$ is a connected subsurface with geodesic boundary, ∂S is incompressible, and $\mathcal{D}(S, M)$ is large.

Lemma 7.3 *Suppose $\lambda \subset S$ is a lamination, there is a meridian μ that does not intersect λ transversely, and that if S is a closed surface of genus 2 then μ is nonseparating. Then λ is strongly commensurable to a Hausdorff limits of meridians on S.*

The proof of Lemma 7.3 uses some ideas that the first author developed with Sebastian Hensel, whom we thank for his contribution.

Proof We may assume that λ is a finite union of minimal components. It suffices to assume μ is not a leaf of λ, as long as we prove the conclusion both for such a λ and for $\lambda \cup \mu$.

Assume first that μ is nonseparating in S. Let (c_i) be a sequence of simple closed curves on S that Hausdorff-converges to λ. One can do this by constructing for each component $\lambda_0 \subset \lambda$ a simple closed geodesic approximating λ_0, by taking an arc that runs along a leaf of λ_0 for a long time, and then closing it up the next time it passes closest to its initial endpoint in the correct direction. Let α be a simple closed curve on S that intersects μ once, and intersects all the components of λ. For each k, let γ_i^k be the geodesic homotopic to the "band sum" of μ and $T_{c_i}^k(\alpha)$, where T_{c_i} is the twist around c_i and a band sum of two curves intersecting once is the boundary of a regular neighbourhood of their union. Note that γ_i^k is a meridian for all i, k. If (k_i) is a sequence that increases quickly enough, $(\gamma_i^{k_i})$ converges to a lamination strongly commensurable to λ. And if we pick a meridian β on S that intersects both μ and λ, then $T_\mu^i \circ T_{\gamma_i^{k_i}}^i(\beta)$ Hausdorff converges to a lamination strongly commensurable to $\lambda \cup \mu$.

Now suppose μ is separating. We claim that there is another separating meridian in S that is disjoint from μ. Let m be a maximal multicurve of meridians in S that contains μ as a component. Since $\mathcal{D}(S, M)$ is large, $m \neq \mu$. If m has a separating component distinct from μ, we are done. So, suppose we have a nonseparating component $m_0 \subset m$. We can make a (separating) band sum of m_0 that is disjoint from μ unless m_0 lies in a punctured torus component of $S \setminus \mu$. So, we assume the latter is true. Since $\mathcal{D}(S, M)$ is large, it cannot be that $m = \mu \cup m_0$, since then all meridians are disjoint from m_0. So, there is another component m_1 of m, which we can assume is also nonseparating. This m_1 must lie on the opposite side of μ from m_0, and as before we're done unless the component of $S \setminus \mu$ containing m_1 is also a punctured torus. But in this case, S is a genus-two surface contrary to assumption.

Let $T \subset S \setminus \mu$ be a component that contains a nonperipheral separating meridian, which we call μ'. Let V be the other component. Write $\lambda = \lambda_T \cup \lambda_V$, where $\lambda_T \subset T$ and $\lambda_V \subset S \setminus T$. Let C be the compression body with exterior boundary equal to the component of ∂M that contains S, that one obtains by compressing the meridian μ.

We claim that there are sequences of simple closed curves (α_i), (β_i) in T such that

- (α_i) and (β_i) both Hausdorff converge to a geodesic lamination strongly commensurable to λ_T, and
- for all i, α_i and β_i bound an essential annulus in C.

To construct these sequences, start by picking a simple closed curve α in T such that α and each component of λ_T together fill T. Let β be a simple closed curve on T such that α, β, μ bound a pair of pants in T. In C, we can compress the boundary component μ of this pair of pants, so α, β bound an annulus in C. Moreover, this annulus is essential, since otherwise α, β bound an annulus in T, implying T is torus with the one boundary component μ, contradicting the fact that there is a separating nonperipheral meridian in T. Then find a sequence (c_i) of simple closed curves in T that Hausdorff converge to a geodesic lamination strongly commensurable to λ_T, take k_i to be a fast increasing sequence and set $\alpha_i = T_{c_i}^{k_i}(\alpha)$ and $\beta_i := T_{c_i}^{k_i}(\beta)$. Since α fills with every component of λ_T, the curve β intersects every component of λ_T. It follows that (α_i) and (β_i) Hausdorff converge to a geodesic lamination strongly commensurable to λ_T. And since each c_i is nonperipheral in T, each component of c_i bounds an annulus in C with a curve on the interior boundary of C, so the twist T_{c_i} extends to C, implying that α_i, β_i bound an annulus in C as desired above.

Now let C' be the compression body obtained by compressing both μ and μ', so $C \subset C' \subset M$. Note that since both curves are separating and are disjoint, Proposition 3.1 says that $\mathcal{D}(S, C')$ is large, so we can pick a meridian $m \in \mathcal{D}(S, C')$ that intersects μ and every component of λ. Fix a sequence of geodesic multicurves (d_i) in V that Hausdorff converges to λ_V. As with the twists T_{c_i} in C, the twists T_{d_i} extend to C'. And the compositions $T_{\alpha_i} \circ T_{\beta_i}^{-1}$ extend to C' because the curves bound annuli in $C \subset C'$. We then define $\gamma_i := (T_{\alpha_i} \circ T_{\beta_i}^{-1})^{k_i} \circ T_{d_i}^{k_i}(m)$ for some fast increasing $k_i \to \infty$. These γ_i are all meridians and Hausdorff converge to a lamination strongly commensurable to λ. To obtain $\lambda \cup \mu$ instead, hit γ_i with high powers of twists around μ. □

Here is a more powerful version of Lemma 7.3. The idea of the proof is more or less the same, but more complicated.

Proposition 7.4 (promoting Hausdorff limits) *Suppose that ν, η are disjoint geodesic laminations on S that are finite unions of minimal components. Suppose also that no component of ν is a meridian.*

Let X be the union of the subsurfaces with geodesic boundary that are filled by the components of ν. Suppose that there are disjoint, nonhomotopic meridians μ, μ' on S that are disjoint from η, and a sequence of homeomorphisms

$$f_i : S \to S, \quad f_i|_{S \setminus \mathrm{int}(X)} = \mathrm{id},$$

such that $\mu_i := f_i(\mu)$ and $\mu_i' := f_i(\mu')$ are both sequences of meridians that Hausdorff converge to laminations strongly commensurable to ν. Then $\nu \cup \eta$ is strongly commensurable to a Hausdorff limit of meridians in S.

Before proving the proposition, we record the following application.

Corollary 7.5 *Suppose that λ is a geodesic lamination on S that is a finite union of minimal components. If either*

- *some component $\nu \subset \lambda$ that is not a simple closed curve is an intrinsic limit of meridians,*
- *there are (possibly equal) components $\lambda_\pm \subset \lambda$, neither of which is a simple closed curve, and where each fills a component of the horizontal boundary (possibly the same component if $\partial_h B$ is connected) of*

an essential interval bundle

$$(B, \partial_H B) \hookrightarrow (M, S),$$

where $\lambda_\pm$ *are essentially homotopic through* B, *and where there is a compression arc* α *for* B *that is disjoint from* λ,

then λ *is strongly commensurable to a Hausdorff limit of meridians.*

Proof Suppose some component $\nu \subset \lambda$ that is not a simple closed curve is an intrinsic limit of meridians. Setting $X := S(\lambda)$ we can take (μ_i) to be any sequence of meridians in X that Hausdorff converges to a lamination strongly commensurable to ν. Moreover, since ν fills X and is a limit of meridians, the disc set $\mathcal{D}(X, M)$ is large, so for each i there is some meridian μ_i' disjoint from μ_i. Since there are only finitely many topological types of pairs of disjoint curves in X up to the pure mapping class group of X, after passing to a subsequence we can assume that all μ_i, μ_i' are of the form in the proposition. The desired conclusion follows.

In the second case, we let X be the subsurface with geodesic boundary obtained by tightening $\partial_H B$ and set $\nu = \lambda_- \cup \lambda_+$. Write the interval bundle as $\pi : B \to Y$, where Y is a compact surface with boundary. We can assume without loss of generality that α is a strict compression arc, i.e., that it is homotopic rel endpoints to a fibre $\pi^{-1}(y)$, $y \in \partial Y$. Note that since $\lambda_\pm$ are not simple closed curves, Y is not an annulus or Möbius band.

Since $\lambda_\pm$ are essentially homotopic through B, Fact 2.16 says that if our reference hyperbolic metrics are chosen appropriately, we have that $\lambda_- \cup \lambda_+ = (\pi|_{\partial_H B})^{-1}(\bar{\lambda})$ for some geodesic lamination $\bar{\lambda}$ on Y. Since $\lambda_\pm$ together fill $\partial_H B$, the lamination $\bar{\lambda}$ is minimal and fills Y. So in particular, it has no closed, one-sided leaves, and therefore if we pick a nonzero transverse measure on $\bar{\lambda}$, we have that $\bar{\lambda}$ is the projective limit of a sequence of two-sided nonperipheral simple closed curves (c_i) in Y, by Theorem 1.2 of [20]. Homotope the c_i on Y to be based simple loops at $y \in \partial Y$, let $m(c_i)$ be the associated compressible curves on S constructed in Claim 2.11, and let μ_i be the geodesic meridians on S in their homotopy classes. Then (μ_i) Hausdorff converges to a lamination strongly commensurable to $\lambda_- \cup \lambda_+$. After passing to a subsequence, we can assume that all the c_i differ by pure homeomorphisms of Y, in which case the meridians μ_i are as required in Proposition 7.4, for some μ, f_i. Note that since our compression arc is assumed to be disjoint from λ, all the μ_i are disjoint from $\eta := \lambda \setminus \lambda_\pm$, and hence so is our μ. We create disjoint meridians μ_i' similarly, by taking some c_i' on Y disjoint from c_i, and letting μ_i' be the geodesic meridian homotopic to $m(c_i')$. It then follows from Proposition 7.4 that λ is strongly commensurable to a Hausdorff limit of meridians as desired. □

We now prove the proposition.

Proof of Proposition 7.4 Assume that μ, μ' are disjoint meridian on S that are disjoint from η, that $f_i : S \to S$ are homeomorphisms that are the identity outside of X, and that $\mu_i := f_i(\mu)$ and $\mu_i' := f_i(\mu')$ are sequences of meridians that Hausdorff converge to laminations strongly commensurable to ν.

We want to show that $\nu \cup \eta$ is strongly commensurable to a Hausdorff limit of meridians on S. We now basically repeat the argument in Lemma 7.3, so the reader should make sure that they understand that argument before continuing here.

Suppose μ (say) is nonseparating in S. Choose a simple closed curve α on S that intersects μ once and intersects essentially each component of η. Let $\alpha_i := f_i(\alpha)$, and note that α_i intersects μ_i once, and also intersects essentially each component of η. Let (c_i) be a geodesic multicurve that Hausdorff converges to η, and let γ_i^k be the geodesic homotopic to the "band sum"

$$B(\mu_i, T_{c_i}^k(\alpha_i)) = T_{c_i}^k(B(\mu_i, \alpha_i)) = T_{c_i}^k \circ f_i(B(\mu, \alpha)), \tag{4}$$

where here $B(\cdot,\cdot)$ takes in two simple closed curves that intersect once and returns the boundary of the regular neighbourhood of their union. If one of the inputs in a band sum is a meridian, then so is the output, so γ_i^k is a meridian for all i, k. The given equalities are true at least for large i. The first equality holds because μ is disjoint from η, $f_i = \mathrm{id}$ on the subsurfaces filled by the components of η, and therefore μ_i is disjoint from c_i for large i. The second equality is obvious from the definitions of μ_i, α_i.

Let (k_i) be a fast increasing sequence. After passing to a subsequence, we can assume that $(\gamma_i^{k_i})$ Hausdorff converges to a lamination λ. *We claim that λ is strongly commensurable to $\nu \cup \eta$.*

First, using the second term in (4), if k_i is huge with respect to i, then c_i is contained in a small neighbourhood of $\gamma_i^{k_i}$, and so since (c_i) Hausdorff converges to η, *we have $\lambda \supset \eta$.*

We claim that each $\gamma_i^{k_i}$ essentially intersects each component $X_0 \subset X$. If not, then from the third term in (4) it follows that $B(\mu, \alpha)$ is disjoint from X_0. But μ essentially intersects X_0, since otherwise the Hausdorff limit of the μ_i will not contain the associated component $\nu_0 \subset \nu$. So, μ, α and X_0 all lie in the punctured torus $T \subset S$ bounded by $B(\mu, \alpha)$. But since α intersects every component of η, we have that η intersects T as well, in a collection of arcs disjoint from μ. Since X_0 is disjoint from η, $X_0 = \mu$, so $\nu_0 = \mu$ is a meridian, contrary to our standing assumption.

It now follows that the Hausdorff limit λ essentially intersects each component of X. Since $\gamma_i^{k_i}$ is disjoint from μ_i and (μ_i) Hausdorff converges to a lamination containing ν, The laminations λ, ν cannot intersect transversely. Since each component of X is filled by a component of ν, *we have $\lambda \supset \nu$.*

Finally, if Y is the union of all the subsurfaces with geodesic boundary that are filled by the components of η, then as $f_i = \mathrm{id}$ outside X and $c_i \subset Y$ for large i, the intersection of $\gamma_i^{k_i}$ with $S \setminus (X \cup Y)$ is properly homotopic to the intersection of $B(\mu, \alpha)$, which is independent of i. It follows that $\lambda \setminus (\nu \cup \eta)$ *is a finite collection of nonclosed leaves*, so we are done.

We can now assume that both μ, μ' are separating, so that μ_i, μ_i' are also separating for all i. Let $T_i \subset S \setminus \mu_i$ be the component containing μ_i', and let V_i be the other component. Note that T_i is not a punctured torus, since it contains a nonperipheral separating curve. Since $\partial T_i \cap \eta = \varnothing$, we have

$$\eta = \eta_T \sqcup \eta_V,$$

where the first term is the intersection of η with T_i, and the second term is defined similarly. Note that since $f_i = \mathrm{id}$ on $S \setminus X$, all the μ_i induce the same two-element partition of the components of $S \setminus X$, so at least after passing to a subsequence the decomposition of η above is actually independent of i, which is why we have omitted the i in the notation.

Let C be the compression body whose exterior boundary is the component of ∂M containing S, and which is obtained by compressing the curve μ. Let C' be similarly obtained by compressing both μ and μ', so that $C \subset C' \subset M$. Since C' admits two nonhomotopic disjoint separating meridians, the disc set $\mathcal{D}(S, C')$ is large by Proposition 3.1, so we can pick a meridian $m \in \mathcal{D}(S, C')$ that intersects every component of $\nu \cup \eta$, as well as μ, μ'. Let $C_i \subset C_i' \subset M$ be the compression bodies obtained by compressing μ_i, μ_i'. Then f_i extends to a map $C' \to C_i'$, implying that $m_i := f_i(m)$ is a meridian in C_i'.

As in the proof of Lemma 7.3, we can pick sequences (α_i), (β_i) of simple closed curves in T_i such that (α_i) and (β_i) both Hausdorff converge to η_T, and where α_i, β_i bound an essential annulus in C_i for all i. As in the lemma, $T_{\alpha_i} \circ T_{\beta_i}^{-1}$ extends to C_i'. Let (c_i) be a sequence of multicurves in V_i that Hausdorff converges to η_V. Each component of c_i bounds an annulus in C_i' with a curve on the interior boundary of C_i', and hence the multitwist T_{c_i} extends to a homeomorphism of C_i'. For any given k, set

$$\gamma_i^k := (T_{\alpha_k} \circ T_{\beta_k}^{-1})^k \circ T_{c_k}^k(m_i).$$

We claim that for fast increasing k_i, the curves $\gamma_i^{k_i}$ Hausdorff converge to a lamination that is strongly commensurable to $\nu \cup \eta$ as desired. This is proved using the same types of arguments we employed in the nonseparating case above. In particular, recall that m was selected to intersect all components of $\nu \cup \eta$. Since f_i is supported on subsurfaces filled by components of ν, all the $m_i = f_i(m)$ intersect all components of $\nu \cup \eta$, and hence for large k_i they intersect $\alpha_{k_i}, \beta_{k_i}$. So, $\gamma_i^{k_i}$ is twisted many times around $\alpha_{k_i}, \beta_{k_i}$, and hence its Hausdorff limit contains ν. Similarly, the m_i intersect c_{k_i} for large i. Since c_k lies in V_k, it is disjoint from $\alpha_k \subset T_k$ and $\beta_k \subset T_k$, and thus the Hausdorff limit of $\gamma_i^{k_i}$ contains η. Finally, the Hausdorff limit has no other minimal components because α_k, β_k, c_k are contained in subsurfaces filled by components of $\nu \cup \eta$, and $m_i = f_i(m)$ is constant outside this subsurfaces. □

We can now start the proof of the theorem.

Proof of Theorem 7.2 Suppose that $\lambda \subset S$ is a lamination and $(\star)$ does not hold, so that it is not the case that S is a genus-two surface and λ is disjoint from a separating meridian μ, but intersects the two nonseparating meridians disjoint from μ. We want to show that λ is strongly commensurable to a Hausdorff limit of meridians if and only if it is strongly commensurable to a lamination containing a homoclinic leaf, which happens if and only if either

(1) λ is disjoint from a meridian,

(2) some component of λ is an intrinsic limit of meridians, or

(3) there is an essential (possibly nontrivial) interval bundle $B \subset M$ over a compact surface Y that is not an annulus or Möbius band, and there are components $\lambda_\pm \subset \lambda$ that each fill a component of $\partial_H B$, such that $\lambda_\pm$ are essentially homotopic through B, as in Section 2.9, and there is a compression arc α for B that is disjoint from λ.

Hausdorff limit $\Longrightarrow$ homoclinic leaf Suppose first that λ is strongly commensurable to a Hausdorff limit of meridians λ'. Then by [42, Theorem B.1], there is a homoclinic leaf $h \subset \lambda'$, so λ is strongly commensurable to a lamination with a homoclinic leaf as desired.

Homoclinic leaf $\Longrightarrow$ (1), (2) or (3) We now assume we have a homoclinic leaf h in some lamination strongly commensurable to λ.

The two ends of h limit onto (possibly equal) components $\lambda_\pm \subset \lambda$. If one of $S(\lambda_\pm)$ has compressible boundary, there is a meridian disjoint from λ, so we are in case (1) and are done. So, $\partial S(\lambda_\pm)$ is incompressible, and we're in the situation of Theorem 6.1 and Corollary 6.2. We now break into cases.

If one of $\lambda_\pm$ is an intrinsic limit of meridians, we're in case (2) and are done. If we're in case (3) of Theorem 6.1 and Corollary 6.2, we're in case (3) of the theorem and are done, unless the given interval bundle $B \to Y$ is over an annulus or Möbius band. But in that case, letting c be a boundary component of Y, we can consider the geodesic meridian μ on S homotopic to the $m(c)$ constructed in Claim 2.11, using the compressing arc given by Corollary 6.2. This μ is disjoint from λ, so we're in case (1) of the theorem.

Finally suppose that the two ends of h are asymptotic on S, so that $\lambda_- = \lambda_+$. Let's separate further into the cases (i) and (ii) in Corollary 6.2. In case (i), using the notation of the corollary, the curve $c \cup h([-s, s])$ is a meridian disjoint from λ. So, we're in case (1) of the theorem. In case (ii), let T be a neighbourhood of $h \cup \lambda_\pm$ that is either a punctured torus or a pair of pants, depending on whether the two ends of h limit onto opposite sides of $\lambda_\pm$, or onto the same side. Because we're in case (ii), there is a meridian in T. Hence, whether T is a pair of pants or a punctured torus, one of the boundary components of T is a meridian, and is disjoint from λ so we're done.

(1), (2) or (3) $\Longrightarrow$ Hausdorff limit Suppose (1), (2) or (3) holds. We want to show λ is strongly commensurable to a Hausdorff limit of meridians. If λ is disjoint from a meridian, then we're done by Lemma 7.3. If a component of λ is an intrinsic limit of meridians, we're done by the first part of Corollary 7.5. In case (3) above, we're done by the second part of Corollary 7.5. □

8 Extending partial pseudo-Anosovs to compression bodies

Let M be a compression body with exterior boundary Σ. Let $S \subset \Sigma$ be an essential subsurface such that ∂S is incompressible. In this section, we prove:

Theorem 8.1 (extending partial pseudo-Anosovs) *Suppose that $f : \Sigma \to \Sigma$ is a partial pseudo-Anosov supported on S. Then f has a power that extends to a nontrivial subcompression body of (M, S) if and only if the attracting lamination of f is a projective limit of meridians that lie in S.*

When $S = \Sigma$, this is a theorem of Biringer, Johnson and Minsky [2]. The proof of Theorem 8.1 is basically the same as their proof, but we need to go through it anyway, to note the places that parabolics appear, and to deal with the fact that we are looking at subcompression bodies of (M, S) rather than of M. Also, before starting on the bulk of the proof in Section 8.2, we isolate part of the argument into a separate purely topological subsection, Section 8.1. This separation of the argument into distinct topological and geometric parts makes it more understandable than the original version, we think.

8.1 Dynamics on the space of marked compression bodies

Let Σ be a closed, orientable surface, and let $S \subset \Sigma$ be an essential subsurface. The *space of marked S-compression bodies* is defined to be

$$\mathrm{CBod}(S) = \{(C, h : \Sigma \to \partial_+ C)\}/\sim,$$

where here C is a compression body, h is a homeomorphism, and

- the multicurve $h(\partial S) \subset \partial_+ C$ is incompressible,
- there is a multicurve m on S such that $h(m)$ is a cut system for C, i.e., $h(m)$ bounds a collection of disks that cut C into balls and trivial interval bundles over the interior boundary components.

We declare $(C_i, h_i : S \to \partial_+ C_i)$, $i = 1, 2$, to be equivalent (written $\sim$ as above) if there is a homeomorphism $\phi : C_1 \to C_2$ that respects the boundary markings: that is, $\phi \circ h_1$ and h_2 are homotopic maps $S \to \partial_+ C_2$.

We write $(C_1, h_1) \subset (C_2, h_2)$ if there is an embedding $\phi : (C_1, \partial_+ C_1) \hookrightarrow (C_2, \partial_+ C_2)$ that respects the boundary markings. This gives a partial ordering on $\mathrm{CBod}(S)$. We often identify Σ with $\partial_+ C$ instead of specifying the boundary marking, and simply write C for an element of $\mathrm{CBod}(S)$. So $\mathrm{CBod}(S)$ is the set of all compression bodies with exterior boundary Σ that one obtains by compressing curves in S (without compressing boundary curves) up to the obvious equivalence.

A marked S-compression body (C, h) has a *disk set* $\mathcal{D}(C) \subset \mathcal{C}(S)$, where a simple closed curve $\gamma \in \mathcal{C}(S)$ lies in the disc set if $h(\gamma)$ is a meridian in C. In fact, the disk set $\mathcal{D}(C)$ determines (C, h) up to equivalence, say by an argument similar to the last paragraph of the proof of Fact 2.4. The set $\mathrm{CBod}(S)$ can then be identified with the "set of all disk sets" in $\mathcal{C}(S)$. It then inherits a topology as a subset of the power set $\mathcal{P}(\mathcal{C}(S))$, wherein $D_n \to D$ if and only if for every $c \in \mathcal{C}(S)$, we have either $c \in D$ and $c \in D_n$ for all large n, or $c \notin D$ and $c \notin D_n$ for all large n.

Lemma 8.2 *If $C_n \to C$ in $\mathrm{CBod}(S)$, then $C \subset C_n$ for large n.*

Proof Suppose that C is obtained by compressing a finite set Γ of disjoint simple closed curves on S. For large n, we have $\Gamma \subset \mathcal{D}(C_n)$, so $\mathcal{C} \subset C_n$. □

Lemma 8.3 $\mathrm{CBod}(S)$ *is compact.*

Proof As $\mathcal{P}(\mathcal{C}(S))$ is compact, we want to show that $\mathrm{CBod}(S)$ is closed. Suppose C_n is a sequence of marked compression bodies with disk sets

$$D_n = \mathcal{D}(C_n) \subset \mathcal{C}(S),$$

and that $D_n \to D \subset \mathcal{C}(S)$. Let Γ be a maximal set of disjoint, pairwise nonhomotopic elements of D. Compressing Γ yields a marked compression body C. Since Γ is finite, $\Gamma \subset D_n$ for large n, so $\mathcal{D}(C) \subset D_n$. Thus, $\mathcal{D}(C) \subset D$.

It therefore suffices to show $D \subset \mathcal{D}(C)$. Suppose this is not the case, and pick $\beta \in D \setminus \mathcal{D}(C)$ such that the intersection number of β and Γ is minimal. By maximality of Γ, this intersection number cannot be zero. Since $\beta \in D$, if n is large we have $\beta \in \mathcal{D}(C_n)$. By an outermost disk argument, if $\gamma \in \Gamma$ is a component that intersects β, there is an arc $c \subset \gamma$ with endpoints on β and interior disjoint from β, that is homotopic rel endpoints in C_n to the two arcs $b_1, b_2 \subset \beta$ into which β is cut by ∂c. Passing to a subsequence, we can assume that c, b_1, b_2 are independent of n. Then $c \cup b_1$ and $c \cup b_2$ are both meridians in C_n for all large n, and hence lie in D. Since they intersect Γ fewer times than β does, they lie in $\mathcal{D}(C)$. But then β (which is a band sum of the two curves) also lies in $\mathcal{D}(C)$, a contradiction. □

Let $f : \Sigma \to \Sigma$ be a homeomorphism with $f = \mathrm{id}$ on $\Sigma \setminus S$. Then f acts on the space $\mathrm{CBod}(S)$ by $f \cdot (C, h) = (C, h \circ f^{-1})$. When we regard marked S-compression bodies as compression bodies with exterior boundary *equal* to Σ, we'll just write C and $f(C)$ for a marked compression body and its image. Note that $f(C) = C$ if and only if f extends to a homeomorphism of C.

Fixing $M \in \mathrm{CBod}(S)$ and f as above, let $\mathcal{A}$ be the set of accumulation points in $\mathrm{CBod}(S)$ of the f-orbit of M, and let

$$\mathcal{A}_{\min} = \{C \in \mathcal{A} \mid \nexists D \in \mathcal{A} \text{ such that } D \subsetneq C\}$$

be the subset consisting of all minimal elements of $\mathcal{A}$.

Theorem 8.4 (existence of maximal subcompression body) *The set $\mathcal{A}_{\min}$ is a finite f-orbit that contains a single element C_f such that $C_f \subset M$.*

Moreover, C_f is the unique maximal element of $\mathrm{CBod}(S)$ *such that $C_f \subset M$ and a power of f extends to C_f.*

This result was proved in [2] when $S = \Sigma$. Our proof follows the same general lines, but is topological instead of hyperbolic geometric.

We proceed with a series of lemmas.

Lemma 8.5 *The set $\mathcal{A}_{\min}$ is nonempty, finite and f-invariant.*

Proof The set $\mathcal{A}$ is nonempty, since $\mathrm{CBod}(S)$ is compact. This implies that $\mathcal{A}_{\min}$ is nonempty, for example, since the "height" of a compression body is nonnegative and decreases under strict containment; see Section 3 of [3].

By Lemma 8.2, $\mathcal{A}_{\min}$ is discrete. But $\mathcal{A}_{\min}$ is closed in $\mathcal{A}$, which is closed in $\mathrm{CBod}(S)$, which is compact. So, $\mathcal{A}_{\min}$ is compact, and must be finite. Finally, $\mathcal{A}_{\min}$ is f-invariant since $\mathcal{A}$ is and the f-action respects containment. □

Lemma 8.6 *Suppose that for $i = 1, 2$ we have $C_i \in \mathrm{CBod}(S)$ with $C_i \subset M$, and that $f^i(C_1) = C_1$ while $f^j(C_2) = C_2$. Then there is an element $C \in \mathcal{A}_{\min}$ such that $C_1, C_2 \subset C \subset M$.*

Proof Every element of $\mathcal{A}$ is the image under a power of f of an accumulation point of the sequence $f^{nij}(M)$, so since $\mathcal{A}_{\min}$ is f-invariant there is some $C' \in \mathcal{A}_{\min}$ to which a subsequence of $f^{nij}(M)$ limits. As $C_1, C_2 \subset f^{nij}(M)$ for all n, we must have $C_1, C_2 \in C'$.

By Lemma 8.2, there is some n such that $f^{nij}(M) \supset C'$. Then $C := f^{-nij}(C') \in \mathcal{A}_{\min}$ is contained in M and must contain C_1, C_2 as well. □

Lemma 8.7 *There is a unique element $C_f \in \mathcal{A}_{\min}$ that is contained in M, and $A_{\min}$ is an f-orbit.*

Proof Applying the previous lemma to two copies of the trivial compression body $\Sigma \times I$ shows that $\mathcal{A}_{\min}$ has an element that is contained in M.

So, suppose that $C, D \in \mathcal{A}_{\min}$ are both contained in M. By the previous lemma, there is another element of $\mathcal{A}_{\min}$ that contains them both, which contradicts the minimality assumption unless $C = D$. Therefore, there is a unique element $C_f \in \mathcal{A}_{\min}$ that is contained in M.

To show that $\mathcal{A}_{\min}$ is an f-orbit, suppose that $C \in \mathcal{A}_{\min}$. Since C is an accumulation point, there is some n such that $f^n(M) \supset C$. Then $f^{-n}(C) \subset M$, implying that $f^{-n}(C) = C_f$ by uniqueness. □

This finishes the proof of Theorem 8.4, since Lemma 8.5 shows that a power of f extends to C_f and Lemmas 8.6 and 8.7 imply that any subcompression body of M to which a power of f extends is contained in C_f.

8.2 The proof of Theorem 8.1

Let $S \subset \Sigma = \partial M_+$ be a compact essential subsurface, with ∂S incompressible in M, and let $f : \Sigma \to \Sigma$ be a pseudo-Anosov map on S.

The "only if" direction of the theorem is trivial. Namely, suppose that some power f^k extends to a nontrivial subcompression body C of (M, S). Pick a meridian $m \subset S$ for C. Then $(f^k(m))$ is a sequence of meridians in M that lie in S, and converge to the attracting lamination of f.

For the "if" direction of the theorem, assume that no nonzero power of f extends to a nontrivial subcompression body of (M, S). We must show that the attracting lamination λ^+ is not in the limit set $\Lambda(S, M)$. The argument is similar to the proof of the main theorem in [2]. As such, we will sketch the argument in places and refer to [2] for details.

Consider the sequence $M_n = f^{-n}(M)$ of marked S-compression bodies, where we consider the exterior boundary of each M_n as identified with the surface Σ. Fix a base point $[X] \in \mathcal{T}(\Sigma)$ and give the interior of each M_n a geometrically finite hyperbolic metric such that the end adjacent to the exterior boundary $\Sigma = \partial_+ M$ is convex cocompact, and when its conformal boundary is identified with Σ, the conformal structure is $[X]$. Let

$$\rho_n : \pi_1 \Sigma \to \mathrm{PSL}_2 \mathbb{C}, \quad N_n := \mathbb{H}^3 / \rho_n(\pi_1 \Sigma),$$

be a representation (unique up to conjugacy) uniformising the interior of M_n and compatible with our markings, in the sense that ρ_n is the composition of the map $\pi_1 \Sigma \to \pi_1 M_n \cong \pi_1 N_n$ induced by inclusion and a faithful uniformising representation of $\pi_1 N_n$. Note that the kernel of ρ_n is

$$\ker(\rho_n) = f_*^{-n}(\ker(\pi_1 \Sigma \to \pi_1 M)).$$

By Theorem 8.4 and the assumption that no power of f extends to a nontrivial subcompression body of M, the only minimal accumulation point of (M_n) in $\mathrm{CBod}(S)$ is the trivial compression body. So in particular, we can choose a subsequence (M_{n_j}) that converges to the trivial compression body. By the compactness of generalised Bers slices (see [2, Theorem 4.3]), we may assume after appropriate conjugations and passing to a further subsequence that (ρ_{n_j}) converges algebraically to a representation

$$\rho_\infty : \pi_1\Sigma \to \mathrm{PSL}_2\,\mathbb{C}, \quad N_\infty := \mathbb{H}^3/\rho_\infty(\pi_1\Sigma),$$

and that N_∞ can be identified with the interior of a compression body M_∞ with exterior boundary Σ in such a way that the end of N_∞ adjacent to Σ is convex cocompact with conformal structure $[X]$ and the representation ρ_∞ is compatible with the marking in the same way as before.

The disk set $\mathcal{D}(S, M_\infty)$ consists of all simple closed curves on S represented by elements $\gamma \in \pi_1\Sigma$ with $\rho_\infty(\gamma) = 1$. By Chuckrow's theorem (see [2, Lemma 2.11]), $\rho_\infty(\gamma) = 1$ if and only if $\rho_{n_j}(\gamma) = 1$ for all sufficiently large i. Since (M_{n_j}) converges to the trivial compression body in the topology of $\mathrm{CBod}(S)$, it follows that the surface $S \subset \Sigma = \partial_+ M_\infty$ is incompressible in M_∞.

Claim 8.8 *The repelling lamination λ^- of f is unrealisable in N_∞.*

Proof The proof is almost identical to that of [2, Lemma 6.2], so we offer a sketch and we refer the reader to their paper for details.

Fixing an M-meridian $\gamma \subset S$, the sequence $f^{-n_j}(\gamma)$ converges in the Hausdorff topology to a lamination λ_M that is the union of λ^- and finitely many leaves spiralling onto it. It suffices by [7, Theorem 2.3] to show that λ_M is unrealisable. So, hoping for a contradiction, assume λ_M *is* realisable; then λ_M is carried by a train track τ that maps nearly straightly into N_∞ (see [2]).

By algebraic convergence, τ also maps nearly straightly into N_{n_j} when j is large. Since $f^{-n_j}(\gamma) \to \lambda_M$, for large j the curve $f^{-n_j}(\gamma)$ is carried by τ. This implies that $f^{-n_j}(\gamma)$ is geodesically realisable in N_{n_j} for large j, contradicting the fact that it is homotopically trivial. □

By work of Thurston [58, Proposition 9.7.1], the π_1-injective surface $S \subset \partial_+ M_\infty$ is isotopic into a degenerate end of N_∞ with ending lamination λ^-. In particular, the peripheral curves of S represent cusps in N_∞ and every nonperipheral curve on S has hyperbolic type in N_∞. Any pair of disjoint nonperipheral simple closed curves on S can then be realised geodesically by a pleated surfaces $S \to N_\infty$ in the given homotopy class, and Thurston's compactness of pleated surfaces (see [50, Lemma 6.13]) implies the following.

Lemma 8.9 (compare with [2, Lemma 6.3]) *Let $\alpha \subset S$ be a simple closed curve. Then for every k, there is some K such that for any other simple closed curve β in S, we have*

$$d_{\mathcal{C}(S)}(\alpha, \beta) \le k \implies d_{N_\infty}(\alpha_\infty, \beta_\infty) \le K,$$

where α and β_∞ are the geodesics in N_∞ in the homotopy classes of α and β.

Hoping for a contradiction, suppose now that $\lambda^+ \in \Lambda(S, M)$. When regarded as an element of $\partial_\infty \mathcal{C}(S)$, the support of λ^+ is then an accumulation point of $\mathcal{D}(S, M) \subset \mathcal{C}(S)$. If $\alpha \in \mathcal{C}(S)$, then for $n = 1, 2, \ldots$

the sequence $(f^n(\alpha))$ is a quasigeodesic path that limits to $\lambda^+ \in \partial_\infty \mathcal{C}(S)$; see [47]. Since $\mathcal{D}(S, M)$ is a quasiconvex subset (see Theorem 2.3, due to Masur and Minsky), there is a constant C and for each n a meridian $\gamma_n \in \mathcal{D}(S, M)$ with

$$d_{\mathcal{C}(S)}(f^n(\alpha), \gamma_n) \le C.$$

Translating the points $f^n(\alpha)$ and γ_n by f^{-n}, this becomes

$$d_{\mathcal{C}(S)}(\alpha, f^{-n}(\mathcal{D}(S, M)) \le C.$$

By Lemma 8.9, an element $\gamma_{n_j} \in f^{-n_j}(\mathcal{D}(S, M))$ at distance at most C from α in $\mathcal{C}(S)$ can be geodesically realised in some fixed compact subset $A \subset N_\infty$. Algebraic convergence implies that for sufficiently large j this geodesic can be pulled back and tightened to a geodesic in N_{n_j}. But by construction, γ_{n_j} is a meridian in M_{n_j}, so it cannot possibly be realised geodesically in N_{n_j}, which is a contradiction.

9 Extending reducible maps to compression bodies

We present here a generalisation of [2, Theorem 1.1] that characterises which (possibly reducible) mapping classes of the boundary of a 3-manifold M have powers that extend to subcompression bodies.

In what follows, let M be a compression body with exterior boundary $\partial_+ M$. Let $S \subset \partial_+ M$ be an essential subsurface such that ∂S is incompressible. Let $f : \partial_+ M \to \partial_+ M$ be a homeomorphism that is "supported" in S, meaning that $f = \mathrm{id}$ on $\partial_+ M \setminus S$.

Definition 9.1 We say f is *pure* if there are disjoint, compact, essential f-invariant subsurfaces $S_i \subset S$, $i = 1, \dots, n$, such that $f = \mathrm{id}$ on $S_{\mathrm{id}} := S \setminus \bigcup_i S_i$, and where for each i, if we set $f_i := f|_{S_i}$, then either

(1) S_i is an annulus and f_i is a power of a Dehn twist, or

(2) f_i is a pseudo-Anosov map on S_i.

It follows from the Nielsen–Thurston classification, see [21], that every f has a power that is isotopic to a pure homeomorphism.

When f is pure, with associated restrictions $f_i : S_i \to S_i$ as above, we define a geodesic lamination $\lambda = \bigcup_i \lambda_i$ on S, where $\lambda_i \subset S_i$ as follows. If f_i is pseudo-Anosov, we let λ_i be the support of the attracting lamination of f_i. If f_i is a Dehn twist, we let λ_i be the core curve of the annulus S_i. So defined, λ is called the *attracting lamination* of the pure homeomorphism f.

Theorem 9.2 *Suppose that $S \subset \partial_+ M$ is an essential subsurface such that the multicurve ∂S is incompressible. Let $f : \partial_+ M \to \partial_+ M$ be a pure homeomorphism supported in S. Then f has a power that extends to a nontrivial subcompression body of (M, S) if and only if either*

(1) *there is a meridian in S_{id},*

(2) *for some i, the map f_i has a power that extends to a nontrivial subcompression body of (M, S_i), or*

(3) *there are (possibly equal) indices i, j such that S_i, S_j bound an essential interval bundle B in M, such that some power of $f|_{S_i \cup S_j}$ extends to B, and there is a compression arc α for B whose interior lies in S_{id}.*

Note that by Theorem 8.1, if f_i is pseudo-Anosov then (2) above is equivalent to the condition that λ_i is an intrinsic limit of meridians.

Recall from Section 2.7 that a "subcompression body of (M, S)" is a compression body obtained from $\partial_+ M$ by compressing some meridian multicurve in S. In (3), the condition that a power of $f|_{S_i \cup S_j}$ extends to B is easier to check. Indeed, if $\sigma : \partial_H B \to \partial_H B$ is the canonical involution, as defined in Section 2.5, then by Fact 2.7 we have that $f^k|_{\partial_H B}$ extends to B exactly when $\sigma \circ f_i^k$ is isotopic to f_j^k. When B is a twisted interval bundle, f_i, f_j are both pseudo-Anosovs on $\partial_H B$ and this means that as mapping classes we have $f_j = g \circ f_i$ for some finite order g commuting with both f_i, f_j; see, e.g., McCarthy's thesis [51]. When B is a trivial interval bundle, σ identifies S_i and S_j, and we have similarly that $f_j = g \circ \sigma(f_i)$ for some g commuting with both.

Proof of Theorem 9.2 Let's start with the "if" direction, since that's easier. If there is a meridian in S_{id}, then f extends to the compression body obtained by compressing that meridian. Suppose (2) holds, so that some power f_i^k extends to a nontrivial subcompression body C of (M, S_i). Then f also extends to C, since all the S_j, where $j \neq i$, bound trivial interval bundles with subsurfaces of the interior boundary of C. So we're done.

The only interesting case is if (3) holds, so that some S_i, S_j bound an essential interval bundle B in M such that some power of $f^k|_{S_i \cup S_j}$ extends to B, and there is a compression arc α for B whose interior lies in S_{id}. Here, let $S' \subset S$ be the smallest essential subsurface containing S_i, S_j and α; so, S' is obtained from a regular neighbourhood of the union of these three subsets of S by capping off any inessential boundary components with discs. Let C be the characteristic compression body of the pair (M, S'), as defined in Fact 2.4.

We claim that f^k extends to C. To see this, note that we can construct C as follows. For concreteness, first assume that the boundary components of S_i, S_j that contain the endpoints of α bound an annulus $A \supset \alpha$ on S. Then $S' = S_i \cup S_j \cup A$, the annulus A is parallel in M to component $A' \subset \mathrm{Fr}(B)$ that is an annulus with the same boundary curves as A, and C is the union of the interval bundle B, the solid torus bounded by A, A', and a trivial interval bundle over $\partial_+ M \setminus S'$. We can then extend f^k to C by letting it be the given extension of $f^k|_{S_i \cup S_j}$ on B, the identity on the solid torus, and the obvious fibre preserving extension of $f^k|_{\partial_+ M \setminus S'}$ to the adjacent interval bundle. The case that the boundary components of S_i, S_j that contain the endpoints of α do not bound an annulus on S is similar, except that instead of the solid torus above we take a thickened disk bounded by a rectangular neighbourhood of α on $S \setminus (S_i \cup S_j)$, and a rectangular neighbourhood of the homotopic arc on the frontier of B.

We now work on the "only if" direction. Passing to a power, suppose that f extends to a nontrivial subcompression body C of (M, S). We may assume that there is no proper, f-invariant essential subsurface $S' \subset S$ such that $f|_{S'}$ extends to a nontrivial subcompression body of (M, S'). If there were such a subsurface S', we could replace S by a minimal such S', therefore reducing the argument to the minimal case we are assuming we are in above.

If $f = \mathrm{id}$ on S, the fact that there is a nontrivial subcompression body of (M, S) means there is a meridian in $S = S_{\mathrm{id}}$, so we're in case (1) and are done. This case may seems silly, but observe that if f is some complicated pure homeomorphism where there's a meridian in S_{id}, the associated "minimal" case that we pass to in the previous paragraph is where S is an annular neighbourhood of some such meridian, and $f = \mathrm{id}$ on S.

Assume from now on that f is not the identity map of S.

We claim that every meridian $m \in \mathcal{D}(C, S)$ intersects every component of λ. Indeed, suppose some λ_i is disjoint from some such m and let S' be the component of $S \setminus S_i$ containing m. Since f extends to C and $S' \subset \partial C_+$ is f^k-invariant, f extends to the characteristic subcompression body C' of the pair (C, S'), defined by starting with $\partial_+ M$ and compressing all meridians of C that lie in S'; see Fact 2.4. Since m is a meridian in C', this C' is nontrivial, which contradicts the minimality assumption in the first paragraph.

Pick a meridian $m \in \mathcal{D}(C, S)$. Since m intersects all components of λ, the sequence of meridians $m_i := f^i(m)$ Hausdorff converges to a lamination λ' strongly commensurable to λ. Applying Theorem 7.2 to the pair (C, S) and using that all meridians in C intersect all components of λ, we have that either

- some component λ_a is an intrinsic limit of meridians lying in $S(\lambda_a)$, in which case Theorem 8.1 (applied to $f_a : S_a \to S_a$) says we're in case (2), or
- there are indices a, b such that S_a, S_b bound an essential interval bundle $B \subset C$, where λ_a, λ_b are essentially homotopic in B, and where there is a compression arc $\alpha \subset S$ for B that is disjoint from λ, and hence can be isotoped so that its interior lies in S_{id}.

Let's assume we're in the last case, since otherwise we're done. We want to show that some power of $f_a \cup f_b : \partial_H B \to \partial_H B$ extends to B.

First, suppose that B is a twisted interval bundle, so that $S_a = S_b$, $f_a = f_b$, $\lambda_a = \lambda_b$. Using just the index a from now on, if σ is the canonical involution of B, then Fact 2.16 implies that $\sigma(\lambda_a)$ is isotopic to λ_a. Let $A \subset \mathcal{T}(S)$ be the axis of f_a on the Teichmüller space $\mathcal{T}(S)$. By Theorem 12.1 of [22] and Theorem 2 of [45], we have that A, $\sigma(A)$ are asymptotic, so as they are both pseudo-Anosov axes they must be equal by discreteness of the action of the mapping class group. Since σ has finite order, it then fixes A pointwise. Now the group $\Gamma = \langle \sigma, f_a \rangle \subset \mathrm{Mod}(\partial_H B)$ is isomorphic to the direct product of a finite group fixing A pointwise and a cyclic group of pseudo-Anosovs, so for some positive k we have $\sigma \circ f_a^k = f_a^k$ in $\mathrm{Mod}(\partial_H B)$. By Theorem 3 of [4] we may isotope f_a^k, σ so that they commute, while preserving that $\sigma^2 = \mathrm{id}$; we can then alter the bundle map $\pi : B \to Y$ so that the new σ is still the canonical involution. It follows that f_a^k is a lift to $\partial_H B$ of a pseudo-Anosov map $g : Y \to Y$, and hence f_a^k extends to B as desired, see Fact 2.7.

Next, assume B is a trivial interval bundle, with canonical involution σ that switches S_a, S_b. As in the previous paragraph, we have that $\sigma(\lambda_b)$ is isotopic to λ_a, so $\Gamma = \langle f_a, \sigma \circ f_b \circ \sigma^{-1} \rangle \subset \mathrm{Mod}(S_a)$ is a direct product $\Gamma = F \times \langle \phi \rangle$ of a finite group F and a cyclic group generated by a pseudo-Anosov ϕ, where if we quotient by F then f_a and $\sigma \circ f_b \circ \sigma^{-1}$ both project to positive powers of ϕ. It suffices to show that they project to the *same* positive power of ϕ, for then we are done by the same argument as in the previous paragraph. For this, recall that all meridians of C intersect S_a, S_b, so these surfaces are "holes" for the disk set of C, as discussed in [49]. So with $m_i = f^i(m)$ the sequence of meridians in C constructed above, Lemma 12.20 of [49] says that for each i, the distance in the arc complex of S_a between

$$m_i \cap S_a = f_a^i(m \cap S_a) \quad \text{and} \quad \sigma(m_i \cap S_b) = (\sigma \circ f_b \circ \sigma^{-1})^i \sigma(m \cap S_b)$$

is at most 6. However, if f_a and $\sigma \circ f_b \circ \sigma^{-1}$ project to different positive powers of ϕ, their stable translation lengths on the arc complex of S_a are different, which is a contradiction. □

Acknowledgements

The authors would like to thank Jeff Brock, Juan Souto and Sebastian Hensel for useful conversations, and also the referee, who helped improve the accuracy and readability of the paper. Biringer was partially supported by NSF grant DMS-1308678.

References

[1] **I Agol**, *Tameness of hyperbolic* 3*-manifolds* (2004) arXiv math/0405568

[2] **I Biringer**, **J Johnson**, **Y Minsky**, *Extending pseudo-Anosov maps into compression bodies*, J. Topol. 6:4 (2013) 1019–1042 MR

[3] **I Biringer**, **N G Vlamis**, *Automorphisms of the compression body graph*, J. Lond. Math. Soc., II. Ser. 95:1 (2017) 94–114

[4] **J S Birman**, **H M Hilden**, *On isotopies of homeomorphisms of Riemann surfaces*, Ann. of Math. (2) 97 (1973) 424–439 MR

[5] **F Bonahon**, *Cobordism of automorphisms of surfaces*, Ann. Sci. École Norm. Sup. (4) 16:2 (1983) 237–270 MR

[6] **B H Bowditch**, *Relatively hyperbolic groups*, Internat. J. Algebra Comput. 22:3 (2012) art. id. 1250016 MR

[7] **J F Brock**, *Continuity of Thurston's length function*, Geom. Funct. Anal. 10:4 (2000) 741–797 MR

[8] **J F Brock**, *Iteration of mapping classes and limits of hyperbolic* 3*-manifolds*, Invent. Math. 143:3 (2001) 523–570 MR

[9] **J F Brock**, **R D Canary**, **Y N Minsky**, *The classification of Kleinian surface groups, II: The ending lamination conjecture*, Ann. of Math. (2) 176:1 (2012) 1–149 MR

[10] **J Brock**, **Y Minsky**, **H Namazi**, **J Souto**, *Bounded combinatorics and uniform models for hyperbolic* 3*-manifolds*, J. Topol. 9:2 (2016) 451–501 MR

[11] **D Calegari**, **D Gabai**, *Shrinkwrapping and the taming of hyperbolic* 3*-manifolds*, J. Amer. Math. Soc. 19:2 (2006) 385–446 MR

[12] **R D Canary**, *Algebraic convergence of Schottky groups*, Trans. Amer. Math. Soc. 337:1 (1993) 235–258 MR

[13] **R D Canary**, *Ends of hyperbolic* 3*-manifolds*, J. Amer. Math. Soc. 6:1 (1993) 1–35 MR

[14] **R D Canary**, *A covering theorem for hyperbolic* 3*-manifolds and its applications*, Topology 35:3 (1996) 751–778 MR

[15] **J W Cannon**, **C D Feustel**, *Essential embeddings of annuli and Möbius bands in* 3*-manifolds*, Trans. Amer. Math. Soc. 215 (1976) 219–239 MR

[16] **A J Casson**, **S A Bleiler**, *Automorphisms of surfaces after Nielsen and Thurston*, London Mathematical Society Student Texts 9, Cambridge Univ. Press (1988) MR

[17] **A J Casson**, **C M Gordon**, *Reducing Heegaard splittings*, Topology Appl. 27:3 (1987) 275–283 MR

[18] **A J Casson**, **D D Long**, *Algorithmic compression of surface automorphisms*, Invent. Math. 81:2 (1985) 295–303 MR

[19] **M T Clay**, **C J Leininger**, **J Mangahas**, *The geometry of right-angled Artin subgroups of mapping class groups*, Groups Geom. Dyn. 6:2 (2012) 249–278 MR

[20] **V Erlandsson**, **M Gendulphe**, **I Pasquinelli**, **J Souto**, *Mapping class group orbit closures for non-orientable surfaces*, Geom. Funct. Anal. 33:3 (2023) 637–693 MR

[21] **B Farb**, **D Margalit**, *A primer on mapping class groups*, Princeton Mathematical Series 49, Princeton Univ. Press (2012) MR

[22] **A Fathi**, **F Laudenbach**, **V Poenaru**, *Travaux de Thurston sur les surfaces*: *séminaire Orsay*, Astérisque 66-67, Soc. Math. France, Paris (1979) MR

[23] **V Gadre**, *The limit set of the handlebody set has measure zero* (2011) arXiv 1105.2980

[24] **U Hamenstädt**, *Train tracks and the Gromov boundary of the complex of curves*, from "Spaces of Kleinian groups" (Y N Minsky, M Sakuma, C Series, editors), London Math. Soc. Lecture Note Ser. 329, Cambridge Univ. Press (2006) 187–207 MR

[25] **U Hamenstädt**, *Hyperbolic relatively hyperbolic graphs and disk graphs*, Groups Geom. Dyn. 10:1 (2016) 365–405 MR

[26] **U Hamenstädt**, *Asymptotic dimension and the disk graph, I*, J. Topol. 12:2 (2019) 658–673 MR

[27] **U Hamenstädt**, *Asymptotic dimension and the disk graph, II*, J. Topol. 12:3 (2019) 675–684 MR

[28] **G C Hruska**, *Relative hyperbolicity and relative quasiconvexity for countable groups*, Algebr. Geom. Topol. 10:3 (2010) 1807–1856 MR

[29] **W Jaco**, *Lectures on three-manifold topology*, CBMS Regional Conference Series in Mathematics 43, Amer. Math. Soc., Providence, RI (1980) MR

[30] **W H Jaco**, **P B Shalen**, *Seifert fibered spaces in* 3*-manifolds*, Mem. Amer. Math. Soc. 220, Amer. Math. Soc., Providence, RI (1979) MR

[31] **W Jeon**, **I Kim**, **K Ohshika**, **C Lecuire**, *Primitive stable representations of free Kleinian groups*, Israel J. Math. 199:2 (2014) 841–866 MR

[32] **K Johannson**, *Homotopy equivalences of* 3*-manifolds with boundaries*, Lecture Notes in Mathematics 761, Springer (1979) MR

[33] **M Kapovich**, *Hyperbolic manifolds and discrete groups*, Progress in Mathematics 183, Birkhäuser, Boston, MA (2001) MR

[34] **S P Kerckhoff**, *The measure of the limit set of the handlebody group*, Topology 29:1 (1990) 27–40 MR

[35] **I Kim**, **C Lecuire**, **K Ohshika**, *Convergence of freely decomposable Kleinian groups*, Invent. Math. 204:1 (2016) 83–131 MR

[36] **I Kim**, **M Lee**, *Separable-stable representations of a compression body*, Topology Appl. 206 (2016) 171–184 MR

[37] **E Klarreich**, *The boundary at infinity of the curve complex and the relative Teichmüller space*, Groups Geom. Dyn. 16:2 (2022) 705–723 MR

[38] **G Kleineidam**, **J Souto**, *Algebraic convergence of function groups*, Comment. Math. Helv. 77:2 (2002) 244–269 MR

[39] **M Lackenby**, *Attaching handlebodies to* 3*-manifolds*, Geom. Topol. 6 (2002) 889–904 MR

[40] **C Lecuire**, *Structures hyperboliques convexes sur les variétés de dimension* 3, PhD thesis, École normale supérieure de Lyon (2004) Available at https://theses.fr/2004ENSL0273

[41] **C Lecuire**, *An extension of the Masur domain*, from "Spaces of Kleinian groups" (Y N Minsky, M Sakuma, C Series, editors), London Math. Soc. Lecture Note Ser. 329, Cambridge Univ. Press (2006) 49–73 MR

[42] **C Lecuire**, *Plissage des variétés hyperboliques de dimension* 3, Invent. Math. 164:1 (2006) 85–141 MR

[43] **D D Long**, *Discs in compression bodies*, Pacific J. Math. 122:1 (1986) 129–146 MR

[44] **J F Manning**, *The Bowditch boundary of* $(G, \mathcal{H})$ *when* G *is hyperbolic* (2015) arXiv 1504.03630

[45] **H Masur**, *Uniquely ergodic quadratic differentials*, Comment. Math. Helv. 55:2 (1980) 255–266 MR

[46] **H Masur**, *Measured foliations and handlebodies*, Ergodic Theory Dynam. Systems 6:1 (1986) 99–116 MR

[47] **H A Masur**, **Y N Minsky**, *Geometry of the complex of curves, I*: *Hyperbolicity*, Invent. Math. 138:1 (1999) 103–149 MR

[48] **H A Masur**, **Y N Minsky**, *Quasiconvexity in the curve complex*, from "In the tradition of Ahlfors and Bers, III" (W Abikoff, A Haas, editors), Contemp. Math. 355, Amer. Math. Soc., Providence, RI (2004) 309–320 MR

[49] **H Masur**, **S Schleimer**, *The geometry of the disk complex*, J. Amer. Math. Soc. 26:1 (2013) 1–62 MR

[50] **K Matsuzaki**, **M Taniguchi**, *Hyperbolic manifolds and Kleinian groups*, Clarendon, New York (1998) MR

[51] **J D McCarthy**, *Normalizers and centralizers of pseudo-anosov mapping classes*, preprint (1982) Available at https://users.math.msu.edu/users/mccarthy/publications/normcent.pdf

[52] **Y Minsky**, *The classification of Kleinian surface groups, I: Models and bounds*, Ann. of Math. (2) 171:1 (2010) 1–107 MR

[53] **M Mj**, *Cannon–Thurston maps for Kleinian groups*, Forum Math. Pi 5 (2017) art. id. e1 MR

[54] **J W Morgan**, *On Thurston's uniformization theorem for three-dimensional manifolds*, from "The Smith conjecture" (New York, 1979) (J W Morgan, H Bass, editors), Pure Appl. Math. 112, Academic, Orlando, FL (1984) 37–125 MR

[55] **H Namazi**, **J Souto**, *Heegaard splittings and pseudo-Anosov maps*, Geom. Funct. Anal. 19:4 (2009) 1195–1228 MR

[56] **J-P Otal**, *Courants géodesique et surfaces*, PhD thesis, Université Paris-Sud, Orsay (1989) Available at https://theses.fr/1989PA112051

[57] **P Scott**, *Strong annulus and torus theorems and the enclosing property of characteristic submanifolds of 3-manifolds*, Quart. J. Math. Oxford Ser. (2) 35:140 (1984) 485–506 MR

[58] **W P Thurston**, *The geometry and topology of three-manifolds*, lecture notes, Princeton University (1979) Available at https://url.msp.org/gt3m

[59] **W P Thurston**, *Hyperbolic Structures on 3-manifolds, III: Deformations of 3-manifolds with incompressible boundary* (1998) arXiv math/9801058

[60] **G S Walsh**, *The bumping set and the characteristic submanifold*, Algebr. Geom. Topol. 14:1 (2014) 283–297 MR

IAN BIRINGER ian.biringer@bc.edu
Department of Mathematics, Boston College, Chestnut Hill, MA, United States

CYRIL LECUIRE cyril.lecuire@ens-lyon.fr
ENS de Lyon site Monod, UMPA UMR 5669 CNRS, Lyon, France

Received: June 6, 2024 Revised: February 16, 2025

ALGEBRAIC & GEOMETRIC TOPOLOGY

msp.org/agt

EDITORS

PRINCIPAL ACADEMIC EDITORS

John Etnyre	Vesna Stojanoska
etnyre@math.gatech.edu	vesna@illinois.edu
Georgia Institute of Technology	University of Illinois at Urbana-Champaign

BOARD OF EDITORS

Julie Bergner	University of Virginia jeb2md@eservices.virginia.edu	Daniel Isaksen	Wayne State University isaksen@math.wayne.edu
Steven Boyer	Université du Québec à Montréal cohf@math.rochester.edu	Thomas Koberda	University of Virginia thomas.koberda@virginia.edu
Tara E Brendle	University of Glasgow tara.brendle@glasgow.ac.uk	Markus Land	JGU Mainz mland@uni-mainz.de
Indira Chatterji	CNRS & Univ. Côte d'Azur (Nice) indira.chatterji@math.cnrs.fr	Christine Lescop	Université Joseph Fourier lescop@ujf-grenoble.fr
Octav Cornea	Université' de Montreal cornea@dms.umontreal.ca	Norihiko Minami	OCAMI (Osaka Central Adv. Math. Inst.) norihikominami@gmail.com
Alexander Dranishnikov	University of Florida dranish@math.ufl.edu	Andrés Navas	Universidad de Santiago de Chile andres.navas@usach.cl
Tobias Ekholm	Uppsala University, Sweden tobias.ekholm@math.uu.se	Jessica S Purcell	Monash University jessica.purcell@monash.edu
Mario Eudave-Muñoz	Univ. Nacional Autónoma de México mario@matem.unam.mx	Birgit Richter	Universität Hamburg birgit.richter@uni-hamburg.de
David Futer	Temple University dfuter@temple.edu	Jérôme Scherer	École Polytech. Féd. de Lausanne jerome.scherer@epfl.ch
John Greenlees	University of Warwick john.greenlees@warwick.ac.uk	Zoltán Szabó	Princeton University szabo@math.princeton.edu
Matthew Hedden	Michigan State University mhedden@math.msu.edu	Maggy Tomova	University of Iowa maggy-tomova@uiowa.edu
Kristen Hendricks	Rutgers University kristen.hendricks@rutgers.edu	Daniel T Wise	McGill University, Canada daniel.wise@mcgill.ca
Hans-Werner Henn	Université Louis Pasteur henn@math.u-strasbg.fr	Lior Yanovski	Hebrew University of Jerusalem lior.yanovski@gmail.com
Kathryn Hess	École Polytechnique Féd. de Lausanne kathryn.hess@epfl.ch		

See inside back cover or msp.org/agt for submission instructions.

The subscription price for 2026 is US $795/year for the electronic version, and $1170/year (+$80, if shipping outside the US) for print and electronic. Subscriptions, requests for back issues and changes of subscriber address should be sent to MSP. Algebraic & Geometric Topology is indexed by Mathematical Reviews, Zentralblatt MATH, Current Mathematical Publications and the Science Citation Index.

Algebraic & Geometric Topology (ISSN 1472-2747 printed, 1472-2739 electronic) is published 9 times per year and continuously online, by Mathematical Sciences Publishers, 2000 Allston Way # 59, Berkeley, CA 94701-4004. Periodical rate postage paid at Oakland, CA 94615-9651, and additional mailing offices. POSTMASTER: send address changes to Mathematical Sciences Publishers, 2000 Allston Way # 59, Berkeley, CA 94701-4004.

AGT peer review and production are managed by EditFlow® from MSP.

PUBLISHED BY

mathematical sciences publishers

nonprofit scientific publishing

https://msp.org/

© 2026 Mathematical Sciences Publishers

ALGEBRAIC & GEOMETRIC TOPOLOGY

Volume 26 Issue 5 (pages 1597–1963) 2026