

THE GEOMETRY OF TOTALLY GEODESIC SUBVARIETIES OF MODULI SPACES OF RIEMANN SURFACES

FRANCISCO ARANA-HERRERA AND ALEX WRIGHT

ABSTRACT. We prove a semisimplicity result for the boundary, in the corresponding Deligne–Mumford compactification, of a totally geodesic subvariety of a moduli space of Riemann surfaces. At the level of Teichmüller space, this semisimplicity theorem gives that each component of the boundary is a product of simple factors, each of which behaves metrically like a diagonal embedding. Building on this result, we also show that the associated totally geodesic submanifolds of Teichmüller space and orbifold fundamental groups are hierarchically hyperbolic.

The proof intertwines in a novel way results and perspectives originating in dynamics, algebraic geometry, geometric group theory, and both classical and modern Teichmüller theory. It establishes both new rigidity and new flexibility for totally geodesic submanifolds and their associated varieties and orbifold fundamental groups and provides a rich set of new tools for the study of these objects.

CONTENTS

1. Introduction	2
2. Cylinders via horocycle and geodesic flows	9
3. The boundary is totally geodesic	16
4. The boundary is semisimple	23
5. The boundary is smooth and irreducible	27
6. First consequences of semisimplicity	29
7. From Teichmüller spaces to non-annular curve graphs	31
8. From Teichmüller spaces to annular curve graphs	33
9. Hierarchical hyperbolicity	35
10. The proof of Theorem 1.6 and open questions	42
Appendix A. Annular projections using cylinders (after Rafi)	43
Appendix B. Electrified Teichmüller geodesics	46
References	51

1. INTRODUCTION

1.1. Initial context. Let $\mathcal{T}_{g,n}$ be the Teichmüller space of genus g Riemann surfaces with n punctures or marked points, and let

$$\pi : \mathcal{T}_{g,n} \rightarrow \mathcal{M}_{g,n}$$

be the map to the associated moduli space of Riemann surfaces $\mathcal{M}_{g,n}$. We also use the same notation for the associated maps on the bundles of quadratic differentials.

Teichmüller space admits several metrics of interest, but here we consider exclusively the Teichmüller metric. This metric reflects the modular nature of Teichmüller space, giving a concrete measure of how non-conformal a map between two Riemann surfaces must be, and also the intrinsic complex geometry of Teichmüller space as a complex manifold, being equal to the Koyabashi metric [Roy71]. It is in a sense as inhomogeneous as possible [Roy71] and its fine geometry is often mysterious.

Teichmüller discs, also known as complex geodesics, are holomorphic isometric embeddings of the hyperbolic plane into Teichmüller space. There is a unique Teichmüller disc through any pair of distinct points, and one might say that the mystery of Teichmüller and moduli spaces can sometimes be pierced with Teichmüller discs.

Here we study totally geodesic submanifolds of Teichmüller space. We say a complex submanifold N of Teichmüller space is totally geodesic if N contains the Teichmüller disc¹ through every pair of distinct points in N .

The one complex dimensional totally geodesic submanifolds are exactly the Teichmüller discs. We call a complex manifold or variety higher dimensional if it has dimension greater than one.

Starting with smaller Teichmüller spaces and taking covering constructions in the sense of [MMW17, Section 6], one can obtain higher dimensional totally geodesic submanifolds of larger Teichmüller spaces, which we call trivial. Recently, non-trivial, higher dimensional examples were discovered for the first time [MMW17, EMMW20].

A totally geodesic submanifold $N \subseteq \mathcal{T}_{g,n}$ is called algebraic if $\pi(N) \subseteq \mathcal{M}_{g,n}$ is a closed algebraic variety; in this case $\pi(N)$ is called a totally geodesic subvariety. The non-trivial examples mentioned above are all algebraic, and were discovered via their associated totally geodesic subvarieties. Later, the second author established the following [Wri20].

Theorem 1.1 (Wright). *Every higher dimensional totally geodesic submanifold is algebraic. There are only finitely many higher dimensional totally geodesic subvarieties in each moduli space.*

For a typical Teichmüller disc $N \subseteq \mathcal{T}_{g,n}$, its projection $\pi(N)$ is dense in $\mathcal{M}_{g,n}$, and each moduli space $\mathcal{M}_{g,n}$ of complex dimension greater than one has infinitely many totally geodesic subvarieties of dimension one, typically known as Teichmüller curves. So the previous theorem establishes a very strong contrast between the cases of dimension one and higher dimension.

Already in [MMW17] it was suggested that the example constructed there might be seen as “a new type of Teichmüller space, on an equal footing with $\mathcal{T}_{g,n}$ ”. Indeed, we believe that that higher dimensional totally geodesic submanifolds are as

¹Equivalently, the (real) Teichmüller geodesic.

complicated as, and behave very much like, whole Teichmüller spaces. We elaborate on this as follows:

Metaconjecture 1.2. *Many results known for Teichmüller spaces also hold for higher dimensional totally geodesic submanifolds; many results known for moduli spaces also hold for higher dimensional totally geodesic subvarieties; and many results known for mapping class groups also hold for the orbifold fundamental groups of higher dimensional totally geodesic subvarieties.*

Theorem 1.5 below can be seen as an instance of this metaconjecture.

1.2. New results. In this paper we make use of augmented Teichmüller space $\overline{\mathcal{T}}_{g,n}$, which is obtained from Teichmüller space by adding nodal surfaces where collections of simple closed curves have been pinched. A more descriptive but less common name for augmented Teichmüller space is the Deligne–Mumford bordification of Teichmüller space, and indeed the quotient of $\overline{\mathcal{T}}_{g,n}$ by the corresponding mapping class group is the Deligne–Mumford compactification $\overline{\mathcal{M}}_{g,n}$ of moduli space. See [Bai06, Section 5.2] for an expository introduction.

The Deligne–Mumford bordification is stratified and each stratum is a product of Teichmüller spaces. A complex submanifold L of a product of Teichmüller spaces is called simple if its projection to each factor is an isometric embedding whose image is a totally geodesic submanifold; see the next paragraph for details. A complex submanifold L of a product of Teichmüller spaces is called semisimple if it is a product of simple factors.

In the definition of *simple* above we use the sup metric on the product of Teichmüller spaces. The isometric embedding condition can be rephrased without referencing a metric on the product by saying that given two points $(X_1, \dots, X_k)$ and $(Y_1, \dots, Y_k)$ in L we have $d(X_i, Y_i) = d(Y_j, X_j)$ for all $i, j \in \{1, \dots, k\}$. In this sense, the isometric embedding condition indicates that L looks metrically like a diagonally embedded copy of a subset of $\mathcal{T}_{g',n'}$ in $(\mathcal{T}_{g',n'})^k$.

Our first main result is the following, where $\overline{N} \subseteq \overline{\mathcal{T}}_{g,n}$ denotes the closure of $N \subseteq \mathcal{T}_{g,n}$ in the Deligne–Mumford bordification.


Theorem 1.3. *Let $N \subseteq \mathcal{T}_{g,n}$ be an algebraic totally geodesic submanifold of Teichmüller space and let L be the intersection of $\overline{N}$ with a stratum of $\overline{\mathcal{T}}_{g,n}$. Then L is semisimple and algebraic.*

The main point of Theorem 1.3 is that L is semisimple; the algebraicity is not hard and is included for completeness. Implicit in Theorem 1.3 is that L is a connected complex manifold, or in other words it is irreducible and smooth.

Example 1.4. There is a trivial totally geodesic submanifold $N \subseteq \mathcal{T}_3$ that can be defined as the fixed point set of an involution in the mapping class group, and which consists of unbranched double covers of genus two surfaces. Consider such a cover $X \rightarrow Y$ and let α be a separating simple closed curve on Y whose preimage on X has two components; see Figure 1. Pinching this preimage multi-curve gives rise to a component

$$L \subseteq \mathcal{T}_{1,1} \times \mathcal{T}_{1,1} \times \mathcal{T}_{1,2}$$

of $\overline{N}$. If $L_1 \subseteq \mathcal{T}_{1,1} \times \mathcal{T}_{1,1}$ is the diagonal and $L_2 \subseteq \mathcal{T}_{1,2}$ is the set of $(Z, \{p, q\})$ where the two marked points $\{p, q\}$ are exchanged by an appropriate involution of Z , we have $L = L_1 \times L_2$. Since L_1 and L_2 are each simple, L is semisimple.

FIGURE 1. The cover $X \rightarrow Y$ in Example 1.4.

Our second main result is the following, whose proof uses Theorem 1.3.

Theorem 1.5. *Let $N \subseteq \mathcal{T}_{g,n}$ be an algebraic totally geodesic submanifold of Teichmüller space. Then N is a hierarchically hyperbolic space and the associated orbifold fundamental group is a hierarchically hyperbolic group.*

By “the associated orbifold fundamental group” we mean the stabilizer of N in the mapping class group, or equivalently the orbifold fundamental group of the normalization of $\pi(N) \subseteq \mathcal{M}_{g,n}$.

Roughly speaking, Theorem 1.5 indicates that we can use curve graphs of subsurfaces as a coarse coordinate system, allowing us to study these objects combinatorially. To do this we only use subsurfaces that appear as nodal subsurfaces in $\bar{N}$. We define two subsurfaces to be equivalent if they occur in the same simple factor, and, roughly speaking, we use one subsurface from each equivalence class. In N the shape of an active subsurface roughly determines that of all equivalent subsurfaces, and, in contrast, the shapes of appropriate nonequivalent subsurfaces are independent. This gives both new rigidity phenomenon (equivalent subsurfaces) and new flexibility (independent subsurfaces).

Theorem 1.5 also encodes a large amount of information and structure on which strata of the Deligne-Mumford bordification $\bar{N}$ intersects.

Independent work. Benirschke, Dozier, and Rached [BDR24] have independently and simultaneously obtained a version of Theorem 1.3 (but without smoothness or irreducibility). Their analysis involves significant use of Beltrami differentials, which do not appear in our work, and conversely most of our main tools do not appear in their work. There is however some overlap, in particular at a key step of our analysis in Section 4.2.

Overview of motivation. Theorem 1.3 implies that totally geodesic subvarieties have remarkably understandable closures in the corresponding Deligne-Mumford compactification. This could prove to be very useful for the dynamical and algebro-geometric study of these subvarieties.

Theorem 1.5 implies that a number of previous results for Teichmüller spaces, some mentioned below, now apply to clarify the coarse geometry of totally geodesic submanifolds.

Theorem 1.5 also establishes the objects in question as potentially clarifying examples in the theory of subgroups of mapping class groups and in the theory of hierarchical hyperbolicity. For instance, the fundamental groups of higher dimensional totally geodesic subvarieties may serve as prototypical examples in the search for a useful general definition of geometrically finite subgroups of mapping class groups, beyond the special case of parabolically geometrically finite subgroups, and some of the ideas behind Theorem 1.5 may provide a prototype for understanding

certain nice subsets of hierarchically hyperbolic spaces which are not hierarchically quasi-convex. See [DDLS24, Definition 1.10] and [RST23] for context.

There is also a significant potential for our results to provide new tools in the study of the classification problem of higher dimensional totally geodesic subvarieties. We will comment on this in more detail below, after discussing the relationship to invariant subvarieties of differentials and saying more about the known examples.

Invariant subvarieties of differentials. There is a $GL(2, \mathbb{R})$ action on the bundle of quadratic differentials (non-zero, holomorphic away from marked points, and with at most simple poles at marked points) over $\mathcal{M}_{g,n}$, and deep results imply that its orbit closures are complex algebraic varieties [EMM15, EM18, Fil16]. Let ρ be the natural forgetful map defined by $\rho(X, q) := X$. As we recall in Remark 3.6 below, if $\mathcal{M}$ is a closed invariant subvariety of quadratic differentials, then $\rho(\mathcal{M})$ is a totally geodesic subvariety if

$$\dim(\mathcal{M}) = 2 \dim \rho(\mathcal{M}),$$

and all totally geodesic subvarieties arise in this way.

In general there seems to be no reason to expect $\dim(\mathcal{M}) = 2 \dim \rho(\mathcal{M})$ when $\dim(\mathcal{M}) > 2$. All that is known at present is that if $\mathcal{M}$ intersects the principle stratum of quadratic differentials this condition is automatically satisfied [KW22, Proposition 5.1], and if $\mathcal{M}$ is not transverse to the isoperiodic foliation this condition is never satisfied [Wri20, Theorem 1.3].

Recent examples. Recently, a collection of six new invariant subvarieties of quadratic differentials were discovered [MMW17, EMMW20], all of dimension four. Of these, three satisfy the condition $\dim(\mathcal{M}) = 2 \dim \rho(\mathcal{M})$ and hence give rise to totally geodesic subvarieties in moduli space. These are complex surfaces in $\mathcal{M}_{1,3}$, $\mathcal{M}_{1,4}$, and $\mathcal{M}_{2,1}$.

Hopes for classification. It is a mystery how many higher dimensional totally geodesic subvarieties of moduli space remain to be discovered. Given how few are known at present, one might hope for a complete classification.

One approach to this problem is as follows: First, classify all invariant subvarieties of quadratic differentials, then, determine which give rise to totally geodesic subvarieties of moduli space. The first step is a major goal in its own right and has seen major progress in recent years; see [McM07, AN20, Api18, EFW18, AW23, Ygo23, Win24] and the surveys [McM23, Wri15b] for a very incomplete selection of results. That being said, the problem of finding a complete classification of all invariant subvarieties of quadratic differentials remains extremely challenging and seems to us unlikely to be completed anytime soon.

This raises the question of if it is possible to approach the classification of totally geodesic subvarieties of moduli space in a new way, making more use of the totally geodesic assumption and not merely the general features of the associated invariant subvariety of quadratic differentials. We believe that the answer to this question is yes, and the results we present in this paper provide new tools for this purpose.

In particular, Theorem 1.3 allows for inductive arguments, but this may well be only the beginning. One might dream that Theorem 1.3 could lead to new insights related to the fact that the Teichmüller metric stores an enormous amount of useful information. Recall for example that complex linear isometric embeddings $Q(Y) \rightarrow Q(X)$ of spaces of quadratic differentials must arise from maps $X \rightarrow Y$

of the Riemann surfaces [GG22, Theorem 1.3] and that isometric embeddings of Teichmüller spaces must arise from covering constructions [BS23].

Putting aside such uncertain dreams, we wish to give a deeper explanation of how our results seem likely to lead to major improvements of existing techniques for the classification problem. Let us start with the Cylinder Deformation Theorem of the second author [Wri15a], which has been generalized in [BDG22], and is now a primary tool in the classification problem. Given an invariant subvariety $\mathcal{M}$, and a differential $(X, q) \in \mathcal{M}$, one defines an equivalence relation on cylinders of q , with two cylinders being equivalent if they are parallel and remain parallel on nearby differentials in $\mathcal{M}$. The Cylinder Deformation Theorem states that equivalence classes of cylinders can be deformed while remaining in $\mathcal{M}$. The rigidity provided by the equivalence relation is in tension to the flexibility provided by the deformations, giving the first indication of the potential usefulness of this theorem.

One of our intermediate results allows us to consider simple closed curves on surfaces rather than cylinders on quadratic differentials, raising the possibility of easier and more topological analogues of arguments involving the Cylinder Deformation Theorem. Let $N \subseteq \mathcal{T}_{g,n}$ be an algebraic totally geodesic submanifold. We define QN to be the set of quadratic differentials generating Teichmüller discs completely contained in N . Note that $\pi(QN)$ is an algebraic variety [EMM15, EM18, Fil16]. It is not hard to show QN is irreducible, and [BS23, Proposition 3.1] shows the harder fact that it is a holomorphic vector bundle over N . We define $Q_{\max}N$ to be the largest stratum of QN ; see Section 2 for details. We say that a simple closed curve is a cylinder at a quadratic differential if it is homotopic to the core curve of a cylinder on that differential.

Theorem 1.6. *Suppose $N \subseteq \mathcal{T}_{g,n}$ is an algebraic totally geodesic submanifold. Let $\mathcal{S}$ be the set of simple closed curves on the marking surface that are cylinders at some point of $Q_{\max}N$ and define two such curves to be equivalent if there is a point of $Q_{\max}N$ where they are equivalent cylinders.*

- (1) *For every $X \in N$ and every $\alpha \in \mathcal{S}$ there is a quadratic differential $(X, q) \in QN$ where α is a cylinder.*
- (2) *If $\alpha \in \mathcal{S}$ is a cylinder at $(X, q) \in QN$, then the equivalent curves in $\mathcal{S}$ are also cylinders at (X, q) and these cylinders are equivalent. The ratios of circumferences and the ratios of moduli of these cylinders depends only on their core curves and not on (X, q) .*
- (3) *For any two simple closed curves $\alpha, \beta \in \mathcal{S}$, there is a quadratic differential in QN where they are both cylinders.*
- (4) *Unions of disjoint equivalence classes in $\mathcal{S}$ are exactly the multi-curves that can be pinched in $\bar{N}$.*
- (5) *For any two equivalence classes of simple closed curves in $\mathcal{S}$, the boundary of the subsurface they fill is a union of equivalence classes of $\mathcal{S}$.*

In addition to this result, our proof of Theorem 1.5 extends the definition of equivalent simple closed curves to a definition of equivalent subsurfaces. The structure consisting of all these curves and subsurfaces and the equivalence relation on them is vaguely reminiscent of the structure of a root system and may be useful for the classification problem.

Additional context on totally geodesic submanifolds. See [Gou21] for a recent survey on totally geodesic submanifolds written to accompany a Bourbaki seminar on the subject.

The three special examples discussed above are remarkable from many points of view, sharing connections with billiards in quadrilaterals and classical algebraic geometry, and have already been the topic of further study [TT20, Che22].

Totally geodesic submanifolds of Teichmüller space which are isometric to covering constructions are known to arise from covering constructions [BS23].

There are no higher dimensional analogues of Teichmüller discs which are symmetric spaces [Ant17c]. The holomorphicity in the definition of Teichmüller discs is essentially redundant [Ant17b]. See also [Ant17a].

The core of Theorem 1.5. A series of papers beginning with seminal work of Masur and Minsky developed the idea that curve graphs can be used as coarse coordinate systems for mapping class groups and Teichmüller spaces; see [MM99, MM00, BKMM12, Raf07a] for some of the key entries in this story. In particular, there is a map from Teichmüller space to the product of the curve graphs of all subsurfaces whose image is characterized by remarkably simple consistency inequalities, each involving only two subsurfaces at a time; the fact that the consistency inequalities characterize the image is called the Consistency Theorem [BKMM12] or the Realization Theorem [BHS19]. The so-called Distance Formula allows one to coarsely compute distances in Teichmüller space using the image of this map.

Roughly speaking, Theorem 1.5 indicates that we can pick out a collection of subsurfaces that are distinguished enough so that the image of N in the corresponding product of curve graphs is characterized by similar consistency inequalities, and large enough so that the Distance Formula still holds. It would be easy to maintain the Distance Formula if we used all subsurfaces, but this would cause the Realization Theorem to fail. Conversely, it would be easy to understand the image if we used only a very small collection of subsurfaces, but using too small of a collection would cause the Distance Formula formula to fail. To prove Theorem 1.5 we must use exactly the right collection, allowing us to coarsely specify points in N using their images in curve graphs uniquely and without unnecessary redundant information.

Additional context on hierarchical hyperbolicity. The relevant structure required to have a version of the Realization Theorem and the Distance Formula was axiomatized in influential work of Behrstock, Hagen, and Sisto [BHS17b, BHS19]. We recommend the survey [Sis19] for an introduction. The resulting structure, called a hierarchically hyperbolic structure, is not generic, but is enormously powerful: Even when given a space with nice non-positively curved behavior, one typically does not expect it to be hierarchically hyperbolic, but if it is hierarchically hyperbolic then one gains a remarkably good understanding of the space. Despite the stringent requirements, many spaces and groups have been shown to be hierarchically hyperbolic [BHS17b, BHS19, BR20, Che20, HMS24, HRSS22, BR22, Vok22, DDLS24, Rus21, Hug22, NQ24, CR24, Kop23].

Other perspectives on hierarchically hyperbolicity exist beyond the coordinate system point of view we highlighted above. For instance, while hyperbolic spaces satisfy that hulls of finite subsets are coarsely trees, hierarchically hyperbolic spaces satisfy that hulls of finite subsets are coarsely CAT(0) cube complexes [BHS21,

DMS20, BHS17a]. Furthermore, hierarchically hyperbolic spaces are hyperbolic up to products of simpler hierarchically hyperbolic spaces [BHS17a].

Consequences of hierarchical hyperbolicity, sometimes with extra restrictions satisfied in our examples, include: bounds on asymptotic dimension [BHS17a], classification of maximal dimensional quasi-flats [BHS21], coarse injectivity [HHP23], semi-hyperbolicity [HHP23, DMS20], and being quasi-cubical [Pet21].

The specific way in which Theorem 1.5 holds has consequences like the fundamental group being undistorted in the mapping class group, which in the case of Teichmüller curves was proven in [Tan21] and in the case of examples arising from covering constructions corresponds to [RS09, Theorem 9.1].

Some comments on the proofs. Our analysis relies on and strengthens a profitable marriage between coarse geometry and results of a more analytic, algebro-geometric, and dynamical flavor. In particular, we make use of theorems of Rafi expressing how the coarse geometry of the Teichmüller metric affects the fine geometry of Teichmüller geodesics. Even our proof of Theorem 1.3, a statement in fine, non-coarse geometry, is intimately linked with coarse geometry.

Our inspiration for why Theorem 1.3 should be true was, in the first place, the structure theorem [CW21, Theorem 1.3] on the boundary of invariant subvarieties of Abelian differentials, and, secondly, the rather natural idea that Theorem 1.5 should not only be true at a coarse level but also be witnessed by fine geometry.

Having proved Theorem 1.3, the main difficulty in proving Theorem 1.5 is obtaining a coarse statement analogous to Theorem 1.3 but for curve graphs. This is surprisingly non-trivial to obtain from first principles but is well within the reach of existing technology. It requires the proof of a result of independent interest that does not seem to have been previously recorded in the literature, namely that the electrification of the real line parameterizing a Teichmüller geodesic along the intervals where curves are short quasi-isometrically embeds into the curve graph of the surface.

Conventions. We emphasize that we work only with complex submanifolds. All dimensions we give are over $\mathbb{C}$.

When we state a result related to the Teichmüller space $\mathcal{T}_{g,n}$, we allow all constants to implicitly depend on g and n without further comment; so for the purposes of such constants one should consider the Teichmüller space $\mathcal{T}_{g,n}$ as fixed.

Acknowledgments. We thank Janusz Adamus, Matt Bainbridge, Dawei Chen, Howard Masur, Jesús Ruiz, Jacob Russell, Edson Sampaio, Saul Schleimer, Carlos Serván, Robert Tang, Sam Taylor, and Scott Wolpert for helpful conversations and comments. We are especially grateful to Kasra Rafi for generously teaching us the results in Appendix A and for clarifying a number of his results to us.

As mentioned above, Fred Benirschke, Ben Dozier, and John Rached have independently obtained a version of the first of our two main results [BDR24]; we thank them for discussing their work with us and for offering helpful comments on our work.

The second author was partially supported by NSF Grants DMS 1856155 and 2142712 and a Sloan Research Fellowship.

2. CYLINDERS VIA HOROCYCLE AND GEODESIC FLOWS

2.1. Main results. This section lays the foundation for the rest of the paper and contains multiple results and constructions that will be used later on.

Let $N \subseteq \mathcal{T}_{g,n}$ be an algebraic totally geodesic submanifold of Teichmüller space. Denote by QN the set of quadratic differentials generating Teichmüller discs completely contained in N (together with the zero differentials on surfaces in N) and let $Q_{\max}N$ denote the largest stratum of QN ; this is well defined since $\pi(QN)$ is an irreducible variety. The set $Q_{\max}N$ is a lift of an invariant subvariety $\pi(Q_{\max}N)$ of quadratic differentials in moduli space.

We say that a simple closed curve (on the marking surface) is the core curve of a cylinder on a quadratic differential if it is freely homotopic to the core curve of that cylinder. We say that a simple closed curve is a cylinder curve for QN if it is the core curve of a cylinder at some point of QN . Note that, since QN is the closure of $Q_{\max}N$, every cylinder curve for QN is the core curve of a cylinder at some point of $Q_{\max}N$. One of the main technical results of this section is the following.

Theorem 2.1. *Let $N \subseteq \mathcal{T}_{g,n}$ be an algebraic totally geodesic submanifold of Teichmüller space. For any cylinder curve α of QN , the subset of $Q_{\max}N$ where α is the core curve of a cylinder is path connected.*

Following [Wri15a], given $q \in Q_{\max}N$, we say two cylinders on q are $Q_{\max}N$ -parallel if they are parallel on every quadratic differential in a sufficiently small neighborhood of q in $Q_{\max}N$. We say that two cylinder curves are N -equivalent if there is a point of $Q_{\max}N$ where they are the core curves of $Q_{\max}N$ -parallel cylinders; if N is clear from context we sometimes omit it from the terminology. The significance of Theorem 2.1 is that it allows us to prove the following result, which will be the foundation for much of our later discussion.

Corollary 2.2. *Let $N \subseteq \mathcal{T}_{g,n}$ be an algebraic totally geodesic submanifold of Teichmüller space and let α, β be N -equivalent cylinder curves. Then, at every point q of QN where α is a cylinder, β is also a cylinder. If $q \in Q_{\max}N$, then the associated cylinders are $Q_{\max}N$ -parallel. Furthermore, the ratios of moduli, heights, and circumferences of the associated cylinders are constant depending only on α, β , and N , and, in the case of moduli, such constant is rational.*

Notice that Corollary 2.2 guarantees the terminology N -equivalent is appropriate in the sense that it defines an equivalence relation among cylinder curves.

Another crucial result of this section is Lemma 2.13, which guarantees that the ratio of hyperbolic lengths of N -equivalent cylinder curves is bounded when the curves have bounded length.

2.2. Deducing the corollary. To deduce Corollary 2.2 from Theorem 2.1 we will make use of the following result.

Lemma 2.3. *Let $N \subseteq \mathcal{T}_{g,n}$ be an algebraic totally geodesic submanifold of Teichmüller space. The ratio of moduli, heights, and circumferences of different $Q_{\max}N$ -parallel cylinders is locally constant, i.e., given a point q of $Q_{\max}N$ and a pair of $Q_{\max}N$ -parallel cylinders on it, on any sufficiently small deformation in $Q_{\max}N$, these ratios do not change. Furthermore, the ratios of moduli are rational.*

Proof. That the ratio of circumferences stays constant is an immediate consequence of the definitions [Wri15a]. For the ratio of the moduli, one has to first recall [Wri20,




FIGURE 2. Horocycle arcs joining pairs of points in Teichmüller space. Arrows represent positive time directions.

Theorem 1.3], which says that $Q_{\max}N$ is the lift of a particular kind of invariant subvariety, more precisely, it has no rel, and then invoke [MW17, Corollary 1.6], which gives the desired result. (See also [AW24, Lemma 6.10] for a different proof of the result powering [MW17, Corollary 1.6].) $\square$

We are now ready to deduce Corollary 2.2 from Theorem 2.1.

Proof of Corollary 2.2. Let α and β be N -equivalent cylinder curves and let q be a point of $Q_{\max}N$ where they are core curves of $Q_{\max}N$ -parallel cylinders. Let $q' \in Q_{\max}N$ be any point where α is a cylinder. Theorem 2.1 gives a path in $Q_{\max}N$ from q to q' such that α is a cylinder at every point on this path. By Lemma 2.3, we get that β must also be a cylinder at every point on this path. This argument also gives that the ratios of moduli, heights, and circumferences only depend on α, β , and N for differentials in $Q_{\max}$.

Because QN is the closure of $Q_{\max}N$ and because all the quantities involved are continuous, the ratios are the same on QN as they are on $Q_{\max}N$. $\square$

2.3. Cylinders via horocycle flow. To prove Theorem 2.1 we will need different methods to produce quadratic differentials with cylinders. We first discuss a method that relies on the Teichmüller horocycle flow

$$u_t = \begin{pmatrix} 1 & t \\ 0 & 1 \end{pmatrix}.$$

Given $X \in \mathcal{T}_{g,n}$, denote by $Q(X)$ the set of quadratic differentials on X .

Lemma 2.4. *For any two distinct points $X, Y \in \mathcal{T}_{g,n}$, there exists a unique, up to scale, $q \in Q(X)$ and a unique time $t > 0$, both depending continuously on X and Y , such that $u_t(X, q) = (Y, q')$ for some quadratic differential q' , and this q generates the Teichmüller disc through X and Y . If both X and Y are in a totally geodesic submanifold N , then $q \in QN$.*

Proof. There exists a unique positive Teichmüller horocycle arc joining X to Y ; see Figure 2. Furthermore, this arc belongs to the unique Teichmüller disc through X and Y . Any quadratic differential q as in the statement generates this Teichmüller disc. $\square$

In the setting of Lemma 2.4 we say that the quadratic differential q generates the horocycle from X to Y . We next recall the following.

Lemma 2.5. *Let $X \in \mathcal{T}_{g,n}$ and let $\alpha = \sum k_i \alpha_i$ be a multi-curve with all k_i positive integers. Let $D_\alpha := \prod D_{\alpha_i}^{k_i}$ be the Dehn twist in α . Then the differential that generates the horocycle from X to $D_\alpha X$ is horizontally periodic and the core curves of the horizontal cylinders are the α_i . If M_i is the modulus of the cylinder with core curve α_i , then $M_i/M_j = k_i/k_j$.*

Proof. For any choice of positive numbers m_i there exists a, unique up to scale, differential q on X which is horizontally periodic with core curves α_i of modulus cm_i for some $c = c(X) > 0$ [Str84, Theorem 21.7]. Here we use $m_i = k_i$. It takes time $1/M$ for horocycle flow to accomplish a Dehn twist on a horizontal cylinder of modulus M , so let $t := k_i/(cm_i) = 1/c$. Then we get that $u_t q$ is a differential on $D_\alpha X$. By Lemma 2.4, q generates the desired horocycle. $\square$

Next, we produce Dehn twists that stabilize a totally geodesic submanifold.

Lemma 2.6. *Let $N \subseteq \mathcal{T}_{g,n}$ a totally geodesic submanifold and let $\{\alpha_1, \dots, \alpha_r\}$ be the core curves of an equivalence class of $Q_{\max} N$ -parallel cylinders. Let $k_1, \dots, k_r$ be positive integers such that the α_i are cylinder core curves on a differential in QN with the corresponding moduli M_i satisfying $M_i/M_k = k_i/k_j$. If $\alpha := \sum k_i \alpha_i$, then $D_\alpha N = N$.*

Morally, this lemma is true because D_α is in the fundamental group of $\pi(N)$ and that fundamental group stabilizes N , but we find it easier to give a careful proof from a slightly different perspective.

Proof. There is an open subset of QN of quadratic differentials where the α_i are core curves of cylinders. The Cylinder Deformation Theorem then gives that $D_\alpha^{\pm 1}$ applied to this subset is contained in QN . By analytic continuation, we deduce $D_\alpha^{\pm 1} QN \subseteq QN$. Thus, $D_\alpha QN = QN$. $\square$

We summarize this discussion as follows.

Corollary 2.7. *Let $N \subseteq \mathcal{T}_{g,n}$ be a totally geodesic submanifold. Suppose that at some point of $Q_{\max} N$ there is an equivalence class of $Q_{\max} N$ -parallel cylinders with core curves $\alpha_1, \dots, \alpha_r$ and moduli $ck_1, \dots, ck_r$ for some $c > 0$ and positive integers k_i . Then, for any $X \in N$, there exists a unique up to scale horizontally periodic quadratic differential q on X with core curves α_i with moduli $c'k_1, \dots, c'k_r$ for some $c' > 0$. This q depends continuously on X and lies in QN .*

Proof. Lemma 2.6 gives a multi-twist D_α stabilizing N . Using this and Lemma 2.5 one obtains the desired differential; it lies in QN by Lemma 2.4. $\square$

2.4. Annular curve graphs. There are two versions of annular curve graphs. One is used in the hierarchically hyperbolic structure of mapping class groups and is quasi-isometric to $\mathbb{R}$. The other one is used in the hierarchically hyperbolic structure of Teichmüller spaces and is isometric to a horoball in the hyperbolic plane.

For the annular region corresponding to a simple closed curve α , the mapping class group version keeps track of the twist along α . We think of the Teichmüller space version as mapping to

$$\{x + iy : y \geq 1\} \subseteq \mathbb{H},$$

with the x coordinate again encoding the twist along α and the y coordinate being equal to $\max(1, 1/\ell_\alpha(X))$, where $\ell_\alpha(X)$ is the hyperbolic length of α at $X \in \mathcal{T}_{g,n}$. In particular, we note that the mapping class group version can be thought of as the x coordinate of the Teichmüller space version.

If desired, we can view $\mathbb{R}$ as the boundary of the horoball with its intrinsic path metric, and consider the projection $x + iy \mapsto x + i$ from the horoball to its boundary. We can think of the mapping class group version of the map to the annular curve graph as the composition of the Teichmüller space version together with this projection. The fact that the projection from a horoball to its boundary is highly non-Lipschitz offers some hint that using the mapping class group version of the annular curve graph when discussing Teichmüller spaces requires extra care.

The mapping class group version has been standard since the work of Masur-Minsky [MM00, Section 2.4] and the Teichmüller space version was foreshadowed in Minsky's product region theorem [Min96], is implicit in formulas of Rafi [Raf07a], and is explicit in subsequent work [EMR17, Dur16]. We suggest [EMR17, Section 3.1] as a reference on this topic.

By default we will use the Teichmüller space version and we will comment when using the mapping class group version.

2.5. Cylinders via geodesic flow. Next, we discuss a coarser method for producing quadratic differentials with cylinders, via the Teichmüller geodesic flow. Let $d_\alpha^{\text{MCG}}(X, Y)$ denote the distance between the subsurface projections of $X, Y \in \mathcal{T}_{g,n}$ in the mapping class group version of the annular curve graph of α and let $\ell_\alpha(X)$ denote the hyperbolic length of the geodesic representative of α on $X \in \mathcal{T}_{g,n}$. Roughly speaking, the following result guarantees the existence of cylinders whenever there is large enough twist.

Lemma 2.8. *There exists a constant $D > 0$ depending only on g and n such that for any $X, Y \in \mathcal{T}_{g,n}$ and every simple closed curve α , if*

$$d_\alpha^{\text{MCG}}(X, Y) > D(1 + \ell_X(\alpha)^{-1} + \ell_Y(\alpha)^{-1}),$$

then the quadratic differential generating the Teichmüller geodesic segment from X to Y has a cylinder with core curve α .

Proof. This can be thought of as a black box coming from work of Rafi, but we briefly outline the structure of the argument.

Denote by ν^+ and ν^- the vertical and horizontal foliations of the quadratic differential q generating the Teichmüller geodesic segment from X to Y . Since the ensuing argument gives a lower bound on the modulus of the cylinder, we may take an arbitrarily small perturbation of X or Y , allowing us to assume that α has non-zero intersection number with the $\nu^\pm$.

As we recall briefly in Appendix A, Rafi has an alternative definition of twisting along a Teichmüller geodesic which makes use of the flat metric of the generating quadratic differential [Raf07a, Section 4]. By [Raf07a, Theorem 4.3], it is possible to compare differences in this alternative twisting to $d_\alpha^{\text{MCG}}(X, Y)$ with an error of the form $O(\ell_X(\alpha)^{-1} + \ell_Y(\alpha)^{-1})$. Hence, in our situation, there is large change in Rafi's alternative twist between X and Y . Formulas in [Raf07a, Section 4], which are recalled in Appendix A, imply the existence of a large modulus cylinder (at some point between X and Y). $\square$

2.6. Hyperbolic lengths. Later we will need to estimate hyperbolic lengths of short curves. We recall some of the necessary background for this now.

Lemma 2.9. *Let q be a horizontally periodic quadratic differential such that the ratios of circumferences of horizontal cylinders is at most M . Then, there exists a uniform bound, depending only on M (and, as always, g and n), for the modulus of the expanding annuli corresponding to the core curves of the horizontal cylinders of q .*

Expanding annuli appear in work of Minsky and Rafi [Min96, Raf07b] and are now a basic tool. Their definition admits slight variations which are not important here, but, for concreteness, let us say we follow [Raf14, Section 3], so the inner boundary is the flat geodesic representative.

Proof. Notice that, for any such expanding annulus of large modulus, there exists an embedded horizontal segment crossing it that is of size proportional to the size of the outer boundary. The circumference of one of the cylinders under consideration is as big as this horizontal segment. On the other hand, the length of the inner boundary is by definition a circumference in the case under consideration. Thus, the desired bound follows from the fact that, as recalled in [Raf14, Section 3], the moduli of expanding annuli can be calculated as the logarithm of the ratio of the lengths of the inner and outer boundaries. $\square$

The ratio assumption in Lemma 2.9 is crucial, as seen by considering a one by one square torus with a small circumference cylinder glued in along a horizontal slit to give a surface in $\mathcal{H}(2)$.

We deduce the following consequence of Lemma 2.9; here we denote the modulus of a cylinder F by $\text{Mod}(F)$.

Corollary 2.10. *There exists $\delta > 0$ (depending only on g and n) such that for all $M > 0$ there exists an $E > 0$ such that the following holds. Let (X, q) be a horizontally periodic quadratic differential such that the ratio of circumferences of horizontal cylinders on q is at most M . Suppose α is the core curve of a horizontal cylinder F on q with $\ell_\alpha(X) \leq \delta$. Then,*

$$1/E \leq \ell_\alpha(X) \cdot \text{Mod}(F) \leq E.$$

Proof. The ratio of the inverse extremal length of α and the sum of $\text{Mod}(F)$ and the moduli of the corresponding expanding annuli is uniformly bounded; see [Raf14, Theorem 3.1, Statement (2)] for a convenient reference and [Min92, Theorem 4.6] for the original. Lemma 2.9 gives that the moduli of the expanding annuli are bounded, and, in particular, if $\delta > 0$ is small enough, this implies that F has large modulus.

With the assumption that $\ell_\alpha(X) \leq \delta$, hyperbolic and extremal lengths have bounded ratios [Mas85, Proposition 1, Corollary 3], so the result follows. $\square$

2.7. Proof of path connectivity. We are now ready to prove Theorem 2.1.

Proof of Theorem 2.1. Suppose α is the core curve of a cylinder C on $(X, q) \in Q_{\max}N$. Denote by $[C]$ the $Q_{\max}N$ -parallelism class of C . Suppose also that α is the core curve of a cylinder D on $(X', q') \in Q_{\max}N$. Our goal is to produce a path connecting q to q' in $Q_{\max}N$ along which α is always a cylinder.

Consider the reparametrized Teichmüller geodesic $\mathcal{G}: [0, 1] \rightarrow N$ from X to X' . After rotating q and q' , we can assume without loss of generality that the cylinders C and D are horizontal. Applying the Teichmüller horocycle flow to q and q' for long enough times yields Riemann surfaces $Y \in N$ and $Y' \in N$ whose projections to the annular curve graph of α are coarsely equal and very far from the projection of any point on $\mathcal{G}$. Furthermore, because the Teichmüller horocycle flow is strongly nondivergent [MW02], it is possible to arrange for $X, Y, X', Y' \in N$ to be ϵ -thick for some $\epsilon := \epsilon(X, Y) > 0$ independent of the constant implicit in the notion of “very far” above.




FIGURE 3. Dotted lines are Teichmüller horocycles and solid straight lines are Teichmüller geodesics.

Let $\mathcal{F}: [0, 1] \rightarrow N$ denote the reparametrized Teichmüller geodesic from X' to Y' . Now, by Corollary 2.7, for every $Y'' \in \mathcal{F}$ we can find a continuously varying horizontally periodic quadratic differential on Y'' whose horizontal cylinder core curves are exactly those corresponding to cylinders in $[C]$. Using the Cylinder Deformation Theorem over these differentials together with Corollary 2.10, it is possible to produce a continuous path $\mathcal{F}': [0, 1] \rightarrow N$ with the same endpoints as $\mathcal{F}$ and such that for every $Y'' \in \mathcal{F}'$ it holds that $\ell_\alpha(Y'') > \epsilon'$ and that the projection of Y'' to the annular curve graph of α is very close to that of X' and Y' . Here ϵ' is a positive constant depending on ϵ and the constant in Corollary 2.10. See Figure 3 for a picture of this construction.

Now let $q_t: [0, 1] \rightarrow QN$ be the continuous one parameter family of quadratic differentials such that $q_t \in QN$ is the initial quadratic differential of the Teichmüller geodesic segment joining $\mathcal{G}(t) \in N$ to $\mathcal{F}'(t) \in N$. Since the Teichmüller horocycle flow preserves Teichmüller disc, q_0 can be obtained as a rotation of q and q_1 can be obtained as a rotation of q' . Originally C was defined as a cylinder on q , but, by abuse of notation, we use the same symbol to denote the corresponding cylinder on q_0 , since q and q_0 only differ by a rotation.

Recall that QN is a vector bundle over N . The non-maximal part $QN \setminus Q_{\max}N$ is the union of the intersections of QN with certain lower dimensional strata. It is thus possible (for example smoothing and using Sard's Theorem) to produce arbitrarily small deformations $\hat{q}_t: [0, 1] \rightarrow Q_{\max}N$ of the path q_t with image contained in $Q_{\max}N$. Since the endpoints of q_t are already in $Q_{\max}N$, we can additionally arrange for $\hat{q}_t$ to agree with q_t at the endpoints.

Now, by Lemma 2.8, assuming the path $\hat{q}_t$ is sufficiently close to the path q_t , we get that α must be a cylinder on $\hat{q}_t$ for all $t \in [0, 1]$. We have thus produced the desired path connecting q and q' in $Q_{\max}N$. $\square$

2.8. Short curves and cylinders. The following lemma and subsequent corollary emphasize the important connection between short curves and cylinders.

Lemma 2.11. *Let $N \subseteq \mathcal{T}_{g,n}$ be an algebraic totally geodesic submanifold. Then, there exists some $\epsilon > 0$ depending only on N such that if $\{\alpha_1, \dots, \alpha_r\}$ is a collection of disjoint simple closed curves that can be made simultaneously of hyperbolic length less than ϵ on some surface in N , then there exists a quadratic differential in QN where they are all core curves of cylinders.*

Proof. A compactness argument using the Deligne–Mumford compactification provides an $\epsilon > 0$ such that if $\{\alpha_1, \dots, \alpha_r\}$ is a collection of disjoint simple closed curves that can be made simultaneously of hyperbolic length less than ϵ on some surface in N , then the boundary of N intersects a stratum of the Deligne–Mumford bordification where all these α_i are pinched to nodes.

Working in moduli space rather than Teichmüller space, considering a holomorphic disc through such a boundary point, and using plumbing coordinates, one can check that N is stabilized by a multi-twist whose support includes the α_i ; see [BDG22, Proposition 3.1] and [Ben23, Proposition 7.6] for details. Corollary 2.7 then gives the result. Alternatively, after applying a large power of this multi-twist to any point in N and considering the corresponding quadratic differential, Lemma 2.8 gives the desired result. $\square$

Corollary 2.12. *Let $N \subseteq \mathcal{T}_{g,n}$ be an algebraic totally geodesic submanifold. Then, the core curves of cylinders of quadratic differentials in QN are exactly the curves that get arbitrarily short on N .*

Proof. The hyperbolic length of the core curve of any cylinder can be made arbitrarily small using the Teichmüller geodesic flow or the Cylinder Deformation Theorem. Conversely, if α is a curve that gets arbitrarily short on N , Lemma 2.11 gives that α is a cylinder curve for QN . $\square$

It will be very useful for our purposes to control the ratios of hyperbolic lengths of short curves, as guaranteed by the following lemma.

Lemma 2.13. *Let $N \subseteq \mathcal{T}_{g,n}$ be an algebraic totally geodesic submanifold. Suppose $\{\alpha_1, \dots, \alpha_r\}$ is an equivalence class of N -equivalent simple closed curves. Then, there exist constants $M_0 > 0$ and $C > 0$ depending only on N such that for every $i, j \in \{1, \dots, r\}$ and every $X \in N$, if $\min(\ell_{\alpha_i}(X), \ell_{\alpha_j}(X)) \leq M_0$, then*

$$1/C \leq \ell_{\alpha_i}(X)/\ell_{\alpha_j}(X) \leq C.$$

Proof. Corollary 2.7 produces a differential on X in QN with horizontal cylinders corresponding to the α_i and whose ratio of moduli do not depend on X . Corollary 2.10 allows the hyperbolic lengths to be estimated in terms of these moduli, keeping in mind that once one of the α_i gets short the corresponding modulus must be large, hence all the moduli must be large and all the other α_j must be short as well. $\square$

2.9. Simultaneous realization. For later use we highlight the following results.

Lemma 2.14. *Let $N \subseteq \mathcal{T}_{g,n}$ be an algebraic totally geodesic submanifold. Let α and β be cylinders at possibly different points of QN . Then there exists $q \in QN$ on which both α and β are cylinders.*

Proof. Fix a point $X \in N$ and find points $Y, Z \in N$ such that

$$\begin{aligned} \ell_\alpha(X) = \ell_\alpha(Y) \quad \text{and} \quad \ell_\beta(X) = \ell_\beta(Z), \\ d_\alpha(X, Y) > K \quad \text{and} \quad d_\beta(X, Z) > K, \end{aligned}$$

where $K > 0$ is a very large constant. This can be done, as in previous proofs, by applying large powers of appropriate multi-twists.

Now, suppose in first instance that

$$d_\alpha(Y, Z) > K/100 \quad \text{and} \quad d_\beta(Y, Z) > K/100.$$

Then, by Lemma 2.8, the quadratic differential generating the Teichmüller geodesic segment from Y to Z proves the desired result.

Thus, it suffices by symmetry to consider the case when $d_\alpha(Y, Z) \leq K/100$. In this case, $d_\alpha(X, Z) > K/2$ by the triangle inequality, so, again by Lemma 2.8, the quadratic differential generating the Teichmüller geodesic segment from X to Z proves the desired result. $\square$

Notice that, in the context of Lemma 2.14, it follows by Corollary 2.2 that all simple closed curves N -equivalent to α or β are core curves of cylinders on q .

Corollary 2.15. *Let $N \subseteq \mathcal{T}_{g,n}$ be an algebraic totally geodesic submanifold. Then, any two N -equivalence classes of simple closed curves are either pairwise disjoint or every curve of one intersects a curve of the other.*

Proof. Pick a curve in each equivalence class. Lemma 2.14 provides a quadratic differential $q \in QN$ where both curves are cylinders. Perturbing q , without loss of generality we can assume that q is in $Q_{\max}N$. Corollary 2.2 gives that all the N -equivalent curves are cylinder curves on q as well. Let $[C]$ and $[D]$ denote the corresponding equivalence classes of cylinders on q .

Using the Cylinder Deformation Theorem, we can deform $[C]$ in such a way that the circumference of every cylinder in $[D]$ that intersects a cylinder in $[C]$ changes. Since ratios of circumference of cylinders in $[D]$ are constant, this gives that if any cylinder of $[D]$ is intersected by a cylinder in $[C]$, then they all are. See [NW14, Proposition 3.2] for details. $\square$

3. THE BOUNDARY IS TOTALLY GEODESIC

3.1. Main definition and theorem. Throughout this section we consider a product of Teichmüller spaces of the form

$$\mathcal{T}_{g_1, n_1} \times \cdots \times \mathcal{T}_{g_s, n_s}.$$

We may sometimes use the Teichmüller metric on this space, which can be defined in the usual way and is equal to the sup of the Teichmüller metrics of the components. We will use π to denote the map to the corresponding product of moduli spaces.

Over these product spaces we consider the corresponding bundles of quadratic differentials holomorphic away from marked points and with at most simple poles at marked points; these differentials have finite area and could be zero on some (or all) components. There is an action of $GL^+(2, \mathbb{R})$ on this bundle which acts as usual on components where the differential is non-zero and leaves unchanged components where the differential is zero. We warn the reader that this action is not continuous at points where the differential is zero on at least one component.

We define a T-disc as the projection to $\mathcal{T}_{g_1, n_1} \times \cdots \times \mathcal{T}_{g_s, n_s}$ of the $GL^+(2, \mathbb{R})$ orbit of one of the quadratic differentials introduced above. This term of course stands for “Teichmüller disc”, but we have chosen to instead write “T-disc” to emphasize this is a non-standard definition given the multi-component setting. The T-disc defined using a quadratic differential q will be called the T-disc generated by q .

Remark 3.1. We highlight a few important aspects of T-discs:

- (1) The T-disc associated to a zero differential is a point.
- (2) The projection to each component $\mathcal{T}_{g_i, n_i}$ of any T-disc is either a point or a Teichmüller disc.
- (3) If $(X_1, \dots, X_s)$ and $(Y_1, \dots, Y_s)$ lie on single T-disc, then, for each pair i, j of coordinates where the disc is not constant, we have

$$d_{\mathcal{T}_{g_i, n_i}}(X_i, Y_i) = d_{\mathcal{T}_{g_j, n_j}}(X_j, Y_j).$$

- (4) Every T-disc is an isometrically and holomorphically embedded copy of the hyperbolic plane.
- (5) Not every isometrically and holomorphically embedded copy of the hyperbolic plane is a T-disc.
- (6) It is not true that there is a T-disc containing every pair of points in a product of Teichmüller spaces.
- (7) If q and q' are equal up to rescaling on each component by a non-zero factor depending on the component, then the T-discs they define are the same. In particular, when the number of components is at least 2, it is possible to have $\mathbb{C}q \neq \mathbb{C}q'$ but for q and q' to generate the same T-disc.

Given an irreducible complex analytic subset

$$L \subseteq \mathcal{T}_{g_1, n_1} \times \cdots \times \mathcal{T}_{g_s, n_s},$$

we define QL to be the set of quadratic differentials that generate T-discs completely contained in L .

Definition 3.2. An irreducible complex analytic subset $L \subseteq \mathcal{T}_{g_1, n_1} \times \cdots \times \mathcal{T}_{g_s, n_s}$ is said to be totally geodesic if there is an irreducible complex analytic subset $Q_{\text{alg}}L$ of the bundle of quadratic differentials over $\mathcal{T}_{g_1, n_1} \times \cdots \times \mathcal{T}_{g_s, n_s}$ such that

- (1) $Q_{\text{alg}}L \subseteq QL$ and $Q_{\text{alg}}L$ is $GL^+(2, \mathbb{R})$ invariant,
- (2) $\pi(Q_{\text{alg}}L) \subseteq Q\mathcal{M}_{g_1, n_1} \times \cdots \times Q\mathcal{M}_{g_s, n_s}$ is a variety, and
- (3) $Q_{\text{alg}}L$ contains a $\dim L$ dimensional vector subspace above every point of L .

Remark 3.3. Recall that a complex analytic subset is a subset that is locally defined near each point of the ambient space as the zero set of a finite collection of holomorphic functions. Such subsets are always closed, and we emphasize that this means $Q_{\text{alg}}L$ is required to be closed.

Remark 3.4. Later, after we understand their remarkable structure, in Lemma 5.1 we will see that totally geodesic analytic subsets are actually complex submanifolds, i.e., have no singularities. So the appearance of complex analytic subsets can be thought of as a technical detail internal to our proofs. The reader wishing to pay attention to this detail may want to keep in mind that there are complex analytic subsets which are topological submanifolds but not complex submanifolds; for example, consider the subset of $\mathbb{C}^2$ defined by $z_1^2 = z_2^3$.

Remark 3.5. The whole product space $\mathcal{T}_{g_1, n_1} \times \cdots \times \mathcal{T}_{g_s, n_s}$ is totally geodesic, but, as remarked above, it does not contain a T-disc through every pair of points. Thus, our definition does not satisfy all the conditions one might initially hope from a generalized definition of totally geodesic.

Remark 3.6. As is discussed in [Gou21, Proposition 1.3] and elsewhere, in the single component case, L is totally geodesic if and only if $\dim QL = 2 \dim L$. Since [BS23, Proposition 3.1] shows that in this case QL is a vector bundle, it is not hard to see that our definition agrees with the usual definition when $s = 1$; see Lemma 5.1 for why L is a complex submanifold and not merely a complex analytic subset.

Remark 3.7. In the single component case one does not always consider zero dimensional totally geodesic submanifolds, i.e., single points, but we note that, for our purposes, a single point will be considered as such.

Remark 3.8. In the multi-component case, QL is often bigger than $Q_{\text{alg}}L$. For example, QL is closed under scaling each component individually by positive real numbers, and need not be closed in the bundle of quadratic differentials. This behavior can already be seen for a diagonal.

Definition 3.2 is easy enough to satisfy that we can verify the following, which is the main result of this section.

Theorem 3.9. *Let $N \subseteq \mathcal{T}_{g,n}$ be an algebraic totally geodesic submanifold. Then any irreducible component of the intersection of its closure $\overline{N}$ with a stratum of the Deligne–Mumford bordification is an algebraic totally geodesic subset.*

Note that Definition 3.2 actually assumes algebraicity, so when in the phrase “algebraic totally geodesic subset” above the word “algebraic” is really just for emphasis and for consistency with other situations where algebraicity is not automatic.

We emphasize that Definition 3.2 should be thought of as a temporary tool, since ultimately in proving Theorem 1.3 we will obtain vastly stronger information.

3.2. The quadratic Hodge bundle. Following the notation of, for example, [CMS23], we consider the quadratic Hodge bundle over the Deligne–Mumford compactification $\overline{\mathcal{M}}_{g,n}$; see also [BCG⁺19b, Sections 3.1, 3.2]. This holomorphic vector bundle parameterizes quadratic differentials with at worst simple poles at the marked points and at worst double poles at the nodes, subject to residue matching condition at nodes; these differentials are allowed to be zero on some components. We also consider the pull-back of this bundle to the Deligne–Mumford bordification, leaving it to the reader to infer from context if we are referring to the lifted or the original version. (The pullback is a vector bundle, see for instance [Hat, Proposition 1.5], but since the bordification is not even locally compact it is not a holomorphic vector bundle).

The residue matching condition is equality of the coefficients of $1/z^2$ in the Laurent expansions of the differential on either side of the node; following the work cited above we call this coefficient the 2-residue. When the differential is viewed as a flat surface, having non-zero and matching 2-residues corresponds to having half-infinite cylinders on either side of the node with core curves that are parallel and of the same length.

Let $N \subseteq \mathcal{T}_{g,n}$ be an algebraic totally geodesic submanifold. Denote by $\overline{QN}$ the closure of QN in the quadratic Hodge bundle. Given two N -equivalent cylinder

curves α and β , denote by $c(\alpha, \beta)$ the circumference of a cylinder with core curve α divided by the circumference of a cylinder with core curve β ; this quantity is well defined by Corollary 2.2.

Lemma 3.10. *Let $N \subseteq \mathcal{T}_{g,n}$ be an algebraic totally geodesic submanifold and α, β be a pair of N -equivalent cylinder curves. Then, for any $q \in \overline{QN}$ where α and β are pinched to nodes n_α, n_β , the 2-residue of q at n_α is equal to $c(\alpha, \beta)$ times the 2-residue of q at n_β .*

Proof. This follows by continuity of the circumference of a cylinder and the fact that 2-residues give the circumferences of the corresponding infinite cylinders. $\square$

Corollary 3.11. *Let $N \subseteq \mathcal{T}_{g,n}$ be an algebraic totally geodesic submanifold and suppose $X \in \overline{N}$ is a point where exactly k classes of N -equivalent curves have been pinched to nodes. Then, the fiber of $\overline{QN}$ at X contains a linear subspace of dimension $d - k$ consisting of finite area differentials. Moreover any finite area differential in the fiber is contained in such a subspace.*

Proof. Recall from [BS23, Proposition 3.1] that QN is a vector bundle of rank $d = \dim N$. Thus, every fiber of $\overline{QN}$ is a union of d -dimensional vector subspaces W of differentials. Given such a W , Lemma 3.10 gives that the subspace W_0 of W where all 2-residues vanish has codimension at most k . This W_0 is exactly the finite area differentials in W . $\square$

We deduce the following corollary from [MW17, CW21].

Corollary 3.12. *Let V be a $GL^+(2, \mathbb{R})$ -invariant subvariety of non-zero quadratic differentials over $\mathcal{M}_{g,n}$. Then the closure of V in the finite area part of the quadratic Hodge bundle is $GL^+(2, \mathbb{R})$ invariant.*

Because the $GL^+(2, \mathbb{R})$ action on the closure of the quadratic Hodge bundle is not continuous, this is a non-trivial statement. Although restricting to the finite area part is not required, we will only need that case here.

Proof. For concreteness of notation we will prove invariance under the one parameter subgroup $\{u_t\}_{t \in \mathbb{R}}$ consisting of unipotent upper triangular matrices. This is sufficient because the closure of V in the finite area part of the quadratic Hodge bundle is easily seen to be invariant under multiplication by complex scalars, and complex scalar multiplication together with $\{u_t\}_{t \in \mathbb{R}}$ generate $GL^+(2, \mathbb{R})$.

Consider a sequence $(X_k, q_k) \in V$ that converges to (X_∞, q_∞) in the finite area part of the quadratic Hodge bundle. We show $u_t(X_\infty, q_\infty)$ is also in the closure of V . Without loss of generality, assume that all (X_k, ω_k) lie in a single stratum and in a single irreducible component of the intersection of V with that stratum.

Let (Y_k, ω_k) be the double cover of (X_k, q_k) . After passing to a subsequence, we can assume (Y_k, ω_k) converges in the Hodge bundle to a finite area limit $(Y_\infty, \omega_\infty)$. Let Y'_∞ (respectively X'_∞) be the union of the components of Y_∞ (respectively X_∞) where ω_∞ (respectively q_∞) is nonzero. Then, $(Y'_\infty, \omega_\infty)$ is a double cover of (X'_∞, q_∞) .

The construction above guarantees (Y_k, ω_k) converges to $(Y'_\infty, \omega_\infty)$ in the “What You See Is What You Get” partial compactification of the stratum. The main result of [MW17, CW21] gives that, for any deformation on the boundary in this partial compactification, a corresponding deformation can be done on (Y_k, ω_k) for

k sufficiently large. The corresponding deformation does not change the vanishing cycles, so, at an intuitive level, one should think that the deformation does not affect the part of Y_k where ω_k is converging to zero. (Results in [MW17, Section 9] relate the topology on the What You See Is What You Get to the topology on the Hodge bundle, and allow vanishing cycles to be related to subsurfaces where ω_k is converging to zero.)

Applying this to the u_t orbit of $(Y'_\infty, \omega_\infty)$, which in period coordinates is given by $\omega_\infty + t\text{Im}(\omega_\infty)$, shows that the path $\omega_k + t\text{Im}(\omega_\infty)$ is contained in the set of double covers of surfaces in V , for k sufficiently large and t sufficiently small. For fixed t small, this can be seen to converge to $u_t(Y_\infty, \omega_\infty)$. For more details, compare to [MW17, Section 9] and [CW21, Section 8]. This shows that for t small $u_t(X_\infty, q_\infty)$ is in the closure of V as desired. Since the closure is closed, proving the result for t sufficiently small depending on the limit point is sufficient. $\square$

3.3. Algebraicity. We will need the following basic algebraicity result.

Lemma 3.13. *Let V be a subvariety of $\mathcal{M}_{g,n}$ and $\hat{V}$ be an irreducible component of its preimage in $\mathcal{T}_{g,n}$. Let B be the closure of V in the Deligne–Mumford bordification intersected with a stratum of the bordification. Then, the image of B in the Deligne–Mumford compactification $\overline{\mathcal{M}}_{g,n}$ is a variety.*

Proof. We will use the well studied construction of Dehn space; see [Ber74] for details and [Bai06, Section 5.2] for a more recent introduction. We follow the notation of the later reference, so S is a set of curves that are pinched to obtain a stratum of the Deligne–Mumford bordification, $\text{Tw}(S)$ is the group generated by Dehn twists along curves of S , and $\text{Stab}(S)$ is the stabilizer of the set S in the whole mapping class group. Note that $\text{Tw}(S)$ is a normal subgroup of $\text{Stab}(S)$.

The Dehn space $\mathcal{D}_{g,n}(S)$ associated to S can be viewed as arising from a two step construction: first take Teichmüller space union the stratum of the Deligne–Mumford bordification corresponding to S and then quotient by $\text{Tw}(S)$. The key results we need are: the map from $\mathcal{D}_{g,n}(S)$ to $\overline{\mathcal{M}}_{g,n}$ is invariant under the induced action of $\text{Stab}(S)/\text{Tw}(S)$; stabilizers for this action are finite; and up to those finite stabilizers the map is a local homeomorphism.

We will also make use of the following topological fact: Given a sequence Z_n in Dehn space converging to a boundary point Z , for any choice of preimages Z'_n of the Z_n in Teichmüller space, the sequence Z'_n converges to the unique point Z' (in the appropriate stratum of the Deligne–Mumford bordification) that maps to Z . (One can keep in mind that the different preimages of Z_n all differ by $\text{Tw}(S)$, and applying elements of $\text{Tw}(S)$ does not effect convergence to points where the curves in S are pinched.)

Let $\tilde{V}$ be the image of $\hat{V}$ in Dehn space, so we have

$$\hat{V} \rightarrow \tilde{V} \rightarrow V.$$

We first note that the topological fact above implies that the boundary of $\hat{V}$ in the stratum corresponding to S covers the boundary of $\tilde{V}$. Thus it suffices to show that the image of the boundary of $\tilde{V}$ is a variety in $\overline{\mathcal{M}}_{g,n}$.

Consider $\mathcal{M}_{g,n}$ union the boundary stratum corresponding to S , and let $\mathcal{U}$ be a neighborhood of the boundary. Up to passing to a finite cover, this $\mathcal{U}$ can be chosen so that it can be parametrized by a product of moduli spaces (one for each component after the curves in S are pinched) and punctured discs (one for each

curve in S). Viewed as a complex analytic space, we can arrange for the intersection $\mathcal{U} \cap V$ to have finitely many irreducible components. The boundary of each of these components is a variety.

Let $\tilde{\mathcal{U}}$ be the component of the preimage of $\mathcal{U}$ in Dehn space whose closure intersects the boundary. This can be taken to be parametrized by a product of Teichmüller spaces (one for each component after the curves in S are pinched) and punctured discs (one for each curve in S).

Because $\tilde{V}$ is an irreducible component of the preimage of V in Dehn space and because the map from Dehn space to $\overline{\mathcal{M}}_{g,n}$ is locally a homeomorphism up to the action of finite groups, the image of $\tilde{V} \cap \tilde{\mathcal{U}}$ in $\mathcal{U}$ is equal to a finite union of components of $V \cap \mathcal{U}$, and the boundary of $\tilde{V}$ maps onto the union of the boundaries of these components. $\square$

We also note two corollaries of the analysis above, in whose proofs we continue to use the notation above.

Corollary 3.14. *Let V be a subvariety of $\mathcal{M}_{g,n}$ and $\hat{V}$ be an irreducible component of its preimage in $\mathcal{T}_{g,n}$. Then, up to the action of the stabilizer of $\hat{V}$ in the mapping class group, there are only finitely many orbits of strata in the Deligne–Mumford bordification that intersect the closure of $\hat{V}$.*

Proof. There are only finitely many mapping class group orbits of strata of the Deligne–Mumford bordification, so it suffices to prove finiteness in each such orbit.

Consider a stratum that intersects the closure of $\hat{V}$. Near that stratum, there must be a subset of $\hat{V}$ that maps onto one of the finitely many components of $\mathcal{U} \cap V$. Given a second stratum in the same mapping class group orbit where one similarly gets a subset mapping onto the same component, there must be a mapping class g taking a germ of one subset to a germ of the other. That mapping class must stabilize $\hat{V}$ because $\hat{V}$ and $g(\hat{V})$ are both irreducible varieties and share a germ, so $\hat{V} = g(\hat{V})$. $\square$

Corollary 3.15. *Let V be a subvariety of $\mathcal{M}_{g,n}$ and $\hat{V}$ be an irreducible component of its preimage in $\mathcal{T}_{g,n}$. Let S be a disjoint collection of curves. Let $\text{Stab}(\hat{V})$ denote the stabilizer of $\hat{V}$ in the mapping class group, and*

$$\text{Tw}_S(\hat{V}) = \text{Stab}(\hat{V}) \cap \text{Tw}(S).$$

Then the map from $\hat{V}/\text{Tw}_S(\hat{V})$ to $\mathcal{T}_{g,n}/\text{Tw}(S)$ is proper.

Proof. The image of this map is $\tilde{V}$. At each point the germ of $\tilde{V}$ is a finite union of irreducible germs. As in the proof of the previous corollary, if two germs of $\hat{V}$ at different points are related by a mapping class, that mapping class must stabilize $\hat{V}$. It follows that any point of $\tilde{V}$ has a compact neighborhood whose preimage in $\hat{V}/\text{Tw}_S(\hat{V})$ is compact. $\square$

3.4. Unpinching. To complete the proof of Theorem 3.9 we also need a result that allows us to unpinch nodes in a given equivalence class. We will deduce this as a corollary of the following result, which is stronger than needed at the moment but will also be useful later.

Lemma 3.16. *Let $N \subseteq \mathcal{T}_{g,n}$ be an algebraic totally geodesic submanifold of Teichmüller space. For any point X in the boundary of $\bar{N}$ it is possible to find a quadratic differential $q \in QN$ with a collection of disjoint equivalence classes of*

cylinders such that replacing each cylinder in this collection with a pair of half-infinite cylinders gives a differential on X .

Proof. Consider a sequence X_k in N converging to X . For each equivalence class of nodes of X we get an associated multi-twist and from Lemma 2.5 we also get horizontally periodic differentials η_k on X_k where the horizontal cylinders correspond to the given equivalence class. By Corollary 2.10, up to rescaling and possibly passing to a subsequence, these converge in the quadratic Hodge bundle to a differential on X with a pair of half-infinite cylinders at each node in the equivalence class.

Taking the sum of one such differential for each equivalence class of nodes we get a differential q_∞ on X with a pair of half-infinite cylinders at each node. (The differentials for each equivalence class should be produced using the same sequence X_k , so they will lie in the same linear space obtained as a limit of fibers of QN over a subsequence of the X_k). With q_∞ thus constructed in the boundary of QN , we can find a sequence $q_k \in QN$ converging to q_∞ .

Apply the Cylinder Deformation Theorem to shrink cylinders on q_k without twisting them to obtain differentials $q'_k \in QN$ where all the cylinders under consideration have modulus bounded above and below. The Deligne-Mumford bordification is not locally compact, and we cannot yet assume that the q'_k converge after passing to a subsequence.

Let S be the set of curves that are pinched on X . The Dehn space of S is locally compact, so we can pass to a subsequence to assume that the images of the q'_k in the Dehn space of S converge; this is possible because all the q'_k can be chosen to live on surfaces in a compact neighborhood around the image of X .

Using Corollary 3.15, we can find $g_k \in \text{Stab}(\hat{V}) \cap \text{Tw}(S)$ so that $g_k q'_k$ converge to a differential q . This q has the desired properties. $\square$

Corollary 3.17. *Let $N \subseteq \mathcal{T}_{g,n}$ be an algebraic totally geodesic submanifold and $X \in \overline{N}$. Then, for any N -equivalence class of nodes on X , there is a nearby point in $\overline{N}$ where these nodes are not pinched but all other nodes remain pinched.*

Proof. This follows by applying cylinder deformations to the differential q produced in the lemma. We can stretch all but one equivalence class out to become pairs of half-infinite cylinders, and stretch the remaining equivalence class to have large but finite modulus. $\square$

3.5. Conclusion of the proof. We can now complete the goal of this section.

Proof of Theorem 3.9. Consider the closure of N in the Deligne-Mumford bordification intersected with one of its strata where exactly k classes of N -equivalent curves have been pinched to nodes. By Lemma 3.13, this intersection covers a variety and is, in particular, a complex analytic subset. Let L be an irreducible component of this set.

Corollary 3.11 gives that over every point of L there is a linear subspace of the fiber of QN of dimension $d - k$ consisting of finite area differentials. Corollary 3.12 gives that this subspace is contained in QL .

We claim that L has dimension at most $d - k$. This is true because Corollary 3.17 guarantees it is contained at least k levels deep in a stratification of the boundary. (We of course expect that $\dim L = d - k$ but the verification of this fact will be postponed to the next section).

Let $Q'_{\text{alg}}L$ be the intersection of $\overline{QN}$ with the relevant part of the boundary (or rather, the union of the irreducible components of this intersection which lie over L). This satisfies all the properties required of $Q_{\text{alg}}L$ except that we have not proven it is irreducible. By Corollary 3.11, each fiber of $Q'_{\text{alg}}L$ is a union of linear subspaces, so in particular it is closed under the $\mathbb{C}$ action which rescales quadratic differentials (equally on all components). We obtain an associated complex analytic set $\mathbb{P}Q'_{\text{alg}}L$ by projectivising fibers, and this still covers a variety. The map $\mathbb{P}Q'_{\text{alg}}L \rightarrow L$ is proper, so the image of any irreducible component is an analytic subset [Lo91, page 290]. Hence irreducibility of L implies that one of the irreducible components of $Q'_{\text{alg}}L$ maps surjectively onto L , and we can define $Q_{\text{alg}}L$ to be that component. $\square$

4. THE BOUNDARY IS SEMISIMPLE

4.1. Main definitions and statement. The following are the two main definitions of this section.

Definition 4.1. Suppose $1 \leq \ell < s$ and

$$L_1 \subseteq \mathcal{T}_{g_1, n_1} \times \cdots \times \mathcal{T}_{g_\ell, n_\ell}, \quad L_2 \subseteq \mathcal{T}_{g_{\ell+1}, n_{\ell+1}} \times \cdots \times \mathcal{T}_{g_s, n_s}$$

are both subsets of products of Teichmüller spaces. Then we say that the set

$$L := L_1 \times L_2 \subseteq \mathcal{T}_{g_1, n_1} \times \cdots \times \mathcal{T}_{g_s, n_s}$$

is the product of L_1 and L_2 . We also say L is a product if this is true after permuting factors, and we use the natural extension of this definition allowing for more than two factors.

Definition 4.2. Following Definition 3.2, suppose

$$L \subseteq \mathcal{T}_{g_1, n_1} \times \cdots \times \mathcal{T}_{g_s, n_s}$$

is a totally geodesic complex analytic set. We say L is simple if the projection to any factor is an isometric embedding. We say L is semisimple if it is a product of simple totally geodesic complex analytic sets.

The main result of this section is the following.

Theorem 4.3. *Every algebraic totally geodesic complex analytic subset*

$$L \subseteq \mathcal{T}_{g_1, n_1} \times \cdots \times \mathcal{T}_{g_s, n_s}$$

is semisimple.

Recall that, according to Definition 3.2, the L here is assumed to be irreducible.

4.2. Primality and the proof of Theorem 4.3. Let

$$L \subseteq \mathcal{T}_{g_1, n_1} \times \cdots \times \mathcal{T}_{g_s, n_s}$$

be a totally geodesic complex analytic set. Fix a $Q_{\text{alg}}L \subseteq QL$ as given by Definition 3.2. We now set up notation to deal with an important technical detail.

Suppose $X \in L$, $V \subseteq QD(X)$ has dimension at least $\dim L$, and $V \subseteq Q_{\text{alg}}L$. We assume V is not contained in a larger subspace of the fiber of $Q_{\text{alg}}L$ above X .

The components of X are indexed by the numbers $\{1, \dots, s\}$. Let $I_Z \subseteq \{1, \dots, s\}$ correspond to the components of X on which all differentials of V are zero and I_R

correspond to the remaining components. Here “ Z ” stands for “zero” and “ R ” for “rest”. Note that $\{1, \dots, s\}$ is the disjoint union of I_Z and I_R . Let

$$Z : \prod_{i \in \{1, \dots, s\}} \mathcal{T}_{g_i, n_i} \rightarrow \prod_{i \in I_Z} \mathcal{T}_{g_i, n_i}, \quad R : \prod_{i \in \{1, \dots, s\}} \mathcal{T}_{g_i, n_i} \rightarrow \prod_{i \in I_R} \mathcal{T}_{g_i, n_i}$$

be the associated maps that forget factors.

Consider the slice of L defined by

$$S = \{X' \in L : Z(X') = Z(X)\}.$$

This is the subset of L of all surfaces which are equal to X in the coordinates where all differentials in V are zero. We also consider the image $R(S)$, which we call the reduced slice, and note that $R : S \rightarrow R(S)$ is a bijection.

Later we will see that all of the objects just introduced do not depend on the choice of X and V , but until then we keep in mind the possibility that these objects, including S , may depend on these choices.

Define QS to be the set of pairs $(X', q') \in Q_{\text{alg}}L$ with $X' \in S$ and q' zero on all components indexed by I_Z . Let $R(QS)$ be the corresponding collection of quadratic differentials on elements of $R(S)$. Again, note that the corresponding forgetful map, which by abuse of notation we denote by $R : QS \rightarrow R(QS)$, is also a bijection.

We emphasize that no information is lost by R in our context because we are restricting to points of L where the components forgotten by R are constant. Furthermore, we are restricting to differentials that are zero on the components that are forgotten. In our context, R has an inverse which simply “adds back” the $Z(X)$ that was forgotten. Our use of slices and the map R is a technical detail required to connect to previous work considering products of strata of differentials, where by definition the differentials are non-zero on all components.

Consider the typical stratum of a differential in V and let $V^{(0)} \subseteq V$ denote the elements of V in this stratum. Note that $V^{(0)}$ is open, connected, and dense in V . Consider the intersection of $R(QS)$ with the associated stratum after applying R . Let $\mathcal{M}$ be an irreducible component of this intersection that contains $R(V^{(0)})$.

All the constructions so far are fully compatible with the maps π from products of Teichmüller spaces to the associated products of moduli spaces (of Riemann surfaces or quadratic differentials), so we get that $\mathcal{M}$ covers a variety $\pi(\mathcal{M})$. By definition (and this is the point of our use of slices and forgetful maps) this $\pi(\mathcal{M})$ lies in a product of strata of *non-zero* quadratic differentials.

Lemma 4.4. *$\mathcal{M}$ and $\pi(\mathcal{M})$ are $GL^+(2, \mathbb{R})$ -invariant.*

Proof. Recall from Definition 3.2 that $Q_{\text{alg}}L$ is $GL^+(2, \mathbb{R})$ -invariant. By definition of the action, it follows that QS is invariant, and hence that $R(QS)$ is invariant as well. The action preserves strata, so it follows that the relevant intersection of $R(QS)$ with a stratum is invariant as well. The action on a stratum is real analytic, so it preserves irreducible components and we conclude that $\mathcal{M}$ and $\pi(\mathcal{M})$ are $GL^+(2, \mathbb{R})$ -invariant. $\square$

Recall from [CW21] that an invariant subvariety of a product of strata of quadratic differentials is called prime if it cannot be written as a product. Furthermore, every invariant subvariety of a product of strata of quadratic differentials can be written uniquely as a product of primes. We will make use of the following structure theorem; see [CW21, Theorem 1.3].

Theorem 4.5 (Chen–Wright). *In any prime invariant subvariety of a product of strata, the absolute periods in one component locally determine the absolute periods in any other component.*

This is stated in [CW21] in the case of Abelian differentials but it follows immediately that it also holds for quadratic differentials with the standard convention that period coordinates, and hence the notion of “absolute periods”, come from considering double covers where the quadratic differential pulls back to the square of an Abelian differential.

It follows that we get a decomposition of $\pi(\mathcal{M})$ as a product of primes and hence a corresponding decomposition of $\mathcal{M}$, which we denote

$$\mathcal{M} = \mathcal{M}_1 \times \cdots \times \mathcal{M}_k.$$

Recall that $\mathcal{M}$ contains $R(V^{(0)})$, which is an open dense subset of $R(V)$. Since V is not contained in a larger subspace, it follows that we get an associated decomposition of V ,

$$V = V_1 \oplus \cdots \oplus V_k,$$

where the V_i can be defined as the differentials of V supported on the same components as $\mathcal{M}_i$. We now make the following key observation:

Lemma 4.6. *Suppose $q' \in V_i$ is zero on one of the components of X corresponding to the factor $\mathcal{M}_i$. Then $q' = 0$.*

Proof. Suppose otherwise. Pick some arbitrary $q \in V^{(0)}$. Adding small multiples of q' to q gives rise to a path in $\mathcal{M}_i$ where multiples of a fixed non-zero differential are added in some components and some other components remained unchanged. This causes area to change in some components but not others. (For example one can consider large multiples instead of small to let a differential go to infinity and see that the area cannot be constant.) Since area can be computed from absolute periods we see that absolute periods change in some components but not all, contradicting the fact that $\mathcal{M}_i$ is prime. $\square$

We now consider a non-standard exponential map

$$\exp : V_1 \oplus \cdots \oplus V_k \rightarrow \mathcal{T}_{g_1, n_1} \times \cdots \times \mathcal{T}_{g_s, n_s}$$

defined as follows. Given

$$v = (q_1, \dots, q_k) \in V_1 \oplus \cdots \oplus V_k$$

let $\exp(v)$ be the result of applying geodesic flow for time $|q_i|$ using the differential q_i for $i = 1, \dots, k$. Geodesic flows in different components commute, so it does not matter what order one applies these flows in or if one applies them simultaneously.

Lemma 4.7. *$\exp$ is continuous.*

Proof. Consider a sequence

$$v^{(j)} = (q_1^{(j)}, \dots, q_k^{(j)}) \in V$$

converging to $v = (q_1, \dots, q_k) \in V$. To show convergence of $\exp(v^{(j)})$ to $\exp(v)$, it suffices to prove convergence in each of the k factors. Restricting to the factor i , we proceed in two cases.

In the first case, suppose that $q_i \neq 0$. Then, Lemma 4.6 gives that q_i is non-zero in all components associated to this factor, and we get that eventually all the

components of $q_i^{(j)}$ are non-zero. The result of flowing for time $|q_i^{(j)}|$ along $q_i^{(j)}$ converges to the result of flowing for time $|q_i|$ along q_i because $|q_i^{(j)}| \rightarrow |q_i|$ and the action is continuous when the set of components where the differential is non-zero is constant.

In the second case, suppose $q_i = 0$. Then $|q_i^{(j)}| \rightarrow 0$ and continuity follows since flowing for amounts of time converging to zero results in no change in the limit. $\square$

Lemma 4.8. *exp is injective.*

Proof. Suppose $\exp(q_1, \dots, q_k) = \exp(q'_1, \dots, q'_k)$. The Teichmüller Uniqueness Theorem gives that q_i and q'_i are equal up to scale in each component. It follows that $q_i = c_i q'_i$ for some $c_i > 0$, since otherwise a linear combination of q_i and q'_i would contradict Lemma 4.6. The distance moved in each component of the i -th factor is by definition $|q_i| = |q'_i|$, so we get that $c_i = 1$. $\square$

Lemma 4.9. *The image of exp is contained in the slice S .*

Proof. Because $\mathcal{M}$ is a product and is invariant under geodesic flow, it follows that $\mathcal{M}$ is invariant under doing different amounts of geodesic flow in each factor $\mathcal{M}_i$. By continuity and density, we get that same statement for the closure of $\mathcal{M}$ in $R(QS)$, and hence also for the associated locus in QS . $\square$

Lemma 4.10. *exp is a homeomorphism onto S and $S = L$. Furthermore, $\dim V = \dim L$.*

Proof. We know that exp is continuous and injective. Since it is also proper, the image of exp is a topological manifold of real dimension $2 \dim V$, where, as always, $\dim$ denotes complex dimension. Invariance of Domain and our assumption that $\dim V \geq \dim L$ gives $\dim V = \dim L$. We will show that it follows from the irreducibility of L that L is equal to the image of this map. Since the image is contained in S and $S \subseteq L$, this will also prove $S = L$.

Indeed, L can be stratified into real submanifolds (actually even complex submanifolds), and irreducibility of L implies that there is only one largest dimensional stratum L^{top} in this stratification, which is furthermore open and dense in L [Whi65, Section 18]. Let $V^{\text{top}} \subseteq \bigoplus_{i=1}^k V_i$ be the preimage of L^{top} under exp. Using Invariance of Domain as well as properness and injectivity of exp, we see that the image of $\exp|_{V^{\text{top}}}$ is both open and closed in L^{top} and hence this image is equal to L^{top} . Since the image of exp is closed and L^{top} is dense in L , we get that the image of exp is equal to L . $\square$

We get the following as an immediate corollary.

Corollary 4.11. *L can be written as a product*

$$L = L_0 \times L_1 \times \cdots \times L_k$$

such that

- (1) L_0 is a subset of the product of Teichmüller spaces indexed by I_Z ,
- (2) L_0 consists of a single point, and
- (3) each L_i , $i > 0$ is a subset of the product of the factors corresponding to $\mathcal{M}_i$.

Furthermore, if $X = (X_0, X_1, \dots, X_k)$ is the point chosen at the beginning of this subsection and

$$Y = (Y_0, Y_1, \dots, Y_k) \in L$$

is arbitrary, then L_i contains a T -disc through X_i and Y_i for each i .

We can now conclude as follows:

Lemma 4.12. I_Z, I_R , the partition of I_R into k subsets associated to the product decomposition $\mathcal{M} = \mathcal{M}_1 \times \dots \times \mathcal{M}_k$, and the L_i do not depend on the choice of X and V .

Proof. We see from the above that I_Z indexes the components where L is constant, so I_Z and its complement I_R do not depend on our choices. We now consider the partition of I_R into k subsets. By Remark 3.1 part (3), L_i cannot be a product in a non-trivial way, and it follows that the decomposition of I_R and the L_i cannot depend on our choices either. $\square$

Proof of Theorem 4.3. Since X was arbitrary and the L_i do not depend on choices, we get that the L_i contain a T -disc through every pair of distinct points. Thus Remark 3.1 part (3) gives the result. A technical point is that if a complex analytic set is a product of sets, then the factors are automatically complex analytic. (For example, if $A' \times B'$ is complex analytic in a complex manifold $A \times B$, note that A' can be expressed in the form $(A' \times B') \cap (A \times \{b_0\})$ for some $b_0 \in B'$, and use that the intersection of analytic sets is analytic). $\square$

5. THE BOUNDARY IS SMOOTH AND IRREDUCIBLE

5.1. Smoothness and the proof of Theorem 1.3. The following ties up a loose end in previous discussions.

Lemma 5.1. *Every totally geodesic complex analytic set L of a product of Teichmüller spaces is a complex submanifold.*

Note that the arguments above show that L is a topological manifold, but these arguments do not show it is a complex submanifold, and the exponential maps lack the smoothness to be immediately useful for this problem [Ree04, Theorem 1].

Proof. By Theorem 4.3 it suffices to prove this when L is simple. For this it is helpful to use the version of the exponential map defined on the tangent space, so $\exp_X(v)$ is the constant speed (but typically not unit speed) geodesic through X with derivative v . By first considering the smooth points of L , as above we see that most and hence all points X of L have a $\dim L$ -dimensional subspace of the ambient tangent space whose image under $\exp_X$ is equal to L . This shows that L is locally well approximated by a linear space at every point of L , as is the case for a submanifold. The result now follows from [BFLS16, Proposition 3.4], keeping in mind that all Finsler metrics are equivalent on compact sets.

(The idea of the argument is to consider (locally) the projection to the approximating linear space. If this projection has degree 1, one can find a holomorphic inverse off of the branching locus, and then extend it to the branching locus using the Removable Singularity Theorem. This would show that L is locally the graph of a holomorphic map, giving the result. On the other hand, if the degree is greater than 1, then as points in the same fiber come together one finds pairs of points such that the most efficient way to join them is to leave L). $\square$

5.2. Irreducibility. The final loose end is the following, which will be important later.

Proposition 5.2. *Suppose $N \subseteq \mathcal{T}_{g,n}$ is algebraic and totally geodesic and let L denote the intersection of $\overline{N}$ with a boundary stratum. Then L is irreducible.*

We will prove some lemmas before we prove this proposition.

Lemma 5.3. *Suppose $X_k, Y_k \in \mathcal{T}_{g,n}$, with $X_k \rightarrow X$ and $Y_k \rightarrow X$ in the Deligne-Mumford bordification. Suppose that X, Y are in the same stratum, that $X_k \neq Y_k$ for all $k \in \mathbb{N}$, and that $\sup_k d(X_k, Y_k) < \infty$. Let Δ_k be the Teichmüller disc through X_k and Y_k . Then, possibly after passing to a subsequence, the discs Δ_k converge to a holomorphic disc in the stratum of X and Y containing X and Y .*

Proof. For the rest of the proof let ρ denote the map that forgets a differential on a surface, so $\rho(X, q) = X$.

Let q_k be a quadratic differential on X_k and $t_k > 0$ be a real number such that doing geodesic flow along q_k for time t_k starting at X_k gives Y_k . Pass to a subsequence so $t_k \rightarrow t$. Also pass to a subsequence to assume there is a meromorphic differential q on X that is non-zero on every component and such that, with appropriate rescaling depending on the component, q_k converges to q (away from nodes). Compare for example to [EKZ14, Theorem 10] and the references therein.

For any matrix $g \in GL(2, \mathbb{R})$ one can check directly using the quasi-conformal topology that $\rho(gq_k)$ converges to $\rho(gq)$, and, in particular, that $\rho(g_t q) = Y$. Let Δ be the image under ρ of the $GL^+(2, \mathbb{R})$ orbit of q . It follows that the Δ_k converge to Δ . One can check that Δ is holomorphic using holomorphic dependence of parameters in the measurable Riemann mapping theorem.

(An alternative approach to this lemma would be to use normal families and the fact that Dehn space is a bounded domain [Ber81, Section 16]). $\square$

Remark 5.4. One can compare the previous proof to [BCG⁺19a, Theorem 1.4], keeping in mind that Lemma 5.3 is much simpler because it does not require any global structural results for a compactification.

Lemma 5.5. *In the setting of Proposition 5.2, given two points $X, Y \in L$, it is possible to find $X_k, Y_k \in N$ such that $X_k \rightarrow X$, $Y_k \rightarrow Y$, and the distance between X_k and Y_k is bounded.*

Proof. Lemma 3.16 gives differentials in QN such that replacing some disjoint equivalence classes of cylinders with pairs of half-infinite cylinders gives differentials on X and Y . We can define X_k and Y_k to be the result of using the Cylinder Deformation Theorem to stretch cylinders (without twisting them) so that a chosen cylinder in each horizontal equivalence class has modulus k . Then all the moduli of the relevant cylinders are the same at X_k and Y_k .

One can then directly build quasi-conformal maps from X_k to Y_k of bounded dilatation which are the identity on a large part of the horizontal cylinders. $\square$

Remark 5.6. For convenience we gave a proof of Lemma 5.5 that is very specific to our situation. There are other approaches that hinge on finding holomorphic disks through the given points in the variety and using plumbing coordinates. One such approach might try to bound distances using the Minsky product region estimates, leaning on asymptotic results on the hyperbolic metric to control lengths and twists for the short curves: see for example [Wol10, page 71] and [Wol90].

Another approach would be to estimate the Teichmüller norm of an infinitesimal change of the plumbing parameter and (roughly speaking) integrate that to get bounded distance. Another approach would be to do something similar to what we did above but without the use of flat geometry. We thank Scott Wolpert for helpful discussions about several of these approaches.

Proof of Proposition 5.2. Lemma 5.3 applied to the points produced by Lemma 5.5 gives that every pair of points in L is contained in a holomorphic disc in L , which shows that L is irreducible. $\square$

5.3. Proof of Theorem 1.3. This concludes the proof of Theorem 1.3 as follows:

Proof of Theorem 1.3. Irreducibility is given by Proposition 5.2. Algebraicity follows from Lemma 3.13, and semisimplicity from Theorem 4.3, with the additional input of Lemma 5.1 guaranteeing that the complex analytic sets that appear are in fact complex submanifolds. $\square$

6. FIRST CONSEQUENCES OF SEMISIMPLICITY

In this section we begin the work of clarifying how some of the different structures previously introduced relate to each other.

6.1. The boundary of a subsurface. If U is a subsurface, we let $\mathcal{T}(U)$ denote the Teichmüller space of the finite type surface obtained by contracting each boundary component of U to a marked point or puncture.

Lemma 6.1. *Let $N \subseteq \mathcal{T}_{g,n}$ be an algebraic totally geodesic submanifold of Teichmüller space. Suppose $F \subseteq \mathcal{T}(U_1) \times \cdots \times \mathcal{T}(U_k)$ is a simple factor of the intersection of $\overline{N}$ with a stratum of the Deligne-Mumford bordification. Furthermore, assume F is not a point. Then:*

- (1) *For all $i \neq j$, every boundary curve of U_i is equivalent to a boundary curve of U_j .*
- (2) *For all cylinder curves in $\mathcal{C}(U_i)$, every equivalent curve lies in one of the $\mathcal{C}(U_j)$, and there is at least one equivalent curve in each $\mathcal{C}(U_j)$.*
- (3) *The boundary of the union of the U_i is a union of equivalence classes.*

Here, saying a curve is in $\mathcal{C}(U_i)$ is just a convenient way of saying it is contained in U_i non-peripherally. A simple factor is by definition one of the L_i , $i \geq 0$ in Corollary 4.11; the requirement that F not be a point exactly excludes L_0 .

Note that assertion (2) is trivially true when F is a point, but that assertion (1) fails, for example, when $N = \mathcal{T}_{g,n}$ and a convenient multi-curve is pinched. In this case, when F is a point, the U_i are the pants components of the complement of the pinched multi-curve. See Figure 4.

Proof. Suppose to the contrary that there is a boundary curve α of U_1 that is not equivalent to any boundary curve of U_2 . By Corollary 3.17, we can unpinch α and its equivalence class, without unpinching any other curves. We get a new boundary stratum where U_2 is still a component and where there is a component U'_1 containing U_1 . Note that U'_1 and U_2 must be in the same simple factor, since otherwise we can take limits and contradict the fact that F is a simple factor. However, a Dehn multi-twist along the equivalence class of α (produced by Lemma 2.6) changes U'_1 without changing U_2 , giving a contradiction. A similar argument proves the second claim after pinching the boundary curves of the U_i .




FIGURE 4. If N is the whole Teichmüller space, pinching the curves drawn gives a simple factor which is a point in $\mathcal{T}(U_1) \times \mathcal{T}(U_2)$ for the trivial reason that $\mathcal{T}(U_1)$ and $\mathcal{T}(U_2)$ are points themselves. In this example each curve is only equivalent to itself.

It remains only to address the third claim, where we will require some ideas from papers on invariant subvarieties. These ideas are isolated in the second of the following two sublemmas, so they can be treated as a black box if desired.

Sublemma 6.2. *Fix an equivalence class B of a boundary curve of one of the U_i . There is a finite area quadratic differential q in the boundary of $Q_{\max}N$ which is horizontally periodic and such that the following hold:*

- (1) *the curves in B are horizontal cylinders on q ,*
- (2) *there is precisely one other equivalence class E of horizontal cylinders on q ,*
- (3) *the pinched curves at q are exactly the other equivalence classes of boundary curves, and*
- (4) *all core curves of cylinders in E are contained non-peripherally in one of the U_i .*

Proof of Sublemma 6.2. Because F is not a point, it contains non constant T-discs. The quadratic differentials generating these T-discs have cylinders, and so we see that the U_i contain cylinder curves non-peripherally. Consider such a cylinder curve and note that the second claim of Lemma 6.1 gives that its equivalence class E consists entirely of curves that are contained non-peripherally in one of the U_i .

Pinch all the equivalence classes of boundary curves of the U_i except B . Note that the curves of B and E must all lie in the same simple factor, again by the second claim of Lemma 6.1.

Lemma 2.6 gives a positive Dehn multi-twist in all the curves of $E \cup F$ that stabilizes the given simple factor. Consider the T-disc in that simple factor that contains a point and its image under this multi-twist, and consider the quadratic differential q generating this T-disc which also generates the horocycle between these two points in this disc. By applying Lemma 2.5 in each component, we see that it is horizontally periodic with horizontal cylinders corresponding to $B \cup E$. $\square$

Sublemma 6.3. *Let C_1 and C_2 be parallel equivalence classes of cylinders on a differential q in the boundary of $Q_{\max}N$. If a cylinder of C_1 is adjacent to a cylinder of C_2 , then every cylinder of C_2 is adjacent to a cylinder of C_1 .*

Here we say two cylinders are adjacent exactly when they have a common saddle connection on their boundary.

Proof of Sublemma 6.3. This follows from an “over-collapsing” technique developed for the classification of invariant subvarieties of differentials. See [Api25, Lemma 4.4] for a precise result sufficient for our needs and see the paragraph after [AW22, Lemma 4.8] for a list of similar arguments which have appeared in the literature, many of which include illustrations. The rough summary of the argument is that one considers an appropriate linear path of cylinder deformations of $\mathbf{C}_1$ which keeps the circumferences the same but decreases the areas of these cylinders to 0. Continuing this path in the stratum, one finds differentials where the moduli of some of the cylinders adjacent to $\mathbf{C}_1$ start to decrease. Since this ratio of moduli of cylinders in $\mathbf{C}_2$ is constant by Lemma 2.3, this gives the result. $\square$

To conclude the proof of the third claim, consider the differential q given by Sublemma 6.2. Let $\mathbf{C}_1 = E$, and let $\mathbf{C}_2 = B$ denote the equivalence class of one of the boundary curves of one of the U_i . Then Sublemma 6.3 gives the result because the curves of E are contained in the U_i . $\square$

6.2. Projection to a subsurface. We also note the following result, which will clarify much of our later discussion.

Lemma 6.4. *Suppose $F \subseteq \mathcal{T}(U_1) \times \cdots \times \mathcal{T}(U_k)$ is a simple factor of the intersection of $\overline{N}$ with a stratum of the Deligne-Mumford bordification. Fix $1 \leq i_0 \leq k$. Consider any non-empty intersection of $\overline{N}$ with a stratum of the Deligne-Mumford bordification which has $\mathcal{T}(U_{i_0})$ as a factor.*

- (1) *Suppose F is not a point. Then this intersection has F as a factor.*
- (2) *Suppose F is a point. Then the projection of this intersection to $\mathcal{T}(U_{i_0})$ is the same as the projection of F to $\mathcal{T}(U_{i_0})$, and so, in particular, a point.*

Proof. In either case, we can consider the component of the boundary of $\overline{N}$ resulting from pinching the equivalence classes of the boundary curves of U_{i_0} . Let F' be the simple factor involving the $\mathcal{T}(U_{i_0})$ coordinate.

Any boundary component of $\overline{N}$ where U_{i_0} appears results from pinching additional equivalence classes, with the restriction that the equivalence classes must contain disjoint curves and cannot contain curves in U_{i_0} .

If F' is a point, then any additional degeneration will not affect the projection to $\mathcal{T}(U_{i_0})$. If F' is not a point, then the product structure of the boundary gives that any additional degenerations which leaves U_{i_0} intact cannot affect F' . $\square$

Remark 6.5. In Figure 4 we can first pinch the curves bounding U_1 and then pinch the curves bounding U_2 to see that a simple factor which is a point can in a certain sense “grow” from involving only U_1 to involving both U_1 and U_2 .

Lemma 6.4 says that, with a minor caveat in the case of simple factors that are points, the factor F is well defined given any of the U_i .

7. FROM TEICHMÜLLER SPACES TO NON-ANNULAR CURVE GRAPHS

Theorem 1.3 provides a decomposition of the boundary of any algebraic totally geodesic submanifold into simple factors. The projection of any such factor into any of its Teichmüller coordinate domains is an isometric embedding whose image is an algebraic totally geodesic submanifold. The main results of this section is the following theorem which provides control over the coarse geometry of these projections and which is particularly useful for the applications we discuss later.

Recall that if U is a surface then $\mathcal{T}(U)$ and $\mathcal{C}(U)$ denote its Teichmüller space and curve complex, respectively.

Theorem 7.1. *Let $N \subseteq \mathcal{T}_{g,n}$ be an algebraic totally geodesic submanifold. Then, there exists $K > 0$ such that if*

$$F \subseteq \mathcal{T}(U_1) \times \cdots \times \mathcal{T}(U_k)$$

is a simple factor of the intersection of $\overline{N}$ with a stratum of the Deligne–Mumford bordification, then the projection maps from the image of F in

$$\mathcal{C}(U_1) \times \cdots \times \mathcal{C}(U_k)$$

to each factor are K -quasi-isometric embeddings.

As a corollary, it is possible to deduce the same statement without the need to pass to the boundary.

Corollary 7.2. *Let $N \subseteq \mathcal{T}_{g,n}$ be an algebraic totally geodesic submanifold. Then, there exists $K > 0$ such that if the subsurfaces $U_1, \dots, U_k$ arise as in Theorem 7.1 from a simple factor of the boundary, then the projection maps from the image of N in*

$$\mathcal{C}(U_1) \times \cdots \times \mathcal{C}(U_k)$$

to each factor are K -quasi-isometric embeddings.

We isolate some of the work of proving the corollary in the following lemma.

Lemma 7.3. *Let $N \subseteq \mathcal{T}_{g,n}$ be an algebraic totally geodesic submanifold. There exists a (small) $\epsilon > 0$ and a (large) $K > 0$ such that the following holds for any $Y \in N$.*

For any collection of equivalence classes of curves, all of hyperbolic length at most ϵ at Y , there is a point Z of the boundary of N where exactly these curves are pinched, and for every subsurface U that is a component of the complement of the union of the curves, the projections of Y and Z to $\mathcal{C}(U)$ are within K of each other.

Proof sketch. We will use compactness of the boundary of $\pi(N)$ in $\overline{\mathcal{M}}_{g,n}$. Each point Z in the boundary has a small neighborhood consisting of points Y which are appropriately close to Z . Taking a finite subcover and lifting gives the result. (One should use neighborhoods that have nice preimages in the bordification.) $\square$

Proof of Corollary 7.2 assuming Theorem 7.1. By Corollary 3.14, there are only finitely many orbits of relevant tuples $U_1, \dots, U_k$, so we do not have to worry about the dependence of K on this tuple.

It suffices to show that for every point X in N there is a corresponding point of F with coarsely the same image in the given product of curve graphs. We first do this when F is not a point.

Let C be the constant from Lemma 2.13 controlling ratios of hyperbolic lengths, and let ϵ be as in Lemma 7.3. We will say that a curve is short if its hyperbolic length is at most ϵ and very short if its hyperbolic length is at most ϵ/C . Thus Lemma 2.13 gives that if a curve is very short, all equivalent curves are short.

Let X be a point of N . Consider a Teichmüller geodesic segment from X to a point Z of N where all boundary curves of $U_1, \dots, U_k$ are very short. Let Y be the first point on this segment such that there exists i such that all boundary curves

of U_i are very short at Y . By definition, at every point on the geodesic before Y , each U_i has a boundary curve with length greater than ϵ/C .

By Rafi's theory of active intervals [Raf07a, Section 3], which we recall in Theorem B.3, the projection to the curve complexes $\mathcal{C}(U_i)$ coarsely does not change along the Teichmüller geodesic $[X, Y]$.

By Lemma 6.1 part (1), all the subsurfaces U_i are bounded by curves in the same finite collection of equivalence classes. As discussed above, it follows that all boundary curves of all U_i are short at Y . We now let Z be the point produced from Y by Lemma 7.3 with the collection of curves being the boundary of the U_i . So the curves pinched at Z are exactly the boundary curves of the U_i . Lemma 7.3 gives that Y and Z have coarsely the same projection to each $\mathcal{C}(U_i)$.

This shows that the image of X in $\mathcal{C}(U_1) \times \cdots \times \mathcal{C}(U_k)$ is coarsely equal to the image of a point Z in the stratum of the Deligne–Mumford bordification where the boundary components of all the U_i have been pinched. This part of the boundary is irreducible by Proposition 5.2. The given F may have originated from an even smaller boundary stratum, i.e., with more pinched curves, but this is not a concern thanks to Lemma 6.4.

If F is a point, we cannot use Lemma 6.1, but the proof can be easily modified. Say U_{i_0} has that all boundary curves are very short at Y . The key observation is that subsurface projection to U_{i_0} cannot change along $[Y, Z]$, because by Lemmas 7.3 and 6.4 it is controlled by a single point of $\mathcal{T}(U_{i_0})$. $\square$

Proof of Theorem 7.1. Notice that Theorem 1.3 gives the corresponding result at the level of Teichmüller spaces. Lemma 6.1 part (2) gives that for any cylinder curve in $\mathcal{C}(U_i)$ there must be at least one equivalent cylinder curve in every other $\mathcal{C}(U_j)$.

Hence Corollary 2.12 and Lemma 2.13 give that when a curve is short on one U_i then there is a comparably short curve in the same equivalence class on every other U_j . Hence the result follows from Theorem B.1 and Corollary B.2, because we can use Theorem B.1 to coarsely compute distance in the curve graph of one component using distance in an electrification, and Corollary B.2 gives that the exact same electrification coarsely computes distance in the curve graphs of the other components. $\square$

8. FROM TEICHMÜLLER SPACES TO ANNULAR CURVE GRAPHS

The goal of this section is to show that the analogue of Corollary 7.2 for annular curve graphs also holds. See Section 2.4 for background. If α is a simple closed curve we let $\mathcal{C}(\alpha)$ denote the Teichmüller space version of the annular curve graph of α , i.e., $\mathcal{C}(\alpha)$ can be viewed as the horoball in the upper half plane $\mathbb{H}$ with imaginary part at least 1. There is a projection map π_α from Teichmüller space to $\mathcal{C}(\alpha)$, where the imaginary part is the max of 1 and 1 over the hyperbolic length of α , and the real part encodes twisting, as measured by taking an appropriate marking and projecting it to the mapping class group version of the annular curve graph.

Theorem 8.1. *Let $N \subseteq \mathcal{T}_{g,n}$ be an algebraic totally geodesic submanifold. Then, there exists $K > 0$ such that if $\{\alpha_1, \dots, \alpha_k\}$ is an N -equivalence class of simple closed curves, then the projection maps from the image of N in*

$$\mathcal{C}(\alpha_1) \times \cdots \times \mathcal{C}(\alpha_k)$$

to each factor are K -quasi-isometric embeddings.

We verify this using the following.

Lemma 8.2. *Let $\mathcal{X}$ be a subset of $\{z : \text{Im}(z) > 1\}^k \subset \mathbb{H}^k$. Assume there are constants $C > 0$ and $c_{j,m} > 0$ such that if*

$$(x_j + iy_j)_{j=1}^k, (x'_j + iy'_j)_{j=1}^k \in \mathcal{X}$$

then for all j, m we have $y_j/y_m < C$ and

$$|x_j - x'_j| - c_{j,m}|x_m - x'_m| \leq C(|y_1| + |y'_1|).$$

Then there is a K depending only on $C, c_{j,m}$ and k such that each of the k -coordinate projections $\mathcal{X} \rightarrow \mathbb{H}$ are K -quasi-isometric embeddings.

The metric on $\mathbb{H}^k$ here is the sup of the usual hyperbolic metrics on the factors. (Since everything here is coarse it wouldn't matter if we instead used the product Riemannian metric.)

Proof. The distance between two points $x + iy, x' + iy' \in \mathbb{H}$ is coarsely

$$\log(\max(y, y')/\min(y, y')) + \max(0, \log(|x - x'|/\max(y, y'))),$$

and the result follows from this. $\square$

Proof of Theorem 8.1. We verify the assumptions of Lemma 8.2. Lemma 2.13 immediately gives the assumption on ratios of y coordinates.

Consider two points $X, X' \in N$. Let D_α denote a multi-twist supported on the set of α_j , as produced by Lemma 2.6. We can pick a power p such that the projections $\pi_{\alpha_1}(D_\alpha^p X')$ and $\pi_{\alpha_1}(X)$ have that their real parts, i.e., their twists, are uniformly close depending only on N . (The uniformity here follows from Corollary 3.14, which gives that there are only finitely many equivalence classes up to the action of the fundamental group of N . See also [MW17, Theorem 5.1] and [CW21, Theorem 5.3] for previous justifications of this.)

It now suffices to prove that for every other j the real parts of $\pi_{\alpha_j}(D_\alpha^p X')$ and $\pi_{\alpha_j}(X)$ must be equal up to an error of size $O(\ell_{\alpha_j}(X)^{-1} + \ell_{\alpha_j}(X')^{-1})$. (The constants $c_{j,m}$ arise from the ratio of twisting done by D_α in different curves.)

Suppose otherwise in order to find a contradiction. We consider the geodesic segment from X to $D_\alpha^p X'$ and perform a more in depth version of the analysis of Lemma 2.8. Although we do not use the main results in Appendix A, we use results of Rafi recalled there, so we recommend first reading the beginning of that appendix. One should also recall that [Raf07a, Theorem 4.3] shows that different notions of twisting at α (namely the x coordinate of π_α and the twist_q of the appendix) are equivalent up to an error of size at most $O(\ell_\alpha(X)^{-1})$.

Let q be the quadratic differential on X generating the geodesic segment from X to $D_\alpha^p X'$. Without loss of generality suppose in order to find a contradiction that the real parts of $\pi_{\alpha_2}(D_\alpha^p X')$ and $\pi_{\alpha_2}(X)$ differ by an extremely large multiple of $\ell_{\alpha_j}(X)^{-1} + \ell_{\alpha_j}(X')^{-1}$. As in the proof of Lemma 2.8, we see that α_2 is a cylinder on q . Moreover, the formulas recalled in the appendix give that, as always up to the error of one over hyperbolic length, the amount of twisting accomplished depends on the modulus. Thus we see that, at some point on the geodesic segment, the α_2 cylinder must have large modulus, but that the α_1 cylinder cannot have large modulus because almost no twisting is accomplished in α_1 ; here we are using that the balanced times, in Rafi's terminology, for α_1 and α_2 are the same because the

corresponding cylinders are parallel. This contradicts Corollary 2.2 which gives the moduli of the α_1 and α_2 cylinders must have fixed ratio. $\square$

9. HIERARCHICAL HYPERBOLICITY

9.1. The axioms. We note in advance that to improve readability for non-experts we will verify a slight modification of the axioms for an HHS. Indeed, we will verify the axioms for an almost HHS. These axioms are identical to those of an HHS, except for a technical change to the orthogonality axiom. It is known that every almost HHS is an HHS, see [ABD21, Appendix], so the very slightly modified axioms are equivalent to the usual ones. Additionally and for the sake of completeness, we sketch a proof of the usual orthogonality axiom at the end of this section.

We will paraphrase the axioms for an almost HHS imprecisely, omitting a number of details; the reader should consult other sources, such as [ABD21], for the precise definitions of these axioms. Our goal here is to briefly introduce the axioms for readers new to hierarchical hyperbolicity and to refresh the memory of those that are already familiar with this concept.

Roughly speaking, a quasi-geodesic metric space X , endowed with a good deal of extra structure specified below, is said to be hierarchically hyperbolic if the following axioms hold with uniform constants:

- (0) **Basic set up.** There is an index set $\mathfrak{S}$, called the set of domains, and for each $W \in \mathfrak{S}$ there is a Gromov hyperbolic metric space $\mathcal{C}W$.
- (1) **Projections.** There are coarsely Lipschitz projections $\pi_W : X \rightarrow 2^{\mathcal{C}W}$. The image $\pi_W(x)$ of each point must have uniformly bounded diameter, so each $\pi_W(x)$ can be considered as a coarse point and there is no harm in thinking of π_W as a map to $\mathcal{C}W$ instead of to subsets of $\mathcal{C}W$. The image of each π_W is required to be quasi-convex (which is automatic when it is coarsely surjective).
- (2) **Nesting.** The set of domains $\mathfrak{S}$ has a reflexive partial order $\sqsubseteq$ called nesting with a unique maximal element. If $V \sqsubset W$, then there is a specified point $\rho_W^V \in \mathcal{C}W$ and a map $\rho_V^W : \mathcal{C}W \rightarrow 2^{\mathcal{C}V}$.
- (3) **Orthogonality.** The set of domains $\mathfrak{S}$ has a symmetric, anti-reflexive relation $\perp$ called orthogonality, such that the following hold:
 - (a) If $V \sqsubseteq W$ and $W \perp U$ then $V \perp U$.
 - (b) A nested pair is never orthogonal.
 - (c) The cardinality of any set of pairwise orthogonal domains is uniformly bounded above.
- (4) **Transversality and consistency.** We say two domains V and W are transverse and write $V \pitchfork W$ if they are not orthogonal and neither is nested in the other. Furthermore, for a uniform constant $\kappa_0 > 0$:
 - (a) If $V \pitchfork W$, then there are points $\rho_W^V \in \mathcal{C}W$ and $\rho_V^W \in \mathcal{C}V$ such that the following Behrstock inequality holds:

$$\min\{d_W(\pi_W(x), \rho_W^V), d_V(\pi_V(x), \rho_V^W)\} \leq \kappa_0.$$

- (b) If $V \sqsubseteq W$, either $\pi_V(x)$ is approximately $\rho_V^W(\pi_W(x))$ or

$$d_W(\pi_W(x), \rho_W^V) \leq \kappa_0.$$

(c) If $U \sqsubseteq V$, then we have

$$d_W(\rho_W^U, \rho_W^V) \leq \kappa_0$$

for all W satisfying one of the following: $V \sqsubseteq W$; or W transverse to V and not orthogonal to U .

- (5) **Finite complexity.** There is a uniform upper bound on the length of properly nested chains.
- (6) **Large links.** There are uniform constants $\lambda > 0$ and $E > 0$ such that for every $x, x' \in X$ there exists a set $\{T_i\}$ of at most

$$\lambda d_W(\pi_W(x), \pi_W(x')) + \lambda$$

domains properly nested in W such that if T is properly nested in W and

$$d_T(\pi_T(x), \pi_T(x')) > E$$

then T is nested in one of the T_i .

- (7) **Bounded geodesic image.** If γ is a geodesic in $\mathcal{C}W$ and $V \sqsubseteq W$, then either $\rho_V^W(\gamma)$ has uniformly bounded diameter, or γ comes close to ρ_V^W .
- (8) **Partial realization.** For a uniform $\alpha > 0$ and any pairwise orthogonal family $\{V_j\}$ and points $p_j \in \pi_{V_j}(X)$, there exists a point $x \in X$ such that

$$d_{V_j}(\pi_{V_j}(x), p_j) \leq \alpha$$

for all j , and such that for each j and each V with $V_j \sqsubset V$ or $V_j \pitchfork V$,

$$d_V(\pi_V(x), \rho_V^{V_j}) \leq \alpha.$$

- (9) **Uniqueness.** For all $\kappa > 0$ there exists $\theta_u > 0$ such that if $d(x, y) \geq \theta_u$ then there exists V with $d_V(\pi_V(x), \pi_V(y)) \geq \kappa$.

9.2. Totally geodesic submanifolds. We now verify the axioms for an algebraic totally geodesic submanifold N , following the order introduced above. Note that we write $d_U(x, y)$ instead of $d_U(\pi_U(x), \pi_U(y))$, leaving implicit that to compute such distance we must first project to $\mathcal{C}U$.

Basic set up. We start by determining what our set of domains will be.

Lemma 9.1. *For any algebraic totally geodesic submanifold N there exists a constant $D > 0$ such that for any subsurface U the projection of N to $\mathcal{C}(U)$ is either unbounded or has diameter at most D . The projection is unbounded if and only if U appears as a component of an unbounded simple factor.*

The requirement that the simple factor be unbounded is equivalent to the requirement that it not be a point. We say U is a component of the simple factor if $\mathcal{T}(U)$ is a factor of the product of Teichmüller spaces in which the simple factor lies.

Proof. Suppose first that U is annular. If the core curve does not get short, then the diameter bound follows from Rafi's theory of active intervals (recalled in Theorem B.3). Otherwise, the core curve gets short and, keeping in mind Lemma 2.6 and Corollary 2.12, we can pinch it and twist it to show unbounded image.

Now take U to be non-annular. We proceed in cases, sometimes making use of Theorem 1.3:

- (1) *Some boundary curve of U is not a cylinder curve.* By Corollary 2.12, this curve cannot get short, so the theory of active intervals gives the diameter bound.
- (2) *All boundary curves of U are cylinder curves, but one of these curves is N -equivalent to a curve α in $\mathcal{CU}$.* The theory of active intervals gives that the projection to $\mathcal{CU}$ can only change when the boundary of U is short, but in that case, by Lemma 2.13, α is short, so the projection is coarsely α , providing a uniform bound on the diameter.
- (3) *It is possible to pinch the boundary of U without pinching a curve in the interior of U .* In this case U appears in some simple factor F . If this factor is a point, then the proof of Corollary 7.2 shows that the projection to $\mathcal{CU}$ has bounded diameter. If not, it contains a Teichmüller disc. Since, by [RS09, Lemma 4.4] and [KMS86], all Teichmüller discs have infinite diameter in the curve graph of the surface, this gives unbounded image. $\square$

Definition 9.2. Given an algebraic totally geodesic submanifold $N \subseteq \mathcal{T}_{g,n}$:

- (1) We call the subsurfaces U for which $\pi_U(N)$ is unbounded active.
- (2) We call two non-annular active subsurfaces equivalent if they occur in the same simple factor of the intersection of $\bar{N}$ with a stratum of the Deligne–Mumford bordification.
- (3) We call two annular active subsurfaces equivalent if their core curves are equivalent.

Note that Lemma 6.4 gives that two equivalent active non-annular subsurfaces will be in the same simple factor of every intersection of $\bar{N}$ with a stratum of the Deligne–Mumford bordification where the two subsurfaces appear. Note also that, by definition, a non-annular subsurface cannot be equivalent to an annular subsurface.

Example 9.3. As remarked in the introduction, the basic example of totally geodesic submanifolds are covering constructions in the sense of [MMW17, Section 6], resulting in a copy of a smaller Teichmüller space sitting inside of a larger Teichmüller space. In this case the active subsurfaces are connected components of lifts of non-pants connected subsurfaces, and the different connected components of such a lift are equivalent.

Let $\mathfrak{S}_N$ be the set of equivalence classes of active subsurfaces of N . For each equivalence class $E = \{U_1, \dots, U_k\}$ in $\mathfrak{S}_N$ we let $\mathcal{C}(E)$ be the image of N in

$$\mathcal{C}(U_1) \times \dots \times \mathcal{C}(U_k).$$

Note that Lemma 9.1 gives that $\mathcal{C}(E)$ is unbounded.

Lemma 9.4. *For any $E \in \mathfrak{S}_N$ the space $\mathcal{C}(E)$ is Gromov hyperbolic.*

Proof. By Corollary 7.2 and Theorem 8.1, the projections to each factor are quasi-isometric embeddings, so it suffices to show that $\pi_{U_1}(N)$ is Gromov hyperbolic.

If U_1 is non-annular this follows from the fact that Teichmüller geodesics map to unparametrized quasi-geodesics [Raf14, Theorem B], which gives quasi-convexity and hence hyperbolicity of the image. If U_1 is annular one can explicitly show that π_U is coarsely surjective, which directly gives hyperbolicity of the image. $\square$

For more on the image of Teichmüller discs in the curve graph see [TW18].

Projections. By definition we have natural maps $\pi_E : N \rightarrow \mathcal{C}(E)$ defined as the restriction of the product map $(\pi_{U_1}, \dots, \pi_{U_k})$ to N . The map π_E is coarsely Lipschitz since the maps π_{U_i} are coarsely Lipschitz.

Nesting. We say $E = \{U_1, \dots, U_k\}$ is nested in $F = \{V_1, \dots, V_\ell\}$ if every U_i is nested in a V_j and every V_j has a U_i nested in it. In this case we define $\rho_E^F(\alpha)$ to be the union of the $\pi_E(X)$ for $X \in N$ with $\pi_F(X)$ close to α .

Notice that π_E is surjective by definition. We define ρ_E^E to be the boundary components of the U_i in the corresponding V_j . Since the boundary components of the U_i can be made simultaneously short, this gives a well defined point in $\mathcal{C}(F)$.

We should compare our definition of $\rho_E^F(\alpha)$ to the usual definition used in the HHS structure of Teichmüller space. Let U and V be subsurfaces with U nested in V . To start, we define ρ_V^U to be the boundary of U , which is a coarse point in $\mathcal{C}(V)$. The map $\rho_V^U : \mathcal{C}V \rightarrow \mathcal{C}^U$ is only really required to be meaningful for points a definite distance from ρ_V^U , but in this context there is a meaningful definition for all curves α that cut U , and one defines $\rho_U^V(\alpha)$ to be the usual subsurface projection of α to U . Our definition of $\rho_E^F(\alpha)$ is compatible with the usual definitions in the following sense.

Lemma 9.5. *Suppose U is nested in V . Consider $\alpha \in \mathcal{C}(V)$ far from ρ_V^U . Then the usual subsurface projection $\rho_U^V(\alpha)$ of α to U lies within bounded distance of set of $\pi_U(X)$ where X ranges over the set of $X \in \mathcal{T}_{g,n}$ where $\pi_V(X)$ is close to α .*

Proof. This follows from axiom (4b) for the usual HHS structure on $\mathcal{T}_{g,n}$. $\square$

Our choice of $\rho_E^F(\alpha)$ was made so that it is immediate from the definition that it lies in $\mathcal{C}(E)$, but Lemma 9.5 shows that we could have defined it in other ways. (Moreover [BHS19, Proposition 1.11] shows that we could have even avoided defining ρ_U^V in this case had we wished to.)

Orthogonality. We say $E = \{U_1, \dots, U_k\}$ is orthogonal to $F = \{V_1, \dots, V_\ell\}$ if every U_i is orthogonal to every V_j .

Transversality and consistency. By definition, $E = \{U_1, \dots, U_k\}$ and $F = \{V_1, \dots, V_\ell\}$ are transverse if they are not orthogonal and neither is nested in the other. To proceed, we need to better understand transverse pairs.

Lemma 9.6. *Let $\{U_1, \dots, U_k\}$ and $\{V_1, \dots, V_\ell\}$ be distinct equivalence classes of active subsurfaces of N . Then, one of the following holds:*

- (1) *Every U_i is nested in some V_j and every V_j has some U_i nested in it; or the same holds with the roles of the U_i 's and V_j 's interchanged.*
- (2) *Every U_i is orthogonal to every V_j .*
- (3) *Every U_i is transverse to some V_j and every V_j is transverse to some U_i .*

Here we use the usual terminology for the HHS structure on Teichmüller space, so U_i is nested in V_j means that either $U_i = V_j$, or U_i is non-annular and is contained up to isotopy in V_j , or U_i is annular with core curve a non-peripheral curve of V_j . Also, U_i is orthogonal to V_j if they are not equal and can be isotoped to be disjoint. As usual, U_i is transverse to V_j if they are neither nested nor orthogonal. More concretely, U_i is transverse to V_j if either their boundaries intersect or each contains a boundary curve of the other non-peripherally.

Proof. We proceed in cases.

Case 1: Both equivalence classes are non-annular. Suppose that the boundary of some U_i intersects the boundary of some V_j . Then, Lemma 6.1 part (3) and Corollary 2.15 give that the boundary of every U_i must intersect the boundary of some V_j and vice versa, so the third conclusion holds.

We can thus assume that all the boundaries do not intersect. Lemma 6.1 part (2) gives that if some U_i contains a boundary curve γ of a V_j non-peripherally, then all U_i contain a boundary curve of a V_j non-peripherally. By symmetry, we also get the claim with the role of the U_i 's and V_j 's swapped.

Now suppose that some U_i contains non-peripherally a boundary component of one of the V_j and that some V_j contains non-peripherally a boundary component of the U_i . In this case the previous paragraph gives that the third conclusion holds.

It thus suffices to consider the case when some U_i contains non-peripherally a boundary component of a V_j , but no V_j contains non-peripherally a boundary component of a U_i . In this case the first conclusion holds, since if U and V are connected subsurfaces with disjoint boundary and V contains a boundary component of U but U does not contain a boundary component of V , then V contains U .

Case 2: $\{U_1, \dots, U_k\}$ is annular and $\{V_1, \dots, V_\ell\}$ is non-annular. In this case each U_i can be thought of as a curve as well as an annular domain. If one of the U_i is contained non-peripherally in one of the V_j , then Lemma 6.1 part (2) gives that that the first conclusion holds. If one of the U_i intersects the boundary of one of the V_j , then Corollary 2.15 and Lemma 6.1 part (3) give that the third conclusion holds.

Case 3: Both equivalence classes are annular. In this case Corollary 2.15 gives that either the second or third conclusion hold. $\square$

It follows from Lemma 9.6 that if $E = \{U_1, \dots, U_k\}$ and $F = \{V_1, \dots, V_\ell\}$ in $\mathfrak{S}_N$ are transverse, then each U_i is intersected by the boundary of some V_j and vice versa. In particular, we can define ρ_F^E using these intersections as in the usual HHS structure on Teichmüller space. Since the boundary components of the U_i can be made simultaneously short, this gives a point of $\mathcal{C}(F)$.

With all relevant definitions complete, we observe that axiom (4a) follows immediately from the corresponding axiom for Teichmüller space. Here one must keep in mind the diagonal like behavior of $\mathcal{C}(E)$ inside of $\prod \mathcal{C}(U_i)$, which gives that two points in $\mathcal{C}(E)$ are close if and only if they are close in any given factor $\mathcal{C}(U_i)$. So in particular, if $d_{U_i}(x, \rho_{U_i}^{V_j})$ is small for a single i, j , then $d_E(x, \rho_E^F)$ is small.

A similar shows (4b), given Lemma 9.5. Axiom (4c) also follows from the corresponding axiom for Teichmüller space because we have defined the relevant ρ 's for N using the corresponding ρ 's for Teichmüller space.

Finite complexity. This follows immediately from the previous definitions.

Large links. We assume E is larger than the constant in Lemma 9.1. Fix an equivalence class $\{W_1, \dots, W_r\} \in \mathfrak{S}_N$ and $x, x' \in \mathcal{T}_{g,n}$. The large links axiom for Teichmüller space gives that we can find a set of at most $\lambda d_{W_1}(x, x') + \lambda$ domains properly nested in W_1 that contain all T properly nested in W_1 with $d_T(x, x') > E$. If these domains are active, then considering their equivalence classes proves large links (in this instance) for N . Otherwise, if one such domain U is not active, then Lemma 9.1 gives that $d_U(x, x')$ is uniformly small, so applying the large links axiom again for Teichmüller space gives a uniformly finite collection of domains properly

nested in U . Iteratively applying large links in this way, we eventually obtain a (possibly empty) uniformly finite set of domains properly nested in U that are active and which serve to prove the large links axiom for N . Compare to the Passing Up Lemma [BHS19, Lemma 2.5].

Bounded geodesic image. This follows directly from the corresponding axiom for Teichmüller space, keeping Lemma 9.5 in mind.

Partial realization. Fix a pairwise orthogonal family of domains in $\mathfrak{S}_N$ and a point in each of the corresponding products of curve graphs. To start, assume none of the domains are annular. By the proof of Corollary 7.2, each point can be realized as a point in the corresponding Teichmüller space after pinching all the boundary curves of every surface in every domain of the family. Using semisimplicity, it follows that the tuple of points can be realized simultaneously on the corresponding pinched surface. Unpinching the nodes we realize the tuple by a point of N where all the corresponding boundary curves are short; see Corollary 3.17. This point proves the partial realization axiom when no annular domains are considered.

The proof when some domains are annular is similar and we only provide a sketch. First pinch and arrange for all non-annular domains to realize the appropriate projections. Then, keeping in mind Lemma 3.16 as well as the proof of Lemma 7.3, one can show that it is possible to unpinch so that each equivalence class of previously pinched curves is approximately any desired small length, and without changing the subsurface projections to the relevant non-annular domains. Then one can apply Dehn multi-twists to adjust the x -coordinates of the image in the curve complexes of the annular domains as required.

Uniqueness. This follows immediately from the uniqueness axiom for Teichmüller space and the uniform bound provided by Lemma 9.1.

Conclusion. This concludes the proof that algebraic totally geodesic submanifolds of Teichmüller space are hierarchically hyperbolic.

9.3. Fundamental groups. Let Γ_N denote the stabilizer of N in the mapping class group. Our goal is now to show that Γ_N is a hierarchically hyperbolic group. We will only sketch this, leaving some details to the reader. Recall that a group is a hierarchically hyperbolic group if it has a suitably nice action on a hierarchically hyperbolic space. This action in particular must be metrically proper and cobounded, and induce an action on the set of domains that is cofinite. See [BHS19, Definition 1.21] for the full definition.

Note that Γ_N is finitely presented because it is the fundamental group of an algebraic variety. To establish that Γ_N is a hierarchically hyperbolic group, we will fix a finite generating set and consider the action of Γ_N on its Cayley graph, which we denote by $\text{Cay}(\Gamma_N)$.

The action of Γ_N on $\text{Cay}(\Gamma_N)$ is certainly metrically proper and cobounded. We will describe a hierarchically hyperbolic structure on $\text{Cay}(\Gamma_N)$ and then complete the proof that Γ_N is a hierarchically hyperbolic group.

Basic set up. The set of domains will be the same as for N , which is to say it will still be $\mathfrak{S}_N$. The set of associated hyperbolic spaces will be almost the same: For E non-annular, we use the same hyperbolic space $\mathcal{C}(E)$. When E is annular, the hyperbolic space we used above is quasi-isometric to a horoball, and we will now replace this horoball with its boundary together with the path metric on its

boundary. That is to say, when E is annular, the new hyperbolic space can be taken to be the real line. We will abuse notation and continue to call this $\mathcal{C}(E)$.

Projections. The projection maps are obtained as follows. Fix $X_0 \in N$. It may be helpful to think of X_0 as being in the thick part of N , but of course every point of N is in the ε -thick part for some $\varepsilon > 0$.

We then define $\pi_E(g)$ to be $\pi_E(gX_0)$. Since gX_0 is uniformly thick, if E is annular, this projection will lie uniformly close to the boundary of the horoball, so this gives a coarsely well defined point even with our new version of annular curve graphs. It will be implicit in our discussion of partial realization that these projections are coarsely surjective.

Nesting, orthogonality, and finite complexity. We are using the same set of domains as for N and we leave the definitions of nesting and orthogonality unchanged.

Transversality and consistency, large links, and bounded geodesic image. This follows immediately from the corresponding statements for N .

Partial realization. Given a collection of orthogonal domains and a point in each of their associated hyperbolic spaces, we use partial realization for N to get an associated point $X \in N$. We then modify this point, as in the proof of Lemma 7.3, to be uniformly thick. We continue to denote the modification by X . Since the Γ_N action on the thick part of N is co-compact, there exists $g \in \Gamma$ such that gX_0 is uniformly close to X , and this g shows that the partial realization holds for the tuple under consideration.

Uniqueness. Consider $g, h \in \Gamma_N$ and suppose g and h are far apart in Γ_N . Because the action of the mapping class group on Teichmüller space is metrically proper, the action of Γ_N on N is metrically proper, so it follows that gX_0 and hX_0 are far apart. We then get uniqueness for Γ_N from uniqueness for N , keeping in mind that the inclusion of the boundary of a horoball into the horoball is metrically proper.

Conclusion. This concludes our sketch that $\text{Cay}(\Gamma_N)$ is hierarchically hyperbolic. The fact that Γ_N acts on N easily implies it acts on $\mathfrak{S}_N$. Since each Γ_N orbit for this action gives a component of the boundary of $\pi(N)$ in $\overline{\mathcal{M}}_{g,n}$, the fact that $\pi(N)$ is a variety gives that there are only finitely many orbits by Corollary 3.14.

9.4. The container axiom. We conclude by commenting a little bit more on the difference between an HHS and an almost HHS. The definition of an almost HHS includes the following as axiom (3c).

The Bounded Orthogonality Axiom: The cardinality of a set of pairwise orthogonal domains is uniformly bounded.

In contrast, the orthogonality axiom of an HHS includes the following.

The Container Axiom: For each domain T , and each domain U nested in T for which there exists a domain V nested in T and orthogonal to U , there exists a domain W properly nested in T which contains all such V .

If the so called container domain W can be additionally chosen to be orthogonal to U then the HHS is said to have clean containers [ABD21, Definition 7.1] and the W can be thought of as the perp of U inside T .

If one uses a slightly larger set of domains than we use above, it is easy to directly verify that the resulting HHS or HHG has clean containers. To do this one only has to make new domains for every orthogonal tuple of old domains; this roughly allows disconnected subsurfaces as well as connected ones. The hyperbolic spaces of the new domains should be viewed as (coarse) points. Given a new domain, which is an orthogonal tuple of the old domains, it has a clean container that can be obtained as follows. Consider the complement of all the boundaries of all the subsurfaces involved. First take the union of the active domains in the complement that do not contain any of the subsurfaces involved. Then add on the annular boundary domains that were not part of the orthogonal tuple.

This has some relevance that we have suppressed above, since [ABD21, Remark A.7] has an omission and in fact there is a mild complication in the equivalence of hierarchically hyperbolicity and almost hierarchically hyperbolicity for HHGs [ABR23, Remark 3.4].

10. THE PROOF OF THEOREM 1.6 AND OPEN QUESTIONS

10.1. About the proof of Theorem 1.6. We can now point out where the various parts of this theorem were proved in the paper.

Proof of Theorem 1.6. Part (1) follows from Corollary 2.7; part (2) from Corollary 2.2; part (3) from Lemma 2.14.

We check part (4) as follows. Given a pinchable multicurve, Corollary 2.12 and Lemma 2.13 show that it is a union of equivalence classes of curves in $\mathcal{S}$. Conversely, given a union of disjoint equivalent classes of curves in $\mathcal{S}$, we can argue as in Corollary 2.7 (using a product of twists given by Lemma 2.6) to find a differential where all the curves are simultaneously cylinders and then pinch these curves as in Corollary 2.12.

The final claim, part (5), is proven as follows. There is nothing to prove if the two equivalence classes of curves are disjoint, so assume they are not disjoint. Let $U_1, \dots, U_k$ be the connected components of the subsurface filled by the two equivalence classes. Corollary 2.15 gives that none of the U_i are annular.

Lemma 2.6 gives multi-twists associated to each of the two equivalence classes. A variant of the Thurston–Veech construction gives a mapping class g which is supported on the union of the U_i 's, which stabilizes N , and which is a pseudo-Anosov on each of the U_i 's.

Fix $X \in N$ and let $Y = g^p X$ for p very large. In particular, $d_V(X, Y)$ is large if V is one of the U_i 's and otherwise it is bounded. It follows from the definition that $\{U_i\}$ is a union of equivalence classes of active subsurfaces. Lemma 9.1 then gives that the boundary of the U_i 's can be pinched. The fact that the boundary of the U_i is a union of equivalence classes of curves is Lemma 6.1 part (3). $\square$

10.2. Open questions. Of the many open questions related to this paper, we highlight first the following:

Question 10.1. Is it true that every totally geodesic submanifold of dimension at least 3 is a covering construction in the sense of [MMW17, Section 6], based off an entire Teichmüller space? (Compare to [AW23, Question 1.3]).

We also highlight just a small number of instances of Metaconjecture 1.2 to emphasize the variety of perspectives that might be interesting here.

Question 10.2. Is there a version of the curve complex for higher dimensional totally geodesic submanifolds that is homotopy equivalent to a wedge of spheres? Are the fundamental groups of totally geodesic subvarieties virtual duality groups? (Compare to [Har86]).

Question 10.3. Are all algebraic totally geodesic submanifolds biholomorphic to bounded domains?

Question 10.4. Can one give nice, explicit presentations for the fundamental groups of the examples in [MMW17, EMMW20]?

We also mention some open questions that are motivated at least in part by comparison to Teichmüller curves.

Question 10.5. Are the surface bundles and surface group extensions associated to totally geodesic subvarieties hierarchically hyperbolic? (Compare to [DDL24]).

Question 10.6. Are algebraic totally geodesic subvarieties rigid? (Compare [McM09] and note the potential relevance of results along the lines of those in [DM21]).

APPENDIX A. ANNULAR PROJECTIONS USING CYLINDERS (AFTER RAFI)

In this appendix we continue to fix a Teichmüller space $\mathcal{T}_{g,n}$, and all the constants we produce continue to implicitly depend on g and n . See Section 2.4 for background on annular curve graphs.

As we recall in the next remark, the usual subsurface projection to annular subsurfaces does not behave as nicely as one would like along Teichmüller geodesics. The purpose of this appendix is to communicate a beautiful idea of Rafi which for some purposes sidesteps this difficulty. We thank Rafi for sharing this with us. The results of this section are only used in Appendix B.

Remark A.1. Recall from Section 2.4 that there are two versions of annular curve graphs. If one uses the projections to the annular curve graphs that are quasi-isometric to $\mathbb{Z}$, Rafi conjectures there is no backtracking, but this conjecture remains open as of the writing of this paper. If, as we do, one uses the version of annular curve graphs that are horoballs, it is well known that there can be backtracking.²

Since there is potential for confusion, we pause briefly to comment on Rafi's results related to this. Rafi has proven a no backtracking result, which, as stated in [Raf14, Theorem B], does not explicitly rule out annular subsurfaces. And indeed Rafi does prove a no backtracking result for some definition of twisting. However the definition of twisting Rafi uses is not what is used in defining maps from Teichmüller space to annular curve graphs, and instead is part of what we use in this appendix. See [Raf07a, Theorem 4.3] for a comparison and [DDM14, Section 2.7] for some related discussion.

²More specifically, the image of a Teichmüller geodesic can make an excursion deep into the horoball $\{z : \text{Im}(z) \geq 1\}$ starting at a point on the boundary $\{z : \text{Im}(z) = 1\}$ and eventually returning to roughly the same point on the boundary. To see this, take for example a flat torus that generates a cobounded geodesic in $\mathcal{T}_1$. Take two copies of this torus, and glue them along an arbitrarily tiny slit of slope 1. The slit curve is hyperbolically short, but if one flows in either direction for a sufficiently long time it is not hyperbolically short. One can use [Raf07a, Theorem 4.3] to verify that all points where the slit curve is not hyperbolically short have roughly the same twisting.

The idea to sidestep this difficulty is to define a variant of subsurface projection for annular subsurfaces. The advantage of this variant is that one simultaneously maintains a version of the uniqueness axiom (axiom (9)) and gets non-backtracking. The disadvantage of this variant is that it does not give a well defined function on all of Teichmüller space. Instead, for each Teichmüller geodesic, one gets a projection defined on this Teichmüller geodesic, which makes use in an essential way of the quadratic differentials generating Teichmüller geodesic. Equivalently, one gets a projection defined on the bundle of non-zero quadratic differentials.

Definition A.2. For any quadratic differential q and any simple closed curve α , define $\pi_\alpha(q)$ to be the point in the horoball $\text{Im}(z) \geq 1$ in the upper half plane $\mathbb{H}$ as follows. If q does not have a cylinder with core curve α , set $\pi_\alpha(q) = 0 + i$. If q does have a cylinder with core curve α , set

$$\pi_\alpha(q) = \text{twist}_q(\alpha) + \max(1, \text{mod}_q(\alpha))i,$$

where typically $\text{twist}_q(\alpha)$ is the twisting at α between the vertical foliation of q and a flat perpendicular to the core curve of α , as defined in [Raf07a, Section 4], and $\text{mod}_q(\alpha)$ is the modulus of the cylinder with core curve α . In the exceptional situation when the cylinder is vertical, we define $\text{twist}_q(\alpha) = 0$.

We warn that special care should be taken with horizontal and vertical cylinders. It is safest to use this definition along one Teichmüller geodesic at a time, in which case it is easy to understand.

If q_t is a Teichmüller geodesic and α is not vertical or horizontal, then [Raf07a, Equation (16)] and [Raf14, Equation(4)] give that there are constants T_α (the total twisting along the geodesic) and t_α (the balanced time) such that

$$\text{twist}_q(\alpha) \stackrel{+}{\asymp} \frac{T_\alpha e^{2(t-t_\alpha)}}{4 \cosh^2(t-t_\alpha) - 2}$$

and

$$\text{mod}_q(\alpha) \stackrel{*}{\asymp} \frac{T_\alpha}{\cosh^2(t-t_\alpha)} \stackrel{*}{\asymp} \frac{T_\alpha}{4 \cosh^2(t-t_\alpha) - 2}.$$

Here, as in [Raf14], $A \stackrel{+}{\asymp} B$ means $|A - B|$ is bounded above by a constant only depending on g and n , and $A \stackrel{*}{\asymp} B$ means A/B is bounded above and below by constants only depending on g and n .

Lemma A.3. *The path*

$$\frac{e^{2t} + i}{e^{2t} + e^{-2t}}$$

is geodesic of speed 2 in the upper half plane $\mathbb{H}$.

This path is illustrated in Figure 5.

Proof. Consider the vertical geodesic $t \mapsto ie^{2t}$ and apply the Möbius transformation $z \mapsto z/(z+1)$. (We thank Saul Schleimer for pointing out this proof to us.) $\square$

Corollary A.4. *If q_t is a Teichmüller geodesic, then $\pi_\alpha(q_t)$ is an unparametrized quasi-geodesic.*

Proof. In the special case when α is vertical or horizontal, this follows easily because the twist is constant and an even simpler formula governs the modulus. The typical case follows from Lemma A.3 by using the hyperbolic isometry $z \mapsto T_\alpha z$.




FIGURE 5. The path of Lemma A.3.

A technical point is that the max in the formula for π_α does not cause problems; a priori one might worry that significant change to the real part might take place during an interval when the imaginary part is constantly equal to 1. For this one should note that, with the formula in Lemma A.3, after t gets negative enough to have that the imaginary part $1/(e^{2t} + e^{-2t})$ has size at most $\delta = 1/T_\alpha$, the x coordinate is within $e^{2t}\delta < \delta$ of its limiting value. A similar observation holds when t is large and positive. $\square$

In the next lemma we give a uniqueness constant θ_u following the usual notation in the uniqueness axiom; the subscript is just a reminder that this relates to uniqueness.

Proposition A.5. *Fix a Teichmüller space $\mathcal{T}_{g,n}$. For all κ there exists θ_u such that if $X, Y \in \mathcal{T}_{g,n}$ have $d(X, Y) \geq \theta_u$ then either*

- (1) *there exists a non-annular V with $d_V(X, Y) \geq \kappa$, or*
- (2) *if q_X and q_Y are the differentials such that $g_T(X, q_X) = (Y, q_Y)$ for some $T \in \mathbb{R}$, then there is a curve α with*

$$d_{\mathbb{H}}(\pi_\alpha(q_X), \pi_\alpha(q_Y)) \geq \kappa.$$

Before giving a proof we recall a standard lemma, which can be derived from Rafi's thick-thin decomposition (with specific ingredients including [Raf07b, Theorem 4] and [Raf14, Theorem 3.1]) or from more recent technology such as [BCG⁺18] (using that q is close to a limit differential with no horizontal nodes and hence more than one level).

Lemma A.6. *For any D and any M , there exists an ε such that if $q \in Q\mathcal{T}_{g,n}$ has no cylinder of modulus greater than D and has a curve of hyperbolic length less than ε , then there is a subsurface U of q bounded by expanding annuli of modulus at least M .*

Here the expanding annuli are taken to be contained in U . This U is a small subsurface, and the expanding annuli mediate between the small scale of U and larger scale adjacent to U .

Proof of Proposition A.5. The cases of $\mathcal{T}_{1,1}$ and $\mathcal{T}_{0,4}$ are straightforward to analyze, either directly or using the ideas below, so we can induct on $\dim \mathcal{T}_{g,n}$ and assume that the complex dimension of $\mathcal{T}_{g,n}$ is at least two and the result is known for smaller dimensional Teichmüller spaces.

Consider the path $g_t(X, q_X)$, $t \in [0, T]$. If at some point this has a cylinder of large modulus with core curve α , then we can directly check that (2) holds using Corollary A.4 and the related formulas, which together give that the modified

annular projection must change when the modulus is large, and that this change cannot be undone. So suppose there is an upper bound on the modulus of cylinders along this path.

If the path stays in the thick part, it makes uniform progress in the main curve graph [RS09, Lemma 4.4]. So it suffices to consider the case when a curve gets very short. In this case Lemma A.6 and the results of [Raf14, Section 4] on projection of a quadratic differential to subsurfaces give that there is a long subinterval where the path follows a geodesic in a smaller dimensional Teichmüller space. The result then follows from the induction hypothesis, keeping in mind Rafi's no backtracking result [Raf14, Theorem B]. $\square$

APPENDIX B. ELECTRIFIED TEICHMÜLLER GEODESICS

In this appendix we continue to fix a Teichmüller space $\mathcal{T}_{g,n}$, and all the constants we produce continue to implicitly depend on g and n . Given a Teichmüller geodesic

$$X_t, t \in \mathbb{R}$$

and a constant $D > 0$, for each (essential, non-peripheral) simple closed curve α we define $I_\alpha \subset \mathbb{R}$ to be the convex hull of the $t \in \mathbb{R}$ such that α has length at most D at X_t . We define $\mathcal{E}$ to be the electrification of $\mathbb{R}$ along these intervals. The purpose of this section is to prove the following. Since this result does not seem to have been previously recorded in the literature we provide a proof here, making use in particular of Rafi's ideas from the previous appendix.

Theorem B.1. *For any $D > 0$ there exists K such that for any Teichmüller geodesic $X_t, t \in \mathbb{R}$ the natural map*

$$\pi : \mathcal{E} \rightarrow \mathcal{C}(\Sigma)$$

is a K -quasi-isometric embedding.

Here Σ is the surface of genus g with n punctures. Note that the fact that π is coarsely Lipschitz follows from the standard fact that the subset of $\mathcal{T}_{g,n}$ where a curve α has length at most D maps to a subset of $\mathcal{C}(\Sigma)$ that is contained in a ball centered at α of radius a function of D . That standard fact is implicit in the result of Masur and Minsky which proves that the curve graph is equal to a similar electrification of Teichmüller space [MM99, Lemma 7.1].

When we apply Theorem B.1 it is helpful to also have the following.

Corollary B.2. *For any $D_2 \geq D_1 > 0$ there exists K such that the following holds. Consider a Teichmüller geodesic $X_t, t \in \mathbb{R}$. For each α , let $I_\alpha \subset \mathbb{R}$ be an interval such that*

$$t \in I_\alpha \implies \ell_\alpha(X_t) \leq D_2, \quad \text{and} \quad t \notin I_\alpha \implies \ell_\alpha(X_t) \geq D_1.$$

Let $\mathcal{E}$ be the electrification of $\mathbb{R}$ along these intervals. Then the natural map

$$\pi : \mathcal{E} \rightarrow \mathcal{C}(\Sigma)$$

is a K -quasi-isometric embedding.

Proof of Corollary B.2 assuming Theorem B.1. Let $\mathcal{E}_1$ be the electrification of $\mathbb{R}$ used in Theorem B.1 with $D = D_1$, and let $\pi_1 : \mathcal{E}_1 \rightarrow \mathcal{C}(\Sigma)$ be the associated map to the curve graph. Theorem B.1 and the discussion above gives that there is a constant K depending only on D_1 and D_2 such that π_1 is a K -quasi-isometric embedding and π is K -Lipschitz.




FIGURE 6. The proof of Corollary B.2.

We also have a 1-Lipschitz inclusion $\iota_1 : \mathcal{E}_1 \rightarrow \mathcal{E}$, and $\pi_1 = \pi \circ \iota_1$. The situation is summarized in Figure 6. The result follows since ι_1 is coarsely surjective. $\square$

Previous work. Theorem B.1 is related to work of Rafi and Schleimer [RS09]. It seems natural to expect that one could approach Theorem B.1 using the analysis in [RS09], but we do not follow that approach here.

We note that Theorem B.1 can be used to reprove the main result of [RS09]. Another proof of the main result of [RS09] was given in [Tan12]. See also [Tan17] for a related result, [FS14] for a version for arc graphs instead of curve graphs, and [APT22] for effective results. Some of these results use an “electric distance” on hyperbolic three manifolds and hence have a similarity in flavor to Theorem B.1; see [Bow, Theorem 2.1.4] and the discussion after the statement of [APT22, Theorem 4.1] for more context on this.

See also [Tan21, Theorem 1.4] for a result of a similar flavor to Theorem B.1 which gives information on Teichmüller discs with sufficiently non-trivial Veech groups.

Active intervals. We start by recalling some of Rafi’s theory of active intervals, established in [Raf07a, Section 3]. See also [DM23, Theorem 3.22] and [DDM14, Section 2.5], for example, for summaries of this theory. It gives in particular the following.

Theorem B.3 (Rafi). *There exists ϵ_0 such that for $0 < \epsilon < \epsilon_0$ there exists $\epsilon' \in (0, \epsilon)$ and $M_{\text{active}} > 0$ such that for any proper subsurface U and any Teichmüller geodesic X_t there is an interval I_U of times such that*

- (1) *all components of ∂U have length at most ϵ at every X_t with $t \in I_U$,*
- (2) *at each $X_t, t \notin I_U$ at least one component of ∂U has length at least ϵ' , and*
- (3) *for any s, t in the same component of $\mathbb{R} - I_U$, we have $d_U(X_s, X_t) \leq M_{\text{active}}$.*

This theorem allows the U to be annular, and in this case we use the version of $\mathcal{C}(U)$ that is a horoball in the hyperbolic plane. Note that I_U can be empty or (in unusual cases) half-infinite.

Rafi [Raf14, Theorem B] also proved that the subsurface projection of a Teichmüller geodesic to a non-annular curve graph gives an unparametrized quasi-geodesic (a quasi-geodesic composed with a monotone time reparametrization).

Metrically proper maps. A map $p : S \rightarrow T$ is called metrically proper if for any d_0 there exists a c_0 such that

$$d(p(s_1), p(s_2)) < d_0 \implies d(s_1, s_2) < c_0.$$

Here it suffices to consider d_0 larger than any fixed constant.

Dowdall and Taylor observed that a metrically proper alignment preserving map must be a quasi-isometric embedding [DT17, Lemma 3.4]. The definition of alignment preserving is such that a coarsely Lipschitz map to a Gromov hyperbolic space that sends geodesics to unparametrized quasi-geodesics is alignment preserving.

Reduction to a key proposition. The following will be the key to the proof.

Proposition B.4. *For any L and any D there exists m and B such that if $X, Y \in \mathcal{T}_{g,n}$ have*

$$d_{\Sigma}(X, Y) < L$$

then there are $\alpha_1, \dots, \alpha_m \in \mathcal{C}(\Sigma)$ and subintervals $I_1, \dots, I_m$ of the Teichmüller geodesic $[X, Y]$ from X to Y such that at each point in I_j we have the length of α_j is at most D , and such that each subinterval of the complement of $\cup I_j$ has length at most B .

Before we prove the proposition, we explain why it is sufficient.

Proof of Theorem B.1 assuming Proposition B.4. The map $\mathcal{E}(\mathcal{I}) \rightarrow \mathcal{C}(\Sigma)$ is alignment preserving, so it suffices to prove it is metrically proper, and that follows immediately from the key proposition. $\square$

A uniformly finite collection of subsurfaces. We now give a construction reminiscent of [DM23] that captures some of the behavior of initial and terminal parts of a Teichmüller geodesic segment.

Since the next lemma is technical, we motivate it by briefly hinting at how it will be applied. Our goal is now to prove Proposition B.4. The theory of active intervals gives that, given a subsurface U , during the interval of times in which a Teichmüller geodesic segment is far from its initial and terminal values in $\mathcal{C}(U)$ the boundary of $\mathcal{C}(U)$ will be short. During the time in which this boundary is short, we do not require any further information.

Lemma B.5. *There exists a $C > 0$ such that for any M there exists an N such that for any $X, Y \in \mathcal{T}_{g,n}$ there exists a set $\mathcal{S} = \mathcal{S}(X, Y)$ of subsurfaces with the following properties.*

- (1) $\Sigma \in \mathcal{S}$.
- (2) $\mathcal{S}$ has cardinality at most N .
- (3) Suppose $U \in \mathcal{S}$ and V is properly nested in U . Suppose there is large progress in V , in the sense that

$$d_V(X, Y) > C.$$

Suppose also that this progress happens in the initial or terminal part of the Teichmüller geodesic as measured in $\mathcal{C}(U)$, in the sense that

$$\min(d_U(\rho_U^V, X), d_U(\rho_U^V, Y)) < M.$$

Then there exists $V' \in \mathcal{S}$ properly nested in U with V nested in V' .

Here ρ_U^V continues to denote the boundary of V viewed as a coarse point of $\mathcal{C}(U)$. A consequence of the final condition is that for any surface V not in $\mathcal{S}$ where large progress occurs, there is a bigger subsurface U in $\mathcal{S}$ which proves that this progress does not happen in the initial or terminal part of the Teichmüller geodesic, as can be seen by looking at a minimal element in $\mathcal{S}$ that contains V .

The proof will use the large links axiom, which gives constants E, λ such that if X and Y are points of Teichmüller space and U is a subsurface then there is a collection $Q = Q_U(X, Y)$ of at most $\lambda d_U(X, Y) + \lambda$ subsurfaces properly nested in U such that if V is properly nested in U and $d_V(X, Y) > E$ then V is nested in one of the surfaces in Q .

Proof. The proof is via an iterative construction. The number of stages in the construction will be bounded by the maximum length of a properly nested chain of subsurfaces, so to prove uniform finiteness it suffices to prove that the number of subsurfaces added to $\mathcal{S}$ in each stage is uniformly bounded. At the first stage we set $\mathcal{S} = \{\Sigma\}$.

At each subsequent stage, we consider all the subsurfaces added to $\mathcal{S}$ in the previous step that are not annular. Let U be such a subsurface, and note that since U is non-annular the map to $\mathcal{C}(U)$ gives an un-parametrized quasi-geodesic. Let H be such that for any Teichmüller geodesic and any non-annular curve graph U there is a monotone reparametrization of the geodesic that stays within distance H of a geodesic in $\mathcal{C}(U)$.

If it exists, let X_U be the first point on the Teichmüller geodesic from X to Y with

$$d_U(X, X_U) > M + H + 10.$$

Similarly let Y_U be the last point with

$$d_U(Y_U, Y) > M + H + 10.$$

If these points do not exist, we will work with the whole geodesic from X to Y . If they do exist, we'll separately work with the segment from X to X_U and also the segment from Y to Y_U .

We now apply the large links axiom to these segments (of which there are one or two per U added in the last stage) with subsurface U , and add all the resulting subsurfaces to $\mathcal{S}$ and proceed to the next stage.

The result is a uniformly finite collection $\mathcal{S}$ of subsurfaces. Now, consider $U \in \mathcal{S}$, and V properly nested in U and assume

$$d_V(X, Y) > C \quad \text{and} \quad d_U(\rho_U^V, X) < M.$$

We will specify C soon. Note that the case $d_U(\rho_U^V, Y) < M$ follows by symmetry from this case.

If X_U is not defined above set $X_U = Y$. It suffices to prove that when C large enough

$$d_V(X, X_U) > E,$$

since this will give that in our construction a subsurface properly nested in U and in which V is nested is added to $\mathcal{S}$ at the stage when the large links axiom is applied to the segment $[X, X_U]$ with subsurface U .

To this end note that the geodesic from $\pi_U(X_U)$ to $\pi_U(Y)$ consists entirely of curves that cut V . Thus

$$\rho_V^U(\pi_U(X_U), \pi_U(Y)) \leq c_{BGI},$$

where c_{BGI} is the constant in the Bounded Geodesic Image Theorem. Hence, keeping in mind that the constants required for the consistency conditions in this context are actually quite small, we get that

$$d_V(X_U, Y) \leq c_{BGI} + 100.$$

So, picking

$$C = E + c_{BGI} + 200,$$

ensures that $d_V(X, X_V) > E$ as desired. $\square$

The proof. We can now prove the key proposition.

Proof of Proposition B.4. Let K be a constant given by [LR11, Theorem A], so that the hyperbolic length of a curve at any point on a Teichmüller geodesic segment is at most K times the maximum of its lengths at the endpoints of the segment.

By Theorem B.3 there is a constant M_{active} such that projection to a subsurface can change by at most M_{active} during intervals where at least one boundary component has length at least D/K . We can easily arrange for this also to apply to the modified annular projections of Appendix A, since the explicit estimates for the modified projections show in particular that progress only happens with the modulus is large, and large modulus implies the core curve is hyperbolically short.

Let $\mathcal{S}$ be as given in Lemma B.5, with any $M > \max(L, M_{\text{active}})$. We can assume that the C produced by that lemma is bigger than L , since the result with a bigger C is weaker than the result with a smaller C .

In general Lemma B.5 gives for any surface U not in $\mathcal{S}$ where large progress occurs, there is a bigger subsurface in $\mathcal{S}$ which proves that this progress does not happen in the initial or terminal part of the Teichmüller geodesic. By taking $M > L$, we ensure that that larger subsurface is never Σ itself.

Let $\alpha_1, \dots, \alpha_m$ be the boundary curves of the subsurfaces in $\mathcal{S}$ together with the set of all curves that are short at X or Y . The number m of such curves depends only on M (and, as always, g and n).

Let I_i be the interval from the first time α_i has length less than D/K to the last time it has length D/K . The main result of [LR11] gives that α_i has length at most D at every point of I_i .

Let X', Y' be in the same component of $[X, Y]$ minus the segments corresponding to the I_i . By Proposition A.5, to conclude this proof it suffices to show that X' and Y' are bounded distance in all non-annular curve graphs and in the modified annular projections of all curves. We will do this by contradiction.

Suppose that U is a non-annular subsurface with

$$d_U(X', Y') > C + M_{\text{active}} + 100.$$

We immediately get that U cannot be in $\mathcal{S}$. It follows that U is properly nested in a proper subsurface V in $\mathcal{S}$, and that the progress in U is not initial or terminal as measured in $\mathcal{C}(V)$. It follows that the active interval for U is contained in the active interval for V , giving a contradiction to the fact that no α_i is short between X' and Y' .

Now suppose that X' and Y' are far apart as measured using the modified annular projection of Appendix A for some curve α . It follows from Corollary A.4 that X and Y are far apart using the same modified annular projection. Rafi has proven that when α is not short, the regular and modified annular projections agree [Raf07a, Theorem 4.3], so we either get that X and Y are far apart using the usual annular projections or we get that α is short at X or Y .

Case 1: α is short at X or Y . Then $\alpha = \alpha_i$ for some i , and we get that α is not short along $[X', Y']$. This contradicts the fact that the geodesic makes progress in the modified annular projection during this time.

Case 2: X and Y are far apart using the usual annular projections and α is not short at X or Y . Then we know from Lemma B.5 that α must be nested in a proper subsurface V which is contained in $\mathcal{S}$, and that in $\mathcal{C}(V)$ we have that α lies far from X and far from Y . The progress of the modified projection thus happens far from X or Y in $\mathcal{C}(V)$, showing that the interval when the modified projection changes is contained in the active interval for V , again giving a contradiction. $\square$

REFERENCES

- [ABD21] Carolyn Abbott, Jason Behrstock, and Matthew Gentry Durham, *Largest acylindrical actions and stability in hierarchically hyperbolic groups*, Trans. Amer. Math. Soc. Ser. B **8** (2021), 66–104, With an appendix by Daniel Berlyne and Jacob Russell.
- [ABR23] Carolyn Abbott, Jason Behrstock, and Jacob Russell, *Structure invariant properties of the hierarchically hyperbolic boundary*, 2023, arXiv:2208.07930.
- [AN20] David Auricino and Duc-Manh Nguyen, *Rank 2 affine manifolds in genus 3*, J. Differential Geom. **116** (2020), no. 2, 205–280.
- [Ant17a] Stergios M. Antonakoudis, *The complex geometry of Teichmüller spaces and bounded symmetric domains*, In the tradition of Ahlfors-Bers. VII, Contemp. Math., vol. 696, Amer. Math. Soc., Providence, RI, 2017, pp. 1–16.
- [Ant17b] ———, *Isometric disks are holomorphic*, Invent. Math. **207** (2017), no. 3, 1289–1299.
- [Ant17c] ———, *Teichmüller spaces and bounded symmetric domains do not mix isometrically*, Geom. Funct. Anal. **27** (2017), no. 3, 453–465.
- [Api18] Paul Apisa, *$GL_2\mathbb{R}$ orbit closures in hyperelliptic components of strata*, Duke Math. J. **167** (2018), no. 4, 679–742.
- [Api25] ———, *Invariant Subvarieties of Minimal Homological Dimension, Zero Lyapunov Exponents, and Monodromy*, Geom. Funct. Anal. **35** (2025), no. 1, 113–160.
- [APT22] Tarik Aougab, Priyam Patel, and Samuel J. Taylor, *Covers of surfaces, Kleinian groups and the curve complex*, J. Topol. **15** (2022), no. 4, 1833–1863.
- [AW22] Paul Apisa and Alex Wright, *Generalizations of the Eierlegende-Wollmilchsau*, Camb. J. Math. **10** (2022), no. 4, 859–933.
- [AW23] ———, *High rank invariant subvarieties*, Ann. of Math. (2) **198** (2023), no. 2, 657–726.
- [AW24] ———, *Reconstructing orbit closures from their boundaries*, Mem. Amer. Math. Soc. **298** (2024), no. 1487, v+141.
- [Bai06] Matthew Charles Bainbridge, *Euler characteristics of Teichmüller curves in genus two*, ProQuest LLC, Ann Arbor, MI, 2006, Thesis (Ph.D.)—Harvard University.
- [BCG⁺18] Matt Bainbridge, Dawei Chen, Quentin Gendron, Samuel Grushevsky, and Martin Möller, *Compactification of strata of Abelian differentials*, Duke Math. J. **167** (2018), no. 12, 2347–2416.
- [BCG⁺19a] Matt Bainbridge, Dawei Chen, Quentin Gendron, Samuel Grushevsky, and Martin Möller, *The moduli space of multi-scale differentials*, 2019, arXiv:1910.13492.
- [BCG⁺19b] Matt Bainbridge, Dawei Chen, Quentin Gendron, Samuel Grushevsky, and Martin Möller, *Strata of k -differentials*, Algebr. Geom. **6** (2019), no. 2, 196–233.
- [BDG22] Frederik Benirschke, Benjamin Dozier, and Samuel Grushevsky, *Equations of linear subvarieties of strata of differentials*, Geom. Topol. **26** (2022), no. 6, 2773–2830.
- [BDR24] Frederik Benirschke, Benjamin Dozier, and John Rached, *The boundary of a totally geodesic subvariety of moduli space*, 2024, arXiv:2412.06765.
- [Ben23] Frederik Benirschke, *The boundary of linear subvarieties*, J. Eur. Math. Soc. (JEMS) **25** (2023), no. 11, 4521–4582.
- [Ber74] Lipman Bers, *On spaces of Riemann surfaces with nodes*, Bull. Amer. Math. Soc. **80** (1974), 1219–1222.
- [Ber81] ———, *Finite-dimensional Teichmüller spaces and generalizations*, Bull. Amer. Math. Soc. (N.S.) **5** (1981), no. 2, 131–172.
- [BFLS16] L. Birbrair, A. Fernandes, D. T. Lê, and J. E. Sampaio, *Lipschitz regular complex algebraic sets are smooth*, Proc. Amer. Math. Soc. **144** (2016), no. 3, 983–987.

- [BHS17a] Jason Behrstock, Mark F. Hagen, and Alessandro Sisto, *Asymptotic dimension and small-cancellation for hierarchically hyperbolic spaces and groups*, Proc. Lond. Math. Soc. (3) **114** (2017), no. 5, 890–926.
- [BHS17b] ———, *Hierarchically hyperbolic spaces, I: Curve complexes for cubical groups*, Geom. Topol. **21** (2017), no. 3, 1731–1804.
- [BHS19] Jason Behrstock, Mark Hagen, and Alessandro Sisto, *Hierarchically hyperbolic spaces II: Combination theorems and the distance formula*, Pacific J. Math. **299** (2019), no. 2, 257–338.
- [BHS21] Jason Behrstock, Mark F. Hagen, and Alessandro Sisto, *Quasiflats in hierarchically hyperbolic spaces*, Duke Math. J. **170** (2021), no. 5, 909–996.
- [BKMM12] Jason Behrstock, Bruce Kleiner, Yair Minsky, and Lee Mosher, *Geometry and rigidity of mapping class groups*, Geom. Topol. **16** (2012), no. 2, 781–888.
- [Bow] Brian Bowditch, *The Ending Lamination Theorem*, <https://homepages.warwick.ac.uk/~masgak/papers/elt-fourth.pdf>.
- [BR20] Federico Berlai and Bruno Robbio, *A refined combination theorem for hierarchically hyperbolic groups*, Groups Geom. Dyn. **14** (2020), no. 4, 1127–1203.
- [BR22] Daniel Berlyne and Jacob Russell, *Hierarchical hyperbolicity of graph products*, Groups Geom. Dyn. **16** (2022), no. 2, 523–580.
- [BS23] Frederik Benirschke and Carlos A. Serván, *Isometric embeddings of Teichmüller spaces are covering constructions*, 2023, arXiv:2305.04153.
- [Che20] Marissa Chesser, *Stable subgroups of the Genus Two Handlebody Group*, 2020, arXiv:2009.05067.
- [Che22] Dawei Chen, *Dynamical invariants and intersection theory on the flex and gothic loci*, Eur. J. Math. **8** (2022), no. suppl. 1, S42–S52.
- [CMS23] Dawei Chen, Martin Möller, and Adrien Sauvaget, *Masur-Veech volumes and intersection theory: the principal strata of quadratic differentials*, Duke Math. J. **172** (2023), no. 9, 1735–1779.
- [CR24] Aaron Calderon and Jacob Russell, *Hierarchical hyperbolicity of admissible curve graphs and the boundary of marked strata*, 2024, arXiv:2409.06798.
- [CW21] Dawei Chen and Alex Wright, *The WYSIWYG compactification*, J. Lond. Math. Soc. (2) **103** (2021), no. 2, 490–515.
- [DDL24] Spencer Dowdall, Matthew G. Durham, Christopher J. Leininger, and Alessandro Sisto, *Extensions of Veech groups II: Hierarchical hyperbolicity and quasi-isometric rigidity*, Comment. Math. Helv. **99** (2024), no. 1, 149–228.
- [DDM14] Spencer Dowdall, Moon Duchin, and Howard Masur, *Statistical hyperbolicity in Teichmüller space*, Geom. Funct. Anal. **24** (2014), no. 3, 748–795.
- [DM21] Georgios Daskalopoulos and Chikako Mese, *Infinite energy maps and rigidity*, 2021, arXiv:2112.13961.
- [DM23] Spencer Dowdall and Howard Masur, *Lattice point counting for finite-order mapping classes*, 2023, arXiv:2305.03721.
- [DMS20] Matthew G. Durham, Yair N. Minsky, and Alessandro Sisto, *Stable cubulations, bicomings, and barycenters*, 2020, arXiv:2009.13647.
- [DT17] Spencer Dowdall and Samuel J. Taylor, *The co-surface graph and the geometry of hyperbolic free group extensions*, J. Topol. **10** (2017), no. 2, 447–482.
- [Dur16] Matthew Gentry Durham, *The augmented marking complex of a surface*, J. Lond. Math. Soc. (2) **94** (2016), no. 3, 933–969.
- [EFW18] Alex Eskin, Simion Filip, and Alex Wright, *The algebraic hull of the Kontsevich-Zorich cocycle*, Ann. of Math. (2) **188** (2018), no. 1, 281–313.
- [EKZ14] Alex Eskin, Maxim Kontsevich, and Anton Zorich, *Sum of Lyapunov exponents of the Hodge bundle with respect to the Teichmüller geodesic flow*, Publ. Math. Inst. Hautes Études Sci. **120** (2014), 207–333.
- [EM18] Alex Eskin and Maryam Mirzakhani, *Invariant and stationary measures for the $SL(2, \mathbb{R})$ action on moduli space*, Publ. Math. Inst. Hautes Études Sci. **127** (2018), 95–324.
- [EMM15] Alex Eskin, Maryam Mirzakhani, and Amir Mohammadi, *Isolation, equidistribution, and orbit closures for the $SL(2, \mathbb{R})$ action on moduli space*, Ann. of Math. (2) **182** (2015), no. 2, 673–721.

- [EMMW20] Alex Eskin, Curtis T. McMullen, Ronen E. Mukamel, and Alex Wright, *Billiards, quadrilaterals and moduli spaces*, J. Amer. Math. Soc. **33** (2020), no. 4, 1039–1086.
- [EMR17] Alex Eskin, Howard Masur, and Kasra Rafi, *Large-scale rank of Teichmüller space*, Duke Math. J. **166** (2017), no. 8, 1517–1572.
- [Fil16] Simion Filip, *Splitting mixed Hodge structures over affine invariant manifolds*, Ann. of Math. (2) **183** (2016), no. 2, 681–713.
- [FS14] David Futer and Saul Schleimer, *Cusp geometry of fibered 3-manifolds*, Amer. J. Math. **136** (2014), no. 2, 309–356.
- [GG22] Dmitri Gekhtman and Mark Greenfield, *Isometric submersions of Teichmüller spaces are forgetful*, Israel J. Math. **247** (2022), no. 2, 499–517.
- [Gou21] Élise Goujard, *Sous-variétés totalement géodésiques des espaces de modules de Riemann [d’après Eskin, McMullen, Mukamel, Wright]*, Astérisque (2021), no. 430, Séminaire Bourbaki. Vol. 2019/2021. Exposés 1166–1180, Exp. No. 1178, 407–424.
- [Har86] John L. Harer, *The virtual cohomological dimension of the mapping class group of an orientable surface*, Invent. Math. **84** (1986), no. 1, 157–176.
- [Hat] Allen Hatcher, *Vector Bundles and K-Theory*, <https://pi.math.cornell.edu/~hatcher/VBKT/VB.pdf>.
- [HHP23] Thomas Haettel, Nima Hoda, and Harry Petyt, *Coarse injectivity, hierarchical hyperbolicity and semihyperbolicity*, Geom. Topol. **27** (2023), no. 4, 1587–1633.
- [HMS24] Mark Hagen, Alexandre Martin, and Alessandro Sisto, *Extra-large type Artin groups are hierarchically hyperbolic*, Math. Ann. **388** (2024), no. 1, 867–938.
- [HRSS22] Mark Hagen, Jacob Russell, Alessandro Sisto, and Davide Spriano, *Equivariant hierarchically hyperbolic structures for 3-manifold groups via quasimorphisms*, 2022, arXiv:2206.12244.
- [Hug22] Sam Hughes, *Lattices in a product of trees, hierarchically hyperbolic groups and virtual torsion-freeness*, Bull. Lond. Math. Soc. **54** (2022), no. 4, 1413–1419.
- [KMS86] Steven Kerckhoff, Howard Masur, and John Smillie, *Ergodicity of billiard flows and quadratic differentials*, Ann. of Math. (2) **124** (1986), no. 2, 293–311.
- [Kop23] Michael Kopreski, *Multiarc and curve graphs are hierarchically hyperbolic*, 2023, arXiv:2311.04356.
- [KW22] Jeremy Kahn and Alex Wright, *Hodge and Teichmüller*, J. Mod. Dyn. **18** (2022), 149–160.
- [Lo91] Stanisław Łojasiewicz, *Introduction to complex analytic geometry*, Birkhäuser Verlag, Basel, 1991, Translated from the Polish by Maciej Klimek.
- [LR11] Anna Lenzhen and Kasra Rafi, *Length of a curve is quasi-convex along a Teichmüller geodesic*, J. Differential Geom. **88** (2011), no. 2, 267–295.
- [Mas85] Bernard Maskit, *Comparison of hyperbolic and extremal lengths*, Ann. Acad. Sci. Fenn. Ser. A I Math. **10** (1985), 381–386.
- [McM07] Curtis T. McMullen, *Dynamics of $SL_2(\mathbb{R})$ over moduli space in genus two*, Ann. of Math. (2) **165** (2007), no. 2, 397–456.
- [McM09] ———, *Rigidity of Teichmüller curves*, Math. Res. Lett. **16** (2009), no. 4, 647–649.
- [McM23] ———, *Billiards and Teichmüller curves*, Bull. Amer. Math. Soc. (N.S.) **60** (2023), no. 2, 195–250.
- [Min92] Yair N. Minsky, *Harmonic maps, length, and energy in Teichmüller space*, J. Differential Geom. **35** (1992), no. 1, 151–217.
- [Min96] ———, *Extremal length estimates and product regions in Teichmüller space*, Duke Math. J. **83** (1996), no. 2, 249–286.
- [MM99] Howard A. Masur and Yair N. Minsky, *Geometry of the complex of curves. I. Hyperbolicity*, Invent. Math. **138** (1999), no. 1, 103–149.
- [MM00] H. A. Masur and Y. N. Minsky, *Geometry of the complex of curves. II. Hierarchical structure*, Geom. Funct. Anal. **10** (2000), no. 4, 902–974.
- [MMW17] Curtis T. McMullen, Ronen E. Mukamel, and Alex Wright, *Cubic curves and totally geodesic subvarieties of moduli space*, Ann. of Math. (2) **185** (2017), no. 3, 957–990.
- [MW02] Yair Minsky and Barak Weiss, *Nondivergence of horocyclic flows on moduli space*, J. Reine Angew. Math. **552** (2002), 131–177.
- [MW17] Maryam Mirzakhani and Alex Wright, *The boundary of an affine invariant submanifold*, Invent. Math. **209** (2017), no. 3, 927–984.

- [NQ24] Hoang Thanh Nguyen and Yulan Qing, *Sublinearly Morse boundary of CAT(0) admissible groups*, J. Group Theory **27** (2024), no. 4, 857–897.
- [NW14] Duc-Manh Nguyen and Alex Wright, *Non-Veech surfaces in $\mathcal{H}^{\text{hyp}}(4)$ are generic*, Geom. Funct. Anal. **24** (2014), no. 4, 1316–1335.
- [Pet21] Harry Petyt, *Mapping class groups are quasicubical*, 2021, arXiv:2112.10681.
- [Raf07a] Kasra Rafi, *A combinatorial model for the Teichmüller metric*, Geom. Funct. Anal. **17** (2007), no. 3, 936–959.
- [Raf07b] ———, *Thick-thin decomposition for quadratic differentials*, Math. Res. Lett. **14** (2007), no. 2, 333–341.
- [Raf14] ———, *Hyperbolicity in Teichmüller space*, Geom. Topol. **18** (2014), no. 5, 3025–3053.
- [Ree04] Mary Rees, *Teichmüller distance is not $C^{2+\epsilon}$* , Proc. London Math. Soc. (3) **88** (2004), no. 1, 114–134.
- [Roy71] H. L. Royden, *Automorphisms and isometries of Teichmüller space*, Advances in the Theory of Riemann Surfaces (Proc. Conf., Stony Brook, N.Y., 1969), Ann. of Math. Stud., No. 66, Princeton Univ. Press, Princeton, NJ, 1971, pp. 369–383.
- [RS09] Kasra Rafi and Saul Schleimer, *Covers and the curve complex*, Geom. Topol. **13** (2009), no. 4, 2141–2162.
- [RST23] Jacob Russell, Davide Spriano, and Hung Cong Tran, *Convexity in hierarchically hyperbolic spaces*, Algebr. Geom. Topol. **23** (2023), no. 3, 1167–1248.
- [Rus21] Jacob Russell, *Extensions of multicurve stabilizers are hierarchically hyperbolic*, 2021, arXiv:2107.14116.
- [Sis19] Alessandro Sisto, *What is a hierarchically hyperbolic space?*, Beyond hyperbolicity, London Math. Soc. Lecture Note Ser., vol. 454, Cambridge Univ. Press, Cambridge, 2019, pp. 117–148.
- [Str84] Kurt Strebel, *Quadratic differentials*, Ergebnisse der Mathematik und ihrer Grenzgebiete (3) [Results in Mathematics and Related Areas (3)], vol. 5, Springer-Verlag, Berlin, 1984.
- [Tan12] Robert Tang, *The curve complex and covers via hyperbolic 3-manifolds*, Geom. Dedicata **161** (2012), 233–237.
- [Tan17] ———, *Projections in the curve complex arising from covering maps*, Pacific J. Math. **291** (2017), no. 1, 213–239.
- [Tan21] ———, *Affine diffeomorphism groups are undistorted*, J. Lond. Math. Soc. (2) **104** (2021), no. 2, 747–769.
- [TT20] David Torres-Teigell, *Masur-Veech volume of the gothic locus*, J. Lond. Math. Soc. (2) **102** (2020), no. 1, 405–436.
- [TW18] Robert Tang and Richard C. H. Webb, *Shadows of Teichmüller discs in the curve graph*, Int. Math. Res. Not. IMRN (2018), no. 11, 3301–3341.
- [Vok22] Kate M. Vokes, *Hierarchical hyperbolicity of graphs of multicurves*, Algebr. Geom. Topol. **22** (2022), no. 1, 113–151.
- [Whi65] Hassler Whitney, *Tangents to an analytic variety*, Ann. of Math. (2) **81** (1965), 496–549.
- [Win24] Karl Winsor, *Saturated orbit closures in the Hodge bundle*, Int. Math. Res. Not. IMRN (2024), no. 6, 4710–4737.
- [Wol90] Scott A. Wolpert, *The hyperbolic metric and the geometry of the universal curve*, J. Differential Geom. **31** (1990), no. 2, 417–472.
- [Wol10] ———, *Families of Riemann surfaces and Weil-Petersson geometry*, CBMS Regional Conference Series in Mathematics, vol. 113, Published for the Conference Board of the Mathematical Sciences, Washington, DC; by the American Mathematical Society, Providence, RI, 2010.
- [Wri15a] Alex Wright, *Cylinder deformations in orbit closures of translation surfaces*, Geom. Topol. **19** (2015), no. 1, 413–438.
- [Wri15b] ———, *Translation surfaces and their orbit closures: an introduction for a broad audience*, EMS Surv. Math. Sci. **2** (2015), no. 1, 63–108.
- [Wri20] ———, *Totally geodesic submanifolds of Teichmüller space*, J. Differential Geom. **115** (2020), no. 3, 565–575.
- [Ygo23] Florent Ygouf, *Nonarithmetic affine invariant orbifolds in $\mathcal{H}^{\text{odd}}(2, 2)$ and $\mathcal{H}(3, 1)$* , Geom. Dedicata **217** (2023), no. 4, Paper No. 72, 24.