

SKINNING MAPS ARE FINITE-TO-ONE

DAVID DUMAS

1. INTRODUCTION

Skinning maps were introduced by William Thurston in the proof of the Geometrization Theorem for Haken 3-manifolds (see [Ota]). At a key step in the proof one has a compact 3-manifold M with nonempty boundary whose interior admits a hyperbolic structure. The interplay between deformations of the hyperbolic structure and the topology of M and ∂M determines a holomorphic map of Teichmüller spaces, the *skinning map*

$$\sigma_M : \mathcal{T}(\partial_0 M) \rightarrow \mathcal{T}(\overline{\partial_0 M}),$$

where $\partial_0 M$ is the union of the non-torus boundary components and $\overline{\partial_0 M}$ denotes the boundary with the opposite orientation. The problem of finding a hyperbolic structure on a related closed manifold is solved by showing that the composition of σ_M with a certain isometry $\tau : \mathcal{T}(\overline{\partial_0 M}) \rightarrow \mathcal{T}(\partial_0 M)$ has a fixed point.

Thurston's original approach to the fixed point problem involved extending the skinning map of an acylindrical manifold to a continuous map $\hat{\sigma}_M : \text{AH}(M) \rightarrow \mathcal{T}(\overline{\partial_0 M})$ defined on the compact space $\text{AH}(M)$ of hyperbolic structures on M . (A proof of this Bounded Image Theorem can be found in [Ken, Sec. 9].) McMullen provided an alternate approach based on an analytic study of the differential of the skinning map [McM1] [McM2]. In each case there are additional complications when M has essential cylinders.

More recently, Kent studied the diameter of the image of the skinning map (in cases when it is finite), producing examples where this diameter is arbitrarily large or small and relating the diameter to hyperbolic volume and the depth of an embedded collar around the boundary [Ken]. However, beyond the contraction and boundedness properties used to solve the fixed point problem—and the result of Kent and the author that skinning maps are never constant [DK3]—little is known about skinning maps in general.

Our main theorem concerns the fibers of skinning maps:

Theorem A. *Skinning maps are finite-to-one. That is, let M be a compact oriented 3-manifold whose boundary is nonempty and not a union of tori. Suppose that the interior of M admits a complete hyperbolic structure without accidental parabolics, so that M has an associated skinning map σ_M . Then for each $X \in \mathcal{T}(\overline{\partial_0 M})$, the preimage $\sigma_M^{-1}(X)$ is finite.*

Date: June 25, 2015 (revised). March 1, 2012 (original).

Work partially supported by the National Science Foundation.

As a holomorphic map with finite fibers, it follows from Theorem A that skinning maps are open (answering a question in [DK3]) and locally bi-holomorphic away from the analytic hypersurface defined by the vanishing of the Jacobian determinant. Thus our results give strong nondegeneracy properties for all skinning maps.

Our proof of Theorem A does not bound the size of the finite set $\sigma_M^{-1}(X)$; instead, we show that each fiber of the skinning map is both compact and discrete. In particular it is not clear if the number of preimages of a point is uniformly bounded over $\mathcal{T}(\overline{\partial_0 M})$.

In case M is acylindrical, we also show that Thurston's extension of the skinning map to $\text{AH}(M)$ is finite-to-one; this result appears as Theorem 10.1 below.

The intersection problem. To study the fibers of the skinning map we translate the problem to one of intersections of certain subvarieties of the $\text{SL}_2\mathbb{C}$ -character variety of $\partial_0 M$. The same reduction to an intersection problem is used in [DK3]. The relevant subvarieties are:

- The *extension variety* \mathcal{E}_M , which is the smallest closed algebraic subvariety containing the characters of all homomorphisms $\pi_1(\partial_0 M) \rightarrow \text{SL}_2\mathbb{C}$ that can be extended to $\pi_1 M \rightarrow \text{SL}_2\mathbb{C}$.
- The *holonomy variety* \mathcal{H}_X , which is the analytic subvariety consisting of characters of the $\text{SL}_2\mathbb{C}$ -lifts of holonomy representations of \mathbb{CP}^1 -structures on $\overline{\partial_0 M}$ compatible with the complex structure $X \in \mathcal{T}(\overline{\partial_0 M})$.

Precise definitions of these objects are provided in Section 6, with additional details of the disconnected boundary case in Section 9.

The main theorem is derived from the following result about the intersections of holonomy and extension varieties, the proof of which occupies most of the paper.

Theorem B (Intersection Theorem). *Let M be an oriented 3-manifold with nonempty boundary that is not a union of tori. Let X be a marked Riemann surface structure on $\partial_0 M$. Then the intersection $\mathcal{H}_X \cap \mathcal{E}_M$ is a discrete subset of the character variety.*

This theorem applies in a more general setting than the specific intersection problem arising from skinning maps. For example, while skinning maps are defined for manifolds with incompressible boundary, such a hypothesis is not needed in Theorem B.

While the theorem above involves an oriented manifold M , the set \mathcal{E}_M is independent of the orientation. Thus we also obtain discreteness of intersections $\mathcal{H}_{\overline{X}} \cap \mathcal{E}_M$ where \overline{X} is a Riemann surface structure on $\partial_0 M$ that induces an orientation opposite that of the boundary orientation of M .

Steps to the intersection theorem. Our study of $\mathcal{H}_X \cap \mathcal{E}_M$ is based on the parameterization of the irreducible components of \mathcal{H}_X by the vector

space $Q(X)$ of holomorphic quadratic differentials; this parameterization is the *holonomy map* of \mathbb{CP}^1 -structures, denoted by “hol”. The overall strategy is to show that the preimage of \mathcal{E}_M , i.e. the set

$$\mathcal{V}_M = \text{hol}^{-1}(\mathcal{E}_M),$$

is a complex analytic subvariety of $Q(X)$ that is subject to certain constraints on its behavior at infinity, and ultimately to show that only a discrete set can satisfy these.

We now sketch the main steps of the argument and state some intermediate results of independent interest. In this sketch we restrict attention to the case of a 3-manifold M with *connected* boundary S . Let $X \in \mathcal{T}(S)$ be a marked Riemann surface structure on the boundary.

Step 1. Construction of an isotropic cone in the space of measured foliations.

The defining property of this cone in $\mathcal{MF}(S)$ is that it determines which quadratic differentials $\phi \in Q(X)$ have dual trees T_ϕ that admit “nice” equivariant maps into trees on which $\pi_1 M$ acts by isometries. Because of the way this cone is used in a later step of the argument, here we must consider actions of $\pi_1 M$ not only \mathbb{R} -trees but also on Λ -trees, where Λ is a more general ordered abelian group. And while an isometric embedding of T_ϕ into an \mathbb{R} -tree on which $\pi_1 M$ acts is the prototypical example of a nice map, we must also consider the *straight maps* defined in [D] and certain partially-defined maps arising from non-isometric trees with the same length function.

In the case of straight maps, we show:

Theorem C. *There is an isotropic piecewise linear cone $\mathcal{L}_{M,X} \subset \mathcal{MF}(S)$ with the following property: If $\pi_1 M$ acts on a Λ -tree T by isometries, and if $\phi \in Q(X)$ is a quadratic differential whose dual tree T_ϕ admits an equivariant straight map $T_\phi \rightarrow T$, then the horizontal foliation of ϕ lies in $\mathcal{L}_{M,X}$.*

This result is stated precisely in Theorem 4.1 below, and Theorem 4.4 presents a further refinement that is used in the proof of the main theorem.

Step 2. Limit points of \mathcal{V}_M have foliations in the isotropic cone.

In [D] we analyze the large-scale behavior of the holonomy map, showing that straight maps arise naturally when comparing limits in the Morgan-Shalen compactification—which are represented by actions of $\pi_1 S$ on \mathbb{R} -trees—to the dual trees of limit quadratic differentials. The main result can be summarized as follows:

Theorem. *If a divergent sequence in $Q(X)$ can be rescaled to have limit ϕ , then any Morgan-Shalen limit of the associated holonomy representations corresponds to an \mathbb{R} -tree T that admits an equivariant straight map $T_\phi \rightarrow T$.*

Here rescaling of the divergent sequence uses the action of \mathbb{R}^+ on $Q(X)$. The precise limit result we use is stated in Theorem 6.4, and other related results and discussion can be found in [D].

When we restrict attention to the subset $\mathcal{V}_M \subset Q(X)$, the associated holonomy representations lie in \mathcal{E}_M and so they arise as compositions of $\pi_1 M$ -representations with the map $i_* : \pi_1 S \rightarrow \pi_1 M$ induced by the inclusion of S as the boundary of M . (Strictly speaking, this describes a Zariski open subset of \mathcal{E}_M .) We think of these as representations that “extend” from $i_*(\pi_1 S)$ to the larger group $\pi_1 M$.

Note that *a priori* the passage from a $\pi_1 S$ -representation to a $\pi_1 M$ -representation could radically change the geometry of the associated action on \mathbb{H}^3 , as measured for example by the diameter of the orbit of a finite generating set. This possibility, combined with the need to keep track of the limiting $\pi_1 S$ - and $\pi_1 M$ -dynamics simultaneously, requires us to consider Λ -trees more general than \mathbb{R} -trees.

Using the valuation constructions of Morgan-Shalen, we show that there is a similar extension property for the trees obtained as limits points of \mathcal{E}_M , or more precisely, for their length functions (Theorem 6.3). In the generic case of a non-abelian action, the combination of this construction with the holonomy limit theorem above gives an \mathbb{R}^n -tree \widehat{T} (where $\mathbb{R}^n = \Lambda$ is given the lexicographical order) on which $\pi_1 M$ acts by isometries and a straight map

$$T_\phi \rightarrow \widehat{T}$$

where ϕ is the rescaled limit of a divergent sequence in \mathcal{V}_M . These satisfy the hypotheses of Theorem C, so the horizontal foliation of ϕ lies in $\mathcal{L}_{M,X}$. Limit points of \mathcal{E}_M that correspond to abelian actions introduce minor additional complications that are handled by Theorem 4.4.

Step 3. The foliation map $\mathcal{F} : Q(X) \rightarrow \mathcal{MF}(S)$ is symplectic.

Hubbard and Masur showed that the foliation map $\mathcal{F} : Q(X) \rightarrow \mathcal{MF}(S)$ is a homeomorphism. In order to use the isotropic cone $\mathcal{L}_{M,X}$ to understand the set \mathcal{V}_M , we analyze the relation between the foliation map, the complex structure of $Q(X)$, and the symplectic structure of $\mathcal{MF}(S)$.

We introduce a natural Kähler structure on $Q(X)$ corresponding to the Weil-Petersson-type hermitian pairing

$$\langle \psi_1, \psi_2 \rangle_\phi = \int_X \frac{\psi_1 \bar{\psi}_2}{4|\phi|}$$

Here we have a base point $\phi \in Q(X)$ and the quadratic differentials $\psi_i \in T_\phi Q(X) \simeq Q(X)$ are considered as tangent vectors. This integral pairing is not smooth, nor even well-defined for all tangent vectors, due to singularities of the integrand coming from higher-order zeros of ϕ . However we show that the pairing does give a well-defined Kähler structure relative to a stratification of $Q(X)$.

We show that the underlying symplectic structure of this Kähler metric is the one pulled back from $\mathcal{MF}(S)$ by the foliation map:

Theorem D. *For any $X \in \mathcal{T}(S)$, the map $\mathcal{F} : Q(X) \rightarrow \mathcal{MF}(S)$ is a real-analytic stratified symplectomorphism, where $Q(X)$ is given the symplectic structure coming from the pairing $\langle \psi_1, \psi_2 \rangle_\phi$ and where $\mathcal{MF}(S)$ has the Thurston symplectic form.*

The lack of a smooth structure on $\mathcal{MF}(S)$ means that the regularity aspect of this result must be interpreted carefully. We show that for any point $\phi \in Q(X)$ there is a neighborhood in its stratum and a train track chart containing $\mathcal{F}(\phi)$ in which the local expression of the foliation map is real-analytic and symplectic. The details are given in Theorem 5.8.

Step 4. Analytic sets with totally real limits are discrete.

In a Kähler manifold, an isotropic submanifold is totally real. While the piecewise linear cone $\mathcal{L}_{M,X}$ is not globally a manifold, Theorem D allows us to describe $\mathcal{F}^{-1}(\mathcal{L}_{M,X})$ locally in a stratum of $Q(X)$ in terms of totally real, real-analytic submanifolds. Since limit points of \mathcal{V}_M correspond to elements of $\mathcal{F}^{-1}(\mathcal{L}_{M,X})$, this gives a kind of “totally real” constraint on the behavior of \mathcal{V}_M at infinity.

To formulate this constraint we consider the set $\partial_{\mathbb{R}}\mathcal{V}_M \subset S^{2N-1}$ of points in the unit sphere of $Q(X) \simeq \mathbb{C}^N$ that can be obtained as \mathbb{R}^+ -rescaled limits of divergent sequences in \mathcal{V}_M . Projecting this set through the Hopf fibration $S^{2N-1} \rightarrow \mathbb{C}\mathbb{P}^{N-1}$ we obtain the set $\partial_{\mathbb{C}}\mathcal{V}_M$ of boundary points of \mathcal{V}_M in the complex projective compactification of $Q(X)$. Using the results of steps 1–3 we show (in Theorem 7.2) that:

- (i) In a neighborhood of some point, $\partial_{\mathbb{C}}\mathcal{V}_M$ is contained in a totally real manifold, and
- (ii) The intersection of $\partial_{\mathbb{R}}\mathcal{V}_M$ with a fiber of the Hopf map has empty interior.

Using extension and parameterization results from analytic geometry it is not hard to show that among analytic subvarieties of $Q(X)$, only a discrete subset of can have both of these properties. Condition (i) forces any analytic curve in \mathcal{V}_M to extend to an analytic curve in a neighborhood of some boundary point $p \in \mathbb{C}\mathbb{P}^{N-1}$. Within this extension there is a generically a circle of directions in which to approach the boundary point, some arc of which is realized by the original curve. Analyzing the correspondence between this circle and the Hopf fiber over p , one finds that $\partial_{\mathbb{R}}\mathcal{V}_M$ contains an open arc of this fiber, violating condition (ii).

This contradiction shows that \mathcal{V}_M contains no analytic curves, making it a discrete set. The intersection theorem follows.

Applications and complements. The construction of the isotropic cone in Theorem C was inspired by the work of Floyd on the space of boundary curves of incompressible, ∂ -incompressible surfaces in 3-manifolds [Flo]. Indeed, in the incompressible boundary case, lifting such a surface to the universal cover and considering dual trees in the boundary and in the 3-manifold gives rise to an isometric embedding of \mathbb{Z} -trees. Using Theorem C

we recover Floyd’s result in this case. This connection is explained in more detail in Section 4.5, where we also note that the same “cancellation” phenomenon is at the core of both arguments.

Since Theorem D provides an interpretation of Thurston’s symplectic form in terms of Riemannian and Kähler geometry of a (stratified) smooth manifold, we hope that it might allow new tools to be applied to problems involving the space of measured foliations. In Section 5.7 we describe work of Mirzakhani in this direction, where it is shown that a certain function connected to $\text{Mod}(S)$ -orbit counting problems in Teichmüller space is constant.

As a possible extension of Theorem B one might ask whether \mathcal{H}_X and \mathcal{E}_M always intersect transversely, or equivalently, whether skinning maps are always immersions. In the slightly more general setting of manifolds with rank-1 cusps a negative answer was recently given by Gaster [Gas]. In addition to Gaster’s example, numerical experiments conducted jointly with Richard Kent suggest critical points for some other manifolds with rank-1 cusps [DK1].

Nevertheless it would be interesting to understand the discreteness of the intersection $\mathcal{H}_X \cap \mathcal{E}_M$ through local or differential properties rather than the compactifications and asymptotic arguments used here. The availability of rich geometric structure on the character variety and its compatibility with the subvarieties in question offers some hope in this direction; for example, both \mathcal{H}_X and \mathcal{E}_M are Lagrangian with respect to the complex symplectic structure of the character variety (see [Kaw] [Lab, Sec. 7.2] [KS, Sec. 1.7]).

Structure of the paper. Section 2 recalls some definitions and basic results related to Λ -trees, measured foliations, and Teichmüller theory.

Sections 3–7 contain the proofs of the main theorems in the case of a 3-manifold with connected boundary. Working in this setting avoids some cumbersome notation and other issues related to disconnected spaces, while all essential features of the argument are present. Sections 3–4 are devoted to the isotropic cone construction, Section 5 introduces the stratified Kähler structure, and Sections 6–7 combine these with the results of [D] to prove Theorem A in the connected boundary case.

In Section 9 we adapt the definitions and results of the previous sections as necessary to handle a manifold with disconnected boundary, possibly including torus components, completing the proof of Theorem A.

Finally, in Section 10 we prove an analogue of Theorem A for the extended skinning map of an acylindrical 3-manifold.

Acknowledgments. A collaboration with Richard Kent (in [DK2] [DK3]) provided essential inspiration in the early stages of this project. The author is grateful for this and for helpful conversations with Athanase Papadopoulos, Kasra Rafi, and Peter Shalen. The author also thanks Maryam Mirzakhani for allowing the inclusion of Theorem 5.10 and the anonymous referees for many helpful comments and suggestions.

2. PRELIMINARIES

2.1. Ordered abelian groups. An *ordered abelian group* is a pair $(\Lambda, <)$ consisting of an abelian group Λ and a translation-invariant total order $<$ on Λ . We often consider the order to be implicit and denote an ordered abelian group by Λ alone. Note that an order on Λ also induces an order on any subgroup of Λ .

The positive subset of an ordered abelian group Λ is the set $\Lambda^+ = \{g \in \Lambda \mid g > 0\}$. If $x \in \Lambda$ is nonzero, then exactly one of $x, -x$ lies in Λ^+ , and we denote this element by $|x|$.

If $x, y \in \Lambda$, we say that y is *infinitely larger* than x if $n|x| < |y|$ for all $n \in \mathbb{N}$. If neither of x, y is infinitely larger than the other, then x and y are *archimedean equivalent*. When the set of archimedean equivalence classes of nonzero elements is finite, the number of such classes is the *rank* of Λ . In what follows we consider *only* ordered abelian groups of finite rank.

A subgroup $\Lambda' \subset \Lambda$ is *convex* if whenever $g, k \in \Lambda'$, $h \in \Lambda$, and $g < h < k$ we have $h \in \Lambda'$. The convex subgroups of a given group are ordered by inclusion. A convex subgroup is a union of archimedean equivalence classes and is uniquely determined by the largest archimedean equivalence class that it contains (which exists, since the rank is finite). In this way the convex subgroups of a given ordered abelian group are in one-to-one order-preserving correspondence with its archimedean equivalence classes.

2.2. Embeddings and left inverses. If we equip \mathbb{R}^n with the lexicographical order, then the inclusion $\mathbb{R} \hookrightarrow \mathbb{R}^n$ as one of the factors is order-preserving. This inclusion has a left inverse $\mathbb{R}^n \rightarrow \mathbb{R}$ given by projecting onto the factor. This projection is of course a homomorphism but it is *not* order-preserving.

Similarly, the following lemma shows that an order-preserving embedding of \mathbb{R} into any ordered abelian group of finite rank has a left inverse; this is used in Section 3.6.

Lemma 2.1. *If Λ is an ordered abelian group of finite rank and $\mathfrak{o} : \mathbb{R} \hookrightarrow \Lambda$ is an order-preserving homomorphism, then \mathfrak{o} has a left inverse. That is, there is a homomorphism $\varphi : \Lambda \rightarrow \mathbb{R}$ such that $\varphi \circ \mathfrak{o} = \text{Id}$.*

To construct a left inverse we use the following structural result for ordered abelian groups; part (i) is the *Hahn embedding theorem* (see e.g. [Gra] [KK, Sec. II.2]):

Theorem 2.2. *Let Λ be an ordered abelian group of rank n .*

- (i) *There exists an order-preserving embedding $\Lambda \hookrightarrow \mathbb{R}^n$ where \mathbb{R}^n is given the lexicographical order.*
- (ii) *If $n = 1$, the order-preserving embedding $\Lambda \hookrightarrow \mathbb{R}$ is unique up to multiplication by a positive constant.*

□

Proof of Lemma 2.1. By Theorem 2.2, it suffices to consider the case of an order-preserving embedding $\mathfrak{o} : \mathbb{R} \rightarrow \mathbb{R}^n$ where \mathbb{R}^n has the lexicographical

order. Since the only order-preserving self-homomorphisms of the additive group \mathbb{R} are multiplication by positive constants, it is enough to find a homomorphism $\varphi : \mathbb{R}^n \rightarrow \mathbb{R}$ such that $\varphi \circ \mathfrak{o}$ is order-preserving and fixes a point.

Let $a = \mathfrak{o}(1)$, and write $a = (a_1, \dots, a_n)$. Since $a > 0$ in \mathbb{R}^n , the first nonzero element of the tuple a is positive. Let k be the index of this element, i.e. $k = \min\{i \mid a_i > 0\}$.

We claim that for any $t \in \mathbb{R}$, the image $\mathfrak{o}(t)$ has the form $(0, \dots, 0, b_k, \dots, b_n)$. If not, then after possibly replacing t by $-t$ we have $t \in \mathbb{R}_+$ such that $\mathfrak{o}(t)$ is infinitely larger than $\mathfrak{o}(1)$. The existence of a positive integer n such that $t < n$ shows that this contradicts the order-preserving property of \mathfrak{o} .

Similarly, we find that if $t > 0$ then $b = \mathfrak{o}(t)$ satisfies $b_k > 0$: The order-preserving property of \mathfrak{o} implies that $b_k \geq 0$, so the only possibility to rule out is $b_k = 0$. But if $b_k = 0$ then $\mathfrak{o}(1)$ is infinitely larger than $\mathfrak{o}(t)$, yet there is a positive integer n such that $1 < nt$, a contradiction.

Now define $\varphi : \mathbb{R}^n \rightarrow \mathbb{R}$ by

$$\varphi(x_1, \dots, x_n) = \frac{x_k}{a_k}.$$

This is a homomorphism satisfying $\varphi(\mathfrak{o}(1)) = 1$, and the properties of \mathfrak{o} derived above show that $\varphi \circ \mathfrak{o}$ is order-preserving, as desired. \square

We now consider the properties of embeddings $\Lambda \rightarrow \mathbb{R}^n$, such as those provided by Theorem 2.2, with respect to convex subgroups. First, a proper convex subgroup maps into \mathbb{R}^{n-1} :

Lemma 2.3. *Let $F : \Lambda \rightarrow \mathbb{R}^n$ be an order-preserving embedding, where Λ has rank n . If $\Lambda' \subset \Lambda$ is a proper convex subgroup, then $F(\Lambda') \subset \{(0, x_2, \dots, x_n)\}$.*

Proof. Since $\Lambda' \neq \Lambda$, the convex subgroup Λ' does not contain the largest Archimedean equivalence class of Λ . Thus there exists a positive element $g \in \Lambda_+$ such that $h < g$ for all $h \in \Lambda'$.

Suppose that there exists $h \in \Lambda'$ such that $F(h) = (a_1, a_2, \dots, a_n)$ with $a_1 \neq 0$. Then we have $F(kh) = ka_1 > F(g)$ for some $k \in \mathbb{Z}$. This contradicts the order-preserving property of F , so no such h exists and $F(\Lambda')$ has the desired form. \square

Building on the previous result, the following lemma shows that in some cases the embeddings given by Hahn's theorem behave functorially with respect to rank-1 subgroups. This result is used in Section 6.4.

Lemma 2.4. *Let Λ be an ordered abelian group of finite rank and $\Lambda' \subset \Lambda$ a subgroup contained in the minimal nontrivial convex subgroup of Λ . Then*

there is a commutative diagram of order-preserving embeddings

$$\begin{array}{ccc} \Lambda' & \longrightarrow & \Lambda \\ \downarrow & & \downarrow \\ \mathbb{R} & \longrightarrow & \mathbb{R}^n \end{array}$$

where $i_n(x) = (0, \dots, 0, x)$ and n is the rank of Λ .

Proof. Let $\Lambda_1 \subset \Lambda_2 \subset \dots \subset \Lambda_n = \Lambda$ be the convex subgroups of Λ . We can assume that $\Lambda' = \Lambda_1$ since all other cases are handled by restricting the maps from this one.

We are given the inclusion $i : \Lambda_1 \rightarrow \Lambda$ and the Hahn embedding theorem provides an order-preserving embedding $F : \Lambda \rightarrow \mathbb{R}^n$. Applying Lemma 2.3 to each step in the chain of convex subgroups of Λ , we find that for all $g \in \Lambda_1$ we have

$$F(g) = (0, \dots, 0, f(g))$$

and the induced map $f : \Lambda_1 \rightarrow \mathbb{R}$ is order-preserving. Since $F \circ i = i_n \circ f$ by construction, these maps complete the commutative diagram. \square

2.3. Λ -metric spaces and Λ -trees. We refer to the book [Chi3] for general background on Λ -metric spaces and Λ -trees. Here we recall the essential definitions and fix notation.

As before let Λ denote an ordered abelian group. A Λ -metric space is a pair (M, d) where M is a set and $d : M \times M \rightarrow \Lambda$ is a function which satisfies the usual axioms for the distance function of a metric space. In particular an \mathbb{R} -metric space (where \mathbb{R} has the standard order) is the usual notion of a metric space.

An isometric embedding of one Λ -metric space into another is defined in the natural way. Generalizing this, let (M, d) be a Λ -metric space and (M', d') a Λ' -metric space. An *isometric embedding* of M into M' is a pair (f, \mathfrak{o}) consisting of a map $f : M \rightarrow M'$ and an order-preserving homomorphism $\mathfrak{o} : \Lambda \rightarrow \Lambda'$ such that

$$d'(f(x), f(y)) = \mathfrak{o}(d(x, y)) \text{ for all } x, y \in M.$$

More generally we say $f : M \rightarrow M'$ is an isometric embedding if there exists an order-preserving homomorphism \mathfrak{o} such that the pair (f, \mathfrak{o}) satisfy this condition.

An ordered abelian group Λ is an example of a Λ -metric space, with metric $d(g, h) = |g - h|$. An isometric embedding of the subspace $[g, h] := \{k \in \Lambda \mid g \leq k \leq h\} \subset \Lambda$ into a Λ -metric space is a *segment*. A Λ -metric space is *geodesic* if any pair of points can be joined by a segment.

A Λ -tree is a Λ -metric space (T, d) satisfying three conditions:

- (T, d) is geodesic,
- If two segments in T share an endpoint but have no other intersection points, then their union is a segment, and

- If two segments in T share an endpoint, then their intersection is a segment (or a point).

The notion of δ -hyperbolicity for metric spaces generalizes naturally to Λ -metric spaces, where now $\delta \in \Lambda$, $\delta \geq 0$. In terms of this generalization, any Λ -tree is 0-hyperbolic. (The converse holds under mild additional assumptions on the space.) The 0-hyperbolicity condition has various equivalent characterizations, but the one we will use in the sequel is the following condition on 4-tuples of points:

Lemma 2.5 (0-hyperbolicity of Λ -trees). *If (T, d) is a Λ -tree, then for all $x, y, z, t \in T$ we have*

$$d(x, y) + d(z, t) \leq \max(d(x, z) + d(y, t), d(x, t) + d(y, z)).$$

□

For a proof and further discussion see [Chi3, Lem. 1.2.6 and Lem. 2.1.6]. By permuting a given 4-tuple x, y, z, t and considering the inequality of this lemma, we obtain the following corollary (see [Chi3, p. 35]):

Lemma 2.6 (Four points in a Λ -tree). *Let (T, d) be a Λ -tree and $x, y, z, t \in T$. Then among the three sums*

$$d(x, y) + d(z, t), \quad d(x, z) + d(y, t), \quad d(x, t) + d(y, z),$$

two are equal, and these two are not less than the third.

□

Given a Λ -tree, there are natural constructions that associate trees to certain subgroups or extensions of Λ ; in what follows we require two such operations. First, let (T, d) be a Λ -tree and $\Lambda' \subset \Lambda$ a convex subgroup. For any $x \in T$ we can consider the subset $T_{\Lambda', x} = \{y \in T \mid d(x, y) \in \Lambda'\}$. Then the restriction of d to $T_{\Lambda', x}$ takes values in Λ' , and this gives $T_{\Lambda', x}$ the structure of a Λ' -tree [MoSh1, Prop. II.1.14].

Second, suppose that $\sigma : \Lambda \rightarrow \Lambda'$ is an order-preserving homomorphism and that (T, d) is a Λ -tree. Then there is a natural *base change* construction that produces a Λ' -tree $\Lambda' \otimes_{\Lambda} T$ and an isometric embedding $T \rightarrow \Lambda' \otimes_{\Lambda} T$ with respect to σ (see [Chi3, Thm. 4.7] for details). Roughly speaking, if one views T as a union of segments, each identified with some interval $[g, h] \subset \Lambda$, then $\Lambda' \otimes_{\Lambda} T$ is obtained by replacing each such segment with $[\sigma(g), \sigma(h)] \subset \Lambda'$.

2.4. Group actions on Λ -trees and length functions. Every isometry of a Λ -tree is either *elliptic*, *hyperbolic*, or an *inversion*; see [Chi3, Sec. 3.1] for a detailed discussion of this classification. Elliptic isometries are those with fixed points, while hyperbolic isometries have an invariant axis (identified with a subgroup of Λ) on which they act as a translation. An inversion is an isometry that has no fixed point but which induces an elliptic isometry after an index-2 base change; permitting such base change allows us to make the standing assumption that *isometric group actions on Λ -trees that we consider are without inversions*.

The *translation length* $\ell(g)$ of an isometry $g : T \rightarrow T$ of a Λ -tree is defined as

$$\ell(g) = \begin{cases} 0 & \text{if } g \text{ is elliptic,} \\ |t| & \text{if } g \text{ is hyperbolic and acts on its axis as } h \mapsto h + t. \end{cases}$$

Note that $\ell(g) \in \Lambda^+ \cup \{0\}$. It can be shown that the translation length is also given by $\ell(g) = \min\{d(x, g(x)) \mid x \in T\}$.

When a group G acts on a Λ -tree by isometries, taking the translation length of each element of G defines a function $\ell : G \rightarrow \Lambda^+ \cup \{0\}$, the *translation length function* (or briefly, the *length function*) of the action.

When the translation length function takes values in a convex subgroup, one can extract a subtree whose distance function takes values in the same subgroup:

Lemma 2.7. *Let G act on a Λ -tree T with length function ℓ . If $\Lambda' \subset \Lambda$ is a convex subgroup and $\ell(G) \subset \Lambda'$ then there is a Λ' -tree $T' \subset T$ that is invariant under G and such that ℓ is also the length function of the induced action of G on T' .*

This lemma is implicit in the proof of Theorem 3.7 in [Mor1], which uses the structure theory of actions developed in [MoSh1]. For the convenience of the reader, we reproduce the argument here.

Proof. Because Λ' is convex, there is an induced order on the quotient group Λ/Λ' . Define an equivalence relation on T where $x \sim y$ if $d(x, y) \in \Lambda'$. Then the quotient $T_0 = T/\sim$ is a (Λ/Λ') -tree, and each fiber of the projection $T \rightarrow T_0$ is a Λ' -tree. The action of G on T induces an action on T_0 whose length function is the composition of ℓ with the map $\Lambda \rightarrow \Lambda/\Lambda'$, which is identically zero since $\ell(G) \subset \Lambda'$. It follows that the action of G on T_0 has a global fixed point [MoSh1, Prop. II.2.15], and thus G acts on the fiber of T over T_0 , which is a Λ' -tree T' . By [MoSh1, Prop. II.2.12] the length function of the action of G on Λ' is ℓ . \square

2.5. Measured foliations and train tracks. Let $\mathcal{MF}(S)$ denote the space of measured foliations on a compact oriented surface S of genus g . Then $\mathcal{MF}(S)$ is a piecewise linear manifold which is homeomorphic to \mathbb{R}^{6g-6} . A point $[\nu] \in \mathcal{MF}(S)$ is an equivalence class up to Whitehead moves of a singular foliation ν on S equipped with a transverse measure of full support. For detailed discussion of measured foliations and of the space $\mathcal{MF}(S)$ see [FLP].

Piecewise linear charts of $\mathcal{MF}(S)$ correspond to sets of measured foliations that are carried by a train track; we will now discuss the construction of these charts in some detail. While this material is well-known to experts, most standard references that discuss train track charts use the equivalent language of measured laminations, whereas our primary interest in foliations

arising from quadratic differentials makes the direct consideration of foliations preferable. Additional details of the carrying construction from this perspective can be found in [Pap] [Mos].

A *train track* on S is a C^1 embedded graph in which all edges incident on a given vertex share a tangent line at that point. Vertices of the train track are called *switches* and its edges are *branches*. We consider only *generic* train tracks in which each switch has three incident edges, two *incoming* and one *outgoing*, such that the union of any incoming edge and the outgoing edge forms a C^1 curve.

The complement of a train track is a finite union of subsurfaces with cusps on their boundaries. In order to give a piecewise linear chart of $\mathcal{MF}(S)$, each complementary disk must have at least three cusps on its boundary and each complementary annulus must have at least one cusp. We will always require this of the train tracks we consider.

If τ is such a train track, let $W(\tau)$ denote the vector space of real-valued functions w on its set of edges that obey the relation $w(a) + w(b) = w(c)$ for any switch with incoming edges $\{a, b\}$ and outgoing edge c . This *switch relation* ensures that w determines a signed transverse measure, or *weight*, on the embedded train track. Within $W(\tau)$ there is the finite-sided convex cone of nonnegative weight functions, denoted $\mathcal{MF}(\tau)$. It is this cone which forms a chart for $\mathcal{MF}(S)$.

A measured foliation is *carried* by the train track τ if the foliation can be cut open near singularities and along saddle connections and then moved by an isotopy so that all of the leaves lie in an arbitrarily small open neighborhood of τ and are nearly parallel to its branches, as depicted in Figure 1. Here “cutting open” refers to the procedure of replacing a union of leaf segments and saddle connections coming out of singularities with a subsurface with cusps on its boundary. The result of cutting open a measured foliation is a *partial measured foliation* in which there are *non-foliated regions*, each of which has a union of leaf segments of the original foliation as a spine.

A measured foliation ν determines a weight on any train track that carries it, as follows: For each branch $e \subset \tau$ choose a *tie* r_e , a short closed arc that crosses e transversely at an interior point and which is otherwise disjoint from τ . Now select an open neighborhood U of τ that intersects each tie r_e in a connected open interval. Let ν' be a partial measured foliation associated to ν that has been isotoped to lie in U and to be transverse to each tie. Note that each tie r_e then has endpoints in non-foliated regions of ν' , since the endpoints of r_e lie outside U . Let $w(e)$ be the total transverse measure of r_e with respect to ν' .

The resulting function w lies in $\mathcal{MF}(\tau)$ and regarding this construction as a map $\nu \mapsto w$ gives a one-to-one correspondence between equivalence classes of measured foliations that are carried by τ and the convex cone $\mathcal{MF}(\tau)$. Furthermore, these cones in train track weight spaces form the charts of a piecewise linear atlas on $\mathcal{MF}(S)$.

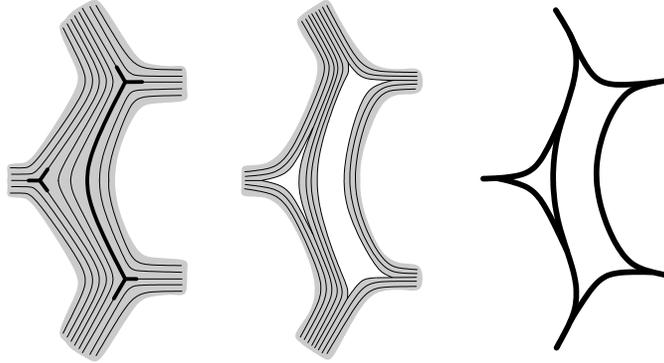


Figure 1. *Local picture of carrying: Cutting a measured foliation along saddle connections and singular leaf segments and inflating these to cusped subsurfaces allows an isotopy into a small neighborhood of the train track.*

2.6. The symplectic structure of $\mathcal{MF}(S)$. The orientation of S induces a natural antisymmetric bilinear map $\omega_{\text{Th}} : W(\tau) \times W(\tau) \rightarrow \mathbb{R}$ on the space of weights on a train track τ . This *Thurston form* can be defined as follows (compare [PH, Sec. 3.2] [Bon, Sec. 3]): For each switch $v \in \tau$, let a_v, b_v be its incoming edges and c_v its outgoing edge, where a_v, b_v are ordered so that intersecting $\{a_v, b_v, c_v\}$ with a small circle around v gives a positively oriented triple. Then we define

$$(2.1) \quad \omega_{\text{Th}}(w_1, w_2) = \frac{1}{2} \sum_v \det \begin{pmatrix} w_1(a_v) & w_1(b_v) \\ w_2(a_v) & w_2(b_v) \end{pmatrix}.$$

If τ defines a chart of $\mathcal{MF}(S)$ then this form is nondegenerate, and the induced symplectic forms on train track charts $\mathcal{MF}(\tau)$ are compatible. This gives $\mathcal{MF}(S)$ the structure of a piecewise linear symplectic manifold.

The Thurston form can also be interpreted as a homological intersection number. If τ can be consistently oriented then each weight function w defines a 1-cycle $c_w = \sum_e w(e)\vec{e}$, where \vec{e} denotes the singular 1-simplex defined by the oriented edge e of τ . In terms of these cycles, we have $\omega_{\text{Th}}(w_1, w_2) = c_{w_1} \cdot c_{w_2}$. For a general train track, there is a branched double cover $\widehat{S} \rightarrow S$ (with branching locus disjoint from τ) such that the preimage $\widehat{\tau} \subset \widehat{S}$ is orientable. Lifting weight functions we obtain cycles $\widehat{c}_{w_i} \in H^1(\widehat{S}, \mathbb{R})$ such that

$$(2.2) \quad \omega_{\text{Th}}(w_1, w_2) = \frac{1}{2} (\widehat{c}_{w_1} \cdot \widehat{c}_{w_2}).$$

Note that if \overline{S} denotes the opposite orientation of the surface S , then there is a natural identification between measured foliation spaces $\mathcal{MF}(S) \simeq \mathcal{MF}(\overline{S})$, but this identification does *not* respect the Thurston symplectic forms. Rather, in corresponding local charts we have $\omega_{\text{Th}}^S = -\omega_{\text{Th}}^{\overline{S}}$.

2.7. Dual trees. Let ν be a measured foliation on S and $\tilde{\nu}$ its lift to the universal cover \tilde{S} . There is a pseudo-metric d on \tilde{S} where $d(x, y)$ is the minimum $\tilde{\nu}$ -transverse measure of a path connecting x to y . The quotient metric space $T_\nu := \tilde{S}/d^{-1}(0)$ is an \mathbb{R} -tree (see [Bow] [MS2]). The action of $\pi_1 S$ on \tilde{S} by deck transformations determines an action on T_ν by isometries. The dual tree of the zero foliation $0 \in \mathcal{MF}(S)$ is a point.

This pseudo-metric construction can be applied to the partial measured foliation ν' obtained by cutting ν open along leaf segments from singularities, as when ν is carried by a train track τ . The result is a tree naturally isometric to T_ν , which we identify with T_ν from now on. Non-foliated regions of $\tilde{\nu}'$ are collapsed to points in this quotient, so in particular each complementary region of the lift $\tilde{\tau}$ has a well-defined image point T_ν .

Similarly, the lift of a tie r_e of τ to the universal cover projects to a geodesic segment in T_ν of length $w(e)$; the endpoints of this segment are the projections of the two complementary regions adjacent to the lift of the edge e .

To summarize, we have the following relation between carrying and dual trees:

Proposition 2.8. *Let ν be a measured foliation carried by the train track τ with associated weight function w . Let $\tilde{\tau}$ denote the lift of τ to the universal cover. If A, B are complementary regions of $\tilde{\tau}$ that are adjacent along an edge \tilde{e} of $\tilde{\tau}$, and if a, b are the associated points in T_ν , then we have*

$$w(e) = d(a, b)$$

where d is the distance function of T_ν . □

2.8. Teichmüller space and quadratic differentials. Let $\mathcal{T}(S)$ be the Teichmüller space of marked isomorphism classes of complex structures on S compatible with its orientation. For any $X \in \mathcal{T}(S)$ we denote by $Q(X)$ the set of holomorphic quadratic differentials on X , a complex vector space of dimension $3g - 3$.

Associated to $\phi \in Q(X)$ we have the following structures on X :

- The flat metric $|\phi|$, which has cone singularities at the zeros of ϕ ,
- The measured foliation $\mathcal{F}(\phi)$ whose leaves integrate the distribution $\ker(\operatorname{Im}(\sqrt{\phi}))$, with transverse measure given by $|\operatorname{Im}(\sqrt{\phi})|$, and
- The dual tree $T_\phi := T_{\mathcal{F}(\phi)}$ and the $\pi_1 S$ -equivariant map $\pi : \tilde{X} \rightarrow T_\phi$ that collapses leaves of the lifted foliation $\widetilde{\mathcal{F}(\phi)}$ to points of T_ϕ .

The dual tree construction is homogeneous with respect to the action of \mathbb{R}^+ on $Q(X)$ in the sense that for any $c \in \mathbb{R}^+$ we have

$$T_{c\phi} = c^{1/2}T_\phi$$

where the right hand side represents the metric space obtained from T_ϕ by multiplying its distance function by $c^{1/2}$.

Note that the point $0 \in Q(X)$ is a degenerate case in which there is no corresponding flat metric, and by convention $\mathcal{F}(0)$ is the empty measured foliation whose dual tree is a point.

We say that a $|\phi|$ -geodesic is *nonsingular* if its interior is disjoint from the zeros of ϕ . Choosing a local coordinate z in which $\phi = dz^2$ (a *natural coordinate* for ϕ), a nonsingular $|\phi|$ -geodesic segment I becomes a line segment in the z -plane. The vertical variation of this segment in \mathbb{C} (i.e. $|\operatorname{Im}(z_2 - z_1)|$, where z_i are the endpoints) is the *height* of I .

Note that leaves of the foliation $\mathcal{F}(\phi)$ are geodesics of the $|\phi|$ -metric. Conversely, a nonsingular $|\phi|$ -geodesic I is either a leaf of $\mathcal{F}(\phi)$ or it is transverse to $\mathcal{F}(\phi)$. In the latter case, the height h of I is equal to its $\mathcal{F}(\phi)$ -transverse measure, and any lift of I to \tilde{X} projects homeomorphically to a geodesic segment in T_ϕ of length h .

3. THE ISOTROPIC CONE: EMBEDDINGS

The goal of this section is to establish the following result relating 3-manifold actions on Λ -trees and measured foliations:

Theorem 3.1. *Let M be a 3-manifold with connected boundary S . There exists an isotropic piecewise linear cone $\mathcal{L}_M \subset \mathcal{MF}(S)$ with the following property: If ν is a measured foliation on S whose dual tree embeds isometrically and $\pi_1 S$ -equivariantly into a Λ -tree T equipped with an isometric action of $\pi_1 M$, then $[\nu] \in \mathcal{L}_M$.*

Here a *piecewise linear cone* refers to a closed \mathbb{R}^+ -invariant subset of $\mathcal{MF}(S)$ whose intersection with any train track chart $\mathcal{MF}(\tau)$ is a finite union of finite-sided convex cones in linear subspaces of $W(\tau)$. Such a cone is *isotropic* if the linear spaces can be chosen to be isotropic with respect to the Thurston symplectic form. Since transition maps between these charts are piecewise linear and symplectic, it suffices to check these conditions in any covering of the set by train track charts.

The first step in the proof of Theorem 3.1 will be to use the foliation and embedding to construct a weight function on the 1-skeleton of a triangulation of M . We begin with some generalities about train tracks, triangulations, and weight functions.

3.1. Complexes and weight functions. Let Δ be a simplicial complex, and let $\Delta^{(k)}$ denote the its set of k -simplices. Given an abelian group G , define the space of G -valued weights on Δ as the G -module consisting of functions $\Delta^{(1)} \rightarrow G$; we denote this space by

$$W(\Delta, G) := G^{\Delta^{(1)}}.$$

The case $G = \mathbb{R}$ will be of primary interest and so we abbreviate $W(G) := W(G, \mathbb{R})$. If $w \in W(\Delta, G)$ and $e \in \Delta^{(1)}$ we say that $w(e)$ is the *weight* of e with respect to w .

A homomorphism of groups $\varphi : G \rightarrow G'$ induces a homomorphism of weight spaces $\varphi_* : W(\Delta, G) \rightarrow W(\Delta, G')$, and an inclusion of simplicial complexes $i : (M, \Delta) \rightarrow (M', \Delta')$ induces a G -linear restriction map $i^* : W(\Delta', G) \rightarrow W(\Delta, G)$. These functorial operations commute, i.e. $\varphi_* \circ i^* = i^* \circ \varphi_*$.

A map f from $\Delta^{(0)}$ to a Λ -metric space induces a Λ -valued weight on Δ in a natural way: For each $e \in \Delta^{(1)}$ we define the weight to be the distance between the f -images of its endpoints. We write w_f for the weight function defined in this way.

This construction has a natural extension to equivariant maps on regular covers. Suppose that $\tilde{\Delta}, \Delta$ are simplicial complexes such that there is a regular covering $\pi : \tilde{\Delta} \rightarrow \Delta$ which is also a simplicial map. Suppose also that the deck group Γ of this covering acts isometrically on a Λ -metric space E . Then if $f : \tilde{\Delta} \rightarrow E$ is a Γ -equivariant map, the resulting weight function $\tilde{w}_f \in W(\tilde{\Delta}, \Lambda)$ is also Γ -invariant, hence it descends to a weight function $w_f \in W(\Delta, \Lambda)$ on the base of the covering.

3.2. Extending triangulations and maps. We will now consider the space of weight functions as defined above in cases where the complex Δ is a triangulation of a 2- or 3-manifold, possibly with boundary.

For example, let τ be a maximal, generic train track on a surface S . Then there is a triangulation Δ_τ of S dual to the embedded trivalent graph underlying τ . Each triangle of Δ_τ contains one switch of the train track, each edge of Δ_τ corresponds to an edge of τ , and each vertex of Δ_τ corresponds to a complementary region of τ . The correspondence between edges gives a natural (linear) embedding

$$W(\tau) \hookrightarrow W(\Delta_\tau).$$

Now suppose that ν is a measured foliation on S that is carried by the train track τ , so we consider the class $[\nu]$ as an element of $\mathcal{MF}(\tau) \subset W(\tau)$. Let $\tilde{\tau}$ denote the lift of τ to the universal cover \tilde{S} . As in Section 2.7, the carrying relationship between ν and τ gives a map from complementary regions of τ to the dual tree T_ν . In terms of the dual triangulation $\Delta := \Delta_\tau$, this is a map

$$f : \tilde{\Delta}^{(0)} \rightarrow T_\nu,$$

and it is immediate from the definitions above and Proposition 2.8 that the associated weight function $w_f \in W(\Delta)$ is the image of $[\nu]$ under the embedding $W(\tau) \hookrightarrow W(\Delta)$.

Let us further assume that, as in the hypotheses of Theorem 3.1, there is an equivariant isometric embedding of T_ν into a Λ -tree T equipped with an action of $\pi_1 M$, where M is a 3-manifold with $\partial M = S$. Using this embedding we can consider the map f constructed above as taking values in T . We extend the triangulation Δ of S to a triangulation Δ_M of M , and

the map f to a $\pi_1 M$ -equivariant map

$$F : \tilde{\Delta}_M^{(0)} \rightarrow T.$$

Such an extension can be constructed by choosing a fundamental domain V for the $\pi_1 M$ -action on $\tilde{\Delta}_M^{(0)}$ and mapping elements of $V \setminus \tilde{\Delta}_M^{(0)}$ to arbitrary points in T . Combining these with the values of f on $\tilde{\Delta}_M^{(0)}$ and the free action of $\pi_1 M$ gives a unique equivariant extension to all of $\tilde{\Delta}_M^{(0)}$.

Associated to the map F is the weight function $w_F \in W(\Delta_M, \Lambda)$. By construction, its values on the edges of Δ are the coordinates of $[\nu]$ relative to the train track chart of τ , considered as elements of Λ using the embedding $\mathfrak{o} : \mathbb{R} \rightarrow \Lambda$ that is implicit in the isometric map $T_\nu \rightarrow T$.

We record the constructions of this paragraph in the following proposition.

Proposition 3.2. *Let ν be a measured foliation on $S = \partial M$ carried by a maximal generic train track τ , and let Δ_M be a triangulation of M extending the dual triangulation of τ . Suppose that there exists a Λ -tree T equipped with an isometric action of $\pi_1 M$ and an equivariant isometric embedding*

$$h : T_\nu \rightarrow T,$$

relative to an order-preserving embedding $\mathfrak{o} : \mathbb{R} \rightarrow \Lambda$. Then there exists a weight function $w \in W(\Delta_M, \Lambda)$ with the following properties:

- (i) *The weight w is induced by an equivariant map $F : \tilde{\Delta}_M^{(0)} \rightarrow T$*
- (ii) *The restriction of w to Δ_τ is the image of $[\nu] \in W(\tau)$ under the natural inclusion $W(\tau) \hookrightarrow W(\Delta_\tau) \xrightarrow{\mathfrak{o}^*} W(\Delta_\tau, \Lambda)$. \square*

3.3. Triangle forms and the symplectic structure. In [PP], Penner and Papadopoulos relate the Thurston symplectic structure of $\mathcal{MF}(S)$ for a punctured surface S to a certain linear 2-form on the space of weights on a “null-gon track” dual to an ideal triangulation of S . In this section we discuss a related construction for a triangulation of a compact surface dual to a train track.

Let σ be an oriented triangle with edges e, f, g (cyclically ordered according to the orientation). Let de, df, dg denote the corresponding linear functionals on $\mathbb{R}^{\{e, f, g\}}$, which evaluate a function on the given edge. We call the alternating 2-form

$$\omega_\sigma := -\frac{1}{2} (de \wedge df + df \wedge dg + dg \wedge de)$$

the *triangle form* associated with σ . Note that if $-\sigma$ represents the triangle with the opposite orientation, then $\omega_{-\sigma} = -\omega_\sigma$.

Given a triangulation Δ of a compact oriented 2-manifold S , the triangle form corresponding to any $\sigma \in \Delta^{(2)}$ (with its induced orientation) is naturally an element of $\Lambda^2 W(\Delta)^*$. Denote the sum of these by

$$(3.1) \quad \omega_\Delta = \sum_{\sigma \in \Delta^{(2)}} \omega_\sigma.$$

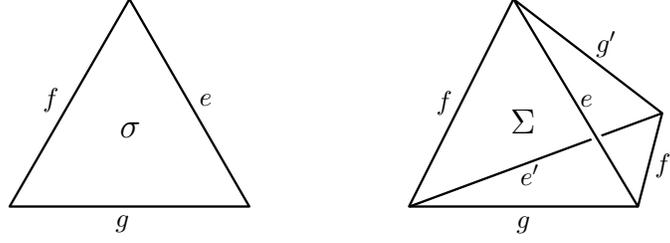


Figure 2. Labeled edges of an oriented 2-simplex σ and an oriented 3-simplex Σ . The 2-forms ω_σ and ω_Σ are defined in terms of these labels.

This 2-form on $W(\Delta)$ is an analogue of the Thurston symplectic form, in a manner made precise by the following:

Lemma 3.3. *If τ is a maximal generic train track on S with dual triangulation $\Delta = \Delta_\tau$, then the Thurston form on $W(\tau)$ is the pullback of ω_Δ by the natural inclusion $W(\tau) \rightarrow W(\Delta)$.*

Proof. By direct calculation: Both the Thurston form and ω_Δ are given as a sum of 2-forms, one for each triangle of Δ (equivalently, switch of τ). The image of $W(\tau)$ in $W(\Delta)$ is cut out by imposing a switch condition for each triangle $\sigma \in \Delta^{(2)}$, which for an appropriate labeling of the sides as e, f, g can be written as

$$de + df = dg.$$

On the subspace defined by this constraint the triangle form pulls back to

$$-\frac{1}{2}(de \wedge df + df \wedge dg + dg \wedge de) = \frac{1}{2}de \wedge df$$

which is the associated summand in the Thurston form (2.1). \square

3.4. Tetrahedron forms. Let Σ be an oriented 3-simplex. Call a pair of edges of Σ *opposite* if they do not share a vertex. Label the edges of Σ as e, f, g, e', f', g' so that the following conditions are satisfied:

- (i) The pairs $\{e, e'\}$, $\{f, f'\}$, and $\{g, g'\}$ are opposite.
- (ii) The ordering e, f, g gives the oriented boundary of one of the faces of Σ .

An example of such a labeling is shown in Figure 2.

Define the *tetrahedron form* $\Omega_\Sigma \in \wedge^2 W(\Sigma)^*$ as

$$\Omega_\Sigma = -\frac{1}{2}(d(e+e') \wedge d(f+f') + d(f+f') \wedge d(g+g') + d(g+g') \wedge d(e+e'))$$

Here we abbreviate $d(e+e') = de + de'$ and similarly for the other edges. It is easy to check that this 2-form does not depend on the labeling (as long as it satisfies the conditions above). As in the case of triangle forms, Ω_Σ is naturally a 2-form on the space of weights for any oriented simplicial complex containing Σ .

A simple calculation using the definition above gives the following:

Lemma 3.4. *The tetrahedron form is equal to the sum of the triangle forms of its oriented boundary faces, i.e.*

$$\Omega_\Sigma = \sum_{\sigma \in \partial\Sigma} \omega_\sigma.$$

□

Now consider a triangulation Δ_M of an oriented 3-manifold with boundary S , and let Δ_S denote the induced triangulation of the boundary. Denote the sum of the tetrahedron forms by

$$\Omega_{\Delta_M} = \sum_{\Sigma \in \Delta_M^{(3)}} \Omega_\Sigma \in \wedge^2 W(\Delta_M)^*.$$

In fact, due to cancellation in this sum, the 2-form defined above “lives” on the boundary:

Lemma 3.5. *The form Ω_{Δ_M} is equal to the pullback of ω_{Δ_S} under the restriction map $W(\Delta_M) \rightarrow W(\Delta_S)$.*

Proof. By Lemma 3.4 we have

$$\Omega_{\Delta_M} = \sum_{\Sigma \in \Delta_M^{(3)}} \sum_{\sigma \in \partial\Sigma} \omega_\sigma.$$

In this sum, each interior triangle of Δ_M appears twice (once with each orientation) and so these terms cancel. The remaining terms are the elements of $\Delta_S^{(2)}$ with the boundary orientation, so we are left with the sum (3.1) defining ω_{Δ_S} . The result is the pullback of ω_{Δ_S} by the restriction map because in the formula above, we are considering ω_σ as an element of $\wedge^2 W(\Delta_M)^*$ rather than $\wedge^2 W(\Delta_S)^*$. □

3.5. The four-point condition. Given four points in a Λ -tree, Lemma 2.6 implies that there is always a labeling $\{p, q, r, s\}$ of these points such that the distance function satisfies

$$(3.2) \quad d(p, q) + d(r, s) = d(p, s) + d(r, q).$$

We call this the *weak four-point condition* to distinguish it from the stronger four-point condition of Lemma 2.6 which also involves an inequality.

If we think of p, q, r, s as labeling the vertices of a 3-simplex Σ , then the pairwise distances give a weight function $w : \Sigma^{(1)} \rightarrow \Lambda$. Condition (3.2) is equivalent to the existence of opposite edge pairs $\{e, e'\}, \{f, f'\} \subset \Sigma^{(1)}$ such that

$$(3.3) \quad w(e) + w(e') = w(f) + w(f').$$

Given a simplicial complex Δ , let $W_4(\Delta, \Lambda)$ denote the set of Λ -valued weights such that in each 3-simplex of Δ there exist opposite edge pairs so that (3.3) is satisfied.

The following basic properties of $W_4(\Delta, \Lambda)$ follow immediately from the definition of this set (and the relation between the four-point condition and 4-tuples in Λ -trees):

Lemma 3.6.

- (i) *The set $W_4(\Delta, \Lambda)$ is a finite union of subspaces (i.e. Λ -submodules) of $W(\Delta, \Lambda)$; each subspace corresponds to choosing opposite edge pairs in each of the 3-simplices of Δ .*
- (ii) *If $\varphi : \Lambda \rightarrow \Lambda'$ is a homomorphism, then we have $\varphi_*(W_4(\Delta, \Lambda)) \subset W_4(\Delta, \Lambda')$.*
- (iii) *If $f : \tilde{\Delta}^{(0)} \rightarrow T$ is an equivariant map to a Λ -tree, then $w_f \in W_4(\Delta, \Lambda)$*

□

Ultimately, the isotropic condition in Theorem 3.1 arises from the following property of the set $W_4(\Delta) = W_4(\Delta, \mathbb{R})$:

Lemma 3.7. *Let M be an oriented 3-manifold and Δ_M a triangulation. Then $W_4(\Delta_M)$ is a finite union of Ω_{Δ_M} -isotropic subspaces of $W(\Delta_M)$.*

Proof. Let $V \subset W(\Delta_M)$ be one of the subspaces comprising $W_4(\Delta_M)$ (as in Lemma 3.6.(i)). Then for each $\Sigma \in \Delta_M^{(3)}$ we have opposite edge pairs $\{e, e'\}$ and $\{f, f'\}$ such that (3.3) holds, or equivalently, on the subspace V the equation

$$d(e + e') = d(f + f')$$

is satisfied. Substituting this into the definition of the tetrahedron form Ω_Σ gives zero. Since Ω_{Δ_M} is the sum of these forms, the subspace V is isotropic. □

3.6. Construction of the isotropic cone. We now combine the results on triangulations, weight functions, and the symplectic structure of $\mathcal{MF}(S)$ with the constructions of Proposition 3.2 to prove Theorem 3.1.

Proof of Theorem 3.1. Let Υ be a finite set of maximal, generic train tracks such that any measured foliation on S is carried by one of them. For each $\tau \in \Upsilon$, let Δ_M^τ be an extension of Δ_τ to a triangulation of M .

Define

$$\mathcal{L}_\tau = i^*(W_4(\Delta_M^\tau)) \cap \mathcal{ML}(\tau)$$

where $i^* : W(\Delta_M^\tau) \rightarrow W(\Delta_\tau)$ is the restriction map (i.e. the map that restricts a weight to the edges that lie on S). That is, an element of \mathcal{L}_τ is a measured foliation carried by τ whose associated weight function on Δ_τ can be extended to Δ_M^τ in such a way that it satisfies the weak 4-point condition in each simplex.

Let $\mathcal{L}_M = \bigcup_{\tau \in \Upsilon} \mathcal{L}_\tau$. By Lemmas 3.3 and 3.7, the set \mathcal{L}_M is an isotropic piecewise linear cone in $\mathcal{MF}(S)$. We need only show that for ν and $T_\nu \hookrightarrow T$ as in the statement of the Theorem we have $[\nu] \in \mathcal{L}_M$.

Given such ν and $T_\nu \hookrightarrow T$, let $\tau \in \Upsilon$ carry ν and abbreviate $\Delta_M = \Delta_M^\tau$. Let $F : \tilde{\Delta}_M^{(0)} \rightarrow T$ and $w = w_F \in W(\Delta_M, \Lambda)$ be the map and associated weight function given by Proposition 3.2. By Lemma 3.6.(iii), we have $w \in W_4(\Delta_M, \Lambda)$.

Let $\varphi : \Lambda \rightarrow \mathbb{R}$ be a left inverse to the inclusion $\mathfrak{o} : \mathbb{R} \rightarrow \Lambda$ associated with the isometric embedding $T_\nu \hookrightarrow T$; such a map exists by Lemma 2.1. Then $\varphi_* : W(\Delta_\tau, \Lambda) \rightarrow W(\Delta_\tau)$ is correspondingly a left inverse to $\mathfrak{o}_* : W(\Delta_\tau) \hookrightarrow W(\Delta_\tau, \Lambda)$. Since by Proposition 3.2.(ii) we have that $i^*(w) \in W(\Delta_\tau, \Lambda)$ is the image of $[\nu] \in \mathcal{ML}(\tau)$ under this inclusion, it follows that $\varphi_*(i^*(w)) = i^*(\varphi_*(w))$ also represents $[\nu]$.

By Lemma 3.6.(ii) we have $\varphi_*(w) \in W_4(\Delta_M)$, so we have shown that $[\nu] \in i^*(W_4(\Delta_M^\tau)) \cap \mathcal{ML}(\tau) = \mathcal{L}_\tau \subset \mathcal{L}_M$, as desired. \square

4. THE ISOTROPIC CONE: STRAIGHT MAPS AND LENGTH FUNCTIONS

In this section we introduce refinements of Theorem 3.1 that will be used in the proof of the main theorem. These refinements replace with isometric embedding hypothesis of Theorem 3.1 with weaker conditions relating the trees carrying actions of $\pi_1 S$ and $\pi_1 M$.

4.1. Straight maps. We first recall (and generalize) the notion of a straight map, which is a certain type of morphism of trees.

Let $X \in \mathcal{T}(S)$ be a marked Riemann surface structure on S and $\phi \in Q(X)$ a holomorphic quadratic differential. Recall that there is a dual \mathbb{R} -tree T_ϕ and projection $\pi : \tilde{X} \rightarrow T_\phi$, and that nonsingular $|\phi|$ -geodesics in \tilde{X} project to geodesics in T_ϕ . Let \mathcal{S}_ϕ denote the set of all geodesics in T_ϕ that arise in this way (including both segments and complete geodesics).

Let T be an \mathbb{R} -tree. Following [D], we say that a map $f : T_\phi \rightarrow T$ is *straight* if it is an isometric embedding when restricted to any element of \mathcal{S}_ϕ . Thus, for example, an isometric embedding of T_ϕ is straight map, but the converse does not hold (see e.g. [D, Lem. 6.5]).

Note that straightness of a map $T_\phi \rightarrow T$ depends on the differential ϕ and not just on the isometry type of the dual tree; Figure 3 shows an example of differentials with isometric dual trees but distinct notions of straightness.

More generally, if T is a Λ -tree, we say that a map $f : T_\phi \rightarrow T$ is straight if there is an order-preserving map $\mathfrak{o} : \mathbb{R} \rightarrow \Lambda$ such that the restriction of f to each element of \mathcal{S}_ϕ is an isometric embedding with respect to \mathfrak{o} . As in the case of \mathbb{R} -trees, isometric embeddings (now in the sense of section 2.3) are examples of straight maps.

For the degenerate case $\phi = 0$, we make the convention that any map of the point T_0 to a Λ -tree is straight.

4.2. Isotropic cone for straight maps. In the following generalization of Theorem 3.1 we fix a Riemann surface structure on $S = \partial M$ and consider straight maps instead of isometric embeddings.

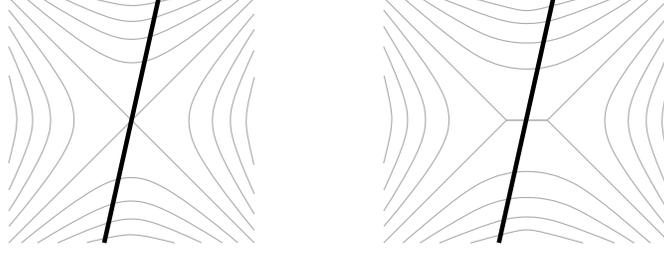


Figure 3. Quadratic differentials with isometric dual trees may induce different notions of straight mapping: The local foliation pictures shown here have isometric leaf spaces, but the indicated path is required to map isometrically by a straight map in one case (right) but not in the other (left).

Theorem 4.1. *Let M be an oriented 3-manifold with connected boundary S , and let $X \in \mathcal{T}(S)$ be a marked Riemann surface structure on S . There exists an isotropic cone $\mathcal{L}_{M,X} \subset \mathcal{MF}(S)$ with the following property: If $\phi \in \mathcal{Q}(X)$ is a holomorphic quadratic differential such that there exists a Λ -tree T equipped with an isometric action of $\pi_1 M$ and a $\pi_1 S$ -equivariant straight map*

$$h : T_\phi \rightarrow T$$

then $[\mathcal{F}(\phi)] \in \mathcal{L}_{M,X}$.

In the proof of Theorem 3.1, the assumption that the map $T_\phi \rightarrow T$ is isometric embedding was only used through its role in the construction of Section 3.2: A train track carrying ν gives a map $f : \tilde{\Delta}_\tau^{(0)} \rightarrow T_\nu$ whose associated weight function represents $[\nu]$, and since $h : T_\nu \rightarrow T$ is an isometric embedding, the composition $h \circ f$ has the same associated weight.

Attempting to reproduce this with the weaker hypotheses of Theorem 4.1, we can again choose a train track τ carrying $\mathcal{F}(\phi)$ and construct a map $f : \tilde{\Delta}_\tau^{(0)} \rightarrow T_\phi$. We would then like to compose f with the straight map $h : T_\phi \rightarrow T$ without changing the associated weight function. This will hold if the segments in T_ϕ corresponding to the ties of τ are mapped isometrically by h , so it is enough to know that they correspond to nonsingular $|\phi|$ -geodesic segments in \tilde{X} . To summarize, we have:

Proposition 4.2. *Let τ be a train track that carries $\mathcal{F}(\phi)$ such that each tie of $\tilde{\tau}$ corresponds to a nonsingular $|\phi|$ -geodesic segment in \tilde{X} . Let Δ_M be a triangulation of M extending the dual triangulation of τ . Suppose that there exists a Λ -tree T equipped with an isometric action of $\pi_1 M$ and a $\pi_1 S$ -equivariant straight map*

$$h : T_\phi \rightarrow T,$$

relative to an order-preserving embedding $\mathfrak{o} : \mathbb{R} \rightarrow \Lambda$. Then there exists a weight function $w \in W(\Delta_M, \Lambda)$ satisfying conditions (i)–(ii) of Proposition 3.2. \square

Therefore, while we used an arbitrary finite collection of train track charts covering $\mathcal{MF}(S)$ in the previous section, we now have a stronger condition that the carrying train track must satisfy. The existence of a suitable finite collection of train tracks that cover $Q(X)$ is given by:

Lemma 4.3. *For each nonzero $\phi \in Q(X)$ there exists a triangulation Δ of X and a maximal train track τ such that:*

- (i) *The vertices Δ are zeros of ϕ , the edges are saddle connections of ϕ , and the triangulation Δ is dual to the train track τ in the sense of Section 3.1,*
- (ii) *The foliation $\mathcal{F}(\phi)$ is carried by τ in such a way each edge e of Δ becomes a tie of the corresponding edge of τ ; in particular,*
- (iii) *The ϕ -heights of the edges of Δ give the weight function on τ representing $[\mathcal{F}(\phi)]$.*

Furthermore, there is a finite set of pairs (Δ, τ) such that the triangulation and train track constructed above can always be chosen to be isotopic to an element of this set.

Proof. Consider a Delaunay triangulation Δ of the singular Euclidean surface $(X, |\phi|)$ with vertices at the zeros of ϕ , as in [MaSm, Sec. 4]. Such a triangulation has nonsingular $|\phi|$ -geodesic segments as edges and is defined by the condition that each triangle has a circumcircle (with respect to the singular Euclidean structure $|\phi|$) which is “empty”, i.e. has no zeros of ϕ in its interior. There are only finitely many Delaunay triangulations of a given singular Euclidean surface, and for generic X and ϕ there is a unique one.

First suppose that this triangulation Δ has no horizontal edges. Each triangle has two “vertically short” edges whose heights sum to that of the third edge, and we construct a train track τ by placing a switch in each triangle so that the incoming branches at the switch are dual to the short edges of the triangle (as shown in Figure 4). The complementary regions of τ are disk neighborhoods of the vertices of the triangulation, so τ is maximal. Thus Δ, τ satisfy condition (i).

After cutting along leaf segments near singularities, an isotopy pushes the leaves of $\mathcal{F}(\phi)$ into a small neighborhood of the train track. Throughout this isotopy the image of an edge e of the triangulation in the dual tree remains the same, and so it corresponds to the tie r_e of the train track. The length of the image segment in T_ϕ is the height of the geodesic edge, so properties (ii)–(iii) follow.

It remains to consider the possibility that the Delaunay triangulation has horizontal edges. In this case we can still form a dual train track but it is not clear whether the dual to a horizontal edge should be incoming or outgoing at the switch in a given triangle. To determine this, we consider a slight deformation of ϕ to a quadratic differential ϕ' with the same zero structure but no horizontal saddle connections. (A generic deformation preserving the multiplicities of zeros will have this property.) For a small enough deformation, the same combinatorial triangulation can be realized geodesically for

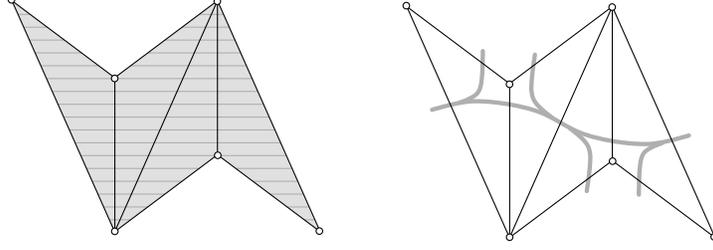


Figure 4. A geodesic triangulation for a quadratic differential with vertices at the zeros and the associated train track carrying the measured foliation.

ϕ' , and the heights of the previously horizontal edges determine how to form switches for τ .

Finally we show that only finitely many isotopy classes of pairs (Δ, τ) arise from this construction. In fact, it suffices to consider Δ alone since filling in the train track τ involves only finitely many choices (incoming and outgoing edges for each switch).

The construction of Δ is independent of scaling ϕ so we assume that $|\phi|$ has unit area, i.e. $\|\phi\| = 1$ where $\|\cdot\|$ is the L^1 norm. The resulting family of metrics (the unit sphere in $Q(X)$) is compact, and in particular the diameters of these spaces are uniformly bounded. By [MaSm, Thm. 4.4] this diameter bound also gives an upper bound, R , on the length of each edge of the Delaunay triangulation. The number of zeros of $\tilde{\phi}$ in a ball of $|\phi|$ -radius R in \tilde{X} is uniformly bounded (again by compactness of the family of metrics, and the fixed number of zeros of ϕ on X), so the edges that appear in Δ belong to finitely many isotopy classes of arcs between pairs of zeros. Thus, up to isotopy, only finitely many triangulations can be constructed of these arcs. \square

With these preliminaries in place, it is straightforward to generalize the proof Theorem 3.1:

Proof of Theorem 4.1. Let Υ_X denote the finite set of train tracks given by Lemma 4.3, and extend each dual triangulation Δ_τ to a triangulation Δ_M^τ of M . Define

$$\mathcal{L}_{M,X} = \bigcup_{\tau \in \Upsilon_X} \mathcal{L}_\tau,$$

where $\mathcal{L}_\tau = i^*(W_4(\Delta_M^\tau)) \cap \mathcal{ML}(\tau)$. As before, Lemmas 3.3 and 3.7 show that this set is an isotropic cone in $\mathcal{MF}(S)$.

If $\phi \in Q(X)$ and $h : T_\phi \rightarrow T$ is a straight map as in the statement of the Theorem, then by Lemma 4.3 and Proposition 4.2 we have a train track $\tau \in \Upsilon_X$ and weight function $w \in W(\Delta_M^\tau, \Lambda)$ such that $i^*(\varphi_*(w)) \in W(\Delta_\tau)$ represents $[\mathcal{F}(\phi)]$. Here we retain the notation of the previous section, i.e. φ denotes a left inverse of $\sigma : \mathbb{R} \rightarrow \Lambda$ and $i^* : W(\Delta_M^\tau) \rightarrow W(\Delta_\tau)$ restricts a weight function to the boundary triangulation.

By Proposition 4.2, the weight w is associated to a map $\tilde{\Delta}_M^{\tau(0)} \rightarrow T$. As in the proof of Theorem 3.1, this implies $w \in W_4(\Delta_M^\tau, \Lambda)$ and therefore $i^*(\varphi_*(w)) \in \mathcal{L}_\tau$. We conclude $[\mathcal{F}(\phi)] \in \mathcal{L}_{M,X}$. \square

4.3. Isotropic cone for length functions. In this section we introduce a further refinement to the isotropic cone construction that addresses special properties of abelian actions of groups on \mathbb{R} -trees (which are described below).

Keeping the notations M, S, X of the previous section, suppose that we have $\phi \in Q(X)$ and a $\pi_1 S$ -equivariant straight map $h : T_\phi \rightarrow T$ as in Theorem 4.1. Then the image $h(T_\phi) \subset T$ is naturally an \mathbb{R} -tree preserved by $\pi_1 S$. Let $\ell : \pi_1 S \rightarrow \mathbb{R}$ denote the translation length function of this action and write $T_\ell = h(T_\phi)$. Then T_ℓ is the intermediate step in a factorization of h as a straight map followed by an isometric embedding:

$$T_\phi \xrightarrow{\text{straight}} T_\ell \xleftarrow{\text{embed}} T$$

Theorem 4.1 shows that this situation forces $[\mathcal{F}(\phi)]$ to lie in an isotropic cone.

The generalization we now consider is to replace T_ℓ with a pair of trees T_ℓ, T'_ℓ on which $\pi_1 S$ acts with the same length function ℓ —we say these actions are *isospectral*. We suppose that one of these is the image of a straight map while the other isometrically embeds in a Λ -tree T with a $\pi_1 M$ action. From this weaker connection between T_ϕ and T , i.e.

$$T_\phi \xrightarrow{\text{straight}} T_\ell \xleftarrow{\text{isospectral}} T'_\ell \xleftarrow{\text{embed}} T$$

we can still conclude $[\mathcal{F}(\phi)] \in \mathcal{L}_{M,X}$. The following theorem makes this precise.

Theorem 4.4. *Let T be a Λ -tree on which $\pi_1 M$ acts. Let T_ℓ, T'_ℓ be \mathbb{R} -trees on which $\pi_1 S$ acts minimally with length function ℓ . Let $\phi \in Q(X)$ be a holomorphic quadratic differential such that there exists a $\pi_1 S$ -equivariant straight map $h : T_\phi \rightarrow T_\ell$ and a $\pi_1 S$ -equivariant isometric embedding $k : T'_\ell \rightarrow T$. Then $[\mathcal{F}(\phi)] \in \mathcal{L}_{M,X}$.*

Evidently this theorem would follow directly from Theorem 4.1 if the isospectrality condition implied the existence of an isometry $T_\ell \simeq T'_\ell$, for this isometry would allow h and k to be composed, giving a straight map $T_\phi \rightarrow T$. This approach works for some length functions but not for others, so before giving the proof we discuss the relevant dichotomy.

4.4. Abelian and non-abelian actions. Recall that an isometric action of a group Γ on an \mathbb{R} -tree is called *abelian* if the associated length function has the form $\ell(g) = |\chi(g)|$ where $\chi : \Gamma \rightarrow \mathbb{R}$ is a homomorphism; otherwise, the action (or length function) is called *non-abelian*. We have the following fundamental result of Culler and Morgan:

Theorem 4.5 ([CM]). *Let T_ℓ, T'_ℓ be \mathbb{R} -trees equipped with minimal, isospectral actions of a group Γ . If the length function ℓ is non-abelian, then there is an equivariant isometry $T_\ell \rightarrow T'_\ell$.*

As remarked above, this shows that the conclusion of Theorem 4.4 follows from Theorem 4.1 whenever the length function is non-abelian. Thus we assume from now on that $\ell = |\chi| : \pi_1 S \rightarrow \mathbb{R}$ is an abelian length function. In this case there may be many non-isometric trees on which $\pi_1 S$ acts with this length function [Br].

An *end* of an \mathbb{R} -tree is an equivalence class of rays, where two rays are equivalent if their intersection is a ray. An abelian action of $\pi_1 S$ on an \mathbb{R} -tree T_ℓ has a fixed end (see [CM, Cor. 2.3] or [AB, Thm. 7.5]). The fixed end has an associated *Busemann function* $\beta : T_\ell \rightarrow \mathbb{R}$, that intertwines the action of $\pi_1 S$ on T_ℓ with the translation action on \mathbb{R} induced by χ [AB, Thm. 7.6]. Furthermore the function β is unique up to adding a constant. Here we use the term “Busemann function” following e.g. [Lev2, Sec. 2], which is consistent with its use in the theory of metric spaces of non-positive curvature [BH, Sec. II.8]; the same object is called an *end map* and discussed in [Chi3, Sec. 2.3], while the above-cited result in [AB] simply calls the map α .

The following result from [D, Sec. 6] shows that composition with a Busemann function preserves straightness of maps from T_ϕ :

Lemma 4.6. *Let T_ℓ be an \mathbb{R} -tree equipped with an abelian action of $\pi_1 S$ by isometries, and let $\beta : T \rightarrow \mathbb{R}$ denote a Busemann function of a fixed end. If $h : T_\phi \rightarrow T_\ell$ is an equivariant straight map, then $\beta \circ h : T_\phi \rightarrow \mathbb{R}$ is also straight. \square*

Effectively this result will allow us to replace T_ℓ with \mathbb{R} in the hypotheses of Theorem 4.4 since any straight map can be composed with a Busemann function, preserving straightness and without changing the length function.

In the proof of Theorem 4.1, the straightness of $h : T_\phi \rightarrow T$ was only used to conclude that the map is isometric when applied to the endpoints of each segment in T_ϕ that corresponds to one of the ϕ -geodesic edges of a triangulation of X (furnished by Lemma 4.3). To generalize the proof to the situation of Theorem 4.4, it will therefore suffice to show that if there exists a straight map $T_\phi \rightarrow T_\ell$, then there also exists a partially-defined map $T_\phi \dashrightarrow T'_\ell$ that is “locally straight” in that it is an isometry when applied to the endpoints of any of these segments. Since these segments arise from lifting the finite set of edges of a triangulation of X , they lie in finitely many $\pi_1 S$ -equivalence classes. Thus, Theorem 4.4 is reduced to:

Theorem 4.7. *Let T_ℓ and T'_ℓ be \mathbb{R} -trees on which $\pi_1 S$ acts minimally and isospectrally, with abelian length function ℓ , and suppose $h : T_\phi \rightarrow T_\ell$ is an equivariant straight map, for some $\phi \in Q(X)$.*

Let $\mathcal{S} \subset \mathcal{S}_\phi$ be a set of segments in T_ϕ that arise from nonsingular ϕ -geodesic segments in \tilde{X} and suppose \mathcal{S} contains only finitely many $\pi_1 S$ -equivalence classes. Let $E \subset T_\phi$ be the set of endpoints of elements of \mathcal{S} .

Then there exists an equivariant map $h' : E \rightarrow T'_\ell$ such that for any segment $J \in \mathcal{S}$ with endpoints x, y , we have

$$d(h(x), h(y)) = d(h'(x), h'(y)).$$

The proof will depend on properties of a certain endomorphism of the tree T'_ℓ related to the end fixed by $\pi_1 S$.

Given an \mathbb{R} -tree T and an end e , for any $x \in T$ and $s \geq 0$ let $P_s(x)$ denote the point on the ray from x to e such that $d(x, P_s(x)) = s$. Then $P_s : T \rightarrow T$ is a weakly contracting map, and if $\pi_1 S$ acts on T fixing e , then P_s is $\pi_1 S$ -equivariant. We call P_s the *pushing map* of distance s for the end e .

Lemma 4.8. *Let T be an \mathbb{R} -tree and let $\beta : T \rightarrow \mathbb{R}$ denote a Busemann function of an end e . Then for any $p, q \in T$ there exists $s_0 \geq 0$ such that $d(P_s(p), P_s(q)) = |\beta(p) - \beta(q)|$ for all $s \geq s_0$, where $P_s : T \rightarrow T$ is the pushing map for the end e .*

Proof. The ray from p to e and the ray from q to e overlap in a ray from o to e , where o is a point on the geodesic segment from p to q , and the Busemann function satisfies $|\beta(p) - \beta(q)| = |d(p, o) - d(q, o)|$. Let $r(t)$ parameterize the ray from o to e , $t \geq 0$. Then for $s \geq \max(d(p, o), d(q, o))$ we have $P_s(p) = r(s - d(p, o))$ and $P_s(q) = r(s - d(q, o))$. Since r is an isometry onto its image, we have $d(P_s(p), P_s(q)) = |d(p, o) - d(q, o)|$. \square

Using the pushing map we can now give the

Proof of Theorem 4.7. Enlarging \mathcal{S} if necessary, we can take this set and its set of endpoints E to be $\pi_1 S$ -invariant.

Let $E_0 \subset E$ be a finite subset containing exactly one point from each $\pi_1 S$ -orbit in E . Let β, β' be Busemann functions of the fixed ends of $\pi_1 S$ acting on T_ℓ, T'_ℓ . For each $x \in E_0$, choose a point $g'(x) \in T'_\ell$ such that $\beta'(g'(x)) = \beta(h(x))$, giving a map $g' : E_0 \rightarrow T'_\ell$. This is possible since the map $\beta' : T' \rightarrow \mathbb{R}$ admits a section, e.g. any complete geodesic $\mathbb{R} \rightarrow T'$ that extends a ray representing the fixed end.

Using the action of $\pi_1 S$ on T_ℓ , we extend g' to an equivariant map $g' : E \rightarrow T'$ which then satisfies $\beta'(g'(x)) = \beta(h(x))$ for all $x \in E$.

For any $s \geq 0$ let $h'_s(x) = P'_s(g'(x))$ where $P'_s : T'_\ell \rightarrow T'_\ell$ is the pushing map for the fixed end of $\pi_1 S$. By Lemma 4.6, for any segment $J \in \mathcal{S}$ with endpoints $x, y \in E$ we have

$$d(h(x), h(y)) = |\beta(h(x)) - \beta(h(y))| = |\beta'(g'(x)) - \beta'(g'(y))|$$

and by Lemma 4.8 there exists $s_J \geq 0$ such that for all $s \geq s_J$ we have

$$d(h'_s(x), h'_s(y)) = |\beta(g'(x)) - \beta(g'(y))| = d(h(x), h(y)).$$

Taking s larger than the maximum of s_J as J ranges over a finite set representing each $\pi_1 S$ -orbit in \mathcal{S} , the above condition holds for each such representative, and by equivariance, for each $J \in \mathcal{S}$. Then $h' = h'_s : E \rightarrow T'_\ell$ is the desired map. \square

As remarked above this completes the proof of Theorem 4.4.

4.5. Application: Floyd's Theorem. In this section we discuss some context for Theorems 3.1, 4.1, and 4.4. The contents of this section are not used in the sequel.

Theorem 3.1 and its refinements can be seen as generalizations of the following theorem of Floyd [Flo]:

Theorem 4.9. *Let M be a compact, irreducible 3-manifold with boundary S . Then the set of boundary curves of two-sided incompressible, ∂ -incompressible surface in M is contained in a finite union of half-dimensional piecewise linear cells in $\mathcal{MF}(S)$.*

Note that the correspondence between measured laminations and measured foliations on surfaces (see e.g. [Lev1]) allows us to consider the boundary of a surface in M as an element of $\mathcal{MF}(S)$. The original statement in [Flo] uses the language of measured laminations.

Floyd's theorem answers a question of Hatcher, who established a similar result for manifolds with torus boundary [Hat]. Hatcher's theorem is often used through its corollary that a knot complement manifold has finitely many boundary slopes.

In both cases the half-dimensional set is constructed as an isotropic cone in the symplectic space $\mathcal{MF}(S)$, and these results can be compared to the more elementary (co)homological version: As a consequence of Poincaré duality, the image of the connecting map

$$H_2(M, \partial M) \xrightarrow{\delta} H_1(\partial M)$$

is isotropic with respect to the intersection pairing. Dually, the image of the map $H^1(M) \rightarrow H^1(\partial M)$ induced by inclusion of the boundary is isotropic for the cup product.

To show the connection with our results, we derive Floyd's theorem from Theorem 3.1 under the additional assumption that the boundary S is incompressible:

Proof of Theorem 4.9 (incompressible boundary case). Let F be an incompressible and ∂ -incompressible surface in M . The preimage \tilde{F} of F in \tilde{M} is a collection of planes, separating \tilde{M} into a countable family of complementary regions. The adjacency graph of these regions, with one vertex for each region and one unit-length edge for each plane, gives an \mathbb{R} -tree (which comes from an underlying \mathbb{Z} -tree) on which $\pi_1 M$ acts by isometries. In this tree, the distance between two vertices is the minimum number of intersections

between \tilde{F} and a path between the corresponding complementary regions in \tilde{M} .

Similarly, the boundary curves ∂F lift to a collection of lines separating \tilde{S} and give a dual tree $T_{\partial F}$ on which $\pi_1 S$ acts by isometries. The equivalence between laminations and foliations allows us to identify ∂F with measured foliation class in $\mathcal{MF}(S)$; under this correspondence, $T_{\partial F}$ becomes the dual tree of that measured foliation (in the sense of Section 2.7).

Since the boundary is incompressible, the inclusion $S \hookrightarrow M$ lifts to $\tilde{S} \rightarrow \tilde{M}$ which induces a map $T_{\partial F} \rightarrow T_F$. This map of trees is an isometric embedding: It is weakly contracting, since minimizing the number of intersections of a path in \tilde{S} with $\tilde{\partial F}$ is a more constrained problem than minimizing intersections of a path in \tilde{M} with \tilde{F} . However, if an isotopy of such a path in \tilde{M} were to decrease the number of intersections with \tilde{F} (i.e. if the map $T_{\partial F} \rightarrow T_F$ strictly contracted any distance), then putting the isotopy in general position relative to \tilde{F} would reveal a boundary compression of F . Since F is ∂ -incompressible, this is a contradiction.

Applying Theorem 3.1 to $T_{\partial F} \rightarrow T_F$ we conclude that $\partial F \in \mathcal{L}_M$. Since \mathcal{L}_M is an isotropic piecewise linear cone, the desired conclusion follows. \square

Comparing Floyd and Hatcher's proofs with that of Theorem 3.1 shows that the same "cancellation" phenomenon is at work in both cases. Briefly, the connection is as follows. Floyd and Hatcher analyze weight functions on *branched surfaces* that carry all of the incompressible, ∂ -incompressible surfaces in M . Weights on a branched surface satisfy a linear condition at each singular vertex. When the Thurston form is applied to a pair of weights on the boundary train track of a branched surface, these vertex conditions lead to pairwise cancellation of terms in the Thurston form, giving an isotropic space of boundary weights.

The finite set of branched surfaces that are used in this argument come from a construction of Floyd-Oertel [FO] which is based on normal surface theory and a triangulation of the 3-manifold. In this way, the weight conditions at the singular vertices of a branched surface are dual to the weak 4-point condition (3.2) in each 3-simplex that defines the cone \mathcal{L}_M in our approach, and the role of the spaces $W_4(\Delta_M)$ in the proof of Theorem 3.1 is analogous to that of the space of boundary weights of a branched surface in the arguments of Floyd and Hatcher.

5. THE KÄHLER STRUCTURE OF $Q(X)$

The goal of this section is to introduce a Kähler metric on $Q(X)$ and then to show that the foliation map $\mathcal{F} : Q(X) \rightarrow \mathcal{MF}(S)$ identifies the underlying symplectic space with the Thurston symplectic structure on $\mathcal{MF}(S)$. The Kähler metric we construct has singularities but we show that it is smooth relative to a stratification of $Q(X)$.

5.1. The stratification. Let Z be a manifold. A *stratification* of Z is a locally finite collection of locally closed submanifolds $\{Z_i \mid i \in I\}$ of Z , the *strata*, indexed by a set I such that

- (1) $Z = \bigcup_{i \in I} Z_i$
- (2) $Z_i \cap \overline{Z_j} \neq \emptyset$ if and only if $Z_i \subset \overline{Z_j}$

These conditions induce a partial order on I where $i \leq j$ if $Z_i \subset \overline{Z_j}$. A stratification of a complex manifold Z is a *complex analytic stratification* if the closure and boundary of each stratum (i.e. $\overline{Z_i}$ and $\overline{Z_i} \setminus Z_i$) are complex analytic sets.

Let $\mathcal{Q}(S)$ denote the space of holomorphic quadratic differentials on marked Riemann surfaces diffeomorphic to S , i.e. the set of all pairs (X, ϕ) where $X \in \mathcal{T}(S)$ and $\phi \in Q(X)$. This space is a vector bundle over $\mathcal{T}(S)$ isomorphic to the cotangent bundle $T^*\mathcal{T}(S)$. Let $s_0 : \mathcal{T}(S) \rightarrow \mathcal{Q}(S)$ denote the zero section.

There is a natural complex analytic stratification of $\mathcal{Q}(S)$ according to the numbers and types of zeros of the quadratic differential (see [Vee2] [MaSm]). Specifically, let the *symbol* of a nonzero quadratic differential ϕ be the pair (\mathbf{n}, ϵ) where $\mathbf{n} = (n_1, \dots, n_k)$ is the list of multiplicities (in weakly decreasing order) of the zeros of ϕ , and where $\epsilon = \pm 1$ according to whether ϕ is the square of a holomorphic 1-form ($\epsilon = 1$) or not ($\epsilon = -1$). Thus we have $\sum_i n_i = 4g - 4$ and there are finitely many possible symbols; we denote the set of all such symbols by \mathcal{S} .

Given $\pi \in \mathcal{S}$ let $\mathcal{Q}(S, \pi)$ denote the set of quadratic differentials with symbol π . This set is a manifold, with local charts described below (in Section 5.4). The stratification of $\mathcal{Q}(S)$ is formed by the sets $\mathcal{Q}(S, \pi)$ and the zero section $s_0(\mathcal{T}(S))$.

There is a related stratification of a fiber $Q(X)$ with the following properties:

Lemma 5.1. *For each $X \in \mathcal{T}(S)$, there exists a complex analytic stratification $\{Q_i(X)\}$ of $Q(X)$ such that:*

- (i) *Each stratum is a connected and \mathbb{C}^* -invariant.*
- (ii) *The symbol is constant on each stratum $Q_i(X)$, and*
- (iii) *If $q \in Q_i(X)$ and $v \in T_q Q_i(X)$, then the meromorphic function v/q has at most simple poles.*

Proof. A complex analytic stratification can always be refined so that a given complex analytic subset becomes a union of strata (see [Whi] [GM, Thm. 1.6, p. 43]), and a further refinement can be taken so that the strata are connected. Here *refinement* refers to changing the stratification in such a way that each new stratum is entirely contained in one of the old strata.

Applying this to the stratification of $\mathcal{Q}(S)$ discussed above and the closed subvariety $Q(X)$ we obtain a stratification of $\mathcal{Q}(S)$ such that the symbol is constant on each stratum and so that $Q(X)$ is a union of strata. In particular there is an induced stratification $\{Q_i(X)\}$ of $Q(X)$ satisfying (ii).

The original stratification of $\mathcal{Q}(S)$ is \mathbb{C}^* -invariant, and the strata of the refinement can be constructed using finitely many operations that preserve this invariance (i.e. boolean operations and passage from a complex analytic set to its singular locus or to an irreducible component), so property (i) also follows.

Thus the proof is completed by the following lemma, which shows that property (iii) is a consequence of property (ii). \square

Lemma 5.2. *Let $M \subset Q(X)$ be a submanifold on which the symbol is constant. Then for any $(q, \dot{q}) \in TM$, the function \dot{q}/q has at most simple poles on X .*

Proof. Let q_t be a smooth family of quadratic differentials in M with $q_0 = q$ and with tangent vector \dot{q} at $t = 0$.

Let $p \in X$ be a zero of q of order $k > 0$, and choose a local coordinate z in which $z(p) = 0$ and $q = z^k dz^2$. Since q_t has the same symbol as q for small t , in a neighborhood of p we can write

$$q_t = \alpha_t^*(z^k dz^2)$$

where α_t is a smooth family of holomorphic functions defined on $\{|z| < \epsilon\}$ and $\alpha_0(z) = z$. This is equivalent to the statement that the family of polynomial differentials $(z^k + a_{k-2}z^{k-2} + \dots + a_0)dz^2$ is a universal deformation of $z^k dz^2$ (see [HM, Prop. 3.1]). Since $\alpha_t^*(z^k dz^2) = \alpha_t(z)^k (\alpha_t'(z))^2 dz^2$, a calculation gives

$$\dot{q} = z^{k-1} (k\dot{\alpha} + 2z\dot{\alpha}') dz^2$$

and \dot{q} has a zero of order at least $k - 1$ at p . It follows that \dot{q}/q has at most simple poles. \square

5.2. The Kähler form. The vector space $Q(X)$ is a complex manifold with a global parallelization which identifies $T_\phi Q(X) \simeq Q(X)$ for any $\phi \in Q(X)$. We consider the hermitian pairing $\langle \cdot, \cdot \rangle_\phi$ on $T_\phi Q(X)$ defined by

$$(5.1) \quad \langle \psi_1, \psi_2 \rangle_\phi := \int_X \frac{\psi_1 \bar{\psi}_2}{4|\phi|}.$$

Note that in this expression we consider $\psi_1 \bar{\psi}_2 / |\phi|$ as a complex-valued quadratic form on TX , and we integrate the corresponding complexified volume form. With respect to a local choice of a holomorphic 1-form $\sqrt{\phi}$, the integrand can also be written as a

$$\frac{i}{2} \left(\frac{\psi_1}{2\sqrt{\phi}} \wedge \frac{\bar{\psi}_2}{2\sqrt{\phi}} \right).$$

A branched double covering of X can be used to globalize this interpretation (as in the proof of Theorem 5.8 below).

Let g_ϕ and ω_ϕ the real and imaginary parts of this hermitian pairing, i.e.

$$\langle \psi_1, \psi_2 \rangle_\phi = g_\phi(\psi_1, \psi_2) + i \omega_\phi(\psi_1, \psi_2).$$

Similarly, we write $\|\psi\|_\phi^2 = g_\phi(\psi, \psi) = \langle \psi, \psi \rangle_\phi$.

The pairing is not defined for all vectors because the function $\psi_1 \bar{\psi}_2 / |\phi|$ is not necessarily integrable on X . However, it is defined on the strata $Q_i(X)$:

Theorem 5.3. *For each stratum $Q_i(X) \subset Q(X)$ we have:*

- (i) *The pairing $\langle \psi_1, \psi_2 \rangle_\phi$ is well-defined and positive-definite on the tangent bundle $TQ_i(X)$,*
- (ii) *The alternating form ω_ϕ on $TQ_i(X)$ can be expressed as*

$$\omega_\phi = \frac{i}{2} \partial \bar{\partial} N,$$

where $N : Q(X) \rightarrow \mathbb{R}$ is defined by $N(\phi) = \|\phi\|$. In particular ω_ϕ is closed, and thus

- (iii) *The hermitian form $\langle \cdot, \cdot \rangle_\phi$ defines a Kähler structure on $Q_i(X)$.*

A formula for the second derivatives of N equivalent to (ii) above was derived by Royden (see [Roy, Lem. 1]) in the case of the open stratum consisting of differentials with at most simple zeros. Royden also analyzes the failure of N to be twice differentiable in certain directions transverse to the other strata, however the fact that N is C^2 when restricted to a stratum and the analogous derivative formula follow easily by similar methods. We describe the necessary adaptation of his argument below.

Proof. The function $|z|^{-1}$ is integrable in a neighborhood of 0 in \mathbb{C} , so if ψ_1/ϕ has at most simple poles, then $\langle \psi_1, \psi_2 \rangle_\phi$ is finite for all $\psi_2 \in Q(X)$. By part (iii) of Lemma 5.1, this holds for all $(\phi, \psi_1) \in TQ_i(X)$, so the pairing is well-defined there. For any $\psi \neq 0$, the function $\frac{|\psi|^2}{|\phi|}$ is positive except for finitely many zeros, and thus

$$\|\psi\|_\phi^2 = \int_X \frac{|\psi|^2}{4|\phi|} > 0.$$

Now we consider the existence of derivatives of the function $N(\phi) = \int_X |\phi|$. The lack of smoothness of $|\phi|$ at the zeros of ϕ is the only problem: If $K \subset X$ is compact and contains no zeros of ϕ , then $\phi \mapsto \int_K |\phi|$ is C^∞ on a neighborhood of ϕ in $Q(X)$, and its derivatives are obtained by differentiating inside the integral. Thus our strategy will be to determine the resulting formula for $\partial \bar{\partial} N$ away from the zeros and then show that the zeros contribute nothing.

Fix $\phi \in Q_i(X)$ and for $\epsilon > 0$ let X_ϵ denote an open neighborhood of the zero set of ϕ such that each connected component of X_ϵ is a disk containing a single zero of ϕ , and where each such disk admits a local holomorphic coordinate z in which the restriction of ϕ is identified with $z^k dz^2$ on the open disk $\{|z| < \epsilon\} \subset \mathbb{C}$. Here $k \in \mathbb{N}$ depends on the component (and is equal the multiplicity of the zero of ϕ it contains). Such standard disk neighborhoods exist for all ϵ sufficiently small.

Then we can write $N(\phi) = N_0^\epsilon(\phi) + N_1^\epsilon(\phi)$ where:

$$N_0^\epsilon(\phi) = \int_{X_\epsilon} |\phi|$$

$$N_1^\epsilon(\phi) = \int_{X \setminus X_\epsilon} |\phi|$$

As explained above, function $N_1^\epsilon(\phi)$ is smooth on a neighborhood of ϕ in $Q(X)$, so it restricts to a smooth function on $Q_i(X)$. From now on we consider N_0^ϵ and N_1^ϵ as functions on $Q_i(X)$. We claim:

(I) N_0^ϵ is smooth on a neighborhood of ϕ in $Q_i(X)$,

and that at the point ϕ we have

- (II) $\partial\bar{\partial}N_1^\epsilon \rightarrow -2i\omega_\phi$ as $\epsilon \rightarrow 0$, and
 (II) $\partial\bar{\partial}N_0^\epsilon \rightarrow 0$ as $\epsilon \rightarrow 0$.

Since (I) allows the expression $\partial\bar{\partial}N = \partial\bar{\partial}N_1^\epsilon + \partial\bar{\partial}N_0^\epsilon$, taking $\epsilon \rightarrow 0$ and using (II) and (III) gives the desired formula $\frac{i}{2}\partial\bar{\partial}N = \omega_\phi$. Thus the Theorem is reduced to these claims.

We begin with (II), which amounts to differentiating inside the integral defining N_1^ϵ . Let $a : U \rightarrow Q_i(X)$ be a local holomorphic parameterization where $U \subset \mathbb{C}^m$ is a neighborhood of the origin and $a(0) = \phi$. Consider the pointwise norm $|a(\zeta)|$ as a function of ζ . Differentiating $|a| = \sqrt{a\bar{a}}$ we find

$$\frac{\partial^2 |a|}{\partial\zeta_k \partial\bar{\zeta}_l}(0) = \frac{a_k \bar{a}_l}{4|\phi|}$$

where for brevity we have written a_k for $\frac{\partial a}{\partial\zeta_k}(0)$, and this formula is valid at any point where $\phi \neq 0$. Since the vectors a_k form a basis for $T_\phi Q_i(X)$, this equivalent to the statement that for all $\psi_1, \psi_2 \in T_\phi Q_i(X)$ we have $\partial\bar{\partial}n(\psi_1, \psi_2) = (-2i)\frac{1}{4|\phi(z)|} \text{Im}(\psi_1(z)\bar{\psi}_2(z))$ where $n(\psi) = |\psi(z)|$ and z is any point with $\phi(z) \neq 0$. Integrating over the set $X \setminus X_\epsilon$, which is compact and which does not contain any zeros of ϕ , we have the corresponding expression

$$\partial\bar{\partial}N_1^\epsilon(\psi_1, \psi_2) = (-2i) \int_{X \setminus X_\epsilon} \frac{\text{Im}(\psi_1 \bar{\psi}_2)}{4|\phi|}.$$

Since $\psi_1, \psi_2 \in T_\phi Q_i$, the form $\frac{1}{4|\phi|} \text{Im}(\psi_1 \bar{\psi}_2)$ is integrable on X , and by definition its integral over X is $\omega_\phi(\psi_1, \psi_2)$. Since the measure of X_ϵ goes to zero as $\epsilon \rightarrow 0$, claim (II) follows.

Now we consider what happens near the zeros of ϕ . Here it will be essential that we are working in a stratum, so the symbol is constant. This means that the zeros move holomorphically as a function of the point in $Q_i(X)$ and with constant multiplicity. In the same coordinate system where ϕ restricted to a component of X_ϵ becomes $z^k dz^2$ on $\Delta_\epsilon := \{|z| < \epsilon\}$, each ψ in a neighborhood U of ϕ in $Q_i(X)$ can therefore be expressed as

$$e^{h(\psi, z)}(z - u(\psi))^k dz^2$$

where $u : U \rightarrow \Delta_\epsilon$ and $h(\psi, z) : U \times \Delta_\epsilon \rightarrow \mathbb{C}$ are holomorphic functions and $u(\phi) = 0$. Thus $N_0^\epsilon(\psi)$ is the sum of finitely many terms of the form

$$\int_{|z| < \epsilon} |e^{h(\psi, z)}(z - u(\psi))^k| |dz|^2 = \int_{|z| < \epsilon} e^{\operatorname{Re} h(\psi, z)} |z - u(\psi)|^k |dz|^2.$$

Using the change of variable $w = z - u(\psi)$ this integral becomes

$$(5.2) \quad \int_{|w+u(\psi)| < \epsilon} e^{\operatorname{Re} h(\psi, w+u(\psi))} |w|^k |dw|^2.$$

In this form, the integrand is C^∞ as a function of ψ and all of its ψ -derivatives are continuous in w . Furthermore, on the boundary curve $|w + u(\psi)| = \epsilon$ the integrand is smooth in both ψ and w (since for ψ near ϕ this curve avoids the locus $w = 0$ where the integrand may fail to be smooth).

These conditions are exactly what we need to differentiate under the integral in computing derivatives of (5.2), and since N_0^ϵ is a sum of finitely many terms of this form, we find that it is smooth. This establishes (II).

Finally, we must estimate the derivative of N_0^ϵ using (5.2). The second ψ -derivatives of the integral at $\psi = \phi$ split into an interior term (an integral over $|w| < \epsilon$) and boundary terms (integrals over $|w| = \epsilon$). The boundary terms involve up to second derivatives of the boundary curve as a function of ψ and the first partial derivatives of the integrand with respect to both ψ and w . The interior term involves the second ψ -partial derivatives of the integrand. This gives an overall estimate for any second partial derivative of the function N_0^ϵ at ϕ of the form

$$O(A B_{2,0}^{\text{int}} + L B_{1,1}^{\text{bdy}} C_2)$$

where

- $A = \pi\epsilon^2$ is the area of the region of integration at $\psi = \phi$,
- $B_{2,0}^{\text{int}}$ is an interior upper bound (i.e. on $|w| < \epsilon$) for the integrand and its ψ -partial derivatives of order at most 2
- $B_{1,1}^{\text{bdy}}$ is a boundary upper bound (i.e. on $|w| = \epsilon$) for the integrand and its first partial derivatives with respect to ψ and w ,
- $L = 2\pi\epsilon$ is the length of the boundary curve, and
- C_2 is a bound on the ψ -derivatives of the boundary.

Since h and u are holomorphic, their derivatives of any fixed order are uniformly bounded once we take U and ϵ small enough. Examining the integrand of (5.2) we then find $B_{2,0}^{\text{int}} = O(\epsilon^k)$, $B_{1,1}^{\text{bdy}} = O(\epsilon^{k-1})$ (with the w -derivative of the integrand on $|w| = \epsilon$ being the dominant term), and $C_2 = O(1)$. Hence $\partial\bar{\partial}N_0^\epsilon$ is $O(\epsilon^2 \cdot \epsilon^k + \epsilon \cdot \epsilon^{k-1} \cdot 1) = O(\epsilon^k)$ as $\epsilon \rightarrow 0$. Since $k > 0$ this establishes (III). \square

Before proceeding to relate the Kähler structure of Theorem 5.3 to the symplectic structure of $\mathcal{MF}(S)$ we will need to describe convenient local coordinates for both spaces. After discussing suitable period and train track coordinates, we return to the matter of relating these spaces in Section 5.6.

5.3. Double covers and periods. For any $\phi \in Q(X)$ let \tilde{X}_ϕ denote the Riemann surface of the locally-defined one-form $\sqrt{\phi}$ on X , i.e. \tilde{X}_ϕ is a branched double cover of X in which ϕ is canonically expressed as the square of a 1-form (also denoted by $\sqrt{\phi}$). By construction \tilde{X}_ϕ has a holomorphic involution σ such that $\sigma^* \sqrt{\phi} = -\sqrt{\phi}$ and $X = \tilde{X}_\phi / \sigma$.

The one-form $\sqrt{\phi}$ on \tilde{X}_ϕ has *absolute periods* obtained by integration along cycles in $H_1(\tilde{X}_\phi)$ and *relative periods* obtained by integration along cycles in $H_1(\tilde{X}_\phi, \tilde{Z}_\phi)$ where \tilde{Z}_ϕ is the set of zeros of $\sqrt{\phi}$. Since $\sigma^* \sqrt{\phi} = -\sqrt{\phi}$, these integrals vanish for cycles invariant under σ and nontrivial periods are only obtained from cycles in the -1 -eigenspace, which we denote by

$$H_1^-(\phi) := \{c \in H_1(\tilde{X}_\phi, \tilde{Z}_\phi; \mathbb{C}) \mid \sigma_* c = -c\}.$$

Collectively, the periods of $\sqrt{\phi}$ determine its cohomology class as an element of

$$H_-^1(\phi) := \{\theta \in H^1(\tilde{X}_\phi, \tilde{Z}_\phi; \mathbb{C}) \mid \sigma^* \theta = -\theta\}.$$

Note that a saddle connection I of ϕ determines an element $[I] \in H_1^-(\phi)$ by taking the difference of its two lifts to \tilde{X}_ϕ . The result is well-defined up to sign. We can therefore consider relative periods of $\sqrt{\phi}$ along such saddle connections.

Since integration of $\sqrt{\phi}$ gives a local natural coordinate for ϕ , the relative period of a saddle connection is simply its displacement vector in such a coordinate system. In particular, the height of a saddle connection I is given by

$$(5.3) \quad \phi\text{-height}(I) = \left| \operatorname{Im} \int_{[I]} \sqrt{\phi} \right|$$

5.4. Period coordinates for strata. Let $\mathcal{Q}(S, \pi)$ be a stratum in $\mathcal{Q}(S)$. The topological type of the double cover $\tilde{X}_\phi \rightarrow X$ is determined by the symbol of ϕ , so in a small neighborhood U of ϕ in $\mathcal{Q}(S, \pi)$ we can trivialize the family of double covers and (co)homology groups. Thus we can regard each space $H_-^1(\psi)$, where $\psi \in U$, as an instance of a single cohomology space $H_-^1(\pi)$ that is determined by topological information contained in (X, ϕ) ; we let $H_1^-(\pi)$ denote the corresponding trivialization of the family of homology groups. When considering a class $a \in H_1^-(\pi)$ we write a_ψ for a representing cycle in $H_1^-(\psi)$.

Using this local trivialization, the cohomology class of $\sqrt{\phi}$ determines the *relative period map*

$$\operatorname{Per} : U \rightarrow H_-^1(\pi).$$

Explicitly, as a linear function on cycles $a \in H_1^-(\pi)$ the map is given by

$$\operatorname{Per}(\phi)(a) = \int_{a_\phi} \sqrt{\phi}.$$

This map provides local coordinates for strata [Vee2] [Vee1, Sec. 28] [MaSm]:

Theorem 5.4. *The relative period construction gives local biholomorphic coordinates for $\mathcal{Q}(S, \pi)$, i.e. for any sufficiently small open set $U \subset \mathcal{Q}(S, \pi)$ the period map $\text{Per} : U \rightarrow H^1_-(\pi)$ is a diffeomorphism onto an open set. In particular, we have $\dim_{\mathbb{C}} H^1_-(\pi) = \dim_{\mathbb{C}} \mathcal{Q}(S, \pi)$. \square*

While the result above applies to strata in $\mathcal{Q}(S)$, each stratum $Q_i(X)$ of $Q(X)$ is a complex submanifold of $\mathcal{Q}(S, \pi)$ for some $\pi \in \mathcal{S}$, so we have:

Corollary 5.5. *Let $\phi \in Q_i(X)$ be a quadratic differential with symbol π . Then there is an open neighborhood of ϕ in $Q_i(X)$ in which the relative period map to $H^1_-(\pi)$ is biholomorphic onto its image. \square*

Later we will need the following formula for the derivative of the relative period coordinates:

Lemma 5.6 (Douady-Hubbard). *Let $\phi \in Q_i(X)$ and $\psi \in T_{\phi}Q_i(X)$. Then for any $a \in H^1_-(\pi)$ we have*

$$d\text{Per}_{\phi}(\psi)(a) = \int_{a_{\phi}} \frac{\psi}{2\sqrt{\phi}}.$$

\square

The proof in [DH] is for differentials with simple zeros, however, the argument only uses the fact that the period of a saddle connection for a family ϕ_t (where $\phi_0 = \phi$ and $\frac{\partial}{\partial t}\phi|_{t=0} = \psi$) can be expressed in the form

$$\int_{a_t}^{b_t} \sqrt{\phi_t(z)} dz$$

where a_t and b_t are smooth paths traced out by the zeros of ϕ_t as t varies near 0. The assumptions on ϕ, ψ in the Lemma above imply that ψ is tangent to a family of differentials whose zeros have constant multiplicity, and so the same argument applies.

5.5. Adapted train tracks. We now consider train track coordinates for $\mathcal{MF}(S)$ compatible with the relative period construction described above. The following refinement of Lemma 4.3 ensures that we can always choose these coordinates so that the foliation in question lies in the interior of the train track chart.

Lemma 5.7. *For each nonzero $\phi \in Q(X)$ there exists a triangulation Δ of X by saddle connections and a dual maximal train track τ satisfying conditions (i)-(iii) of Lemma 4.3 and such that none of the saddle connections in Δ are horizontal. In particular, the point $[\mathcal{F}(\phi)]$ lies in the interior of the train track chart $\mathcal{MF}(\tau)$.*

Proof. Each face of the train track chart $\mathcal{MF}(\tau)$ is defined by the weight of some branch of the track being zero. In this case the weights are heights of saddle connections, so excluding horizontal edges will result in $\mathcal{F}(\phi)$ being in the interior of the chart.

If the Delaunay triangulation of Lemma 4.3 has no horizontal edges, or if ϕ itself has no horizontal saddle connections, then we are done. Otherwise we must alter the construction of the triangulation to eliminate the horizontal edges. Note that ϕ has only finitely many horizontal saddle connections.

Consider the Teichmüller geodesic (X_t, ϕ_t) determined by $e^{i\theta}\phi$. The Riemann surfaces and quadratic differentials in this family are identified by locally affine maps, so saddle connections of ϕ are also saddle connections of ϕ_t and vice versa. Furthermore, if $\theta \neq \pi$ then horizontal saddle connections of ϕ have ϕ_t -length growing exponentially in t . If we choose $\theta \neq \pi$ so that the geodesic is recurrent in moduli space (a dense set of directions have this property [KW]), then by choosing t large enough we can assume that X_t has bounded $|\phi_t|$ -diameter while the ϕ -horizontal saddle connections are arbitrarily long with respect to $|\phi_t|$.

Since the length of an edge of the Delaunay triangulation is bounded by the diameter of the surface [MaSm, Thm. 4.4], this shows that for large t the ϕ -horizontal saddle connections are not edges of the Delaunay triangulation of ϕ_t . Thus the Delaunay triangulation of ϕ_t gives the desired triangulation by non-horizontal saddle connections of ϕ . \square

5.6. The symplectomorphism. Hubbard and Masur showed that for any $X \in \mathcal{T}(S)$, the foliation map $\mathcal{F} : Q(X) \rightarrow \mathcal{MF}(S)$ is a homeomorphism [HM]. We now show that this map relates the Kähler structure on $Q(X)$ introduced above to the Thurston symplectic structure of $\mathcal{MF}(S)$:

Theorem 5.8. *For any $X \in \mathcal{T}(S)$, the map $\mathcal{F} : Q(X) \rightarrow \mathcal{MF}(S)$ is a real-analytic stratified symplectomorphism. That is:*

- (i) *For any $\phi \in Q_i(X)$ there exists an open neighborhood $U \subset Q_i(X)$ of ϕ in its stratum and a train track coordinate chart $\mathcal{MF}(\tau) \subset \mathcal{MF}(S)$ covering $\mathcal{F}(U)$ so that the restriction*

$$\mathcal{F} : U \rightarrow \mathcal{MF}(\tau)$$

is a real-analytic diffeomorphism onto its image, and

- (ii) *The derivative $d\mathcal{F}_\phi$ defines a symplectic linear map from $T_\phi Q_i(X)$ into $W(\tau) \simeq T_{\mathcal{F}(\phi)}\mathcal{MF}(\tau)$, where $T_\phi Q_i(X)$ is equipped with the symplectic form ω_ϕ and $W(\tau)$ is given the Thurston symplectic form.*

Proof.

- (i) Let $\phi \in Q_i(X)$. Applying Lemma 5.7 we obtain a neighborhood $U \subset Q_i(X)$ of ϕ and a train track τ that carries the horizontal foliation of each $\psi \in U$ by assigning to each branch the height of an associated edge of the ψ -geodesic triangulation.

Lift τ and the dual ϕ -geodesic triangulation Δ to the cover \tilde{X}_ϕ , obtaining a triangulation $\hat{\Delta}$ and double covering of train tracks $\hat{\tau} \rightarrow \tau$. Orient the edges of $\hat{\Delta}$ so that the integral of $\text{Im} \sqrt{\phi}$ over any edge is positive. (This integral is nonzero because the original triangulation

did not have any ϕ -horizontal edges.) Then the integral of $\text{Im } \sqrt{\phi}$ over an edge \widehat{e} is the ϕ -height of the corresponding edge e of Δ .

The covering train tracks and oriented triangulations obtained in this way for other $\psi \in U$ are naturally isotopic to $\widehat{\tau}$ and $\widehat{\Delta}$, so this construction extends throughout U . Thus for all $\psi \in U$ we have realized the weights on τ defining $[\mathcal{F}(\psi)]$ as the imaginary parts of periods of $\sqrt{\psi}$, which by Corollary 5.5 are real-analytic functions.

It remains to show that the derivative of the map to $\mathcal{ML}(\tau)$ is an isomorphism, so that after shrinking U appropriately we have a diffeomorphism onto an open set. However this is a consequence of the proof of (ii) below since the Thurston symplectic form is nondegenerate.

(ii) We need to show that any $\psi_1, \psi_2 \in T_\phi Q_i(X)$ satisfy

$$(5.4) \quad \omega_\phi(\psi_1, \psi_2) = \omega_{\text{Th}}(d\mathcal{F}_\phi(\psi_1), d\mathcal{F}_\phi(\psi_2)).$$

We begin by analyzing the left hand side. Let $\widehat{\psi}_i$ denote the lift of ψ_i to the double cover \widetilde{X}_ϕ and define

$$\theta_i = \frac{\widehat{\psi}_i}{2\sqrt{\phi}} \in \Omega(\widetilde{X}_\phi).$$

These 1-forms are holomorphic because all poles of ψ_i/ϕ are simple and occur at branch points of the covering $\widetilde{X}_\phi \rightarrow X$. Since the 2-form $\frac{i}{2}\theta_1 \wedge \bar{\theta}_2$ is the lift of the integrand of $\langle \psi_1, \psi_2 \rangle_\phi$ to the degree-2 cover \widetilde{X}_ϕ , we have

$$\omega_\phi(\psi_1, \psi_2) = \frac{1}{2} \text{Im} \int_{\widetilde{X}_\phi} \frac{i}{2} \theta_1 \wedge \bar{\theta}_2 = \frac{1}{4} \text{Re} \int_{\widetilde{X}_\phi} \theta_1 \wedge \bar{\theta}_2.$$

For any holomorphic 1-forms θ_i we have

$$\text{Re}(\theta_1 \wedge \bar{\theta}_2) = 2(\text{Re } \theta_1) \wedge (\text{Re } \theta_2) = 2(\text{Im } \theta_1) \wedge (\text{Im } \theta_2)$$

so we can express the integral above as

$$(5.5) \quad \omega_\phi(\psi_1, \psi_2) = \frac{1}{2} \int_{\widetilde{X}_\phi} (\text{Im } \theta_1) \wedge (\text{Im } \theta_2) = \frac{1}{2} [\text{Im } \theta_1] \cdot [\text{Im } \theta_2],$$

where in the last expression $[\alpha]$ denotes the de Rham cohomology class of a closed 1-form α and $[\alpha] \cdot [\beta]$ is the cup product.

Now consider the pairing $\omega_{\text{Th}}(d\mathcal{F}_\phi(\psi_1), d\mathcal{F}_\phi(\psi_2))$. The tangent vector $d\mathcal{F}_\phi(\psi_i) \in W(\tau)$ is a weight function whose value on a branch e is the derivative of the height of the associated edge e' of Δ . The height of an edge is the imaginary part of the period of $\sqrt{\phi}$, so Lemma 5.6 gives a formula for the derivatives of these periods. Namely, after lifting to the covering train track $\widehat{\tau}$ we find that $d\mathcal{F}_\phi(\psi_i)$ corresponds to the weight function $\widehat{w}_i \in W(\widehat{\tau})$ defined by

$$\widehat{w}_i(e) = \int_{\widehat{e}} \text{Im } \theta_i,$$

where e is a branch of τ (identified with its dual edge of Δ) and \widehat{e} is an associated oriented edge of $\widehat{\Delta}$. The orientation of $\widehat{\Delta}$ induces a consistent orientation of $\widehat{\tau}$ so that all intersections of $\widehat{\tau}$ with $\widehat{\Delta}$ become positively oriented. In terms of this orientation, the expression above shows that the de Rham cohomology class $[\text{Im } \theta_i]$ is Poincaré dual to the cycle

$$\widehat{c}_i = \sum_{e \in \widehat{\tau}} w_i(e) \vec{e}.$$

Using the formula (2.2) for the Thurston form as a homological intersection of such cycles and the duality of intersection and cup product, we have

$$\omega_{\text{Th}}(d\mathcal{F}_\phi(\psi_1), d\mathcal{F}_\phi(\psi_2)) = \frac{1}{2} (\widehat{c}_1 \cdot \widehat{c}_2) = \frac{1}{2} [\text{Im } \theta_1] \cdot [\text{Im } \theta_2].$$

With (5.5) this gives the desired equality between symplectic pairings. \square

Remark. The smoothness of the foliation map when restricted to a set of quadratic differentials with constant symbol is implicit in [HM]. Because we consider only tangent vectors to strata in $Q(X)$, the subtle issues that arise from breaking up high-order zeros (and which underlie the failure of differentiability for the full map $Q(X) \rightarrow \mathcal{ML}(S)$) do not arise here.

5.7. Application: The Hubbard-Masur constant. Here we mention an application of Theorem 5.8 that is not used in the sequel. It is immediate from the definition (5.1) that the hermitian form $\langle \psi_1, \psi_2 \rangle_\phi$ is invariant under the action of $S^1 \simeq \{e^{i\theta}\}$ on $Q(X)$ by scalar multiplication:

$$\langle c\psi_1, c\psi_2 \rangle_{c\phi} = \langle \psi_1, \psi_2 \rangle_\phi \quad \text{if } |c| = 1.$$

It follows that this S^1 -action preserves the volume form associated to the stratified Kähler structure on $Q(X)$, and thus the symplectomorphism with $\mathcal{MF}(S)$ gives:

Corollary 5.9. *The action of S^1 on $\mathcal{MF}(S)$ induced by the foliation map $\mathcal{F} : Q(X) \rightarrow \mathcal{MF}(S)$ preserves the volume form associated to the Thurston symplectic structure. \square*

In particular this corollary applies to the *antipodal involution* $i_X : \mathcal{MF}(S) \rightarrow \mathcal{MF}(S)$ which corresponds to multiplication by -1 in $Q(X)$. This map exchanges the vertical and horizontal measured foliations of any quadratic differential on X .

Let $b(X) \subset \mathcal{MF}(S)$ denote the unit ball of the extremal length function on X :

$$b(X) = \{[\nu] \in \mathcal{MF}(S) \mid \text{Ext}_{[\nu]}(X) \leq 1\}.$$

Equivalently $b(X)$ is the image of the L^1 norm ball in $Q(X)$ under the foliation map, so it is invariant under i_X .

Let $\Lambda(X)$ denote the volume of this set with respect to the Thurston symplectic form on $\mathcal{MF}(S)$; this defines the *Hubbard-Masur function* Λ :

$\mathcal{T}(S) \rightarrow \mathbb{R}^+$. This function appears as a coefficient in various counting problems related to the action of the mapping class group $\text{Mod}(S)$ on $\mathcal{T}(S)$ studied in [ABEM].

Using Corollary 5.9, Mirzakhani has shown (personal communication):

Theorem 5.10 (Mirzakhani). *The Hubbard-Masur function is constant. That is, the volume of $b(X)$ depends only on the topological type of S and is independent of the point $X \in \mathcal{T}(S)$.*

The following argument is based on the above-cited communication with Mirzakhani. An analogous statement in a different dynamical context is established in [Yue].

Proof. We will use the antipodal map i_X to show that the derivative of Λ vanishes identically. Since $\mathcal{T}(S)$ is connected it will then follow that Λ is constant.

Let $S(X) = \partial b(X)$ denote the extremal length unit sphere in $\mathcal{MF}(S)$. We think of this as a family of hypersurfaces in $\mathcal{MF}(S)$ parameterized by $X \in \mathcal{T}(S)$.

Fix a point $X_0 \in \mathcal{T}(S)$. For any other point $X \in \mathcal{T}(S)$, both $S(X_0)$ and $S(X)$ intersect each ray in $\mathcal{MF}(S)$ in a single point, so we can consider $S(X)$ as obtained from $S(X_0)$ by scaling each point $[\nu] \in S(X_0)$ by a positive real number

$$(5.6) \quad (\text{Ext}_{[\nu]}(X_0) / \text{Ext}_{[\nu]}(X))^{1/2}.$$

Regarding this expression as a function of $[\nu]$, we have described the spheres $S(X)$ as a family of “radial graphs” over $S(X_0)$. Using this description, the derivative of this family at $X = X_0$ (if it exists) is a vector field along $S(X_0)$ which is *radial*, i.e. it is a pointwise multiple of the vector field $\frac{\partial}{\partial t}$ generating the \mathbb{R}^+ action. Since the scaling function relating $S(X_0)$ to $S(X)$ is a quotient of powers of the extremal length functions, differentiability of this family at $X = X_0$ is a consequence of Gardiner’s formula [Gar], which states that the derivative of extremal length is given by

$$\frac{d}{dt} \text{Ext}_{[\nu]}(X_t) \Big|_{t=0} = 2 \text{Re} \int_{X_0} \mu \mathcal{F}^{-1}([\nu])$$

where X_t is smooth a path in $\mathcal{T}(S)$ and μ is a Beltrami coefficient on X_0 representing $\frac{d}{dt} X_t \Big|_{t=0}$. Differentiating (5.6) using this formula, we find that the derivative of $S(X)$ at $X = X_0$ is the continuous vector field

$$V_\mu([\nu]) = - \left(\text{Re} \int_{X_0} \mu \mathcal{F}^{-1}([\nu]) \right) \frac{\partial}{\partial t}.$$

As above $\frac{\partial}{\partial t}$ is the vector field on $\mathcal{MF}(S)$ generating the \mathbb{R}^+ -action. The derivative of the volume enclosed by $S(X)$ at $X = X_0$ is therefore the integral over $S(X_0)$ of the interior product of this vector field with the

Thurston volume form,

$$d\Lambda_{X_0}(\mu) = \int_{S(X_0)} V_\mu \lrcorner \omega_{\text{Th}}^n$$

where $n = \frac{1}{2} \dim_{\mathbb{R}} \mathcal{MF}(S)$. Since it corresponds to the L^1 norm sphere in $Q(X)$, the sphere $S(X_0)$ is invariant under the antipodal involution. By Corollary 5.9 the volume form ω_{Th}^n is also i_X -invariant. But since $\mathcal{F}^{-1}(i_X([\nu])) = -\mathcal{F}^{-1}([\nu])$, the vector field V_μ is odd under this involution (i.e. $i_X^*(V_\mu) = -V_\mu$) as is the integrand $V_\mu \lrcorner \omega_{\text{Th}}^n$. Since the integral of an odd form over $S(X_0)$ vanishes we have $d\Lambda_{X_0}(\mu) = 0$. \square

6. CHARACTER VARIETIES AND HOLONOMY

6.1. Character varieties. Let G be one of the complex algebraic groups $\text{SL}_2\mathbb{C}$ or $\text{PSL}_2\mathbb{C}$ and let Γ be a finitely generated group. We denote by $\mathcal{R}(\Gamma, G) := \text{Hom}(\Gamma, G)$ the G -representation variety of Γ , which carries an action of G by conjugation. The categorical quotient

$$\mathcal{X}(\Gamma, G) := \mathcal{R}(\Gamma, G) // G$$

is the *character variety*, or more precisely, the variety of characters of representations of Γ in G . See [CS] [MoSh1, Sec. II.4] [HP] for detailed discussion of these spaces. Both $\mathcal{R}(\Gamma, G)$ and $\mathcal{X}(\Gamma, G)$ are affine algebraic varieties defined over \mathbb{Q} . The ring $\mathbb{Q}[\mathcal{X}(\Gamma, \text{SL}_2\mathbb{C})]$ is generated by the *trace functions* $\{t_\gamma\}_{\gamma \in \Gamma}$ which are induced by the conjugation-invariant functions on $\mathcal{R}(\Gamma, G)$ defined by

$$t_\gamma(\rho) = \text{tr}(\rho(\gamma)).$$

Similarly the ring $\mathbb{Q}[\mathcal{X}(\Gamma, \text{PSL}_2\mathbb{C})]$ is generated by the squares of trace functions.

There are two types of natural maps between character varieties that we will use in the sequel. First, the covering map $\text{SL}_2\mathbb{C} \rightarrow \text{PSL}_2\mathbb{C}$ induces a map of character varieties

$$r : \mathcal{X}(\Gamma, \text{SL}_2\mathbb{C}) \rightarrow \mathcal{X}(\Gamma, \text{PSL}_2\mathbb{C}),$$

which is a finite-to-one, proper, and whose image is a union of irreducible components; in fact, the group $H_1(\Gamma, \mathbb{Z}/2\mathbb{Z})$ acts on $\mathcal{X}(\Gamma, \text{SL}_2\mathbb{C})$ by biregular maps, and r is the quotient mapping for this action [MS1, Sec. V.1]. Secondly, if $\phi : \Gamma \rightarrow \Gamma'$ is a group homomorphism, then composing representations with ϕ induces a map of character varieties

$$\phi^* : \mathcal{X}(\Gamma', G) \rightarrow \mathcal{X}(\Gamma, G),$$

which is a regular map.

These constructions are functorial in the sense that the maps r and φ^* fit into a commutative diagram

$$(6.1) \quad \begin{array}{ccc} \mathcal{X}(\Gamma', \mathrm{SL}_2\mathbb{C}) & \xrightarrow{\varphi^*} & \mathcal{X}(\Gamma, \mathrm{SL}_2\mathbb{C}) \\ r \downarrow & & \downarrow r \\ \mathcal{X}(\Gamma', \mathrm{PSL}_2\mathbb{C}) & \xrightarrow{\varphi^*} & \mathcal{X}(\Gamma, \mathrm{PSL}_2\mathbb{C}) \end{array}$$

Since character varieties we consider are for groups of the form $\Gamma = \pi_1 N$ where N is a compact 2- or 3-manifold, we often use the abbreviated notation

$$\mathcal{X}(N, G) := \mathcal{X}(\pi_1 N, G).$$

Note that this algebraic variety does not depend on a choice of orientation for N .

6.2. The Morgan-Shalen compactification. In [MoSh1] a compactification of $\mathcal{X}(\Gamma, \mathrm{SL}_2\mathbb{C})$ is defined by mapping $\mathcal{X}(\Gamma, \mathrm{SL}_2\mathbb{C})$ into the infinite-dimensional projective space $\mathbb{P}(\mathbb{R}^\Gamma) := (\mathbb{R}^\Gamma \setminus \{0\})/\mathbb{R}^+$ by

$$(6.2) \quad [\rho] \mapsto (\log(|t_\gamma(\rho)| + 2))_{\gamma \in \Gamma}.$$

The image of this map is precompact and taking the closure gives the *Morgan-Shalen compactification* of $\mathcal{X}(\Gamma, \mathrm{SL}_2\mathbb{C})$. A boundary point $[\ell]$ of this compactification is a projective equivalence class of functions $\ell : \Gamma \rightarrow \mathbb{R}$, and any function arising this way is the translation length function of an action of Γ on an \mathbb{R} -tree by isometries. These \mathbb{R} -trees are constructed algebraically, using valuations on the function fields of subvarieties of $\mathcal{X}(\Gamma, \mathrm{SL}_2\mathbb{C})$. The intermediate stages of this algebraic construction also involve Λ -trees of higher rank. For later use we will now recall some key steps in their construction.

6.3. Valuation constructions. In what follows we consider irreducible subvarieties $V \subset \mathcal{X}(\Gamma, \mathrm{SL}_2\mathbb{C})$, and k will denote a countable subfield of \mathbb{C} over which V is defined. (For example if $\mathcal{X}(\Gamma, \mathrm{SL}_2\mathbb{C})$ is irreducible we can take $V = \mathcal{X}(\Gamma, \mathrm{SL}_2\mathbb{C})$ and $k = \mathbb{Q}$.) The function field $k(V)$ of such a variety is a finitely generated extension of k , and we consider k -valuations $v : k(V)^* \rightarrow \Lambda$, where Λ is an ordered abelian group. Without loss of generality we can assume that Λ has finite rank [ZS, Ch. 5, Sec. 10]. A valuation is *supported at infinity* if there exists a regular function $f \in k[V]$ with $v(f) < 0$.

Boundary points of the Morgan-Shalen compactification correspond to valuations as follows:

Theorem 6.1 ([MoSh1, Thm. I.3.6]). *If $V \subset \mathcal{X}(\Gamma, \mathrm{SL}_2\mathbb{C})$ is an irreducible subvariety defined over k and $[\ell]$ is a boundary point of V in the Morgan-Shalen compactification, then there exists a valuation $v : k(V)^* \rightarrow \Lambda$ such that*

- (i) v is supported at infinity,
- (ii) If $\Lambda_1 \subset \Lambda$ is the minimal nontrivial convex subgroup, then for each $\gamma \in \Gamma$ either $v(t_\gamma) > 0$ or $v(t_\gamma) \in \Lambda_1$, and

- (iii) *There is an order-preserving embedding $p : \Lambda_1 \rightarrow \mathbb{R}$ such that $\ell(\gamma) = p(\max(-v(t_\gamma), 0))$.* \square

Note that the embedding p is unique up to multiplication by a positive constant (by Theorem 2.2) and that condition (iii) above shows that ℓ can be recovered from the valuation v .

The link between valuations and Λ -trees is given by:

Theorem 6.2 ([MoSh1, Thm. II.4.3 and Lem. II.4.5]). *If $V \subset \mathcal{X}(\Gamma, \mathrm{SL}_2\mathbb{C})$ is an irreducible subvariety defined over k and $v : k(V)^* \rightarrow \Lambda$ is a valuation supported at infinity, then there is an isometric action of Γ on a Λ -tree whose translation length function $\ell : \Gamma \rightarrow \Lambda$ satisfies*

$$\ell(\gamma) = \max(-v(t_\gamma), 0).$$

\square

Remark. The statement of Lemma II.4.5 in [MoSh1] involves only \mathbb{R} -trees, however a Λ -tree satisfying the conditions above is constructed as part of its proof. The lemma also involves an additional condition on the valuation, (equivalent to (ii) of Theorem 6.1 above), but this condition is only used at the final step to produce an \mathbb{R} -tree from the Λ -tree. Additional discussion of the Λ -tree construction underlying Theorem 6.2 can be found in [Mor2, Thm. 16] and [MoSh2, pp. 232–233].

6.4. The extension variety. Let M be a compact 3-manifold with connected boundary S , and let $i_* : \pi_1 S \rightarrow \pi_1 M$ be the map induced by inclusion of the boundary. As discussed above, such a homomorphism induces a map of character varieties

$$i^* : \mathcal{X}(M, G) \rightarrow \mathcal{X}(S, G).$$

We call this the *restriction map*. Since it is a regular map of algebraic varieties, the image of i^* is a constructible set which contains a Zariski open subset of its closure.

Considering the case $G = \mathrm{SL}_2\mathbb{C}$, we denote the closure of the image by

$$\mathcal{E}_M := \overline{i^*(\mathcal{X}(M, \mathrm{SL}_2\mathbb{C}))}^{\mathrm{Zariski}},$$

which we call the *extension variety*, since its generic points are conjugacy classes of representations of $\pi_1 S$ that are trivial on $\ker(i_*)$ and which admit an extension from $i_*(\pi_1 S)$ to its supergroup $\pi_1 M$. Note that \mathcal{E}_M is an algebraic subvariety of $\mathcal{X}(S, \mathrm{SL}_2\mathbb{C})$.

Since points in the Morgan-Shalen boundary of $\mathcal{X}(S, \mathrm{SL}_2\mathbb{C})$ correspond to length functions of actions of $\pi_1 S$ on \mathbb{R} -trees, it is natural to expect that the length functions that arise as boundary points of \mathcal{E}_M would have a similar extension property. We now show that this is true if we allow the extended length function to take values in a higher-rank group, \mathbb{R}^n with the lexicographical order.

Theorem 6.3. *Let $[\ell]$ be a boundary point of \mathcal{E}_M in the Morgan-Shalen compactification of $\mathcal{X}(S, \mathrm{SL}_2\mathbb{C})$. Then $\ell : \pi_1 S \rightarrow \mathbb{R}$ extends to a length*

function of an action of $\pi_1 M$ on a \mathbb{R}^n -tree, i.e. there exists a function $\widehat{\ell} : \pi_1 M \rightarrow \mathbb{R}^n$ such that

- (i) The group $\pi_1 M$ acts isometrically on a \mathbb{R}^n -tree with translation length function $\widehat{\ell}$.
- (ii) For each $\gamma \in \pi_1 S$ we have

$$\widehat{\ell}(i_*(\gamma)) = i_n(\ell(\gamma))$$

where $i_* : \pi_1 S \rightarrow \pi_1 M$ is the map induced by the inclusion of S as the boundary of M and $i_n : \mathbb{R} \rightarrow \mathbb{R}^n$ is the order-preserving inclusion as the last (least significant) factor.

Proof. Since $[\ell]$ is a boundary point of \mathcal{E}_M , it is a boundary point of one of its irreducible components. Let \mathcal{E}_M^0 be such a component, and let \mathcal{X}_M^0 be a corresponding irreducible component of $\mathcal{X}(M, \mathrm{SL}_2\mathbb{C})$ so that $i^*(\mathcal{X}_M^0)$ contains a Zariski open subset of \mathcal{E}_M^0 .

Since $i^* : \mathcal{X}_M^0 \rightarrow \mathcal{E}_M^0$ is dominant, it induces an extension of function fields $k(\mathcal{E}_M^0) \hookrightarrow k(\mathcal{X}_M^0)$, where k is a finite extension of \mathbb{Q} over which \mathcal{E}_M^0 and \mathcal{X}_M^0 are defined. Note that when considering $k(\mathcal{E}_M^0)$ as a subfield of $k(\mathcal{X}_M^0)$, the element of $k(\mathcal{E}_M^0)$ represented by the trace function t_γ , $\gamma \in \pi_1 S$, is identified with element of $k(\mathcal{X}_M^0)$ represented by the trace function $t_{i_*(\gamma)}$.

Let $v : k(\mathcal{E}_M^0)^* \rightarrow \Lambda$ and $p : \Lambda_1 \rightarrow \mathbb{R}$ be the valuation and embedding associated to ℓ by Theorem 6.1. Since $k(\mathcal{X}_M^0)$ is finitely generated over $k(\mathcal{E}_M^0)$, the standard extension theorem for valuations (see [ZS, Thm. 5', p. 13], [MoSh1, Lem. II.4.4]) gives an ordered abelian group Λ' such that $\Lambda \subset \Lambda'$ and $\Lambda' \subset m\Lambda$ for some $m \in \mathbb{N}$, and a valuation

$$v' : k(\mathcal{X}_M^0)^* \rightarrow \Lambda'$$

so that $v'(f) = v(f)$ for any $f \in k(\mathcal{E}_M)$. Since $\Lambda' \subset m\Lambda$ it follows that the minimal convex subgroups satisfy $\Lambda_1 \subset \Lambda'_1$.

By Lemma 2.4, we have a commutative diagram of order-preserving embeddings

$$(6.3) \quad \begin{array}{ccc} \Lambda_1 & \longrightarrow & \Lambda' \\ p \downarrow & & \downarrow F \\ \mathbb{R} & \xrightarrow{i_n} & \mathbb{R}^n \end{array}$$

We can arrange that the left vertical map in this diagram agrees with the embedding $p : \Lambda_1 \rightarrow \mathbb{R}$ considered above; this is possible since there is a unique such embedding up to scale (by Theorem 2.2), and both vertical maps in the diagram can be multiplied by an arbitrary positive constant while preserving commutativity and order.

Applying Theorem 6.2 to v' we obtain a Λ' -tree T' on which $\pi_1 M$ acts by isometries with length function ℓ' . Let $T = T' \otimes_{\Lambda'} \mathbb{R}^n$ be the \mathbb{R}^n -tree associated to T' by the embedding F , and let $\widehat{\ell} : \pi_1 M \rightarrow \mathbb{R}^n$ be its length function. Condition (i) is satisfied by definition.

It remains to verify condition (ii). For any $\gamma \in \pi_1 S$ we have $t_{i_*(\gamma)} \in k(\mathcal{X}_M^0)$ and the length function $\widehat{\ell}$ satisfies:

$$\begin{aligned}
 \widehat{\ell}(i_*(\gamma)) &= F(\ell'(i_*(\gamma))) && \text{by definition of } T \\
 &= F(\max(-v'(t_{i_*(\gamma)}), 0)) && \text{by Theorem 6.2} \\
 (6.4) \quad &= F(\max(-v(t_\gamma), 0)) && \text{since } v' \text{ extends } v \\
 &= i_n(p(\max(-v(t_\gamma), 0))) && \text{by commutativity of (6.3)} \\
 &= i_n(\ell(\gamma)) && \text{by Theorem 6.1}
 \end{aligned}$$

□

Remark. It is natural to ask whether the extended length function $\widehat{\ell}$ of Theorem 6.3 can always be taken to be \mathbb{R} -valued, thus avoiding the introduction of \mathbb{R}^n -trees ($n > 1$). The proof above shows a potential obstruction. For any $\gamma \in \pi_1 S$ the valuation $v(t_\gamma)$ is either positive or it lies in the rank-1 convex subgroup Λ_1 , but it is not clear whether this holds for the extended valuation v' applied to a trace function of an element in $\pi_1 M$. A rank-1 subgroup containing the negative valuations of all trace functions is needed in order to apply the construction of [MoSh1, Sec. II.4] to produce an \mathbb{R} -tree from the Λ -tree while preserving the action of $\pi_1 M$ and the length function.

Since these valuations are associated to boundary points of the Morgan-Shalen compactification, this is effectively a question about comparing the rate of growth of traces in a sequence of $\pi_1 S$ -representations in \mathcal{E}_M to that of an associated sequence of $\pi_1 M$ -representations. Alternatively, in the terminology of [D, Sec. 6], we ask whether the *local scales* of $\mathrm{PSL}_2\mathbb{C}$ -representations of $\pi_1 M$ are comparable (within a uniform multiplicative constant) to those of the restrictions to $i_*(\pi_1 S)$, or if such a uniform comparison is possible for *some* sequence representing any given boundary point.

6.5. The holonomy variety. Let X be a marked Riemann surface structure on S . Here we allow that the complex structure of X induces an orientation opposite that of S , so either $X \in \mathcal{T}(S)$ or $X \in \mathcal{T}(\overline{S})$.

The vector space $Q(X)$ of holomorphic quadratic differentials can be identified with the set of complex projective structures (\mathbb{CP}^1 -structures) on X . Here $0 \in Q(X)$ corresponds to the projective structure induced by the uniformization of X .

Each \mathbb{CP}^1 structure on X has an associated holonomy representation $\pi_1 S \rightarrow \mathrm{PSL}_2\mathbb{C}$, which is well-defined up to conjugacy. Considering the conjugacy class of the holonomy representation as a function of the projective structure gives the *holonomy map*

$$\mathrm{hol} : Q(X) \rightarrow \mathcal{X}(S, \mathrm{PSL}_2\mathbb{C}).$$

This map can be lifted through $r : \mathcal{X}(S, \mathrm{SL}_2\mathbb{C}) \rightarrow \mathcal{X}(S, \mathrm{PSL}_2\mathbb{C})$ in several ways; the set of such lifts is naturally in bijection with the set $\mathrm{Spin}(X)$ of spin structures on X , which is a finite set acted upon simply transitively

by $H_1(S, \mathbb{Z}/2\mathbb{Z})$. For each $\varepsilon \in \text{Spin}(X)$ we denote the corresponding lifted holonomy map by

$$\text{hol}_\varepsilon : Q(X) \rightarrow \mathcal{X}(S, \text{SL}_2\mathbb{C}),$$

so $\text{hol} = r \circ \text{hol}_\varepsilon$. The maps hol and hol_ε are proper holomorphic embeddings [GKM, Thm. 11.4.1].

Define $\mathcal{H}_{X,\varepsilon} := \text{hol}_\varepsilon(Q(X))$, which is therefore a complex analytic subvariety of $\mathcal{X}(S, \text{SL}_2\mathbb{C})$. Taking the union of these subvarieties we obtain the *holonomy variety*

$$\mathcal{H}_X := \bigcup_{\varepsilon \in \text{Spin}(X)} \mathcal{H}_{X,\varepsilon} \subset \mathcal{X}(S, \text{SL}_2\mathbb{C}),$$

an analytic variety with irreducible components $\mathcal{H}_{X,\varepsilon}$. Equivalently, we have $\mathcal{H}_X = r^{-1}(\text{hol}(Q(X)))$.

We will be interested in the limiting behavior of \mathcal{H}_X in the Morgan-Shalen compactification and how this relates to the parameterizations of its components by $Q(X)$. Consider a divergent sequence $\phi_n \in Q(X)$. Since the unit sphere in $Q(X)$ is compact, by passing to a subsequence we can assume that $c_n \phi_n \rightarrow \phi$ as $n \rightarrow \infty$, where $c_n \in \mathbb{R}^+$ is a suitable sequence of scale factors with $c_n \rightarrow 0$. We call a limit point ϕ obtained this way a *projective limit* of $\{\phi_n\}$. Projective limits in $Q(X)$ are related to limits of holonomy representations in $\mathcal{X}(S, \text{SL}_2\mathbb{C})$ as follows:

Theorem 6.4 ([D, Thm. A]). *If $\phi_n \in Q(X)$ is a divergent sequence with projective limit ϕ , then any accumulation point of $\text{hol}_\varepsilon(\phi_n)$ in the Morgan-Shalen boundary is represented by an \mathbb{R} -tree T that admits an equivariant, surjective straight map $T_\phi \rightarrow T$. \square*

6.6. Intersections and the isotropic cone. We now consider the intersection of the holonomy variety and the extension variety. Since $\mathcal{H}_{X,\varepsilon}$ is parameterized by $Q(X)$, the intersection $\mathcal{H}_{X,\varepsilon} \cap \mathcal{E}_M$ is parameterized by a subset of $Q(X)$ which we denote by

$$\mathcal{V}_{M,\varepsilon} = \text{hol}_\varepsilon^{-1}(\mathcal{E}_M) \subset Q(X),$$

and $\mathcal{H}_X \cap \mathcal{E}_M$ is the union of the images of these sets under the respective holonomy maps.

Combining the main results of Section 4 with theorems 6.3–6.4, we have the following characterization of the limit points of $\mathcal{V}_{M,\varepsilon}$:

Theorem 6.5. *Let $\{\phi_n\} \subset \mathcal{V}_{M,\varepsilon}$ be a divergent sequence with projective limit ϕ . Then $[\mathcal{F}(\phi)] \in \mathcal{L}_{M,X}$ where $\mathcal{L}_{M,X}$ is the isotropic cone of Theorem 4.1.*

Note that in this theorem we regard $\mathcal{L}_{M,X}$ as a subset of $\mathcal{MF}(S)$ regardless of whether $X \in \mathcal{T}(S)$ or $X \in \mathcal{T}(\bar{S})$. This is possible since the natural identification $\mathcal{MF}(S) \simeq \mathcal{MF}(\bar{S})$ preserves the property of being an isotropic cone (while changing the sign of the symplectic form).

Proof. Let $[\ell]$ be an accumulation point of $\text{hol}_\varepsilon(\phi_n)$ in the Morgan-Shalen compactification. By Theorem 6.4, there is a surjective, equivariant straight map $T_\phi \rightarrow T$ where T is an \mathbb{R} -tree on which $\pi_1 S$ acts with length function ℓ .

By Theorem 6.3, the length function ℓ extends to a length function $\widehat{\ell} : \pi_1 M \rightarrow \mathbb{R}^n$ of an action of $\pi_1 M$ on a \mathbb{R}^n -tree \widehat{T} . By Lemma 2.7 there is an \mathbb{R} -tree $T' \subset \widehat{T}$ on which $\pi_1 S$ acts with length function $\ell = \widehat{\ell}|_{\pi_1 S}$ (where we use the embedding $i_n : \mathbb{R} \rightarrow \mathbb{R}^n$ as the least significant factor to identify \mathbb{R} with the minimal convex subgroup of \mathbb{R}^n). Note that T and T' are then isospectral \mathbb{R} -trees in the terminology of Section 4.3.

Finally we apply Theorem 4.4 to the isospectral \mathbb{R} -trees T, T' , straight map $T_\phi \rightarrow T$, the inclusion of T' as a subtree of \widehat{T} to conclude that $[\mathcal{F}(\phi)] \in \mathcal{L}_{M,X}$. \square

7. ANALYTIC GEOMETRY IN $Q(X)$

Theorem 6.5 shows that the large-scale behavior of the set $\mathcal{V}_{M,\varepsilon} \subset Q(X)$ is constrained by an isotropic cone in $\mathcal{MF}(S)$. The goal of this section is to complete the proof of Theorem B by showing that only a discrete set can satisfy this constraint.

We begin with some generalities on real and complex limit points of sets in a complex vector space.

7.1. Real and complex boundaries. The vector space \mathbb{C}^n can be compactified to $\mathbb{C}\mathbb{P}^n$ by adjoining a hyperplane at infinity $\mathbb{C}\mathbb{P}^{n-1} = (\mathbb{C}^n \setminus \{0\})/\mathbb{C}^*$. When considering this compactification we regard $\mathbb{C}\mathbb{P}^{n-1}$ as the *complex boundary* of \mathbb{C}^n , and write $\partial_{\mathbb{C}}\mathbb{C}^n = \mathbb{C}\mathbb{P}^{n-1}$.

Given a set $R \subset \mathbb{C}^n$, let $\partial_{\mathbb{C}}R \subset \mathbb{C}\mathbb{P}^{n-1}$ denote its set of accumulation points in complex boundary $\partial_{\mathbb{C}}R = \overline{R} \cap \mathbb{C}\mathbb{P}^{n-1}$ where \overline{R} is the closure of R in $\mathbb{C}\mathbb{P}^n$.

In the real analogue of these constructions we identify the sphere S^{2n-1} with the set of rays from the origin in \mathbb{C}^n ; for any $q \in S^{2n-1}$ let $r_q \subset \mathbb{C}^n$ denote the corresponding open ray. There is a corresponding compactification $\overline{B}^{2n} = \mathbb{C}^n \sqcup S^{2n-1}$ where \overline{B}^{2n} is the closed ball of dimension $2n$; here a sequence z_k converges to $q \in S^{2n-1}$ if it can be rescaled by positive real constants so as to converge to a point in r_q . In this sense S^{2n-1} is the *real boundary* of \mathbb{C}^n and write $\partial_{\mathbb{R}}\mathbb{C}^n = S^{2n-1}$.

Given a set $R \subset \mathbb{C}^n$, let $\partial_{\mathbb{R}}R \subset S^{2n-1}$ denote the accumulation points of R in the boundary of the real compactification \overline{B}^{2n} . This real boundary of R corresponds to a set of open rays in \mathbb{C}^n , and we denote by $\text{Cone}_{\mathbb{R}}(R)$ the union of these rays, i.e.

$$\text{Cone}_{\mathbb{R}}(R) = \bigcup_{q \in \partial_{\mathbb{R}}R} r_q.$$

Equivalently $\text{Cone}_{\mathbb{R}}(R)$ is the set of projective limits of the set R in the sense of Section 6.5.

Mapping a ray in \mathbb{C}^n to the complex line it spans induces the *Hopf fibration* $\Pi : S^{2n-1} \rightarrow \mathbb{C}\mathbb{P}^{n-1}$, which gives S^{2n-1} the structure of a principal S^1 -bundle. Identifying S^{2n-1} with the unit sphere in \mathbb{C}^n , the map Π is the restriction of the quotient map $\widehat{\Pi} : (\mathbb{C}^n \setminus \{0\}) \rightarrow \mathbb{C}\mathbb{P}^{n-1}$, which is holomorphic. It is immediate from the definitions that for any set $R \subset \mathbb{C}^n$ we have $\partial_{\mathbb{C}}R = \Pi(\partial_{\mathbb{R}}R)$.

We call a fiber of Π a *Hopf circle*. For any Hopf circle $C = \Pi^{-1}(p) \subset S^{2n-1}$, the set $\bigcup_{q \in C} r_q = \widehat{\Pi}^{-1}(p) \subset \mathbb{C}^n$ is a punctured complex line, and more generally if $I \subset S^{2n-1}$ is an open arc of a Hopf circle, then $\bigcup_{q \in I} r_q$ is an open sector in a complex line.

The following elementary lemma allows us to recognize totally real submanifolds of $\mathbb{C}\mathbb{P}^{n-1}$ arising as boundaries of cones in \mathbb{C}^n :

Lemma 7.1. *Let $L \subset \mathbb{C}^n \setminus \{0\}$ be an \mathbb{R}^+ -invariant and totally real submanifold of dimension m . Then $\partial_{\mathbb{C}}L$ (i.e. the projection of L to $\mathbb{C}\mathbb{P}^{n-1}$) is an immersed totally real submanifold of $\mathbb{C}\mathbb{P}^{n-1}$ of dimension $m - 1$.*

Proof. By \mathbb{R}^+ -invariance, intersecting L with the unit sphere gives a manifold $L_1 \subset S^{2n-1}$ of dimension $m - 1$ naturally identified with $\partial_{\mathbb{R}}L$, and $\partial_{\mathbb{C}}L = \Pi(L_1)$. Let $x \in L_1$. Because L is totally real, the tangent space to $T_x L_1$ is transverse to $i\mathbb{R} \cdot x = \ker d\Pi_x$, thus $\Pi|_{L_1}$ is an immersion. Since the \mathbb{C} -span of $T_x L_1$ has complex dimension $m - 1$ and is transverse to $\mathbb{C} \cdot x$, the differential $d\widehat{\Pi}$ maps it injectively and complex-linearly to $T_{\Pi(x)}\mathbb{C}\mathbb{P}^{n-1}$. The image of the totally real subspace $T_x L_1$ under such a map is totally real, and the lemma follows. \square

While the real and complex boundary constructions have been described for \mathbb{C}^n , they apply naturally to any finite-dimensional complex vector space; we will apply them to $Q(X) \simeq \mathbb{C}^{3g-3}$.

7.2. Real and complex boundaries of $\mathcal{V}_{M,\varepsilon}$. Using the symplectic properties of the foliation map (Theorem 5.8), we can now translate the properties of $\mathcal{V}_{M,\varepsilon}$ established in Theorem 6.5 into analytic conditions satisfied by its real and complex boundaries:

Theorem 7.2. *Let $\mathcal{V}_{M,\varepsilon} = \text{hol}_{\varepsilon}^{-1}(\mathcal{E}_M) \subset Q(X)$ be the set of holomorphic quadratic differentials corresponding to projective structures with holonomy in the extension variety of the 3-manifold M . Then:*

- (i) $\partial_{\mathbb{C}}\mathcal{V}_{M,\varepsilon}$ is locally contained in a totally real manifold.
That is, either $\partial_{\mathbb{C}}\mathcal{V}_{M,\varepsilon} = \emptyset$ or there exists $p \in \partial_{\mathbb{C}}\mathcal{V}_{M,\varepsilon}$, a neighborhood U of p in $\partial_{\mathbb{C}}Q(X)$, and a totally real, real-analytic submanifold $N \subset U$ such that $(\partial_{\mathbb{C}}\mathcal{V}_{M,\varepsilon} \cap U) \subset N$.
- (ii) $\partial_{\mathbb{R}}\mathcal{V}_{M,\varepsilon}$ does not contain an open arc of any Hopf circle.
That is, for any $p \in \partial_{\mathbb{C}}\mathcal{V}_{M,\varepsilon}$ the intersection $\Pi^{-1}(p) \cap \partial_{\mathbb{R}}\mathcal{V}_{M,\varepsilon}$ has empty interior in the relative topology of $\Pi^{-1}(p) \simeq S^1$.

Proof.

(i) Define

$$\mathcal{C}_{M,\varepsilon} := \text{Cone}_{\mathbb{R}}(\mathcal{V}_{M,\varepsilon}),$$

and suppose that $\partial_{\mathbb{C}}\mathcal{V}_{M,\varepsilon} \neq \emptyset$ so that $\mathcal{C} \neq \emptyset$. Note that $\partial_{\mathbb{R}}\mathcal{V}_{M,\varepsilon} = \partial_{\mathbb{R}}\mathcal{C}_{M,\varepsilon}$ and $\partial_{\mathbb{C}}\mathcal{V}_{M,\varepsilon} = \partial_{\mathbb{C}}\mathcal{C}_{M,\varepsilon}$. By Theorem 6.5 we have $\mathcal{C}_{M,\varepsilon} \subset \mathcal{F}^{-1}(\mathcal{L}_{M,X})$ and so

$$\partial_{\mathbb{C}}\mathcal{V}_{M,\varepsilon} \subset \partial_{\mathbb{C}}\mathcal{F}^{-1}(\mathcal{L}_{M,X}).$$

Thus it will suffice to locally cover $\partial_{\mathbb{C}}\mathcal{F}^{-1}(\mathcal{L}_{M,X})$ by a totally real, real-analytic manifold and to ensure that this set contains a limit point of $\mathcal{C}_{M,\varepsilon}$.

Among the strata $Q_i(X)$ intersected by $\mathcal{C}_{M,\varepsilon}$, let $Q_k(X)$ be a maximal element (one not contained in the boundary of another stratum intersecting $\mathcal{C}_{M,\varepsilon}$). Thus $Q_k(X)$ has an open tubular neighborhood $U_0 \subset Q(X)$ disjoint from all other strata intersecting $\mathcal{C}_{M,\varepsilon}$, so that $\emptyset \neq (\mathcal{C}_{M,\varepsilon} \cap U_0) \subset Q_k(X)$. Furthermore since $Q_k(X)$ is \mathbb{C}^* -invariant, we can choose U_0 to be \mathbb{C}^* -invariant as well. In particular the set of complex lines in U_0 is an open set $\partial_{\mathbb{C}}U_0 \subset \partial_{\mathbb{C}}Q(X)$.

By Theorem 5.8, the intersection $(\mathcal{F}^{-1}(\mathcal{L}_{M,X}) \cap Q_k(X))$ is a semi-analytic set, i.e. it is locally defined by finitely many equations and inequalities of real-analytic functions. Indeed, for each $p \in (\mathcal{C}_{M,\varepsilon} \cap U_0)$ the theorem gives an open neighborhood V of p in $Q_k(X)$ such that $\mathcal{F} : V \rightarrow \mathcal{ML}(\tau)$ is a real-analytic map into a train track chart, and $\mathcal{ML}(\tau) \cap \mathcal{L}_{M,X}$ is a union of convex cones in linear subspaces (thus semianalytic).

Since $(\mathcal{F}^{-1}(\mathcal{L}_{M,X}) \cap Q_k(X))$ is a union of rays in $Q(X)$, its real boundary is the same as its intersection with the unit sphere in $Q(X)$ and is also semianalytic. Thus the set

$$\partial_{\mathbb{C}}(\mathcal{F}^{-1}(\mathcal{L}_{M,X}) \cap Q_k(X)) = \Pi(\partial_{\mathbb{R}}(\mathcal{F}^{-1}(\mathcal{L}_{M,X}) \cap Q_k(X)))$$

is the image of a semianalytic set under a proper real-analytic mapping, i.e. a subanalytic set. Such sets can be stratified by connected, real-analytic, subanalytic manifolds (see [Hir] [Har]), and furthermore stratifications of $\partial_{\mathbb{R}}(\mathcal{F}^{-1}(\mathcal{L}_{M,X}) \cap Q_k(X))$ and $\partial_{\mathbb{C}}(\mathcal{F}^{-1}(\mathcal{L}_{M,X}) \cap Q_k(X))$ can be chosen so that Π maps strata to strata and so that the differential of Π has constant rank on each stratum [Har, Cor. 4.4].

Let $N \subset \partial_{\mathbb{C}}(\mathcal{F}^{-1}(\mathcal{L}_{M,X}) \cap Q_k(X))$ be a stratum maximal among those intersecting $\partial_{\mathbb{C}}\mathcal{V}_{M,\varepsilon}$. Taking a tubular neighborhood U of N in $\partial_{\mathbb{C}}U_0$ gives an open set in $\partial_{\mathbb{C}}Q(X)$ in which $\emptyset \neq (\partial_{\mathbb{C}}\mathcal{V}_{M,\varepsilon} \cap U) \subset (N \cap U)$.

It remains to show that the real-analytic manifold N is totally real. Since Π is surjective, for each $p \in N$ there is a stratum

$$N' \subset \partial_{\mathbb{R}}(\mathcal{F}^{-1}(\mathcal{L}_{M,X}) \cap Q_k(X))$$

such that $\Pi(N')$ contains an open neighborhood of p in N . Thus $N'' := \text{Cone}_{\mathbb{R}}(N')$ is a real-analytic submanifold of $Q(X)$ and $N'' \subset (\mathcal{F}^{-1}(\mathcal{L}_{M,X}) \cap Q_k(X))$.

Theorem 5.8 then implies that N'' locally maps by \mathcal{F} into a finite union of isotropic subspaces of a train track chart $\mathcal{ML}(\tau)$, and that in these local charts \mathcal{F} is a real-analytic, symplectic map. Thus N'' is isotropic with respect to the symplectic structure of $Q_k(X)$ given by Theorem 5.3. Since the symplectic form on $Q_k(X)$ is induced by a Kähler structure, an isotropic manifold is totally real. We have therefore described an open neighborhood of an arbitrary point $p \in N$ as $\Pi(\partial_{\mathbb{R}}N'') = \partial_{\mathbb{C}}N''$ where $N'' \subset Q(X)$ is \mathbb{R}^+ -invariant and totally real. It follows by Lemma 7.1 that N itself is totally real.

- (ii) Suppose on the contrary that $\partial_{\mathbb{R}}\mathcal{V}_{M,\varepsilon}$ contains an open arc of a Hopf circle, or equivalently that the set $\mathcal{C}_{M,\varepsilon}$ contains an open sector D in a complex line in $Q(X)$. Then D lies in a stratum $Q_i(X) \subset Q(X)$, so after possibly shrinking D we can apply Theorem 5.8 as in the previous paragraph to conclude that D is totally real, a contradiction. □

In order to show that $\mathcal{V}_{M,\varepsilon}$ is discrete, we will derive a contradiction from the conditions (i)–(ii) of Theorem 7.2 and the assumption that $\mathcal{V}_{M,\varepsilon}$ contains an analytic curve. The next few paragraphs develop necessary machinery for analyzing the real and complex boundaries of such curves, after which we return to the proof of Theorem B in sections 7.5–7.6.

7.3. Tangent cones and analytic curves. A general reference for the following material is [Chi2]. If E is a subset of \mathbb{C}^n and $p \in \mathbb{C}^n$, the *tangent cone* of E at p is the set $C(E, p) \subset \mathbb{C}^n$ of points of the form

$$\lim_{k \rightarrow \infty} t_k(z_k - p)$$

where $z_k \in E$, $z_k \rightarrow p$, and $t_k \rightarrow 0^+$ as $k \rightarrow \infty$. The following basic properties of the tangent cone follow immediately from the definition:

- (1) $C(E, p)$ is a closed \mathbb{R}^+ -invariant set.
- (2) $C(E, p) \neq 0$ if and only if p lies in the closure \overline{E} .
- (3) For any $E_1, E_2 \subset \mathbb{C}^n$ we have $C(E_1 \cap E_2, p) \subset C(E_1, p) \cap C(E_2, p)$.
- (4) If E is a submanifold in a neighborhood of $p \in E$, then $C(E, p)$ is the tangent space $T_p E$.

It follows from (4) that if E is an analytic curve (i.e. a complex analytic set of dimension 1) and $p \in E$ is a smooth point then $C(E, p)$ is a complex line. More generally, the tangent cone of an analytic curve at any point is a finite union of complex lines.

Furthermore, just as a complex submanifold is locally a graph over its tangent space, in a neighborhood of a point an analytic curve can be parameterized as follows (see [Chi2, Sec. 1.6.1]):

Lemma 7.3. *Let U be an open neighborhood of $0 \in \mathbb{C}^n$ and let $E \subset U$ be an analytic curve containing 0 and irreducible at that point. Suppose that $C(E, 0) = \{(w_1, 0, \dots, 0) \mid w_1 \in \mathbb{C}\}$. Then there exists a natural number m and a holomorphic map $f : (\Delta, 0) \rightarrow (E, 0)$ such that $f(\zeta) = (f_1(\zeta), \dots, f_n(\zeta))$ where*

$$f_1(\zeta) = \zeta^m$$

and for each $k > 1$ we have

$$f_k(\zeta) = \zeta^{m_k} h_k(\zeta)$$

where $m_k > m$ and either $h_k(\zeta) \equiv 0$ or $h_k(\zeta)$ is holomorphic with $h_k(0) \neq 0$. The image $f(\Delta)$ contains $E \cap U'$ for some neighborhood U' of 0 . \square

Of course one can permute coordinates to obtain a similar parameterization when $C(E, 0)$ is any of the coordinate axes in \mathbb{C}^n . Also note that in this lemma the number m is the *multiplicity* of E at p , and $m > 1$ if and only if p is a singular point.

7.4. Analytic curves near a totally real manifold. Two results characterizing the behavior of an analytic curve near a totally real submanifold of \mathbb{C}^n will be essential in the sequel. The first describes the tangent cone of an analytic curve in the complement of a totally real manifold at a boundary point. A (closed) *complex half-line* is a set of the form $\{L(x + iy) \mid y \geq 0\}$ where $L : \mathbb{C} \rightarrow \mathbb{C}^n$ is an injective complex linear map.

Theorem 7.4 (Chirka [Chi1, Prop. 19]). *Let $E \subset (U \setminus M)$ be an analytic curve where $U \subset \mathbb{C}^n$ is open and $M \subset U$ is a closed, totally real submanifold of class C^k for some $k > 1$. Then for any $p \in (\overline{E} \cap M)$ the tangent cone $C(E, p)$ is a nonempty finite union of complex lines and half-lines.*

With additional regularity for the submanifold M , one has the following extension result:

Theorem 7.5 (Chirka [Chi1, Sec. 1], Alexander [Ale]). *Let $E \subset (U \setminus M)$ be an analytic curve where $U \subset \mathbb{C}^n$ is open and $M \subset U$ is a closed, totally real, real-analytic submanifold. Then for any $p \in (\overline{E} \cap M)$, the set E admits an analytic continuation near p , i.e. there exists a neighborhood U'_p of p and an analytic curve $E'_p \subset U'_p$ such that $(E \cap U'_p) \subset E'_p$.*

Further discussion of this result can be found in [Chi2, Sec. 20.5].

7.5. Hopf circles. If $E \subset \mathbb{C}^n$ is an algebraic curve, then $\partial_{\mathbb{C}} E$ is a finite set and $\partial_{\mathbb{R}} E$ is the union of Hopf circles lying over $\partial_{\mathbb{C}} E$. The next theorem establishes a similar property of $\partial_{\mathbb{R}} E$ when the algebraic assumption is replaced by the condition that $\partial_{\mathbb{C}} E$ locally lie in a totally real manifold.

Theorem 7.6. *Let E be an analytic curve in \mathbb{C}^n and suppose that for some $p \in \partial_{\mathbb{C}} E$ there is a neighborhood U of p in $\partial_{\mathbb{C}} \mathbb{C}^n$ and a totally real, real-analytic submanifold N of U such that $(\partial_{\mathbb{C}} E \cap U) \subset N$. Then $\partial_{\mathbb{R}} E$ contains an open arc of a Hopf circle.*

In the proof we will consider \mathbb{C}^n as an affine chart of its compactification $\mathbb{CP}^n = \mathbb{C}^n \cup \partial_{\mathbb{C}}\mathbb{C}^n$, but other affine charts of \mathbb{CP}^n will also be used. In order to distinguish among them, we use the notation \mathbb{C}_z^n for the original affine chart (in which E is an analytic curve) with coordinates z_1, \dots, z_n , and \mathbb{C}_w^n will denote another affine chart with coordinates w_1, \dots, w_n .

Proof. Let V be an open neighborhood of p in \mathbb{CP}^n such that $V \cap \partial_{\mathbb{C}}\mathbb{C}_z^n = U$. After possibly shrinking V and U we can assume that V lies in an affine chart \mathbb{C}_w^n of \mathbb{CP}^n .

By Theorem 7.5 we can assume, after further shrinking V , that $E \cap V$ is a subset of an analytic curve $E' \subset V$. One of the irreducible components of E' at p must intersect E , so let \widehat{E}' denote such a component. Then $\widehat{E} = \widehat{E}' \cap E$ is an analytic curve in $U \cap \mathbb{C}_z^n$ such that $\partial_{\mathbb{R}}\widehat{E} \subset N$, and p is an accumulation point of \widehat{E} . By Theorem 7.4 the set $C(\widehat{E}, p)$ contains a complex half-line H . Since $\widehat{E} \subset \widehat{E}'$, the complex line L containing H is one of the finite set of lines comprising $C(\widehat{E}', p)$.

There are now two cases to consider, based on the relative position of L and the hyperplane $\partial_{\mathbb{C}}\mathbb{C}_z^n$:

- (1) L is transverse to $\partial_{\mathbb{C}}\mathbb{C}_z^n$ (equivalently, it intersects \mathbb{C}_z^n)
- (2) L is contained in $\partial_{\mathbb{C}}\mathbb{C}_z^n$

Intuitively, these two cases correspond to whether the analytic curve E meets the hyperplane at infinity $\partial_{\mathbb{C}}\mathbb{C}_z^n$ transversely or tangentially at p .

Case 1. By changing the affine chart \mathbb{C}_w^n and making an affine change of coordinates in \mathbb{C}_z^n and \mathbb{C}_w^n we put p , L , and H in a standard position; specifically, we suppose that in a homogeneous coordinate system for \mathbb{CP}^n the inclusions of \mathbb{C}_z^n and \mathbb{C}_w^n are given by

$$\begin{aligned} (z_1, \dots, z_n) &\mapsto [z_1 : \dots : z_n : 1], \\ (w_1, \dots, w_n) &\mapsto [w_1 : \dots : w_{n-1} : 1 : w_n], \end{aligned}$$

that L is the w_n -axis in \mathbb{C}_w^n , and that $H \subset L$ is defined by $\text{Im}(w_n) \geq 0$.

Let $f : \Delta \rightarrow \mathbb{C}_w^n$ be a local parameterization of \widehat{E}' at 0 as in Lemma 7.3; note that here the tangent cone of \widehat{E}' is the w_n -axis. Then we have for any $\zeta \neq 0$ that

$$\begin{aligned} f(\zeta) &= (\zeta^{m_1} h_1(\zeta), \dots, \zeta^{m_{n-1}} h_{n-1}(\zeta), \zeta^m) \in \mathbb{C}_w^n \\ &= (\zeta^{m_1-m} h_1(\zeta), \dots, \zeta^{m_{n-1}-m} h_{n-1}(\zeta), \zeta^{-m}) \in \mathbb{C}_z^n. \end{aligned}$$

Since H is the tangent cone of $\widehat{E} \subset \widehat{E}'$ as a subset of \mathbb{C}_w^n , for each $\theta \in [0, \pi]$ we have a sequence $\zeta_k \rightarrow 0$ such that $\lim_{k \rightarrow \infty} \arg(\zeta_k^m) = \theta$ and $f(\zeta_k) \in \widehat{E}$. As a point in \mathbb{C}_z^n , $f(\zeta_k)$ lies on the same ray as $|\zeta_k|^m f(\zeta_k)$ which has coordinates

$$\left(|\zeta_k|^m \zeta_k^{m_1-m} h_1(\zeta_k), \dots, |\zeta_k|^m \zeta_k^{m_{n-1}-m} h_{n-1}(\zeta_k), |\zeta_k|^m \zeta_k^{-m} \right) \in \mathbb{C}_z^n.$$

Since $m_i > 0$ and $h_i(\zeta_k)$ is bounded as $\zeta_k \rightarrow 0$ for each $1 \leq i \leq m-1$, the sequence $|\zeta_k|^m f(\zeta_k) \in \mathbb{C}_z^n$ converges to $(0, \dots, 0, e^{-i\theta}) \in \partial_{\mathbb{R}}E$. Since

$\theta \in [0, \pi]$ was arbitrary, we find that $\partial_{\mathbb{R}}E$ contains half of the Hopf circle containing $(0, \dots, 0, 1)$, completing this case.

Case 2. We begin as before, altering the argument as necessary.

Choose coordinates so that z_i and w_i are related to one another as above but now we take L to be the w_1 -axis and H the subset with $\text{Im}(w_1) \geq 0$. Parameterizing \widehat{E}' and calculating as above we find that

$$\begin{aligned} f(\zeta) &= (\zeta^m, \zeta^{m_2} h_2(\zeta), \dots, \zeta^{m_n} h_n(\zeta)) \in C_w^n \\ &= (\zeta^{m-m_n} h_n(\zeta)^{-1}, \zeta^{m_2-m_n} h_2(\zeta) h_n(\zeta)^{-1}, \dots \\ &\quad \dots, \zeta^{m_{n-1}-m_n} h_{n-1}(\zeta) h_n(\zeta)^{-1}, \zeta^{-m_n} h_n(\zeta)^{-1}) \in C_z^n. \end{aligned}$$

The coordinate expression in \mathbb{C}_z^n is well-defined since $h_n(\zeta) \neq 0$ for ζ in a small punctured neighborhood of zero: Indeed, while Lemma 7.3 includes the possibility that $h_n(\zeta) \equiv 0$, this would mean that $\widehat{E}' \cap \mathbb{C}_z^n = \emptyset$, contradicting the assumption that \widehat{E}' intersects E . By a further linear change of coordinates we can also assume without loss of generality that $h_n(0) = 1$. As before the condition that H is the tangent cone of \widehat{E} gives for each $\theta \in [0, \pi]$ a sequence $\zeta_k \rightarrow 0$ such that $\lim_{k \rightarrow \infty} \arg(\zeta_k^m) = \theta$ and $f(\zeta_k) \in \widehat{E}$. As a point in \mathbb{C}_z^n , $f(\zeta_k)$ lies on the same ray as $|\zeta_k|^{m_n} f(\zeta_k)$ which has coordinates

$$\begin{aligned} (|\zeta_k|^{m_n} \zeta_k^{m-m_n} h_n(\zeta_k)^{-1}, |\zeta_k|^{m_n} \zeta_k^{m_2-m_n} h_2(\zeta_k) h_n(\zeta_k)^{-1}, \\ \dots, |\zeta_k|^{m_n} \zeta_k^{m_{n-1}-m_n} h_{n-1}(\zeta_k) h_n(\zeta_k)^{-1}, |\zeta_k|^{m_n} \zeta_k^{-m_n} h_n(\zeta_k)^{-1}) \in \mathbb{C}_z^n. \end{aligned}$$

As $k \rightarrow \infty$ we find $|\zeta_k|^{m_n} f(\zeta_k) \rightarrow (0, \dots, 0, e^{-(m_n/m)i\theta'}) \in \partial_{\mathbb{R}}E$ for some $\theta' \equiv \theta \pmod{2\pi}$. Allowing θ to vary over $[0, \pi]$ we find that $\partial_{\mathbb{R}}E$ contains an arc of a Hopf circle. \square

7.6. Discreteness. Using the above results on analytic curves near a totally real manifold, the discreteness of $\mathcal{H}_X \cap \mathcal{E}_M$ now follows easily:

Proof of Theorem B (connected boundary). Suppose on the contrary that the intersection $\mathcal{H}_X \cap \mathcal{E}_M$ is not discrete. Since this set is a finite union of analytic subvarieties $\{\mathcal{H}_{X,\varepsilon} \cap \mathcal{E}_M \mid \varepsilon \in \text{Spin}(X)\}$, at least one of these subvarieties is not discrete. Thus there exists some $\varepsilon \in \text{Spin}(X)$ so that $\mathcal{H}_{X,\varepsilon} \cap \mathcal{E}_M$ contains an analytic curve, as does its preimage $\mathcal{V}_{M,\varepsilon} = \text{hol}^{-1}(\mathcal{E}_M)$.

By the maximum principle $\mathcal{V}_{M,\varepsilon}$ is non-compact and $\partial_{\mathbb{C}}\mathcal{V}_{M,\varepsilon} \neq \emptyset$, so part (i) of Theorem 7.2 implies that $\partial_{\mathbb{C}}\mathcal{V}_{M,\varepsilon}$ is locally contained in a real-analytic, totally real manifold. But then Theorem 7.6 gives an open arc of a Hopf circle contained in $\partial_{\mathbb{R}}\mathcal{V}_{M,\varepsilon}$, contradicting part (ii) of Theorem 7.2. \square

8. SKINNING MAPS: THE CONNECTED BOUNDARY CASE

8.1. Hyperbolic structures. Let M be a compact, irreducible, atoroidal 3-manifold with connected incompressible boundary $S = \partial M$ of genus $g \geq 2$. Then the interior M° admits a complete hyperbolic structure by Thurston's Geometrization Theorem for Haken manifolds. Let $\text{AH}(M) \subset \mathcal{X}(M, \text{PSL}_2\mathbb{C})$ denote the set of isometry classes of marked hyperbolic structures on M° .

The closed set $\text{AH}(M)$ lies in the smooth locus of $\mathcal{X}(M, \text{PSL}_2\mathbb{C})$ (see [Kap, Sec. 8.8]), and its interior $\text{GF}(M)$ consists of the convex cocompact hyperbolic structures. The quasiconformal deformation theory of Kleinian groups gives a holomorphic parameterization of $\text{GF}(M)$ by the Teichmüller space $\mathcal{T}(S)$ (see e.g. [Ber2] [Kra]). We denote this *Ahlfors-Bers parameterization* by

$$\begin{aligned} \mathcal{T}(S) &\longrightarrow \text{GF}(M) \subset \mathcal{X}(M, \text{PSL}_2\mathbb{C}) \\ X &\longmapsto \rho_X^M \end{aligned}$$

8.2. Quasi-Fuchsian groups. The set $\text{QF}(S) \subset \mathcal{X}(S, \text{PSL}_2\mathbb{C})$ of characters of quasi-Fuchsian representations is an open subset of the smooth locus in $\mathcal{X}(S, \text{PSL}_2\mathbb{C})$ that has a natural parameterization by the product of Teichmüller spaces $\mathcal{T}(S) \times \mathcal{T}(\bar{S})$. As a set $\text{QF}(S)$ does not depend on the orientation of S , but the orientation *is* used to distinguish the factors in this parameterization. This coordinate system for $\text{QF}(S)$ is a particular case of the Ahlfors-Bers coordinates, since $\text{QF}(S) = \text{GF}(S \times I) \subset \mathcal{X}(S \times I, \text{PSL}_2\mathbb{C}) = \mathcal{X}(S, \text{PSL}_2\mathbb{C})$. We write $Q(X, Y)$ for the point in $\text{QF}(S)$ corresponding to the pair $(X, Y) \in \mathcal{T}(S) \times \mathcal{T}(\bar{S})$.

Given $Y \in \mathcal{T}(\bar{S})$, the *Bers slice* is the subset

$$B_Y = \{Q(X, Y) \mid X \in \mathcal{T}(S)\} \subset \text{QF}(S),$$

that is, B_Y is a “horizontal slice” of the product structure of the quasi-Fuchsian space. Similarly we can define the vertical Bers slices $B_X = \{Q(X, \cdot)\}$ for $X \in \mathcal{T}(S)$.

Note B_Y is naturally in one-to-one correspondence with $\mathcal{T}(S)$. An inequality of Bers shows that the lengths of geodesics in the hyperbolic 3-manifold corresponding to a quasi-Fuchsian group $Q(X, Y)$ have an upper bound in terms of lengths in the uniformization of either X or Y [Ber1, Thm. 3]). Since one of these is fixed in a Bers slice, there are uniform bounds on the traces of any finite set of elements in $\pi_1 S$ over B_Y , and thus:

Lemma 8.1 (Bers). *For each $Y \in \mathcal{T}(\bar{S})$, the Bers slice B_Y is precompact.* \square

Each point in the Bers slice B_Y induces a \mathbb{CP}^1 -structure on Y by taking the quotient of one of the domains of discontinuity of the associated quasi-Fuchsian group. This construction is a local inverse of the holonomy map, i.e. it gives an open subset of $Q(Y)$ (the *Bers embedding*) that maps biholomorphically onto B_Y by hol. In particular we have $B_Y \subset r(\mathcal{H}_Y)$, where $r : \mathcal{X}(S, \text{SL}_2\mathbb{C}) \rightarrow \mathcal{X}(S, \text{PSL}_2\mathbb{C})$ is the map induced by the covering $\text{SL}_2\mathbb{C} \rightarrow \text{PSL}_2\mathbb{C}$. This applies equally to the vertical slices, i.e. $B_X \subset r(\mathcal{H}_X)$ for $X \in \mathcal{T}(S)$.

8.3. The skinning map. The restriction map $i^* : \mathcal{X}(M, \text{PSL}_2\mathbb{C}) \rightarrow \mathcal{X}(S, \text{PSL}_2\mathbb{C})$ sends $\text{GF}(M)$ into $\text{QF}(S)$, and in terms of the Ahlfors-Bers parameterization

it is the identity on one Teichmüller space factor, i.e.

$$i^*(\rho_X^M) = Q(X, \sigma_M(X))$$

which defines a map

$$\sigma_M : \mathcal{T}(S) \rightarrow \mathcal{T}(\bar{S}),$$

the *skinning map of M*.

Fibers of the skinning map are related to sets $r(\mathcal{H}_Y \cap \mathcal{E}_M)$ as follows:

Lemma 8.2. *The preimage $\sigma_M^{-1}(Y)$ is in bijection with a precompact set $F_Y \subset r(\mathcal{H}_Y \cap \mathcal{E}_M)$.*

Proof. Given $Y \in \mathcal{T}(\bar{S})$ we consider the set of quasi-Fuchsian groups $F_Y = i^*(\text{GF}(M)) \cap B_Y$, i.e.

$$F_Y = \{Q(X, Y) \mid \text{There exists } X \in \mathcal{T}(S) \text{ such that } i^*(\rho_X^M) = Q(X, Y)\}.$$

From the definition of the skinning map it is immediate that F_Y is in bijection with the preimage $\sigma_M^{-1}(Y)$ by

$$Q(X, Y) \in F_Y \iff X \in \sigma_M^{-1}(Y),$$

Furthermore $F_Y \subset B_Y \subset r(\mathcal{H}_Y)$, and precompactness of F_Y follows from that of B_Y , so it remains only to show that $F_Y \subset r(\mathcal{E}_M)$.

By definition \mathcal{E}_M contains as a Zariski dense subset $i^*(\mathcal{X}(M, \text{SL}_2\mathbb{C}))$, using the commutative diagram (6.1) for r, i^* we have

$$i^*(r(\mathcal{X}(M, \text{SL}_2\mathbb{C}))) = r(i^*(\mathcal{X}(M, \text{SL}_2\mathbb{C}))) \subset r(\mathcal{E}_M).$$

Since $F_Y \subset i^*(\text{GF}(M))$, it is enough to know that $\text{GF}(M) \subset r(\mathcal{X}(M, \text{SL}_2\mathbb{C}))$, i.e. that the $\text{PSL}_2\mathbb{C}$ -representations arising from hyperbolic structures on M can be lifted to $\text{SL}_2\mathbb{C}$. This is a well-known consequence of the parallelizability of 3-manifolds (see [Cul] or [CS, Thm 3.1.1] for details). \square

Finally, using this lemma we have the

Proof of Theorem A (connected boundary). Suppose on the contrary that $\sigma_M^{-1}(Y)$ is infinite. Then by the previous lemma the set $r(\mathcal{H}_Y \cap \mathcal{E}_M)$ has an infinite and precompact subset, and therefore an accumulation point. Since $\mathcal{H}_Y \cap \mathcal{E}_M$ is discrete by Theorem B and the map $r : \mathcal{X}(S, \text{SL}_2\mathbb{C}) \rightarrow \mathcal{X}(S, \text{PSL}_2\mathbb{C})$ is proper, this is a contradiction. \square

9. DISCONNECTED BOUNDARY AND TORI

The previous sections established the main theorems for a 3-manifold with connected boundary. We now adapt the statements and proofs to the more general case of a compact oriented 3-manifold M whose boundary has at least one connected component that is not a torus.

As in the introduction, we denote by $\partial_0 M$ the union of the non-torus boundary components of M . Let S_1, \dots, S_m denote the connected components of $\partial_0 M$, each equipped with the boundary orientation. Let $\partial_1 M = \partial M \setminus \partial_0 M$ denote the union of the torus boundary components of M .

9.1. Measured foliations and Teichmüller spaces. In several cases we first need to adapt definitions of spaces associated to a surface to the disconnected case. We define the measured foliation space $\mathcal{MF}(\partial_0 M)$ and Teichmüller space $\mathcal{T}(\partial_0 M)$ to be the cartesian products of the spaces corresponding to the connected components S_i , e.g.

$$\mathcal{MF}(\partial_0 M) = \prod_{i=1}^m \mathcal{MF}(S_i).$$

The sum of the symplectic forms of the factors (using the boundary orientation) gives the Thurston symplectic form on $\mathcal{MF}(\partial_0 M)$.

For a point $X = (X_1, \dots, X_m) \in \mathcal{T}(\partial_0 M)$, we denote by $Q(X)$ the direct sum of quadratic differential spaces,

$$Q(X) = \bigoplus_{i=1}^m Q(X_i).$$

The product of foliation maps of the factors gives the homeomorphism $\mathcal{F} : Q(X) \rightarrow \mathcal{MF}(\partial_0 M)$.

For the 3-manifold M and its fundamental group, in some cases we must treat torus boundary components differently. Instead of considering arbitrary isometric actions of $\pi_1 M$ on \mathbb{R} -trees, we restrict attention to actions in which each subgroup of $\pi_1 M$ represented by a component of $\partial_1 M$ —that is, each *boundary torus subgroup*—has a fixed point.

9.2. Isotropic cones. The isotropic cone construction (Theorem 4.4, the main result of sections 3–4) generalizes to the disconnected case as follows:

Theorem 9.1. *For each $X = (X_1, \dots, X_m) \in \mathcal{T}(\partial_0 M)$ there exists an isotropic piecewise linear cone $\mathcal{L}_{M,X} \subset \mathcal{MF}(\partial_0 M)$ with the following property:*

Let T be a Λ -tree on which $\pi_1 M$ acts so that each boundary torus subgroup has a fixed point. Let $\mathfrak{o} : \mathbb{R} \rightarrow \Lambda$ be an order-preserving embedding. For each i with $1 \leq i \leq m$ suppose that we have:

- *A pair T_{ℓ_i}, T'_{ℓ_i} of \mathbb{R} -trees on which $\pi_1 S_i$ acts minimally with length function ℓ_i ,*
- *A holomorphic quadratic differential $\phi_i \in Q(X_i)$,*
- *A $\pi_1 S_i$ -equivariant straight map $T_{\phi_i} \rightarrow T_{\ell_i}$ with respect to \mathfrak{o} , and*
- *A $\pi_1 S_i$ -equivariant isometric embedding $k : T'_{\ell_i} \rightarrow T$.*

Then $[\mathcal{F}(\phi)] \in \mathcal{L}_{M,X}$, where $\phi = (\phi_1, \dots, \phi_m)$.

Proof. First we adapt the definition of the cone $\mathcal{L}_{M,X}$ from the connected case. We choose a finite set of triangulations of M that extend the triangulations of S_i given by Lemma 4.3. For each such triangulation Δ_M we have a space $W_4(\Delta_M, \Lambda)$ of weights satisfying the 4-point condition in each 3-simplex, but we now also require these weights to be identically zero on the edges of each boundary torus.

As in Lemmas 3.5 and 3.7 we find that the restriction of $W_4(\Delta_M, \mathbb{R})$ to the edges of the boundary triangulation gives a finite union of isotropic subspaces for a symplectic form that is the sum of the Thurston forms for the triangulated surfaces S_i and a similar alternating 2-form for the weight space of each boundary torus. Since the weights in $W_4(\Delta_M)$ are identically zero in the torus components, these subspaces are still isotropic when projected to the product of non-torus factors. Thus $W_4(\Delta_M)$ gives an isotropic cone in a product of train-track charts $\mathcal{MF}(\tau_i)$ for the surfaces S_i , and taking the union of these over the finite set of triangulations of M gives the cone $\mathcal{L}_{M,X} \subset \mathcal{MF}(\partial_0 M)$.

Now we show that $[\mathcal{F}(\phi)] \in \mathcal{L}_{M,X}$, or equivalently that the train track coordinates of $[\mathcal{F}(\phi)]$ are obtained by restricting an element of $W_4(\Delta_M)$ to the boundary. Here Δ_M is the triangulation from our finite set in which the edges on S_i can be realized ϕ_i -geodesically.

Using the straight maps $T_{\phi_i} \rightarrow T_{\ell_i}$ and either the isometry $T_{\ell_i} \rightarrow T'_{\ell_i}$ of Theorem 4.5 (in the case of a non-abelian length function) or the partially-defined map $T_{\ell_i} \dashrightarrow T'_{\ell_i}$ of Theorem 4.7 (in the abelian case), we obtain a map from the non-torus boundary vertices, $\tilde{\Delta}_M^{(0)} \cap \partial_0 M$, to the tree T . We extend this over the vertices on the torus boundary components by mapping all vertices in a given boundary torus to a fixed point of the associated subgroup of $\pi_1 M$.

Extending over the remaining (interior) vertices of Δ_M as in Propositions 3.2 and 4.2, we obtain a map $\tilde{\Delta}_M^{(0)} \rightarrow T$ whose associated weight function w lies in $W_4(\Delta_M, \Lambda)$; note that since all vertices of a boundary torus are mapped to a single point of T , the associated weight vanishes on edges of the torus boundary components as required. We push forward by a left inverse of \mathfrak{o} to obtain an element of $W_4(\Delta_M)$ whose values on the non-torus boundary edges give the train track coordinates of $[\mathcal{F}(\phi)]$. Thus $[\mathcal{F}(\phi)] \in \mathcal{L}_{M,X}$. \square

9.3. Kähler structure and symplectomorphism. The results of Section 5 generalize easily to disconnected surfaces by taking products of the spaces, maps, and stratifications considered there.

Specifically, for any $X \in \mathcal{J}(\partial_0 M)$, Lemma 5.1 provides a stratification of $Q(X_i)$. There is an induced *product stratification* of the product space $Q(X) = \bigoplus_i Q(X_i)$ consisting of products of strata in the factors. Note that the origin $\{0\} \in Q(X_i)$ is the minimal stratum in each factor, so $Q(X)$ now has nontrivial (positive-dimensional) strata consisting of quadratic differentials that are zero on one or more of the boundary components.

Similarly we take the product of the stratified Kähler structures on the factors $Q(X_i)$ to obtain a Kähler structure on $Q(X)$, smooth relative to the product stratification. Applying Theorem 5.8 to each factor of the map $\mathcal{F} : Q(X) \rightarrow \mathcal{MF}(\partial_0 M)$ we obtain:

Theorem 9.2. *For any $X \in \mathcal{T}(\partial_0 M)$, the map $\mathcal{F} : Q(X) \rightarrow \mathcal{MF}(\partial_0 M)$ is a real-analytic stratified symplectomorphism. That is, if $Q_k(X)$ is a stratum of $Q(X)$ then:*

- (i) *For any $\phi = (\phi_1, \dots, \phi_m) \in Q_k(X)$ there exists an open neighborhood $U \subset Q_k(X)$ of ϕ and a product of train track coordinate charts $\prod_i \mathcal{MF}(\tau_i) \subset \mathcal{MF}(\partial_0 M)$ covering $\mathcal{F}(U)$ so that the restriction*

$$\mathcal{F} : U \rightarrow \prod_i \mathcal{MF}(\tau_i)$$

is a real-analytic diffeomorphism onto its image, and

- (ii) *The derivative $d\mathcal{F}_\phi$ defines a symplectic linear map from $T_\phi Q_k(X)$ into $\bigoplus_i W(\tau_i)$, where $T_\phi Q_k(X)$ is equipped with the symplectic form $\sum_i \omega_{\phi_i}$ and $\bigoplus_i W(\tau_i)$ is given the Thurston symplectic form.*

□

To make sense of this statement in case the differentials in the stratum $Q_k(X)$ are identically zero in some factor, say $Q(X_i)$, we adopt the convention that τ_i is the empty train track and that $\mathcal{MF}(\tau_i) = W(\tau_i) = \{0\}$ is a point representing the representing the empty foliation on X_i , which is the image of 0 under the map $\mathcal{F} : Q(X_i) \rightarrow \mathcal{MF}(S_i)$.

9.4. Character varieties and extension varieties. We generalize the character variety of a connected surface to $\partial_0 M$ by taking the product of character varieties of components

$$\mathcal{X}(\partial_0 M, G) := \prod_i \mathcal{X}(S_i, G),$$

and similarly for the representation variety $\mathcal{R}(\partial_0 M, G)$. Note that while $\mathcal{R}(\partial_0 M, G)$ can also be described as the representation variety of the free product $\pi_1 S_1 * \dots * \pi_1 S_m$, the character variety of this free product does *not* agree with our definition of $\mathcal{X}(\partial_0 M, G)$. To obtain $\mathcal{X}(\partial_0 M, G)$ from $\mathcal{R}(\partial_0 M, G)$ one must take the quotient of by the action of G^m .

The 3-manifold character variety also requires modification to account for the presence of boundary tori. As is standard when considering complete hyperbolic structures, rather than working with the full character variety of $\pi_1 M$, we consider the subvariety

$$\mathcal{X}(M, \partial_1 M, G) \subset \mathcal{X}(M, G)$$

consisting of characters of representations that map each boundary torus subgroup (that is, the fundamental group of each connected component of $\partial_1 M$) to parabolic elements of G .

The inclusion of each boundary component $S_i \hookrightarrow M$ induces a restriction map $\mathcal{X}(M, \partial_1 M, \mathrm{SL}_2\mathbb{C}) \rightarrow \mathcal{X}(S_i, \mathrm{SL}_2\mathbb{C})$, and taking the product of these we obtain a regular map

$$i^* : \mathcal{X}(M, \partial_1 M, \mathrm{SL}_2\mathbb{C}) \rightarrow \mathcal{X}(\partial_0 M, \mathrm{SL}_2\mathbb{C}),$$

We define the extension variety $\mathcal{E}_M \subset \mathcal{X}(\partial_0 M, \mathrm{SL}_2\mathbb{C})$ as the Zariski closure of the image $i^*(\mathcal{X}(M, \partial_1 M, \mathrm{SL}_2\mathbb{C}))$.

The Morgan-Shalen compactification of $\mathcal{X}(\Gamma, \mathrm{SL}_2\mathbb{C})$ was defined in Section 6.2 using trace functions of elements of Γ . On the product $\mathcal{X}(\partial_0 M, \mathrm{SL}_2\mathbb{C})$ we have trace functions for the elements of each component $\pi_1 S_i$, so the family of all such functions is indexed by the disjoint union

$$H := \pi_1 S_1 \sqcup \cdots \sqcup \pi_1 S_m.$$

To adapt the Morgan-Shalen compactification to this case we map the character variety of $\partial_0 M$ to the projective space $\mathbb{P}(\mathbb{R}^H)$ using formula (6.2) and take its closure. Indeed, a compactification in this generality was already discussed in [MoSh1], where the map to projective space arising from an arbitrary collection of regular functions on an algebraic variety is considered.

A boundary point of the resulting compactification of $\mathcal{X}(\partial_0 M, \mathrm{SL}_2\mathbb{C})$ is therefore an \mathbb{R}^+ -equivalence class $[\ell]$ where $\ell = (\ell_1, \dots, \ell_m)$ is a tuple of functions, $\ell_i : \pi_1 S_i \rightarrow \mathbb{R}$. The factor ℓ_i is either a nontrivial length function of an action of $\pi_1 S_i$ on an \mathbb{R} -tree or is identically zero (which is the length function of the action of $\pi_1 S_i$ on a point).

Having adapted the definitions of its objects suitably, the Theorem 6.3 on extensions of length functions arising from the boundary of \mathcal{E}_M generalizes to:

Theorem 9.3. *Let $[\ell]$ be a boundary point of \mathcal{E}_M in the Morgan-Shalen compactification of $\mathcal{X}(\partial_0 M, \mathrm{SL}_2\mathbb{C})$, where $\ell = (\ell_1, \dots, \ell_m)$. Then there exists a function $\widehat{\ell} : \pi_1 M \rightarrow \mathbb{R}^n$ such that*

- (i) *The group $\pi_1 M$ acts isometrically on a \mathbb{R}^n -tree with length function $\widehat{\ell}$ such that each boundary torus subgroup of $\pi_1 M$ has a fixed point, and*
- (ii) *The function $\widehat{\ell}$ is a simultaneous extension of the functions ℓ_i , i.e. for each $\gamma \in \pi_1 S_i$ we have $\widehat{\ell}(i_*(\gamma)) = i_n(\ell(\gamma))$ where $i_* : \pi_1 S_i \rightarrow \pi_1 M$ is induced by the inclusion of S_i as a boundary component of M and $i_n(x) = (0, \dots, 0, x)$.*

Proof. As in the proof of Theorem 6.3 we have an extension of fields $k(\mathcal{E}_M^0) \rightarrow k(\mathcal{X}_M^0)$ where now $\mathcal{E}_M^0 \subset \mathcal{X}(\partial_0 M, \mathrm{SL}_2\mathbb{C})$ is an irreducible component of \mathcal{E}_M which has $[\ell]$ in its boundary and $\mathcal{X}_M^0 \subset \mathcal{X}(M, \partial_1 M, \mathrm{SL}_2\mathbb{C})$.

The boundary point $[\ell]$ determines a valuation $v : k(\mathcal{E}_M^0) \rightarrow \Lambda$ by the analogue of Theorem 6.1 for subvarieties of the product of character varieties $\mathcal{X}(\partial_0 M, \mathrm{SL}_2\mathbb{C})$; like Theorem 6.1, this case is covered by the more general comparison of valuation- and projectivization-based compactifications of [MoSh1, Thm. I.3.6].

Extending this valuation to the superfield $k(\widehat{\mathcal{X}}_M^0)$ and proceeding as in the proof of Theorem 6.3 then gives a function $\widehat{\ell} : \pi_1 M \rightarrow \mathbb{R}^n$ satisfying condition (ii) above; note that the argument of (6.4) applies to each function $\ell_i : \pi_1 S_i \rightarrow \mathbb{R}$, $1 \leq i \leq m$, giving that $\widehat{\ell}$ is a simultaneous extension.

The resulting length function $\widehat{\ell}$ arises from an action of $\pi_1 M$ on a \mathbb{R}^n -tree T , but we must show that each boundary torus subgroup has a fixed point. Since \mathcal{X}_M^0 is an irreducible subvariety of $\mathcal{X}(M, \partial_1 M, \mathrm{SL}_2\mathbb{C})$, the trace function t_γ of an element γ of a boundary torus subgroup restricts to a constant ± 2 on \mathcal{X}_M^0 (the possible traces of parabolics). Therefore we have $v'(t_\gamma) = 0$, where v' is the valuation of $k(\mathcal{X}_M^0)$ from which T is constructed. By Lemma 2.7 each boundary torus subgroup leaves invariant a subtree of T on which it acts with zero length function, and by [MoSh1, Prop. II.2.15] there is a fixed point. \square

9.5. Holonomy, the isotropic cone, and discreteness. The construction of the holonomy variety in Section 6.5 extends to $\partial_0 M$ by taking products; that is, for $X \in \mathcal{T}(\partial_0 M)$ or $X \in \mathcal{T}(\overline{\partial_0 M})$, where $X = (X_1, \dots, X_m)$, we define

$$\mathcal{H}_X := \bigcup_{\varepsilon \in \mathrm{Spin}(X)} \mathcal{H}_{X,\varepsilon} \quad \text{where} \quad \mathcal{H}_{X,\varepsilon} := \prod_{i=1}^m \mathcal{H}_{X_i,\varepsilon_i}.$$

Here $\varepsilon = (\varepsilon_1, \dots, \varepsilon_m)$ is a tuple of spin structures on the components.

The subvariety of the quadratic differential space corresponding to the intersection $\mathcal{H}_{X,\varepsilon} \cap \mathcal{E}_M$ is

$$\mathcal{V}_{M,\varepsilon} := \mathrm{hol}_\varepsilon^{-1}(\mathcal{E}_M) \subset Q(\partial_0 M).$$

Generalizing Theorem 6.5 we have:

Theorem 9.4. *Let $\phi = (\phi_1, \dots, \phi_m)$ be the projective limit of a divergent sequence in $\mathcal{V}_{M,\varepsilon} \subset Q(\partial_0 M)$. Then $[\mathcal{F}(\phi)] \in \mathcal{L}_{M,X}$ where $\mathcal{L}_{M,X}$ is the isotropic cone of Theorem 9.1.*

Proof. Theorems 9.1 and 9.3 provide the necessary generalizations to adapt the proof of Theorem 6.5 to this situation, except for the existence of the straight maps $T_{\phi_i} \rightarrow T_{\ell_i}$.

In the connected boundary case, a straight map $T_\phi \rightarrow T_\ell$ is given by Theorem 6.4 (i.e. [D, Thm. A]). Both the projective limit ϕ and the representative ℓ of the Morgan-Shalen boundary point are only well-defined up to multiplication by a positive constant in this case, but it suffices to have such a straight map for some pair of representatives ℓ and ϕ .

In the disconnected case, both tuples $\phi = (\phi_1, \dots, \phi_m)$ and $\ell = (\ell_1, \dots, \ell_m)$ represent equivalence classes up to a single multiplicative constant, whereas a factor-wise application of Theorem 6.4 would seem to require a separate multiplicative factor for each connected component of the boundary.

To remedy this we use choose representatives ℓ, ϕ as in [D], where it is shown that for a connected surface S , $X \in \mathcal{T}(S)$, and a divergent sequence $\{\phi_n\} \subset Q(X)$, we can extract a representative ℓ of the Morgan-Shalen limit of $\mathrm{hol}(\phi_n)$ by scaling the functions $\gamma \mapsto \log(|t_\gamma(\mathrm{hol}(\phi_n))| + 2)$ by the factors $\|\phi_n\|^{-1/2}$ and taking the limit $\ell \in \mathbb{R}^{\pi_1 S}$. Furthermore, the straight map

$T_\phi \rightarrow T_\ell$ of Theorem 6.4 is defined for this function ℓ and for the projective limit $\phi \in Q(X)$ satisfying $\|\phi\| = C$, for a universal constant C .

Correspondingly, for the disconnected case we consider $X \in \mathcal{T}(\partial_0 M)$ and a sequence of tuples $(\phi_{1,n}, \dots, \phi_{m,n})$ which converges projectively. Equivalently, the sequence

$$c_n(\phi_{1,n}, \dots, \phi_{m,n})$$

converges in $Q(X)$ as $n \rightarrow \infty$ where

$$c_n = \left(\sum_{i=1}^m \|\phi_{i,n}\| \right)^{-1}.$$

Applying Theorem 6.4 to each factor and using the representative length functions and projective limits discussed above, after passing to a subsequence we obtain straight maps

$$(9.1) \quad T_{\phi_i^{(1)}} \rightarrow T_{\ell_i^{(1)}}$$

where $\phi_i^{(1)} = C \lim_{n \rightarrow \infty} \frac{\phi_{i,n}}{\|\phi_{i,n}\|}$ and $\ell_i^{(1)} : \pi_1 S_i \rightarrow \mathbb{R}$ is the limit as $n \rightarrow \infty$ of the functions

$$\gamma \mapsto \frac{1}{\|\phi_{i,n}\|^{\frac{1}{2}}} \log(|t_\gamma(\text{hol}(\phi_{i,n}))| + 2).$$

Since $0 \leq \|\phi_{i,n}\| \leq c_n^{-1}$ we can take a further subsequence so that for each i the limit

$$r_i = \lim_{n \rightarrow \infty} c_n \|\phi_{i,n}\| \in [0, 1],$$

exists. Then

$$\begin{aligned} \phi &:= (r_1 \phi_1^{(1)}, \dots, r_n \phi_n^{(1)}) \\ &= C \lim_{n \rightarrow \infty} c_n(\phi_{1,n}, \dots, \phi_{m,n}) \in Q(X) \end{aligned}$$

is a projective limit of the sequence of quadratic differentials and

$$\begin{aligned} \ell &:= (r_1^{\frac{1}{2}} \ell_1^{(1)}, \dots, r_n^{\frac{1}{2}} \ell_n^{(1)}) \\ &= \lim_{n \rightarrow \infty} \left(\gamma \mapsto c_n^{1/2} \log(|t_\gamma(\text{hol}(\phi_{i,n}))| + 2) \right)_{\gamma \in \pi_1 S_i, 1 \leq i \leq n} \in \mathbb{R}^H \end{aligned}$$

represents the limit of the holonomy representations in the Morgan-Shalen compactification of $\mathcal{X}(\partial_0 M, \text{SL}_2 \mathbb{C})$. The desired straight maps

$$T_{\phi_i} \rightarrow T_{\ell_i}$$

are therefore given by (9.1) for each i with $r_i \neq 0$, since the trees T_{ϕ_i}, T_{ℓ_i} are obtained from $T_{\phi_i^{(n)}}, T_{\ell_i^{(n)}}$ by multiplying the metrics by $r_i^{1/2}$. In each factor where $r_i = 0$ we have $\phi_i = 0$ and $\ell_i = 0$, so each of T_{ϕ_i} and T_{ℓ_i} is a point, and there is nothing to prove. \square

Using Theorems 9.4 and 9.2 in place of their counterparts for the connected boundary case (Theorems 6.5 and 5.8, respectively), the properties of the real and complex boundaries of $\mathcal{V}_{M,\varepsilon}$ from Theorem 7.2 follow for the general case as well. The argument of Section 7.6 then gives the discreteness of these varieties, completing the proof of Theorem B in the general case.

9.6. Skinning maps. Now suppose that in addition to the hypotheses of Theorem B that the 3-manifold M has incompressible boundary and that its interior admits a complete hyperbolic structure with no accidental parabolics. The space $\text{GF}(M)$ of geometrically finite hyperbolic structures on M is naturally a subset of $\mathcal{X}(M, \partial_1 M, \text{PSL}_2\mathbb{C})$ which is parameterized by the Teichmüller space $\mathcal{T}(\partial_0 M)$; as before we denote this parameterization by $X \mapsto \rho_M^X$ where $X = (X_1, \dots, X_m)$.

The restriction map $i^* : \mathcal{X}(M, \partial_1 M, \text{PSL}_2\mathbb{C}) \rightarrow \mathcal{X}(\partial_0 M, \text{PSL}_2\mathbb{C})$ sends ρ_M^X to a tuple of quasi-Fuchsian groups

$$i^*(\rho_M^X) = (Q(X_1, Y_1), \dots, Q(X_m, Y_m)) \in \prod_{i=1}^m \text{QF}(S_i)$$

and this defines the *skinning map*

$$\begin{aligned} \sigma_M : \mathcal{T}(\partial_0 M) &\rightarrow \mathcal{T}(\overline{\partial_0 M}) \\ (X_1, \dots, X_m) &\mapsto (Y_1, \dots, Y_m). \end{aligned}$$

With this definition in hand, the generalization of the proof of the main theorem from the connected case is straightforward:

Proof of Theorem A (general case). Suppose on the contrary that $\sigma_M^{-1}(Y)$ is infinite. As in Lemma 8.2, it follows from the definition of σ_M that the preimage $\sigma_M^{-1}(Y)$ is in bijection with a subset of the $B_Y \cap r(\mathcal{E}_M)$, where $B_Y = \prod_i B_{Y_i}$ is the product of Bers slices. Here $r : \mathcal{X}(\partial_0 M, \text{SL}_2\mathbb{C}) \rightarrow \mathcal{X}(\partial_0 M, \text{PSL}_2\mathbb{C})$ is the finite-sheeted, proper map induced by $\text{SL}_2\mathbb{C} \rightarrow \text{PSL}_2\mathbb{C}$. Using Theorem B we find that $r(\mathcal{H}_Y \cap \mathcal{E}_M)$ is a discrete set. But since B_Y is a precompact subset of $r(\mathcal{H}_Y)$, the infinite subset we have identified with $\sigma_M^{-1}(Y)$ has an accumulation point, which is a contradiction. \square

10. EXTENDED SKINNING MAPS OF ACYLINDRICAL MANIFOLDS

As mentioned in the introduction, the proof of Theorem A can be adapted to show finiteness of fibers of Thurston's continuous extension of the skinning map $\widehat{\sigma}_M : \text{AH}(M) \rightarrow \mathcal{T}(\overline{\partial_0 M})$ for an acylindrical manifold M . A proof of the existence of this continuous extension by Brock, Kent, and Minsky can be found [Ken, Sec. 9].

The only property of the extension we will use is the following, which can be taken as the definition of $\widehat{\sigma}_M$: If the image of $\rho \in \text{AH}(M)$ under the extended skinning map is

$$\widehat{\sigma}_M(\rho) = (Y_1, \dots, Y_m) \in \mathcal{T}(\overline{\partial_0 M}),$$

and if the image of ρ under the restriction map is

$$i^*(\rho) = (\eta_1, \dots, \eta_m) \in \mathcal{X}(\partial_0 M, \mathrm{PSL}_2\mathbb{C})$$

then η_i is the character of a discrete, faithful $\pi_1 S_i$ -representation which has an invariant disk in its domain of discontinuity for which the quotient Riemann surface is Y_i .

The analogue of Theorem A for $\widehat{\sigma}_M$ is:

Theorem 10.1. *Let M be a compact, oriented, irreducible, acylindrical 3-manifold with incompressible boundary that is not empty and not a union of tori. Let $\widehat{\sigma}_M : \mathrm{AH}(M) \rightarrow \mathcal{T}(\overline{\partial_0 M})$ denote the extension of the skinning map of M . Then for each $Y \in \mathcal{T}(\overline{\partial_0 M})$, the set $\widehat{\sigma}_M^{-1}(Y)$ is finite.*

Proof. First we show that $i^* : \mathcal{X}(M, \partial_1 M, \mathrm{PSL}_2\mathbb{C}) \rightarrow \mathcal{X}(\partial_0 M, \mathrm{PSL}_2\mathbb{C})$ is injective when restricted to $\mathrm{AH}(M)$. Suppose $i^*(\rho) = i^*(\rho')$ with $\rho, \rho' \in \mathrm{AH}(M)$. We claim the hyperbolic structures on M associated to ρ and ρ' are bilipschitz. Since $i^*(\rho) = i^*(\rho')$, these hyperbolic structures are actually *isometric* in a neighborhood of each non-torus end. Smooth bilipschitz maps of the torus ends can be obtained by choosing affine maps between the cusp tori and extending normally. We have therefore defined a smooth bilipschitz equivalence on all of the ends, and any diffeomorphic extension to the remaining compact part is globally bilipschitz. By Sullivan's rigidity theorem [Sul], we conclude $\rho = \rho'$.

Now suppose for contradiction that $\rho^{(n)} \in \mathrm{AH}(M)$ is an infinite sequence of pairwise distinct points satisfying $\widehat{\sigma}_M(\rho^{(n)}) = Y = (Y_1, \dots, Y_m)$ for all n . As in the proof of Theorem A, it is enough to show that this leads to an accumulation point in $r(\mathcal{H}_Y \cap \mathcal{E}_M)$.

Consider the sequence

$$i^*(\rho^{(n)}) = (\eta_1^{(n)}, \dots, \eta_m^{(n)}) \in \prod_i \mathcal{X}(S_i, \mathrm{PSL}_2\mathbb{C}),$$

which is also pairwise distinct since i^* is injective on $\mathrm{AH}(M)$. By the definition of $\widehat{\sigma}_M$, the character $\eta_i^{(n)}$ corresponds to a discrete $\pi_1 S_i$ -representation with an invariant disk in its domain of discontinuity having quotient Riemann surface Y_i . Since this quotient describes a projective structure on Y_i with holonomy $\eta_i^{(n)}$, we find that $i^*(\rho^{(n)}) \in r(\mathcal{H}_Y)$. Since $r(\mathcal{E}_M)$ contains the range of i^* , we in fact have $i^*(\rho^{(n)}) \in r(\mathcal{H}_Y \cap \mathcal{E}_M)$.

While the representations $\eta_i^{(n)}$ are not necessarily quasi-Fuchsian, a generalization of Bers' inequality (e.g. [Ohs, Prop. 2.1]) shows that the fixed quotient Riemann surface Y_i constrains $\eta_i^{(n)}$ to lie in a compact subset of $\mathcal{X}(S_i, \mathrm{PSL}_2\mathbb{C})$. Thus $i^*(\rho^{(n)})$ lies in the product of these compact sets for all n , giving an accumulation point of $r(\mathcal{H}_Y \cap \mathcal{E}_M)$ and the desired contradiction. \square

REFERENCES

- [Ale] H. Alexander. Continuing 1-dimensional analytic sets. *Math. Ann.* **191**(1971), 143–144.
- [AB] R. Alperin and H. Bass. Length functions of group actions on Λ -trees. In *Combinatorial group theory and topology (Alta, Utah, 1984)*, volume 111 of *Ann. of Math. Stud.*, pages 265–378. Princeton Univ. Press, Princeton, NJ, 1987.
- [ABEM] J. Athreya, A. Bufetov, A. Eskin, and M. Mirzakhani. Lattice point asymptotics and volume growth on Teichmüller space. *Duke Math. J.* **161**(2012), 1055–1111.
- [Ber1] L. Bers. On boundaries of Teichmüller spaces and on Kleinian groups. I. *Ann. of Math. (2)* **91**(1970), 570–600.
- [Ber2] L. Bers. Spaces of Kleinian groups. In *Several Complex Variables, I (Proc. Conf., Univ. of Maryland, College Park, Md., 1970)*, pages 9–34. Springer, Berlin, 1970.
- [Bon] F. Bonahon. Shearing hyperbolic surfaces, bending pleated surfaces and Thurston’s symplectic form. *Ann. Fac. Sci. Toulouse Math. (6)* **5**(1996), 233–297.
- [Bow] B. Bowditch. Group actions on trees and dendrons. *Topology* **37**(1998), 1275–1298.
- [BH] M. Bridson and A. Haefliger. *Metric spaces of non-positive curvature*, volume 319 of *Grundlehren der Mathematischen Wissenschaften*. Springer-Verlag, Berlin, 1999.
- [Br] K. Brown. Trees, valuations, and the Bieri-Neumann-Strebel invariant. *Invent. Math.* **90**(1987), 479–504.
- [Chi1] E. Chirka. Regularity of the boundaries of analytic sets. *Mat. Sb. (N.S.)* **117(159)**(1982), 291–336, 431.
- [Chi2] E. Chirka. *Complex analytic sets*, volume 46 of *Mathematics and its Applications (Soviet Series)*. Kluwer Academic Publishers Group, Dordrecht, 1989. Translated from the Russian by R. A. M. Hoksbergen.
- [Chi3] I. Chiswell. *Introduction to Λ -trees*. World Scientific Publishing Co. Inc., River Edge, NJ, 2001.
- [Cul] M. Culler. Lifting representations to covering groups. *Adv. in Math.* **59**(1986), 64–70.
- [CM] M. Culler and J. Morgan. Group actions on \mathbf{R} -trees. *Proc. London Math. Soc. (3)* **55**(1987), 571–604.
- [CS] M. Culler and P. Shalen. Varieties of group representations and splittings of 3-manifolds. *Ann. of Math. (2)* **117**(1983), 109–146.
- [DH] A. Douady and J. Hubbard. On the density of Strebel differentials. *Invent. Math.* **30**(1975), 175–179.
- [D] D. Dumas. Holonomy limits of complex projective structures. Preprint, 2011. [arXiv:1105.5102](https://arxiv.org/abs/1105.5102)
- [DK1] D. Dumas and R. Kent. Experiments with skinning maps. In preparation.
- [DK2] D. Dumas and R. P. Kent. Bers slices are Zariski dense. *J. Topol.* **2**(2009), 373–379.
- [DK3] D. Dumas and R. P. Kent. Slicing, skinning, and grafting. *Amer. J. Math.* **131**(2009), 1419–1429.
- [FLP] A. Fathi, F. Laudenbach, and V. Poenaru. *Travaux de Thurston sur les surfaces*, volume 66 of *Astérisque*. Société Mathématique de France, Paris, 1979. Séminaire Orsay, With an English summary.
- [Flo] W. Floyd. Incompressible surfaces in 3-manifolds: the space of boundary curves. In *Low-dimensional topology and Kleinian groups (Coventry/Durham, 1984)*, volume 112 of *London Math. Soc. Lecture Note Ser.*, pages 131–143. Cambridge Univ. Press, Cambridge, 1986.

- [FO] W. Floyd and U. Oertel. Incompressible surfaces via branched surfaces. *Topology* **23**(1984), 117–125.
- [GKM] D. Gallo, M. Kapovich, and A. Marden. The monodromy groups of Schwarzian equations on closed Riemann surfaces. *Ann. of Math. (2)* **151**(2000), 625–704.
- [Gar] F. Gardiner. Measured foliations and the minimal norm property for quadratic differentials. *Acta Math.* **152**(1984), 57–76.
- [Gas] J. Gaster. A family of non-injective skinning maps with critical points. Preprint, 2012. To appear in *Trans. Amer. Math. Soc.*
- [GM] M. Goresky and R. MacPherson. *Stratified Morse theory*, volume 14 of *Ergebnisse der Mathematik und ihrer Grenzgebiete (3) [Results in Mathematics and Related Areas (3)]*. Springer-Verlag, Berlin, 1988.
- [Gra] K. Gravett. Ordered abelian groups. *Quart. J. Math. Oxford Ser. (2)* **7**(1956), 57–63.
- [Har] R. Hardt. Stratification of real analytic mappings and images. *Invent. Math.* **28**(1975), 193–208.
- [Hat] A. Hatcher. On the boundary curves of incompressible surfaces. *Pacific J. Math.* **99**(1982), 373–377.
- [HP] M. Heusener and J. Porti. The variety of characters in $\mathrm{PSL}_2(\mathbb{C})$. *Bol. Soc. Mat. Mexicana (3)* **10**(2004), 221–237.
- [Hir] H. Hironaka. Subanalytic sets. In *Number theory, algebraic geometry and commutative algebra, in honor of Yasuo Akizuki*, pages 453–493. Kinokuniya, Tokyo, 1973.
- [HM] J. Hubbard and H. Masur. Quadratic differentials and foliations. *Acta Math.* **142**(1979), 221–274.
- [Kap] M. Kapovich. *Hyperbolic manifolds and discrete groups*, volume 183 of *Progress in Mathematics*. Birkhäuser Boston Inc., Boston, MA, 2001.
- [Kaw] S. Kawai. The symplectic nature of the space of projective connections on Riemann surfaces. *Math. Ann.* **305**(1996), 161–182.
- [Ken] R. Kent. Skinning maps. *Duke Math. J.* **151**(2010), 279–336.
- [KW] D. Kleinbock and B. Weiss. Bounded geodesics in moduli space. *Int. Math. Res. Not.* **30**(2004), 1551–1560.
- [KK] A. Kokorin and V. Kopytov. *Fully ordered groups*. Halsted Press [John Wiley & Sons], New York-Toronto, Ont., 1974. Translated from the Russian by D. Louvish.
- [Kra] I. Kra. Deformation spaces. In *A crash course on Kleinian groups (Lectures at a Special Session, Annual Winter Meeting, Amer. Math. Soc., San Francisco, Calif., 1974)*, pages 48–70. Lecture Notes in Math., Vol. 400. Springer, Berlin, 1974.
- [KS] K. Krasnov and J.-M. Schlenker. A symplectic map between hyperbolic and complex Teichmüller theory. *Duke Math. J.* **150**(2009), 331–356.
- [Lab] F. Labourie. *Lectures on representations of surface groups*. Zurich Lectures in Advanced Mathematics. European Mathematical Society (EMS), Zürich, 2013.
- [Lev1] G. Levitt. Foliations and laminations on hyperbolic surfaces. *Topology* **22**(1983), 119–135.
- [Lev2] G. Levitt. \mathbf{R} -trees and the Bieri-Neumann-Strebel invariant. *Publ. Mat.* **38**(1994), 195–202.
- [MaSm] H. Masur and J. Smillie. Hausdorff dimension of sets of nonergodic measured foliations. *Ann. of Math. (2)* **134**(1991), 455–543.
- [McM1] C. McMullen. Amenability, Poincaré series and quasiconformal maps. *Invent. Math.* **97**(1989), 95–127.
- [McM2] C. McMullen. Iteration on Teichmüller space. *Invent. Math.* **99**(1990), 425–454.
- [Mor1] J. Morgan. Group actions on trees and the compactification of the space of classes of $\mathrm{SO}(n, 1)$ -representations. *Topology* **25**(1986), 1–33.

- [Mor2] J. Morgan. Λ -trees and their applications. *Bull. Amer. Math. Soc. (N.S.)* **26**(1992), 87–112.
- [MoSh1] J. Morgan and P. Shalen. Valuations, trees, and degenerations of hyperbolic structures. I. *Ann. of Math. (2)* **120**(1984), 401–476.
- [MoSh2] J. Morgan and P. Shalen. An introduction to compactifying spaces of hyperbolic structures by actions on trees. In *Geometry and topology (College Park, Md., 1983/84)*, volume 1167 of *Lecture Notes in Math.*, pages 228–240. Springer, Berlin, 1985.
- [MS1] J. Morgan and P. Shalen. Degenerations of hyperbolic structures. III. Actions of 3-manifold groups on trees and Thurston’s compactness theorem. *Ann. of Math. (2)* **127**(1988), 457–519.
- [MS2] J. Morgan and P. Shalen. Free actions of surface groups on \mathbf{R} -trees. *Topology* **30**(1991), 143–154.
- [Mos] L. Mosher. Train track expansions of measured foliations. Preprint, 2003. <http://andromeda.rutgers.edu/~mosher>
- [Ohs] K. Ohshika. Limits of geometrically tame Kleinian groups. *Invent. Math.* **99**(1990), 185–203.
- [Ota] J.-P. Otal. Thurston’s hyperbolization of Haken manifolds. In *Surveys in differential geometry, Vol. III (Cambridge, MA, 1996)*, pages 77–194. Int. Press, Boston, MA, 1998.
- [Pap] A. Papadopoulos. Réseaux ferroviaires et courbes simples sur les surfaces. *C. R. Acad. Sci. Paris Sér. I Math.* **297**(1983), 565–568.
- [PP] A. Papadopoulos and R. Penner. The Weil-Petersson symplectic structure at Thurston’s boundary. *Trans. Amer. Math. Soc.* **335**(1993), 891–904.
- [PH] R. Penner and J. Harer. *Combinatorics of train tracks*, volume 125 of *Annals of Mathematics Studies*. Princeton University Press, Princeton, NJ, 1992.
- [Roy] H. Royden. Automorphisms and isometries of Teichmüller space. In *Advances in the Theory of Riemann Surfaces (Proc. Conf., Stony Brook, N.Y., 1969)*, pages 369–383. Ann. of Math. Studies, No. 66. Princeton Univ. Press, Princeton, N.J., 1971.
- [Sul] D. Sullivan. On the ergodic theory at infinity of an arbitrary discrete group of hyperbolic motions. In *Riemann surfaces and related topics: Proceedings of the 1978 Stony Brook Conference (State Univ. New York, Stony Brook, N.Y., 1978)*, volume 97 of *Ann. of Math. Stud.*, pages 465–496. Princeton Univ. Press, Princeton, N.J., 1981.
- [Vee1] W. Veech. The Teichmüller geodesic flow. *Ann. of Math. (2)* **124**(1986), 441–530.
- [Vee2] W. Veech. Moduli spaces of quadratic differentials. *J. Analyse Math.* **55**(1990), 117–171.
- [Whi] H. Whitney. Tangents to an analytic variety. *Ann. of Math. (2)* **81**(1965), 496–549.
- [Yue] C. Yue. Conditional measure and flip invariance of Bowen-Margulis and harmonic measures on manifolds of negative curvature. *Ergodic Theory Dynam. Systems* **15**(1995), 807–811.
- [ZS] O. Zariski and P. Samuel. *Commutative algebra. Vol. II*. Springer-Verlag, New York, 1975. Reprint of the 1960 edition, Graduate Texts in Mathematics, Vol. 29.