Continuity of Thurston's length function

Jeffrey F. Brock

Abstract. A measured lamination μ geodesically realized in a hyperbolic 3-manifold M has a well defined *average length*, due to W. Thurston. For $M \cong S \times \mathbb{R}$ we prove that the function measuring the average length of the maximal realizable sublamination of μ varies bi-continuously in M and μ. Since connected, positive, non-realizable measured laminations arise as zeros of the length function, its continuity suggests new behavioral features of quasi-isometry invariants under limits of hyperbolic 3-manifolds.

Contents

1. Introduction

The geodesic length function is a fundamental tool in the study of deformation spaces of hyperbolic manifolds. In dimension 2, *average length* functions for measured laminations interpolate between positive scalar multiples of geodesic length functions for simple closed geodesics.

This notion also works in dimension 3: in general, one seeks a *pleated surface* X that *realizes* a measured lamination μ geodesically in a hyperbolic 3-manifold M. Then length in M is just length on X. For $M \cong S \times \mathbb{R}$, this notion of length in M has been central to the idea that aspects of the geometry of hyperbolic 3-manifolds are controlled by the geometry of hyperbolic surfaces.

Originally, in his 1986 preprint [**Th4**], Thurston stated a continuity theorem for the length function on the subset of the product of the measured lamination

1991 *Mathematics Subject Classification.* Primary 30F40; Secondary 30F60, 58E10.

Research partially supported by an NSF Postdoctoral Fellowship.

Date: November 18, 1998. *Revised:* October 3, 2000.

space with the deformation space of M where such realizing surfaces exist. This theorem was applied in the proof of the celebrated *double limit theorem*, though its proof was referred to a later paper (which has yet to appear).

We develop new uniform estimates on the geometry of pleated surfaces to obtain continuity result that subsumes Thurston's statement, giving a proof of this statement along the way. To state our results we review terminology.

Let S be an oriented surface, closed for simplicity. Let $\mathcal{ML}(S)$ be the space of measured laminations on S: the completion of the set of isotopy classes of essential simple closed curves γ on S equipped with positive real weights $t \in \mathbb{R}_+$. Let $X \in \mathrm{Teich}(S)$ be a marked hyperbolic Riemann surface. Then $t\gamma$ determines a closed geodesic γ^* on X with real weight t, while a general measured lamination $\mu \in \mathcal{ML}(S)$ determines a *geodesic lamination* $|\mu|$ on X, a closed subset foliated by complete geodesics, equipped with a *transverse measure*. There is a unique continuous function $\mathrm{length}_X \colon \mathcal{ML}(S) \to \mathbb{R}$ so that $\mathrm{length}_X(t\gamma) = t\ell_X(\gamma^*)$, where $\ell_X(.)$ denotes arclength in X.

Let $AH(S)$ denote the space of hyperbolic 3-manifolds M *marked* by homotopy equivalences $f\colon S \to M$ (up to marking preserving isometry); $AH(S)$ carries the compact-open topology (or *algebraic topology*) on the induced homomorphisms $f_*\colon \pi_1(S) \to \mathrm{Isom}^+ \mathbb{H}^3$ up to conjugacy. We say $\mu \in \mathcal{ML}(S)$ is *realizable* in M if there is a hyperbolic surface $X \in \mathrm{Teich}(S)$ and a continuous path-isometric mapping $g\colon X \to M$, consistent with markings on X and M, that maps each geodesic in $|\mu|$ by a local isometry to a geodesic in M. When such an X exists, set $\mathrm{length}_M(\mu) = \mathrm{length}_X(\mu)$. Our first goal is to prove Thurston's original claim (theorem 6.1):

Theorem 1 Length Continuous on Realizables *Let $\mathfrak{R}$ be the subset of $AH(S) \times \mathcal{ML}(S)$ consisting of all pairs (M, μ) of measured laminations μ realizable in marked hyperbolic 3-manifolds M. Then the length function*

$$\mathrm{length}\colon \mathfrak{R} \to \mathbb{R}$$

is continuous.

While the length of μ in M is not well defined when μ is not realizable in M, density of $\mathfrak{R}$ in $AH(S) \times \mathcal{ML}(S)$ allows an extension of $\mathrm{length}_M(\mu)$ to the function $\underline{\mathrm{length}}_M(\mu)$ obtained by taking the $\liminf$ of the infima of lengths of laminations realizable in M in smaller and smaller neighborhoods of μ. Using the proof of theorem 1, together with a train track shortening result of F. Bonahon, we prove the following generalization (theorem 7.1).

Theorem 2 Length Extends Continuously *The function*

$$\underline{\mathrm{length}}\colon AH(S) \times \mathcal{ML}(S) \to \mathbb{R}$$

is continuous.

By work of Thurston and Bonahon, when μ is connected and non-realizable we have $\underline{\mathrm{length}}_M(\mu) = 0$. Theorem 2 shows that in general $\underline{\mathrm{length}}$ is the function

$$(M, \mu) \to \mathrm{length}_M \circ\ \mathrm{R}_M(\mu)$$

where $\mathrm{R}_M \colon \mathcal{ML}(S) \to \mathcal{ML}(S)$ is the projection map assigning to each $\mu \in \mathcal{ML}(S)$ the maximal sublamination $\mathrm{R}_M(\mu) \subset \mu$ that is realizable in M.

We briefly discuss some difficulties in the proofs of theorems 1 and 2 that arise from consideration of Gromov-Hausdorff or *geometric* convergence of 1, 2 and 3-dimensional hyperbolic manifolds.

Problematic Hausdorff limits of simple closed geodesics. One subtlety in the proof arises even when M is fixed: given a convergent sequence $t_i c_i \to \mu$ in $\mathcal{ML}(S)$ of weighted simple closed curves, one may pass to a subsequence so that the underlying closed geodesics c_i converge in the Hausdorff topology to a (necessarily connected) geodesic lamination λ. Any such limit λ contains $|\mu|$. When λ (and hence μ) is realizable, convergence of lengths $\text{length}_M(t_i c_i)$ to $\text{length}_M(\mu)$ follows from a straightforward argument due to Thurston using so-called *nearly-straight train tracks*. When μ is realizable and λ contains non-realizable laminations the question arises: which phenomenon dominates? If, for example, λ contains a parabolic curve η and μ is a weighted closed geodesic, is it more efficient for the geodesic representatives c_i^* in M to travel in the cusp for η or near the closed geodesic $|\mu|^*$ realized in M? These tendencies compete, making the realizations c_i^* delicate to control (see §4).

Problematic geometric limits of hyperbolic manifolds. Another central issue in the proof arises from the range of possible geometric limits of manifolds M_i in an algebraically convergent sequence $M_i \to M$. Indeed, one approach to theorem 1 might employ compactness theorems for pleated surfaces to reduce the theorem to the well-known surface case. It seems likely that such an argument works on the product $GH(S) \times \mathcal{ML}(S)$ where $GH(S)$ denotes the finer topology on $H(S)$ of *strong convergence* ($M_i \to M$ only if in addition (M_i, ω_i) converges to (M, ω) geometrically for compatibly chosen baseframes ω_i). Convergent sequences in $AH(S)$ that do not converge in $GH(S)$, however, are a bountiful source of examples where such arguments run aground.

Consider the following example. On a surface of genus 2, let γ denote the central separating curve and let (α_1, β_1) and (α_2, β_2) denote meridian longitude pairs on the complementary punctured tori. It is possible to construct a convergent sequence

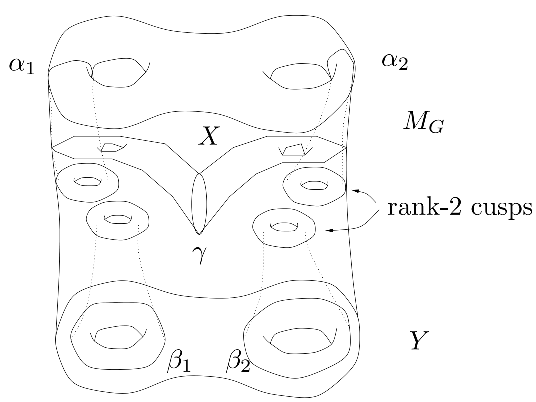

Figure 1. A problematic geometric limit. The algebraic limit M is the cover corresponding to $\pi_1(Y)$.

$\{M_i\} \subset AH(S)$ converging algebraically to M and geometrically to a manifold M_G

covered by M so that

$$M_G \cong S \times \mathbb{R} - \left\{ (\alpha_1 \cup \alpha_2, 1) \bigcup (\beta_1 \cup \beta_2, -1) \right\},$$

and M is the cover corresponding to the inclusion on π_1 of a boundary surface $Y \subset \partial M_G$ with $Y \cong S$. Then there are many families of compatibly marked pleated surfaces $X_i \subset M_i$, so that the triples (M_i, X_i, ω_i) converge geometrically to a triple (M_G, X, ω) for which X lifts in the natural covering $M \to M_G$ to an annulus with core curve γ (see figure 1); the surfaces X_i tend to infinity in Teichmüller space, but always converge geometrically to a possibly smaller surface up to subsequence. For μ_i approximating γ, any such limit X carries insufficient homotopy data from M to control the shape of μ_i sitting on X_i; μ_i could conceivably *unwind* on X_i if the structure on X_i were to become twisted about γ relative to M_i.

The content of theorem 3.6 (uniform relative twisting) is to control this potential unwinding. We show that when surfaces X_i realizing γ in M_i become arbitrarily twisted about γ either $\text{length}_{M_i}(\gamma) \to 0$, or $M_i \to \infty$ in $AH(S)$ (see §3).

Outline of the proof. We have chosen a direct approach that avoids explicit discussion of geometric limits since an explicit description of the spectrum of such limits is still very much under development.

For theorem 1, a diagonal argument gives a reduction to evaluating length on sequences where weighted simple closed curves $t_i c_i \to \mu$ in $\mathcal{ML}(S)$ and $M_i \to M$ in $AH(S)$. Via a projection map, we obtain the "closest" geodesic lamination λ_i to c_i that contains $|\mu|$. We then attempt to control the geodesic representatives of c_i on pleated surfaces X_i realizing λ_i in M_i. When $M_i = M$, this amounts to two challenges:

1. when c_i spirals towards a simple closed geodesic $\gamma \subset |\mu|$, we show it spirals on any pleated surface realizing γ, and
2. when c_i enters a close neighborhood of $|\mu|$ on X deep in a 'spike' of $X - \mu$, we show it does so on any pleated surface realizing μ.

In particular, these properties hold on X_i. When $\{M_i\}$ converges algebraically, we show that the estimates involved hold uniformly.

Then c_i can be forced to run along leaves of λ_i except for short jumps between leaves of λ_i. These short jumps are short in the unit tangent bundle, which allows us to construct a *nearly-straight* 1-complex in τ_i in M_i so that leaves of μ and c_i are simultaneously homotopic into paths on τ_i of small geodesic curvature (see §4 and lemma 5.2 for technical statements). These so-called nearly-straight train tracks τ_i serve to control both the position and the length of c_i in M_i for i sufficiently large.

As for theorem 2, the above argument shows that when μ is not necessarily realizable the limiting lengths are at least the length of the largest realizable sublamination μ^{r} of μ. For the remaining sublamination μ^{nr}, work of Bonahon shows that any sequence approximating μ^{nr} has length tending to zero in M. Combining these techniques and adapting to algebraic convergence, theorem 2 follows.

Applications. It follows from theorem 2 that the zero-locus of $\underline{\text{length}}$ consists of pairs (M, μ) such that either $\mu = 0$ or μ is the union of connected laminations non-realizable in M. Theorem 2, then, has the following application (corollary 7.3).

Corollary Zero-Locus *Let (M_i, μ_i) converge to (M, μ) in $AH(S) \times \mathcal{ML}(S)$ so that the sequence*

$$\left\{ \underline{\text{length}}_{M_i}(\mu_i) \right\}_{i=1}^{\infty}$$

converges to 0. *Then* $\mathrm{R}_M(\mu) = 0$.

In other words, no positive component μ' of μ is realizable in M. Since for any positive connected μ' in the kernel $\ker(\mathrm{R}_M)$ of R_M the support $|\mu'|$ is a quasi-isometry invariant of M, the corollary gives insight into how quasi-isometry invariants behave under algebraic convergence. We take up this question in [**Br1**].

History and references. The length function was originally introduced by Thurston in [**Th1**] and is discussed in [**Th5**]. The statement of theorem 1 appears in [**Th4**, Prop. 3.1] without proof, and $\underline{\text{length}}$ is defined in [**Th4**, §3]. Much work in the surface case is due to Francis Bonahon, who develops an elegant unified theory of Teichmüller space and measured lamination space via *geodesic currents* in [**Bon1**] and [**Bon2**].

F. Bonahon has recently developed a holomorphic analog for the length of a fixed measured lamination μ that generalizes the complex length of a simple closed geodesic, see [**Bon3**]. Bonahon has also shown in [**Bon4**, §7] that given a finer topology (that prevents the type of convergence described in the first problematic example above) measured laminations realized in a fixed manifold M have continuously varying lengths. See also [**Ohs**] who treats the case where no new parabolics arise in the limit M of M_i (an assumption ruling out both problematic examples).

The geometric topology on pleated surfaces is discussed in [**CEG**] and [**Th2**], and geometric limits of the type M_G above are studied in [**BO**], [**Th4**, §7], and [**Br2**]. Nearly-straight train tracks were first introduced by Thurston in [**Th1**, Ch. 8]. Their uses have been developed extensively by F. Bonahon in [**Bon1**] and Y. Minsky in [**Min1**] [**Min2**] to understand *degenerate ends* of hyperbolic 3-manifolds via geodesic currents and harmonic maps.

Plan of the paper. We first discuss background of surfaces and hyperbolic structures. Sections 3 and 4 develop uniform estimates for pleated surfaces, and use these estimates for the construction of uniformly nearly-straight train tracks mentioned above in section 5. Theorem 1 is then proven in section 6 where the train-tracks of section 5 prove sufficiently robust to imply lower semi-continuity of the extended function $\underline{\text{length}}$ defined on the full product $AH(S) \times \mathcal{ML}(S)$. We then adapt the shortening technique of Bonahon [**Bon1**, §5] to the setting of algebraic convergence to show upper semi-continuity of $\underline{\text{length}}$ in section 7.

Acknowledgments. I would like to thank Yair Minsky for discussions of many aspects of this argument, as well as Francis Bonahon for suggesting the application of his shortening process to theorem 7.1. I would also like to thank Dick Canary, Steve Kerckhoff, Bill Thurston, Curt McMullen, and the referee for many helpful suggestions.

2. Preliminaries

Let S be a compact oriented surface of negative Euler characteristic. When S has boundary ∂S, let $\mathrm{int}(S)$ denote its interior $S - \partial S$.

Teich(S). The *Teichmüller space* $\mathrm{Teich}(S)$ of S is the space of isotopy classes of marked hyperbolic structures of finite area on $\mathrm{int}(S)$; that is, isotopy classes of orientation-preserving homeomorphisms $(f\colon \mathrm{int}(S) \to X)$ from $\mathrm{int}(S)$ to the complete, finite-area, hyperbolic surface X. Points $X \in \mathrm{Teich}(S)$ are implicitly marked. We will use the notation $(f, X) \in \mathrm{Teich}(S)$ when explicit reference to the marking is necessary.

The intersection number. Let $\mathcal{S}$ be the set of all essential non-peripheral isotopy classes of simple closed curves on S. The *geometric intersection number*

$$i\colon \mathcal{S} \times \mathcal{S} \to \mathbb{Z}$$

counts the minimal number of intersections of curves γ and η in a pair of isotopy classes $([\gamma], [\eta]) \in \mathcal{S} \times \mathcal{S}$, and takes the value 0 on the diagonal.

$\boldsymbol{\mathcal{ML}(S)}$**.** The *measured lamination space* is the closure of the image of the set $\mathbb{R}_+ \times \mathcal{S}$ under the embedding

$$\iota\colon \mathbb{R}_+ \times \mathcal{S} \to \mathbb{R}_+^{\mathcal{S}}$$

defined by

$$\langle \iota(t\beta) \rangle_\alpha = t i(\alpha, \beta).$$

Let $\mathcal{ML}(S)_+$ denote the positive elements of $\mathcal{ML}(S)$. The function $i(.,.)$ admits a continuous linear extension to

$$i\colon \mathcal{ML}(S) \times \mathcal{ML}(S) \to \mathbb{R}_{\geq 0}$$

(see [**Bon1**, Prop. 4.5]).

$\boldsymbol{\mathcal{GL}(S)}$**.** Let $X \in \mathrm{Teich}(S)$ be a hyperbolic surface. Then a *geodesic lamination* λ on X is a closed subset of X with a decomposition as pairwise disjoint complete simple geodesics called *leaves* of λ. The geodesic laminations $\mathcal{GL}(X)$ are topologized by the Hausdorff topology on closed subsets of X [**CEG**, §4.1] [**Mc**, §2.1]. Since for any other surface $Y \in \mathrm{Teich}(S)$, $\mathcal{GL}(X)$ and $\mathcal{GL}(Y)$ are canonically homeomorphic via the circle at ∞ for $\pi_1(S)$ (see [**Fl**], [**Bon2**]), $\mathcal{GL}(X)$ is canonically associated to the topological surface S. We use the notation $\mathcal{GL}(S)$ and think of a point λ in $\mathcal{GL}(S)$ as specifying a geodesic lamination on any hyperbolic surface. A geodesic lamination $\lambda \in \mathcal{GL}(S)$ on is *maximal* if each component of the complement of its realization on $X \in \mathrm{Teich}(S)$ is the interior of an ideal hyperbolic triangle.

Measured laminations $\mu \in \mathcal{ML}(S)$ are identified with transversely measured geodesic laminations (see [**Th1**, §8] [**Bon2**]) of *compact support*; a measured lamination μ determines a geodesic lamination $|\mu| \in \mathcal{GL}(S)$ whose realization on any $X \in \mathrm{Teich}(S)$ is compact.

The thick-thin decomposition. Given any Riemannian manifold M, the *injectivity radius* $\mathrm{inj}\colon M \to \mathbb{R}$ is the function that measures the radius of the maximal isometrically embedded open ball at each point $x \in M$. Given any $\epsilon > 0$ M admits a decomposition into its ϵ*-thick-part* $M_{\geq \epsilon}$ where $\mathrm{inj} \geq \epsilon$ and its ϵ*-thin-part* $M_{<\epsilon}$ where $\mathrm{inj} < \epsilon$. When M is a complete hyperbolic n-manifold a theorem of Margulis implies there is a constant ϵ_n depending only on the dimension, such that the ϵ_n-thin-part $M_{<\epsilon_n}$ has a standard type.

For $n = 2$, each component of $M_{<\epsilon_2}$ is either an annular neighborhood of a short geodesic, or isometric to a neighborhood of the puncture in a hyperbolic punctured disk. For $n = 3$, each component T of $M_{<\epsilon_3}$ is either

- a *Margulis tube:* $T \cong S^1 \times D^2$ is a solid torus neighborhood of a short geodesic γ in M,
- a *rank-1 cusp:* $T \cong S^1 \times \mathbb{R} \times (0, \infty)$ is the quotient of a horoball by a $\mathbb{Z}$ parabolic subgroup $\mathrm{Isom}^+ \mathbb{H}^3$,

- or a *rank-2 cusp:* $T \cong T^2 \times (0, \infty)$ is the quotient of a horoball by a $\mathbb{Z} \oplus \mathbb{Z}$ parabolic subgroup of $\mathrm{Isom}^+ \mathbb{H}^3$

(see [**BP**, Thm. D.3.3]).

The hyperbolic 3-manifolds of principal interest to us will be described by the following *deformation space*.

AH(S). Let $H(S)$ denote the set of all hyperbolic 3-manifolds M marked by homotopy equivalences $f\colon S \to M$ such that f sends ∂S to cusps of M under the equivalence relation

$$(f\colon S \to M) \quad \sim \quad (g\colon S \to N)$$

if there is an isometry $\phi\colon M \to N$ such that $\phi \circ f$ is homotopic to g. With the compact-open topology on the holonomy representations $[f_*]\colon \pi_1(S) \to \mathrm{Isom}^+(\mathbb{H}^3)$ up to conjugacy, $H(S)$ becomes the *algebraic deformation space* $AH(S)$. As with $\mathrm{Teich}(S)$, manifolds in $AH(S)$ are implicitly marked.

To pin down a specific representation in the conjugacy class of f_* (up to inner automorphism of $\pi_1(S)$), it suffices to make a choice of *baseframe* $\omega \in M$, that is a basepoint, which we denote by $|\omega|$, and an orthonormal frame at the basepoint. Then requiring that the standard frame at the origin in $\mathbb{H}^3$ lie over ω in the covering projection $\mathbb{H}^3 \to M$, determines a representation $\rho \in [f_*]$. Given an algebraically convergent sequence $M_i \to M$ in $AH(S)$, there are by definition baseframes $\omega_i \in M_i$ that determine convergent representations $\rho_i \to \rho$, i.e. so that $\rho_i(g) \to \rho(g)$ in $\mathrm{Isom}^+(\mathbb{H}^3)$ for each $g \in \pi_1(S)$.

A more geometric formulation of the algebraic topology (via compactness of the based hyperbolic manifolds (M_i, ω_i) in the geometric topology, see [**BP**, Ch. E.]) is the following: M_i converges to M in $AH(S)$ if for any compact subset $K \subset M$, there are smooth, marking-preserving homotopy equivalences

$$q_i\colon M \to M_i$$

so that q_i tends to a local isometry on K in the C^∞ topology (see [**Mc**, §3.1]).

Pleated Surfaces. A *pleated surface* $g\colon X \to M$ in a hyperbolic 3-manifold M, is a hyperbolic surface X together with a path-isometry $g\colon X \to M$ (g sends rectifiable arcs in X to rectifiable arcs in M of the same length) with the property that through each $x \in X$ there is a geodesic segment α mapped isometrically to M.

The *pleating locus* $\Pi(g)$ is the set of points on X where g fails to be a local isometry. The pleating locus is a geodesic lamination on X (see [**Th2**, Prop. 5.2]); it is the set where the image of g is bent, or *pleated*.

Given $M = (f\colon S \to M) \in AH(S)$, define the set $\mathcal{PS}(M)$ to be the set of compatibly marked pleated surfaces in M: i.e. the set of pairs (g, X), such that $g\colon X \to M$ is a pleated surface, $X = (\phi, X)$ lies in $\mathrm{Teich}(S)$, and we have $g \circ \phi \simeq f$. As with $\mathrm{Teich}(S)$ we will often suppress reference to marking and write $X \in \mathcal{PS}(M)$ to refer to a compatibly marked pleated surface $g\colon X \to M$.

Realizing laminations. We say the pleated surface $(g, X) \in \mathcal{PS}(M)$ *realizes* the geodesic lamination λ on X if g maps each leaf of λ by a local isometry to M (the intersection $\lambda \cap \Pi(g)$ is again a geodesic lamination). For given M, the subset $\mathcal{R}(M) \subset \mathcal{ML}(S)$ consisting of $\mu \in \mathcal{ML}(S)$ such that $|\mu|$ is realizable in M is dense in $\mathcal{ML}(S)$ (see [**Th1**] [**CEG**]). We say $\mu \in \mathcal{R}(M)$ is *realizable* as a measured lamination; by convention the zero lamination $0 \in \mathcal{ML}(S)$ is trivially realizable.

Given $M \in AH(S)$, let

$$\mathrm{R}_M \colon \mathcal{ML}(S) \to \mathcal{ML}(S)$$

be the projection so that $\mathrm{R}_M(0) = 0$ and for any $\mu \in \mathcal{ML}(S)_+$, $\mathrm{R}_M(\mu)$ is the maximal sublamination of μ that is realizable in M.

If $c \in \mathcal{S}$ satisfies $\mathrm{R}_M(tc) = 0$, $t > 0$, c is called an *accidental parabolic* for M. Note that the image $\mathrm{R}_M(\mathcal{ML}(S)) = \mathcal{R}(M)$ is not in general open when M has accidental parabolics, but it contains an open dense subset of $\mathcal{ML}(S)$ (see [**Th1**, Ch. 8, 9]).

Length functions. Given $X \in \mathrm{Teich}(S)$, there is a unique continuous function

$$\mathrm{length}\colon \mathrm{Teich}(S) \times \mathcal{ML}(S) \to \mathbb{R}$$

written $(X, \mu) \to \mathrm{length}_X(\mu)$, with the property that for any isotopy class $t\gamma \in \mathbb{R}_+ \times \mathcal{S}$ we have $\mathrm{length}_X(t\gamma) = t\ \ell_X(\gamma^*)$ where γ^* is the unique geodesic representative of γ on X (see [**Th4**], [**Bon1**, Prop. 4.5]).

Similarly, given $(f \colon S \to M) \in AH(S)$, if $f_*(\gamma)$ is not parabolic in $\pi_1(M)$, then the geodesic representative γ^* exists and $t\gamma$ has a well defined length

$$\mathrm{length}_M(t\gamma) = t\ \ell_M(\gamma^*).$$

For $\mu \in \mathcal{R}(M)$ realizable by a pleated surface $(g, X) \in \mathcal{PS}(M)$, Thurston extends this notion of length to the *length function* which measures the length of μ in M by setting

$$\mathrm{length}_M(\mu) = \mathrm{length}_X(\mu).$$

While the function $\mathrm{length}_M(\mu)$ is only defined on $\mu \in \mathcal{R}(M)$, one important use of the length function is to detect the non-realizability of a lamination. Indeed, for weighted simple closed curves $t\gamma$, the function $\mathrm{length}_{(.)}(t\gamma)$ extends continuously to a function

$$\underline{\mathrm{length}}_{(.)}(t\gamma)\colon AH(S) \to \mathbb{R}$$

that takes the value 0 on manifolds $M \in AH(S)$ for which γ is a parabolic element of $\pi_1(M)$.

More generally, following Thurston [**Th4**, Cor. 3.5] we define the function

$$\underline{\mathrm{length}}\colon AH(S) \times \mathcal{ML}(S) \to \mathbb{R}$$

on the full product $AH(S) \times \mathcal{ML}(S)$ by setting $\underline{\mathrm{length}}_M(\mu)$ equal to the lim inf of the infima of the lengths of realizable laminations ν over closer and closer neighborhoods of μ. Then zeros of $\underline{\mathrm{length}}$ on $AH(S) \times \mathcal{ML}(S)_+$ are manifold-lamination pairs (M, μ) for which no connected component of μ is realizable in M. As remarked, it follows from [**Th1**, Ch. 9] [**Bon1**, Lem. 5.1] that when μ is connected $\underline{\mathrm{length}}$ is equal to the function

$$(M, \mu) \to \mathrm{length}_M(\mathrm{R}_M(\mu))$$

on $AH(S) \times \mathcal{ML}(S)$.

Train tracks. In [**Th1**], Thurston introduced train tracks as neighborhood systems for measured and geodesic laminations. This theory is developed in detail in [**PH**].

A *train track* τ on a hyperbolic surface X is a 1-complex in X whose edges (called *branches*) are C^1 arcs, and whose vertices (called *switches*) carry the additional data of a tangent line to which all incident edges are mutually tangent at their endpoints: with respect to a tangent vector v at a switch s along such a line l, there are incoming and outgoing branches of τ, and an incoming and outgoing branch form a C^1 arc tangent to v. We require s to have at least one incoming and outgoing branch. Furthermore, we require that the double of each component of $X - \tau$ along the interiors of the branches in its boundary has negative Euler characteristic.

A *train-path* on τ is a monotone C^1 immersion $\rho\colon \mathbb{R} \to X$ (a "bi-infinite" train-path) or $\rho\colon S^1 \to X$ (a "closed" train-path) with image in τ. Given a geodesic lamination λ on a hyperbolic surface X, we say τ *carries* λ if there is a differentiable map $p\colon X \to X$ homotopic to the identity, and non-singular on the tangent spaces to leaves of λ so that $p(\lambda) \subset \tau$. In practice, this is equivalent to the existence of an homotopy of each leaf ℓ of λ (*rel*-ideal endpoints) into τ through smooth arcs to a train-path $\rho\colon \mathbb{R} \to \tau$. We say τ *minimally carries* λ if for each branch $b \subset \tau$ there is a leaf ℓ of λ whose corresponding train-path ρ contains b in its image. Train tracks τ and τ' are *equivalent* if the inclusion of τ is homotopic to the inclusion of τ' through train tracks on X. Equivalent train tracks carry the same laminations.

A train track τ^* in a hyperbolic 3 manifold $(f\colon S \to M) \in AH(S)$ is a train track τ on $X \in \mathrm{Teich}(S)$ together with a smooth map $h\colon X \to M$ compatible with marking so that $h(\tau) = \tau^*$. The map h serves to mark the *realization* τ^* of τ in M with homotopy information from S. We say $\tau^* \subset M$ carries a lamination λ if the train track τ carries λ on X.

When a train track τ carries the support $|\mu|$ of a measured lamination μ we will say it *carries* μ. Then μ assigns a mass $m_\mu(b)$ to each branch b of τ by taking a point $x \in b$ and setting $m_\mu(b)$ total mass of the measure associated to the transverse arc $p^{-1}(x)$. The train track minimally carries μ if it minimally carries $|\mu|$. In this case $m_\mu(b) \neq 0$ for each $b \subset \tau$. Assigning the weight $m_\mu(b)$ to each branch b we obtain a *weighted train track*.

If τ sits on a hyperbolic surface X, where each b has an arclength $\ell_X(b)$, the *track length*

$$\ell_\tau(\mu) = \sum_{b \subset \tau} m_b(\mu)\ell_X(b)$$

bounds $\mathrm{length}_X(\mu)$ from above. Likewise, when τ is realized in the manifold $M \in AH(S)$ with image τ^*, and each branch $b \subset \tau^*$ has arclength $\ell_M(b)$ in M, the track length

$$\ell_{\tau^*}(\mu) = \sum_{b \subset \tau^*} m_b(\mu)\ell_M(b)$$

bounds $\mathrm{length}_M(\mu)$ from above. (See [**Bon1**, §5] for a similar discussion of length and train tracks).

Thurston's Uniform Injectivity theorem. Given any hyperbolic manifold M, a point in its *projective tangent bundle* $\mathbf{P}M$ is determined by a pair (x, γ) of a point $x \in M$ and geodesic γ through x. Let

$$d_{\mathbf{P}}\colon \mathbf{P}M \times \mathbf{P}M \to \mathbb{R}$$

be the natural distance on $\mathbf{P}M$ so that $d_{\mathbf{P}}\left((x_1,\gamma_1),(x_2,\gamma_2)\right)$ is the maximum of the minimal distance $d_M(x_1, x_2)$ from x_1 to x_2 in M and the minimal angle between γ_1 and the parallel transport of γ_2 along the geodesic α joining x_1 to x_2 of length $d_M(x_1, x_2)$.

Let $M = (f\colon S \to M) \in AH(S)$, and let (g, X) lie in $\mathcal{PS}(M)$. Since $g_*(\pi_1(X))$ is an isomorphism, all surfaces $(g, X) \in \mathcal{PS}(M)$ satisfy the technical assumption of *double incompressibility* in the hypotheses of Thurston's *uniform injectivity theorem* which we may thus restate as follows:

Theorem 2.1 (Thurston). Uniform Injectivity *Fix $\epsilon_0 > 0$. Then for every ϵ there is a δ so that if M lies $AH(S)$ and $(g, X) \in \mathcal{PS}(M)$ is a pleated surface realizing $\lambda \in \mathcal{GL}(S)$, then for any two points $x \in \ell_x$ and $y \in \ell_y$ in leaves ℓ_x and ℓ_y of λ with x and y in the thick part $X_{\geq \epsilon_0}$ we have*

$$d_{\mathbf{P}}\left((g(x), g(\ell_x)), (g(y), g(\ell_y))\right) \geq \delta$$

whenever $d_X(x, y) \geq \epsilon$.

In [**Min1**], Y. Minsky obtains the following proportional relationship between ϵ and δ in the above theorem. Note that his statement ([**Min1**, Lem. 2.3]) overlooks a hypothesis which is that δ should be taken sufficiently small. Its uses in [**Min1**] and here, which are similar, satisfy this hypothesis. We give the correct restatement below.

For any arc α in a hyperbolic manifold N, let $\text{length}_N(\alpha)$ denote the length of the geodesic segment homotopic to α *rel*-endpoints.

Theorem 2.2 (Minsky). Proportional Injectivity *There exists $C_{\text{inj}} > 1$ and $\delta_{\text{inj}} > 0$ depending only on S so that for all $\delta < \delta_{\text{inj}}$ the following holds: Let $M \in AH(S)$ and $(g, X) \in \mathcal{PS}(M)$ realize λ as above. Let $x \in \ell_x$ and $y \in \ell_y$ be points on leaves ℓ_x and ℓ_y of λ with an arc $\alpha \subset X$ connecting them whose interior meets neither ℓ_x nor ℓ_y. Then provided* $\text{length}_M(g(\alpha)) < \delta_{\text{inj}}$, *we have*

$$\text{length}_X(\alpha) < C_{\text{inj}} \text{length}_M(g(\alpha)).$$

The constant δ_{inj} is determined by applying the uniform injectivity theorem taking ϵ_2 to be the 2-dimensional Margulis constant and setting $\epsilon = \epsilon_2/2$; then for $\text{length}_M(g(\alpha))$ less than the resulting $\delta = \delta_{\text{inj}}$, the above proportional relationship holds.

To conclude this section, we apply uniform injectivity to give a proof of the following theorem for use later. The theorem seems to be well known but has not been formally published (see [**Th1**, Prop. 8.10.5]).

Theorem 2.3. *Let $M \in AH(S)$, and let $\lambda \in \mathcal{GL}(S)$ be realized by a pleated surface $(g, X) \in \mathcal{PS}(M)$. Let $l \subset X - \lambda$ be a complete geodesic asymptotic to λ in each direction. Then there is a pleated surface $(g', X') \in \mathcal{PS}(M)$ realizing $\lambda \sqcup l$.*

Proof: Assume $\lambda \sqcup l$ is not realizable. Then a lift $\widetilde{g} \colon \widetilde{X} \to \widetilde{M}$ sends a lift $\widetilde{l}$ of l to a bi-infinite path that terminates at a single point $p \in S^2_\infty$ in each direction. If l_1 and l_2 are leaves of λ to which l is asymptotic in each direction, then $\widetilde{l}$ is asymptotic in $\widetilde{X}$ to lifts $\widetilde{l_1}$ and $\widetilde{l_2}$ so that the geodesics $\widetilde{g}(\widetilde{l_1})$ and $\widetilde{g}(\widetilde{l_2})$ each have one end at p.

Since $\widetilde{l}$ is asymptotic to $\widetilde{l_1}$ and $\widetilde{l_2}$ on $\widetilde{X}$, and $\widetilde{g}(\widetilde{l_1})$ and $\widetilde{g}(\widetilde{l_2})$ are asymptotic in $\mathbb{H}^3$, for any $\epsilon > 0$ there are points $x_1 \in \widetilde{l_1}$, $x_2 \in \widetilde{l_2}$, and $x_1', x_2' \in l$ so that

$$d_{\widetilde{X}}(x_1, x_1') < \epsilon, \quad d_{\widetilde{X}}(x_2, x_2') < \epsilon, \quad \text{and} \quad d_{\mathbb{H}^3}(\widetilde{g}(x_1), \widetilde{g}(x_2)) < \epsilon.$$

Thus, there is an arc $\alpha_1 * \beta * \alpha_2$ of length 3ϵ in $\mathbb{H}^3$ where α_1 joins $\widetilde{g}(x_1')$ and $\widetilde{g}(x_1)$ in $\widetilde{g}(\widetilde{(X)})$, β joins $\widetilde{g}(x_1)$ and $\widetilde{g}(x_2)$, and α_2 joins $\widetilde{g}(x_2)$ and $\widetilde{g}(x_2')$ in $\widetilde{g}(\widetilde{X})$. By proportional injectivity, theorem 2.2, if ϵ is chosen sufficiently small β is homotopic to an arc β' in $\widetilde{g}(\widetilde{X})$ of length less than $C\epsilon$.

If $l_0 \subset \widetilde{l}$ is the geodesic segment between x'_1 and x'_2, then the loop

$$\alpha_1 * \beta * \alpha_2 * \widetilde{g}(l_0)$$

bounds a disk in $\mathbb{H}^3$, which, by incompressibility of g, implies that

$$g^{-1}(\alpha_1 * \beta' * \alpha_2) * l_0$$

bounds a disk in $\widetilde{X}$. But l_0 is a geodesic segment in $\widetilde{X}$, which may be made arbitrarily long by choosing ϵ sufficiently small. A long geodesic followed by a short arc cannot bound a disk in $\mathbb{H}^2$, so we have a contradiction

∎

3. Controlling spiraling

Consider a sequence of weighted simple closed curves $\{t_i c_i\} \subset \mathcal{ML}(S)$ that converges to a single weighted simple closed curve $\gamma \in \mathcal{ML}(S)$, with $c_i \neq \gamma$. On a fixed hyperbolic surface, c_i must spiral increasingly about γ as i tends to infinity. An important part of our argument will be to ensure that when γ is realizable in the algebraic limit M of a convergent sequence $M_i \in AH(S)$, the realizations of c_i in M_i spiral increasingly towards the realization of γ in M_i, in an appropriate sense. Due to the phenomena alluded to in the introduction that may arise in the geometric limit of the manifolds M_i, we require a very sensitive measure of this spiraling.

Relative twisting. In [**Min3**] Minsky defines a notion of the relative twisting of a pair of isotopy classes α and β in $\mathcal{S}$ relative to another isotopy class $\gamma \in \mathcal{S}$ which α and β each intersect.

To make the discussion clear, we fix a hyperbolic structure $X \in \mathrm{Teich}(S)$ for reference and pass freely from isotopy classes in $\mathcal{S}$ to their geodesic representatives on X. Let γ be a simple closed geodesic on X and let α and β be distinct simple closed geodesics on X such that $i(\alpha, \gamma)$ and $i(\beta, \gamma)$ are each non-zero. Then following [**Min3**] we define the *relative twisting*

$$\tau_\gamma(\alpha, \beta)$$

of α and β relative to γ to be interval in $\mathbb{Z}$ of width at most 2 as follows: let $G \in \pi_1(X) \subset \mathrm{Isom}^+ \mathbb{H}^2$ be an indivisible hyperbolic element of $\pi_1(X)$ stabilizing a lift $\widetilde{\gamma}$ of γ to the universal cover $\widetilde{X} = \mathbb{H}^2$. Let $\mathrm{fix}(G)$ denote the fixed points of G. Now consider the annular cover X_γ of X corresponding to $\langle G \rangle$. Let

$$\overline{X_\gamma} = \mathbb{H}^2 \cup (S^1_\infty - \mathrm{fix}(G))/\langle G \rangle$$

denote X_γ with its ideal boundary adjoined. Consider lifts $\widetilde{\alpha}$ of α and $\widetilde{\beta}$ of β to $\overline{X_\gamma}$ that cross γ, together with their endpoints $\partial(\widetilde{\alpha})$ and $\partial(\widetilde{\beta})$ at infinity. Let $x(\widetilde{\alpha}, \widetilde{\beta})$ denote the number of intersections of $\widetilde{\alpha}$ and $\widetilde{\beta}$ reckoned positively if $\widetilde{\beta}$ is more *leftward* of γ than $\widetilde{\alpha}$, in the sense that $\widetilde{\beta}$ crosses $\widetilde{\alpha}$ from right to left on $\overline{X_\gamma}$ (this depends only on the orientation on S, not on orientations for the curves γ, α and β). Since any two lifts of α to $\overline{X_\gamma}$ do not intersect, and likewise for β, the range of values for $x(\widetilde{\alpha}, \widetilde{\beta})$ lies in an interval of $\mathbb{Z}$ of width 2 (see figure 2). We let $\tau_\gamma(\alpha, \beta)$ denote this *interval* in $\mathbb{Z}$. Since these quantities are naturally determined by separation properties of pairs of points on the circle at infinity S^1_∞, the interval $\tau_\gamma(\alpha, \beta)$ does not depend on the choice of underlying hyperbolic structure.

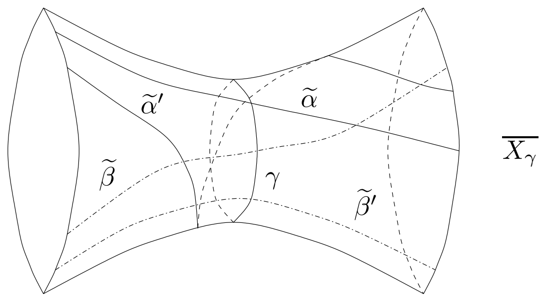

Figure 2. Relative twisting. Here $x(\widetilde{\alpha},\widetilde{\beta})=1$, $x(\widetilde{\alpha}',\widetilde{\beta})=2$, and $x(\widetilde{\alpha},\widetilde{\beta}')=0$. Thus $\tau_\gamma(\alpha,\beta)=[0,2]\subset\mathbb{Z}$.

Our interest will be in coarse properties of $\tau_\gamma(\alpha,\beta)$ so we let $\sigma_\gamma(\alpha,\beta)\in\tau_\gamma(\alpha,\beta)$ be the integer with least absolute value in $\tau_\gamma(\alpha,\beta)$.

The important properties of $\tau_\gamma(\alpha,\beta)$ are easily verified (see [**Min3**, Lem. 1]); we translate them into properties of $\sigma_\tau(\alpha,\beta)$. Given $a,b\in\mathbb{R}$, $Q\in\mathbb{R}^+$, let $a\asymp_Q b$ denote $|a-b|\le Q$.

Lemma 3.1 (Minksy). *Let α, β and δ in $\mathcal{S}$ have non-zero geometric intersection with $\gamma\in\mathcal{S}$ on the oriented hyperbolic surface X. Then we have*

RT1 *Intersection bounds twisting:*

$$|\sigma_\gamma(\alpha,\beta)|\le i(\alpha,\beta)+1$$

RT2 *Quasi-additivity:*

$$\sigma_\gamma(\alpha,\delta)\asymp_5(\sigma_\gamma(\alpha,\beta)+\sigma_\gamma(\beta,\delta))\,.$$

Simplicial hyperbolic surfaces. We employ the technique of continuous families *simplicial hyperbolic surfaces* in hyperbolic 3-manifolds (see [**Th1**, Ch. 8] [**Bon1**] [**Can**, §5]). Let $\mathrm{Sing}_k(S)$ denote the marked singular hyperbolic structures on S, i.e. complete, finite-area, marked surfaces Y that are hyperbolic away from at most $k\ge 0$ cone singularities each with cone angle at least 2π, up to marking preserving isometry. Given $Y\in\mathrm{Sing}_k(S)$ let T be a 'triangulation' of Y with geodesic edges terminating at singularities if Y is singular, and at punctures of Y if Y has cusps (T is a *triangulation of a punctured surface* in the sense of Canary [**Can**, §3]). Let N be a hyperbolic 3-manifold. Then a path-isometry $h\colon Y\to N$ is a *simplicial hyperbolic surface with associated triangulation* T if it is a local isometry on $Y-T$.

Given $M\in AH(S)$, we let $\mathcal{SH}_k(M)$ denote the marking-preserving simplicial hyperbolic surfaces $h\colon Y\to M$ so that $Y\in\mathrm{Sing}_k(S)$. When a simplicial hyperbolic surface $(h\colon Y\to M)\in\mathcal{SH}_k(M)$ maps an edge of T representing $\gamma\in\mathcal{S}$ to its geodesic representative γ^* in M, we say h *realizes* γ. If, in addition, T only has one non-ideal vertex (necessarily lying in γ) we say T and h are *adapted to* γ.

A theorem of Hatcher [**Hat**] guarantees that there is a finite sequence of elementary moves where one edge changes at a time relating any two triangulations adapted to γ, and furthermore that these moves can be chosen to preserve γ. By a theorem of Canary, there is a continuous family $\{(h_t\colon Y_t\to M)\}\subset\mathcal{SH}_{k+1}(M)$ of simplicial hyperbolic surfaces interpolating between simplicial hyperbolic surfaces $(h_0\colon Y_0\to M)$ and $(h_1\colon Y_1\to M)$ in $\mathcal{SH}_k(M)$ whose associated triangulations differ by precisely such an elementary move (see [**Can**, Lem. 5.3], figure 3). The

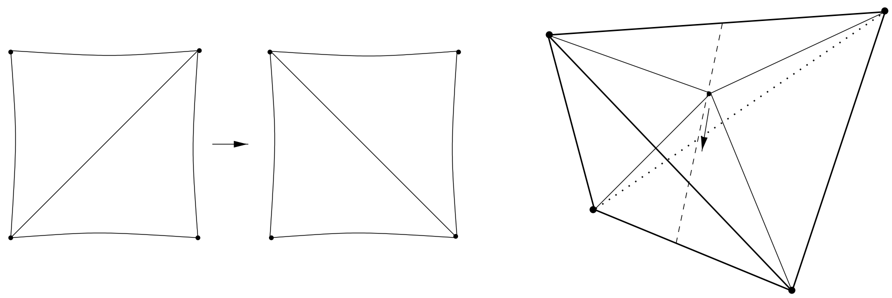

Figure 3. An elementary move on a triangulation is realized by a continuous family of simplicial hyperbolic surfaces with one greater singularity.

associated triangulations for h_t agree with those for h_0 and h_1 except for the edge involved in the elementary move, so we may combine these results to interpolate between simplicial hyperbolic surfaces adapted to γ through surfaces realizing γ:

THEOREM 3.2. *Let $(h_0\colon Y_0 \to M)$ and $(h_1\colon Y_1 \to M)$ in $\mathcal{SH}_1(M)$ be simplicial hyperbolic surfaces adapted to $\gamma \in \mathcal{S}$. Then there is a continuous family $(h_t\colon Y_t \to M)$ in $\mathcal{SH}_2(M)$ so that h_t realizes γ for each t.*

(See [**Min4**, §4] for a similar discussion).

Given $\gamma \in \mathcal{S}$, and a simplicial hyperbolic surface $(h\colon Y \to M) \in \mathcal{SH}_k(M)$, we define $\text{length}_Y(\gamma)$ to be the length of the unique geodesic representative of γ in the singular hyperbolic metric on Y.

Pinched negative curvature. The collar lemma (see e.g. [**Bus**, Thm 4.4.6]) guarantees that for any L there is a K so that two simple closed geodesics of length less than L on a surface $X \in \text{Teich}(S)$ have intersection number less than K; i.e. K does not depend on X. Since, however, simplicial hyperbolic surfaces can have concentrated negative curvature it is straightforward to build sequences of simplicial hyperbolic surfaces on which pairs of curves with length bounded by L have intersection number that grows without bound.

Since each $(h\colon Y \to M) \in \mathcal{SH}_k(M)$ is incompressible (h_* is injective), the non-singular hyperbolic metric on M remedies the situation.

LEMMA 3.3. *For any $L > 0$ there is a $J > 0$ such that for any $M \in AH(S)$ and $(h\colon Y \to M) \in \mathcal{SH}_k(M)$ for which α and β in $\mathcal{S}$ each have length on Y bounded by L, the geometric intersection number $i(\alpha, \beta)$ is bounded by J.*

Proof: Let $(h\colon Y \to M) \in \mathcal{SH}_k(M)$, and let α and β have geodesic representatives of length less than L on Y with images $\widehat{\alpha}$ and $\widehat{\beta}$ under h. Assume that $i(\alpha, \beta) \neq 0$ since the theorem is trivial otherwise.

Fix a point $x \in \widehat{\alpha} \cap \widehat{\beta}$. We establish the following lemma for future reference:

LEMMA 3.4. *For all L there is an ϵ such that the following holds. Let α and β in $\mathcal{S}$ have non-zero geometric intersection. Then if $\widehat{\alpha}$ and $\widehat{\beta}$ are intersecting rectifiable representatives of α and β in $M \in AH(S)$ with lengths $\ell_M(\widehat{\alpha})$ and $\ell_M(\widehat{\beta})$ each less than L, then their union $\widehat{\alpha} \cup \widehat{\beta}$ lies in the ϵ-thick part $M_{\geq \epsilon}$.*

Proof: Let ϵ_3 be the Margulis constant for dimension 3. Since α and β determine non-commuting elements of $\pi_1(S)$, the curves $\widehat{\alpha}$ and $\widehat{\beta}$ cannot lie in the same component of the thin part $M_{<\epsilon_3}$. This means the union $\widehat{\alpha} \cup \widehat{\beta}$ can only penetrate a distance at most L into $M_{<\epsilon_3}$. By the collar lemma for hyperbolic 3-manifolds (see [**BM**, §4]), any sufficiently short closed geodesic γ in M admits an embedded metric tubular neighborhood whose radius depends only on the length of γ, and tends to ∞ as the length of γ tends to 0. It follows that there is an $\epsilon \in (0, \epsilon_3)$ so that $\widehat{\alpha} \cup \widehat{\beta}$ lies in $M_{\geq \epsilon}$.

∎

It follows from lemma 3.4 that there is an embedded ball $B(x, \epsilon)$ of radius ϵ at x in M. Lift the picture to the universal cover $\mathbb{H}^3$ so that x lifts to the origin. Then all covering translations of $B(0, \epsilon)$ are disjoint in $\mathbb{H}^3$. There are $i(\alpha, \beta)$ distinct translates of 0 that are joined to 0 by paths of length bounded by $2L$ that lie in lifts of $\widehat{\alpha}$ and $\widehat{\beta}$ (these paths descend to homotopically distinct closed loops in M based at x). There are thus at least $i(\alpha, \beta)$ disjoint copies $B(0, \epsilon)$ within the ball $B(0, 2L + \epsilon)$ of radius $2L + \epsilon$ of the origin. It follows that there can be at most

$$\frac{\operatorname{vol}(B(0, 2L + \epsilon))}{\operatorname{vol}(B(0, \epsilon))} \tag{3.1}$$

translates of x by $\pi_1(M)$ within the ball $B(0, 2L)$. Setting J equal to the quotient in line 3.1 we have

$$i(\alpha, \beta) < J.$$

∎

Let $X \in \operatorname{Teich}(S)$ be a hyperbolic surface and $\gamma \in \mathcal{S}$. Let

$$\operatorname{short}_\gamma(X) = \left\{ \xi \in \mathcal{S} \;\middle|\; i(\xi, \gamma) > 0 \;\&\; \operatorname{length}_X(\xi) = \min_{\{\nu \in \mathcal{S} \,\mid\, i(\gamma, \nu) > 0\}} \operatorname{length}_X(\nu) \right\}$$

be the "shortest curve(s) on X crossing γ." When $(h \colon Y \to M) \in \mathcal{SH}_k(M)$ is a simplicial hyperbolic surface, define $\operatorname{short}_\gamma(Y)$ similarly. Generically, $\operatorname{short}_\gamma(X)$ will consist of a single isotopy class, but two (or more) isotopy classes can tie to be the shortest.

Lemma 3.5. *Let S be compact surface of negative Euler characteristic. Then for all $\epsilon > 0$ there is a D_ϵ so that if $Y \in \operatorname{Sing}_k(S)$ and $\gamma \in \mathcal{S}$ satisfies $\operatorname{length}_Y(\gamma) > \epsilon$ then any $\alpha \in \operatorname{short}_\gamma(Y)$ satisfies*

$$\operatorname{length}_Y(\alpha) < D_\epsilon.$$

Proof: The general simplicial hyperbolic case ($k \geq 0$) follows from the hyperbolic case ($k = 0$), which we treat first. We differentiate these cases by using $X \in \operatorname{Teich}(S)$ to denote our hyperbolic surface and $Y \in \operatorname{Sing}_k(S)$ to denote our possibly singular hyperbolic surface.

By a theorem of Bers [**Bus**, Thm. 5.2.6] is a constant B_S depending only on S so that any $X \in \operatorname{Teich}(S)$ admits a decomposition into pairs of pants by closed geodesics of length less than B_S. Give X such a decomposition. If $\operatorname{length}_X(\gamma)$ is greater than B_S then it must have non-zero intersection with such a decomposing geodesic. If $\operatorname{length}_X(\gamma) \leq B_S$ then by the collar lemma, there is a w such that the closed metric annular neighborhood $\mathcal{N}_w(\gamma^*)$ of width w on each side with core geodesic γ^* is embedded on X.

Fatten $\mathcal{N}_w(\gamma^*)$ by increasing w to either the width where the two boundary components of $\mathcal{N}_w(\gamma^*)$ become tangent to one another or, the width where one boundary component of $\mathcal{N}_w(\gamma^*)$ becomes self tangent. In the latter case, continue to increase the width of the component of $\mathcal{N}_w(\gamma^*) - \gamma^*$ that is still embedded until its boundary develops a self-tangency or a tangency with the original self-tangent boundary component.

Let w_ϵ denote the larger of the two radii of the sides of the annulus on either side of γ^*. The value of w_ϵ depends only on ϵ and the area of X (which depends only on the genus of S). This process produces a closed annular neighborhood A_γ with embedded interior so that either each boundary component is tangent to itself, or the two boundary components are tangent to each other. Take the shortest route back from the tangency or the two self-tangencies to γ^* and connect by segments of γ^* to form a curve α that has no essential self-intersections. One may check by convexity of $\mathcal{N}_w(\gamma^*)$ that α lies in $\mathbb{S}$ with $i(\alpha, \gamma) \neq 0$. Clearly,

$$\operatorname{length}_Y(\alpha) \leq 4w_\epsilon + 2B_S.$$

Now let $Y \in \operatorname{Sing}_k(S)$ be a singular hyperbolic surface such that $\operatorname{length}_Y(\gamma) > \epsilon$. The induced singular hyperbolic metric on Y has a natural conformal structure. Let g be the hyperbolic metric on Y in the same conformal class as the singular hyperbolic metric s on Y. Since the cone singularities of s have angle at least 2π, a lemma of Ahlfors [**Ah**] guarantees that the metric g is *pointwise* greater than or equal to s.[1] Thus, if $X \in \operatorname{Teich}(S)$ is the surface corresponding to g then we have $\operatorname{length}_X(\gamma) > \epsilon$ as well.

By the above we again have $\operatorname{length}_X(\alpha) \leq 4w_\epsilon + 2B_S$, so by Ahlfors' lemma the length in the singular hyperbolic metric s on Y $\operatorname{length}_Y(\alpha)$ satisfies the same bound. Set $4w_\epsilon + 2B_S$ equal to D_ϵ to prove the theorem.

∎

Uniform twisting. We are now ready to prove the main theorem of this section (a closely related result appears in [**Min4**]). Our particular formulation of the theorem refers to a certain simplicial hyperbolic surface which may at first seem arbitrary. We remark that this surface will be of particular use in the context of algebraically convergent manifolds $M_i \to M$ in $AH(S)$, where it will record the convergent geometry of generators for $\pi_1(M_i)$ (see corollary 3.9).

Theorem 3.6. Uniform Relative Twisting *For any $L > 0$ there is a $K > 0$ depending only on L, S so that the following holds: given $\alpha, \gamma \in \mathbb{S}$ with $i(\alpha, \gamma) \neq 0$, M in $AH(S)$, and a simplicial hyperbolic surface $(h\colon Y \to M) \in \mathcal{SH}_1(M)$ adapted to γ such that $\operatorname{length}_Y(\alpha) < L$, then for any pleated surface $(g, X) \in \mathcal{PS}(M)$ such that $\operatorname{length}_X(\gamma) < L$, any $\alpha_X \in \operatorname{short}_\gamma(X)$ satisfies*

$$|\sigma_\gamma(\alpha_X, \alpha)| < K.$$

Proof: Let $(g, X) \in \mathcal{PS}(M)$ be a pleated surface with $\operatorname{length}_X(\gamma) < L$. The simplicial hyperbolic surface $(h\colon Y \to M)$ realizes γ, so we have $\operatorname{length}_Y(\gamma) = \operatorname{length}_M(\gamma)$, the length of the geodesic representative of γ in M. Thus, α and γ both have length bounded by L on Y. Since $i(\alpha, \gamma) \neq 0$, lemma 3.4 guarantees there is an $\epsilon(L)$, depending only on L, so that $\alpha \cup \gamma$ as they sit on Y lie in $M_{\geq \epsilon(L)/2}$. This means the geodesic representative γ^* of γ in M lies in $M_{\geq \epsilon(L)/2}$ (again, because h realizes γ) and thus $\operatorname{length}_M(\gamma) > \epsilon(L)$.

[1] Thanks to Yair Minsky and Curt McMullen for bringing this to the author's attention.

Let α_X lie in $\text{short}_\gamma(X)$. By lemma 3.5 there is a $D_{\epsilon(L)}$ so that

$$\text{length}_X(\alpha_X) < D_{\epsilon(L)}.$$

By lemma 3.3, there is a J_1 depending only on $\max\{L, D_{\epsilon(L)}\}$ (which depends only on L) such that

$$i(\alpha_X, \gamma) < J_1. \tag{3.2}$$

Form a simplicial hyperbolic surface adapted to γ from α_X as follows: Let x be a basepoint for M on γ^*. Let α'_X and γ' be the images of the geodesic representatives of α_X and γ under the pleated mapping $g \colon X \to M$. Since $\text{length}_M(\gamma) > \epsilon(L)$ there is a homotopy from γ' to γ^* in M all of whose tracks have length bounded by R_L. Let $x' \in \gamma'$ be joined to x by a track of this homotopy. Modify α'_X by a homotopy on X as follows: At each intersection $p \in \alpha'_X \cap \gamma'$, introduce a detour along γ' that takes the shortest route to x' and then returns to p. Using these detours, form a homotopy of α'_X to a curve $\widehat{\alpha_X}$ on X freely homotopic to α_X given as a concatenation of loops based at x' that are disjoint away from x' on X. Extend $\gamma' \cup \widehat{\alpha_X}$ to a triangulation T_X of X with no new vertices. Form a simplicial hyperbolic surface $(\widehat{h} \colon \widehat{X} \to M) \in \mathcal{SH}_1(M)$ by pulling γ' tight to γ^*, pulling T_X tight to a 1-complex $T_{\widehat{X}}$ with vertex at x that contains the closed geodesic γ^*, and using $T_{\widehat{X}}$ as the associated triangulation for $\widehat{h} \colon \widehat{X} \to M$.

We have

$$\ell_X(\widehat{\alpha_X}) < D_{\epsilon(L)} + i(\alpha_X, \gamma) 2L,$$

and thus, since $i(\alpha_X, \gamma) < J_1$ setting $L' = D_{\epsilon(L)} + J_1 2(L + R_L)$ gives

$$\text{length}_{\widehat{X}}(\alpha_X) < L',$$

where L' depends only on L.

Applying theorem 3.2, we have a continuous family $(h_t \colon Y_t \to M)$, $t \in [0,1]$ of simplicial hyperbolic surfaces such that $h_0 = \widehat{h}$ and $h_1 = h$ (from the hypotheses) all of which realize γ^*, i.e. the associated triangulation for h_t has an edge in the isotopy class γ that it maps to γ^*.

Let $\mathcal{B} \subset \mathcal{S}$ be the set of isotopy classes β with $i(\beta, \gamma) \neq 0$ such that

$$\text{length}_{Y_t}(\beta) < D_{\epsilon(L)}$$

for some $t \in [0,1]$. Since γ^* lies in $M_{\geq \epsilon(L)/2}$, and all elements $\beta \in \mathcal{B}$ have a representative intersecting γ^* in M of length less than $D_{\epsilon(L)}$, the constant

$$\mathcal{M}_L = \frac{\text{vol}\, B\left(0, L + D_{\epsilon(L)} + \epsilon(L)/2\right)}{\text{vol}\, B\left(0, \epsilon(L)/2\right)},$$

which depends only on L (cf. lemma 3.3), gives the bound $|\mathcal{B}| < \mathcal{M}_L$, where $|\mathcal{B}|$ denotes the number of elements of $\mathcal{B}$. By lemma 3.5, for each $t \in [0,1]$ there is a $\beta \in \mathcal{B}$ with $\text{length}_{Y_t}(\beta) < D_{\epsilon(L)}$.

Let $U(\beta) \subset [0,1]$ denote the subset

$$U(\beta) = \{t \in [0,1] \mid \text{length}_{Y_t}(\beta) < D_{\epsilon(L)}\}.$$

Since the length of an isotopy class $\beta \in S$ varies continuously on Y_t, the union $\cup_{\beta \in \mathcal{B}} U(\beta)$ gives an open cover of $[0,1]$. If 0 lies in $U(\beta_1)$, $\beta_1 \in \mathcal{B}$, then either $1 \in U(\beta_1)$ or the least upper bound $t(1)$ of $U(\beta_1)$ lies in $U(\beta_2)$ for some $\beta_2 \in \mathcal{B}$, $\beta_2 \neq \beta_1$. Continuing inductively, we find a non-repeating sequence $\{\beta_j\}_{j=1}^k \subset \mathcal{B}$, $k < \mathcal{M}_L$, so that

$$\text{length}_{Y_1}(\beta_k) < D_{\epsilon(L)},$$

and so that if $t(j)$ is the least upper bound of $U(\beta_j)$, $j < k$, then we have,

$$\max\{\text{length}_{Y_{t(j)}}(\beta_j), \text{length}_{Y_{t(j)}}(\beta_{j+1})\} \le D_{\epsilon(L)}. \tag{3.3}$$

Hence, by lemma 3.3 we have the intersection number bound

$$i(\beta_j, \beta_{j+1}) < J_2 \tag{3.4}$$

for $j = 1, \ldots, k-1$, where J_2 depends only on L.

To treat the ends of the sequence, we have

$$\max\{\text{length}_{\widehat{X}}(\beta_1), \text{length}_{\widehat{X}}(\alpha_X)\} < L'$$

on $\widehat{X}$, while on Y we have

$$\max\{\text{length}_Y(\beta_k), \text{length}_Y(\alpha)\} < \max\{L, D_{\epsilon(L)}\}.$$

Thus by lemma 3.3 there are constants, J_3 and J_4 depending only on L so that

$$i(\alpha_X, \beta_1) < J_3 \quad \text{and} \quad i(\beta_k, \alpha) < J_4.$$

Let $J = \max_{i=1,\ldots,4} J_i$.

From lemma 3.1 (**RT1**) it follows that

$$|\sigma_\gamma(\beta_j, \beta_{j+1})| < J + 1$$

for each $j = 1, \ldots, k-1$ and moreover that

$$|\sigma_\gamma(\alpha_X, \beta_1)| < J + 1 \quad \text{and} \quad |\sigma_\gamma(\beta_k, \alpha)| < J + 1.$$

By repeated applications of lemma 3.1 (**RT2**) we have

$$\sigma_\gamma(\alpha_X, \alpha) \quad \asymp_{5k} \quad \sigma_\gamma(\alpha_X, \beta_1) + \sum_{j=1}^{k-1} \sigma_\gamma(\beta_j, \beta_{j+1}) + \sigma_\gamma(\beta_k, \alpha),$$

so by the triangle inequality

$$|\sigma_\gamma(\alpha_X, \alpha)| \quad \le \quad |\sigma_\gamma(\alpha_X, \beta_1)| + \sum_{j=1}^{k-1} |\sigma_\gamma(\beta_j, \beta_{j+1})| + |\sigma_\gamma(\beta_k, \alpha)| + 5k.$$

The bound

$$|\sigma_\gamma(\alpha_X, \alpha)| < (J+1)(k+1) + 5k$$

follows. Since J and k depend only on L, setting $K = (J+1)(k+1) + 5k$ proves the theorem.

∎

We wish to use the metric independent quantity $\sigma_\gamma(\alpha, \beta)$ to control a metric dependent quantity: namely, when simple closed geodesics c and γ cross on $X \in \text{Teich}(S)$, we want to measure the amount c winds around a small metric annular neighborhood of γ.

When a metric $X \in \text{Teich}(S)$ is fixed, the annular cover X_γ corresponding to γ admits a foliation $\mathcal{F}(\gamma)$ by complete geodesics perpendicular to the core geodesic γ. We may compute the relative twisting of c with respect to this foliation

$$\sigma_\gamma^X(c, \mathcal{F}(\gamma))$$

by counting the number of intersections of a leaf l of $\mathcal{F}(\gamma)$ and a lift $\widetilde{c}$ of c to X_γ that intersect least, reckoned positively if $\widetilde{c}$ winds to the right about X_γ and negatively if $\widetilde{c}$ winds to the left. Given ϵ for which the ϵ-collar $\mathcal{N}_\epsilon(\gamma)$ embeds in X, we define the *winding* of c about $\mathcal{N}_\epsilon(\gamma)$

$$\mathcal{W}_{(\epsilon,X)}(c,\gamma) \in \mathbb{Z}$$

by only counting the contributions to $\sigma^X_\gamma(c, \mathcal{F}(\gamma))$ that come from intersections within $\mathcal{N}_\epsilon(\gamma)$ on X_γ. More generally, if $a \subset \mathcal{N}_\epsilon(\gamma)$ is an arc joining boundary components of $\mathcal{N}_\epsilon(\gamma)$, the winding $\mathcal{W}_{(\epsilon,X)}(.,\gamma)$ is defined on the homotopy class $[a]$ of a *rel*-boundary by taking the geodesic representative a^* of a *rel*-endpoints on X, and counting the minimal number of intersections with any leaf of $\mathcal{F}(\gamma)$.

Let $\text{length}_X(\gamma)$ lie in the interval $[\delta, L]$, where $\delta > 0$. The collar lemma (see [**Bus**, Thm. 4.4.6]) guarantees that there is an $\epsilon_L > 0$ so that given any collection $\gamma_1, \ldots, \gamma_p$ of pairwise disjoint simple closed geodesics on X, and $\epsilon < \epsilon_L$ the ϵ-collars $\mathcal{N}_\epsilon(\gamma_1), \ldots, \mathcal{N}_\epsilon(\gamma_p)$ embed disjointly in X. Then we have the following.

LEMMA 3.7. *For any $\epsilon < \epsilon_L$, there is a constant $K(\epsilon, \delta)$ so that if $\alpha_X \in \text{short}_\gamma(X)$, we have*

$$\mathcal{W}_{(\epsilon,X)}(c,\gamma) \asymp_{K(\epsilon,\delta)} \sigma_\gamma(c, \alpha_X).$$

Proof: We first remark that $\mathcal{W}_{(\epsilon,X)}(\alpha_X, \gamma) = 0$ since otherwise every arc of $\alpha_X \cap \mathcal{N}_\epsilon(\gamma)$ winds all the way around the annulus $\mathcal{N}_\epsilon(\gamma)$, and performing a Dehn twist on $\mathcal{N}_\epsilon(\gamma)$ shortens α_X.

To compare the winding of a general curve c about $\mathcal{N}_\epsilon(\gamma)$ with the winding of α_X about $\mathcal{N}_\epsilon(\gamma)$ we apply the following lemma:

LEMMA 3.8. *There is a constant $K_1(\epsilon, \delta) > 0$ so that given any pair of complete geodesics l_1 and l_2 on X_γ, the number of intersections of l_1 and l_2 outside the ϵ-neighborhood $\mathcal{N}_\epsilon(\gamma)$ on X_γ is bounded by $K_1(\epsilon, \delta)$.*

Proof: This is seen by lifting to the universal cover $\mathbb{H}^2$ and considering the number of γ-translates of a lift $\widetilde{l_1}$ of l_1 that can intersect a lift $\widetilde{l_2}$ within $\mathbb{H}^2 - \mathcal{N}_\epsilon(\widetilde{\gamma})$: let Q be the ideal quadrilateral obtained as the convex hull of the ideal endpoints of $\widetilde{\gamma}$ and $\widetilde{l_2}$ on $\mathbb{H}^2$. If $\widetilde{l_1}$ does not cross $\widetilde{\gamma}$ and is not asymptotic to $\widetilde{\gamma}$ then at most two γ-translates of $\widetilde{l_1}$ cross $\widetilde{l_2}$ on $\mathbb{H}^2$.

If on the other hand l_1 either crosses $\widetilde{\gamma}$ or is asymptotic to $\widetilde{\gamma}$, then for each k such that $\gamma^k(\widetilde{l_1})$ crosses $\widetilde{l_2}$ outside of $\mathcal{N}_\epsilon(\widetilde{\gamma})$, the translate $\gamma^k(\widetilde{l_1})$ intersects $\partial\mathcal{N}_\epsilon(\widetilde{\gamma})$ within Q. (see figure 4). The total arc length of the intersection of $\partial\mathcal{N}_\epsilon(\widetilde{\gamma})$ with the interior of Q is bounded by a constant $b_\epsilon > 0$ depending only on ϵ. The translation distance of γ on $\partial\mathcal{N}_\epsilon(\widetilde{\gamma})$ is fixed and greater than δ. Thus, the number of intersections of translates of l_1 by γ that intersect l_2 is bounded by b_ϵ/δ which we set equal to $K_1(\epsilon, \delta)$.

■

Proof of lemma 3.7 continued: It follows from lemma 3.8 that all but at most $K_1(\epsilon, \delta)$ intersections of a lift $\widetilde{c}$ and a leaf l of $\mathcal{F}(\gamma)$ occur within $\mathcal{N}_\epsilon(\gamma)$ on X_γ. Thus, we have that

$$\mathcal{W}_{(\epsilon,X)}(c,\gamma) \asymp_{K_1(\epsilon,\delta)} \sigma^X_\gamma(c, \mathcal{F}(\gamma)),$$

and

$$\sigma^X_\gamma(\mathcal{F}(\gamma), \alpha_X) \asymp_{K_1(\epsilon,\delta)} \mathcal{W}_{(\epsilon,X)}(\alpha_X, \gamma) = 0.$$

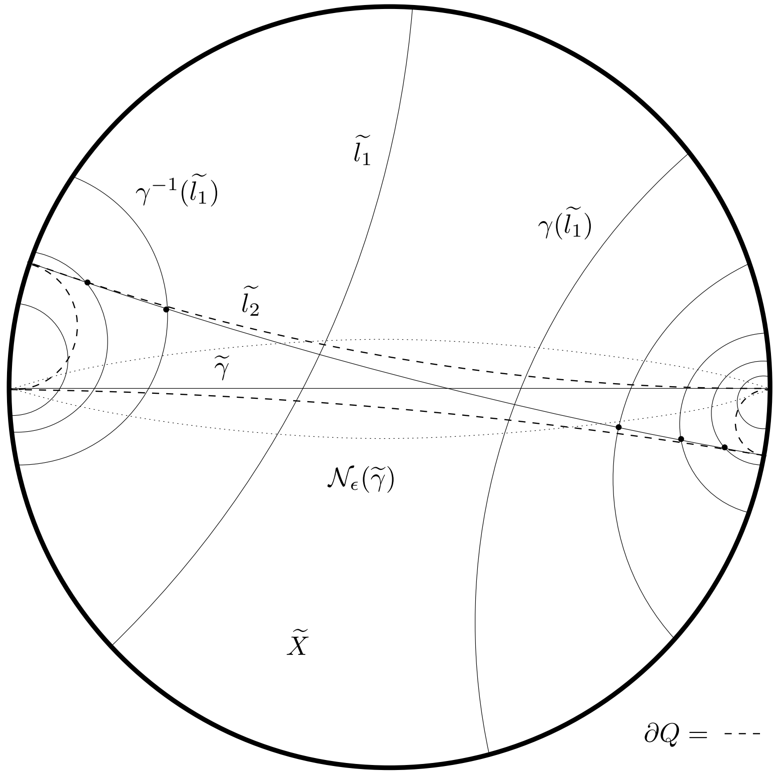

Figure 4. Contributions to relative twisting outside $\mathcal{N}_\epsilon(\gamma)$ are uniformly bounded in terms of ϵ and the length of γ.

From lemma 3.1, we have

$$\sigma^X_\gamma(c, \mathcal{F}(\gamma)) + \sigma^X_\gamma(\mathcal{F}(\gamma), \alpha_X) \asymp_5 \sigma_\gamma(c, \alpha_X).$$

It follows that

$$\mathcal{W}_{(\epsilon,X)}(c,\gamma) \asymp_{2K_1(\epsilon,\delta)+5} \sigma_\gamma(c, \alpha_X).$$

The lemma follows setting $K(\epsilon,\delta) = 2K_1(\epsilon,\delta) + 5$.

∎

We are now ready to apply theorem 3.6 to the setting of algebraically convergent manifolds $M_i \to M$. Recall that we denote by $\mathfrak{R} \subset AH(S) \times \mathcal{ML}(S)$ the set of pairs (M,μ) for which μ is realizable in M.

Corollary 3.9. *Let (M_i,γ) converge to (M,γ) in $\mathfrak{R}$. Let $(g_i, X_i) \in \mathcal{PS}(M_i)$ and $(g, X) \in \mathcal{PS}(M)$ be pleated surfaces all realizing γ, and let $\epsilon > 0$ be sufficiently small so that the metric annular neighborhood $\mathcal{N}_\epsilon(\gamma)$ of γ on X_i is embedded for each i. Then there is a constant $A_\epsilon > 0$ so that*

$$\mathcal{W}_{(\epsilon,X)}(c,\gamma) \asymp_{A_\epsilon} \mathcal{W}_{(\epsilon,X_i)}(c,\gamma)$$

for any $c \in \mathcal{S}$.

Proof: Since $(M_i,\gamma) \to (M,\gamma)$ in $\mathfrak{R}$ there exists $L > \delta > 0$ so that $\text{length}_{M_i}(\gamma)$ lies in the interval $[\delta, L]$ for all i. Taking $\epsilon < \epsilon_L$ ensures that the ϵ-collar $\mathcal{N}_\epsilon(\gamma)$ is embedded on X_i. Let $(\gamma^*)_i$ denote the geodesic representative of γ in M_i, it follows that there is a $\delta' > 0$ depending only on L and δ so that $(\gamma^*)_i$ always lies within $(M_i)_{\geq \delta'}$. Thus we may choose baseframes $\omega_i \in M_i$ on the geodesic representatives $(\gamma^*)_i$ of γ in M_i that determine holonomy representations $\rho_i \colon \pi_1(S) \to \mathrm{Isom}^+(\mathbb{H}^3)$ so that $\rho_i(g)$ converges for each $g \in \pi_1(S)$. As we shall apply theorem 3.6 later, we now choose a reference curve $\alpha \in \mathcal{S}$ so that $i(\gamma, \alpha) \neq 0$.

Realize the isotopy classes of α and γ on X as geodesics with the same names. Let $p \in \gamma$ denote a basepoint for $\pi_1(S)$. As in the proof of theorem 3.6, form a based representative $\widehat{\alpha}$ of α at p once and for all by finding the shortest route along γ from each intersection point $x \in \alpha \cap \gamma$ to p, homotoping α so that all intersections of α and γ occur at p and so that $\alpha - p$ is embedded. Extend $\gamma \cup p \cup \widehat{\alpha}$ to a triangulation T of X with no new vertices, and pull T back to S via the marking on X.

There are simplicial hyperbolic surfaces $(h_i \colon Y_i \to M_i) \in \mathcal{SH}_2(M_i)$ so that

1. each h_i is adapted to γ with associated triangulation T.
2. the vertex $p \in T$ maps to $|\omega_i|$.

If $\widehat{\alpha}$ is represented in $\pi_1(S, p)$ as the composition of edges $g_1 \circ \ldots \circ g_a$ of T, then by algebraic convergence ρ_i converges on each g_b, $b = 1, \ldots, a$, so the translation distance of $\rho_i(g_b)$ at the origin $\widetilde{\omega} \in \mathbb{H}^3$ converges. There is thus an $L' > 0$ so that $\ell_{M_i}(\widehat{\alpha})$ is bounded by L' for all i. Moreover, the length of α on the simplicial hyperbolic surface Y_i remains uniformly bounded by L' throughout the sequence. Let $L_0 = \max\{L, L'\}$.

For each i, the data γ, α, Y_i and M_i together satisfy the hypotheses of theorem 3.6 with length bound L_0, so there is a uniform K (depending only on L_0 (and S)), so that for any sequence of pleated surfaces $(g_i, X_i) \in \mathcal{PS}(M_i)$ realizing γ and curves $\alpha_i \in \mathrm{short}_\gamma(X_i)$, we have the bound

$$|\sigma_\gamma(\alpha_i, \alpha)| < K \tag{3.5}$$

on the relative twisting of α_i and α with respect to γ.

By lemma 3.7 we have

$$\mathcal{W}_{(\epsilon, X_i)}(c, \gamma) \asymp_{K(\epsilon,\delta)} \sigma_\gamma(\alpha_i, c)$$

and

$$\mathcal{W}_{(\epsilon, X)}(c, \gamma) \asymp_{K(\epsilon,\delta)} \sigma_\gamma(\alpha, c)$$

so letting $A_\epsilon = 2K(\epsilon, \delta) + K + 5$ we have

$$\mathcal{W}_{(\epsilon, X)}(c, \gamma) \asymp_{A_\epsilon} \mathcal{W}_{(\epsilon, X_i)}(c, \gamma).$$

∎

Let λ be a geodesic lamination containing a simple closed curve γ. If λ is a limit of simple closed curves $c_i \neq \gamma$, then leaves of λ spiral either to the right or to the left as they approach γ. We say λ is *to the right* of γ or *to the left* of γ respectively.

COROLLARY 3.10. *Let M_i, M, γ, X_i, X, and ϵ be as above. Let $c_i \in \mathcal{S}$, $c_i \neq \gamma$, be simple closed curves converging to a lamination λ in $\mathcal{GL}(S)$ containing γ as a simple closed leaf. Then $\mathcal{W}_{(\epsilon,X)}(c_i, \gamma) \to +\infty$ (resp. $-\infty$) if λ is to the right (resp. left) of γ, and*

$$\lim_{i \to \infty} \frac{\mathcal{W}_{(\epsilon, X_i)}(c_i, \gamma)}{\mathcal{W}_{(\epsilon, X)}(c_i, \gamma)} = 1.$$

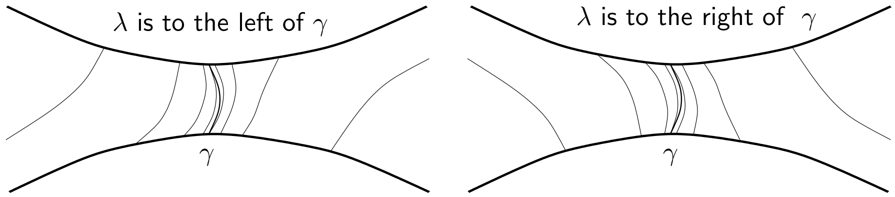

Figure 5. Geodesic laminations to the left and to the right of γ.

Proof: For any $N > 0$ there is an $r > 0$ so that any c_i crossing γ that lies within the r-neighborhood $\mathcal{N}_r(\lambda)$ on X crosses every leaf in the foliation of $\mathcal{N}_\epsilon(\gamma)$ by geodesic arcs orthogonal to γ at least N times. Moreover, any such c_i winds to the right if λ is to the right of γ and to the left if λ is to the left of γ. By Hausdorff convergence, for i sufficiently large c_i lies in $\mathcal{N}_r(\lambda)$ and c_i crosses γ. It follows that $\mathcal{W}_{(\epsilon,X)}(c_i,\gamma)$ tends to $+\infty$ if λ is to the right of γ and $\mathcal{W}_{(\epsilon,X)}(c_i,\gamma)$ tends to $-\infty$ if λ is to the left of γ.

Applying corollary 3.9, we have $\mathcal{W}_{(\epsilon,X)}(c_i,\gamma) \asymp_{A_\epsilon} \mathcal{W}_{(\epsilon,X_i)}(c_i,\gamma)$ for each i; it follows that

$$\lim_{i\to\infty} \frac{\mathcal{W}_{(\epsilon,X_i)}(c_i,\gamma)}{\mathcal{W}_{(\epsilon,X)}(c_i,\gamma)} = 1.$$

∎

4. Constructing nearly straight train tracks

In the next two sections, we develop geometric estimates on *nearly-straight train tracks* which we will use to control the lengths of laminations in varying families of hyperbolic 3-manifolds. This section develops and formalizes basic elements of constructing nearly-straight train tracks. The reader may wish first to peruse sections 6 and 7, in which the main theorems are proven, to motivate the constructions of the next two sections.

Definition 4.1. *Let $\epsilon > 0$ be less than 1. A train track τ in a hyperbolic manifold M^n, $n = 2, 3$, is* ϵ-nearly-straight *if any train path $r\colon \mathbb{R} \to \tau$ lifts to a C^2 embedding $\widetilde{r}\colon \mathbb{R} \to \widetilde{M^n}$ with geodesic curvature less than ϵ.*

A train track $\tau \subset X$ has an *ϵ-nearly-straight realization* τ^* in M^n if τ is equivalent to an ϵ-nearly-straight train track τ^* in M^n.

Features of nearly straight train tracks. From standard hyperbolic geometry, an ϵ-nearly-straight train track τ^* in hyperbolic n-manifold M^n has two important features:

1. For any $\delta_0 < 1$ there is a *tracking constant* $C > 1$ so that for each $\epsilon < \delta_0$, any train path $r\colon \mathbb{R} \to \tau^*$ or $r\colon S^1 \to \tau^*$ lifts to a train path $\widetilde{r}\colon \mathbb{R} \to \mathbf{P}\mathbb{H}^n$, that is smoothly isotopic to a geodesic by an isotopy that moves each point a distance less than $C_{\mathrm{tr}}\epsilon$ in $\mathbf{P}\mathbb{H}^n$. Taking $\delta_0 = 1/2$, let C_{tr} be the corresponding tracking constant.
2. There is a continuous *contraction bound*

$$K\colon [0,1) \to [1,\infty)$$

so that if $\epsilon \in [0,1)$ and $\mu \in \mathcal{ML}(S)$ is carried by τ^*, we have

$$\text{length}_{M^n}(\mu) \geq \frac{1}{K(\epsilon)} \ell_{\tau^*}(\mu)$$

and $K(\epsilon)$ tends to 1 as ϵ tends to 0.

One criterion to determine whether a given train track τ has an ϵ-nearly straight realization is to form its associated *train track graph* $\widehat{\tau}$ by straightening its branches to geodesics *rel*-switches. We call the *external angles* of $\widehat{\tau}$ the angles that can occur as the external angles of the lift to the universal cover of the image of a train-path under this straightening. For any definite length $\ell > 0$ there is a $\delta_\ell > 0$ and a constant $C_{\text{curv}}(\ell)$ so that if the edges of $\widehat{\tau}$ have length at least ℓ, and all external angles are bounded by $\epsilon \in (0, \delta_\ell)$, then τ admits a $(C_{\text{curv}}(\ell)\epsilon)$-nearly-straight realization. We call such a train track graph (ℓ, ϵ)-*nearly-straight* (cf. [**Min1**] [**CEG**, Thm. 4.2.10]).

Horocyclic spike-foliations. We now give a direct construction of nearly straight train tracks carrying a lamination μ. The idea is this: given a lamination on a surface X, we construct a partial foliation of a neighborhood of μ on the surface. The leaves of the lamination are like "train-routes" running along transverse to the leaves of this foliation, which are like "ties" of the train track. Collapsing the leaves of the foliation, we have a graph τ on X that represents our train track. Each leaf of μ naturally determines a train-path on τ, so τ carries μ.

We will develop a substantial amount of geometric control on such foliations in order to control the straightness of the resulting train tracks in the above sense. The straightness will depend on the metric on X, which will later vary. Uniformity will come from

- requiring that $|\mu|$ lie in the thick part $X_{\geq \epsilon_0}$, and
- requiring for any compact leaf γ of $|\mu|$ that $\text{length}_X(\gamma) < L$.

We will introduce various geometric constants in what follows, all of which will depend at most on S, ϵ_0, and L.

To start, let $\lambda \in \mathcal{GL}(S)$ be a geodesic lamination on X with no compact leaves. The completion $\overline{X - \lambda}$ of the path metric on $X - \lambda$ has a finite number of ends where *frontier leaves* of λ, namely leaves in $\partial(\overline{X - \lambda})$, are asymptotic. A *spike* s of $X - \lambda$ is an open neighborhood of such an end that lies to one side of a horocyclic arc α orthogonal to the asymptotic leaves of λ that determine the end. The horocyclic arc α is called the *frontier horocycle* of the spike s, and its arclength $\ell_X(\alpha)$ is called the *width* of the spike s. The number of disjoint spikes is bounded by a constant C_{spike} depending only on S.

Any collection of spikes of $X - \lambda$ with widths less than 1 are pairwise disjoint. Each spike s of $X - \lambda$ admits a foliation by parallel horocycles limiting to the frontier horocycle for s.

Let $s_1, \ldots, s_j$ be spikes of $\overline{X - \lambda}$, one for each end of $\overline{X - \lambda}$, with widths $\epsilon_1, \ldots, \epsilon_j$ all less than 1. Following Thurston [**Th1**, §8.9], the horocyclic foliations of the spikes $s_1, \ldots, s_j$ naturally extend to a foliation $\mathcal{F}_{\text{int}}$ of the interior $\text{int}(\overline{s_1 \cup \ldots \cup s_j})$ on X obtained by extending the union of the horocyclic foliations across λ: the tangent line field along leaves of the foliations of each spike extends continuously to a Lipschitz line field on $\text{int}(\overline{s_1 \cup \ldots \cup s_j})$, which is integrable, and thus $\mathcal{F}_{\text{int}}$ is well defined (see [**Th5**, §3] [**Bon3**, pp. 244]). Denote by $|\mathcal{F}_{\text{int}}| = \text{int}(\overline{s_1 \cup \ldots \cup s_j})$ the support of $\mathcal{F}_{\text{int}}$.

The foliation $\mathcal{F}_{\rm int}$ extends to a decomposition $\mathcal{F}$ of the closure $\overline{s_1 \cup \ldots \cup s_j}$ into *"leaves"* which are extensions of the leaves of $\mathcal{F}_{\rm int}$: If $x \in \partial|\mathcal{F}_{\rm int}|$ lies in a frontier leaf ℓ of λ, we call x an *endpoint* of the leaf l of $\mathcal{F}_{\rm int}$ whose tangent line field limits to the orthogonal line to x at ℓ. Let $\hat{l}$ denote the union of l with its endpoints.[2] Then we obtain the elements of the decomposition $\mathcal{F}$ with the following two steps:

1. If the endpoints of l do not meet the boundary $\partial\alpha$ of any frontier horocycle α, then $\hat{l}$ is a *generic leaf* of $\mathcal{F}$.
2. Each frontier horocycle α_k determines a pair $\hat{l}_k$ and $\hat{l}'_k$ with endpoints at $\partial\alpha_k$. Each connected component of $\cup_k(\overline{\alpha_k} \cup \hat{l}_k \cup \hat{l}'_k)$ is a leaf of $\mathcal{F}$, which we call a *switch leaf* (figure 6).

We call the resulting decomposition of $\overline{s_1 \cup \ldots \cup s_j}$, into generic leaves and switch leaves the *horocyclic spike-foliation* $\mathcal{F}$ for λ. The decomposition $\mathcal{F}$ is almost a foliation of $\overline{s_1 \cup \ldots \cup s_j}$: it is a foliation away from its "corners" at the finite set of endpoints of the frontier horocycles (see figure 6). We denote by $|\mathcal{F}|$ the closed

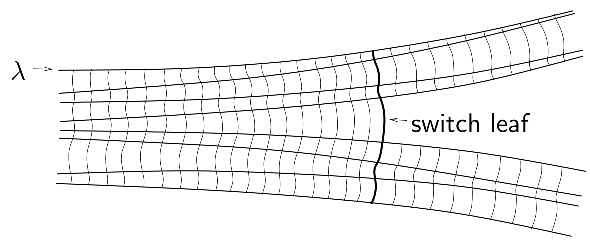

Figure 6. A horocyclic spike-foliation for λ.

set $\overline{s_1 \cup \ldots \cup s_j}$ on X which we call the *support* of the horocyclic spike-foliation $\mathcal{F}$.

We define moves on horocyclic spike-foliations by adjusting the widths of the spikes. Given $r \in \mathbb{R}$, we define the *r-adjustment of $\mathcal{F}$ on the spike s_k* to be the spike-foliation obtained by adjusting the width ϵ_k of the spike s_k to be $e^r\epsilon_k$. The r-adjustment is always well-defined, or *allowable*, provided we constrain r so that $e^r\epsilon_k < 1$. Under an r-adjustment on a spike s_k, the frontier horocycle α_k moves a signed distance r outward along the frontier leaves of λ at its endpoints.

Each switch leaf l of $\mathcal{F}$ contains a frontier horocycle of $\mathcal{F}$ in some spike s (see figure 6). A *transverse branch* b is a closed geodesic segment of a leaf ℓ of λ so that ∂b lies in switch leaves of $\mathcal{F}$ and the interior of b does not intersect any switch leaves of $\mathcal{F}$. Two switch leaves l_1 and l_2 of $\mathcal{F}$ are *adjacent* when they are joined by a transverse branch b. Any two transverse branches joining l_1 and l_2 that meet the same leaves of $\mathcal{F}$ have the same length.

Let $\mathcal{F}_\epsilon(\lambda)$ denote the horocyclic spike-foliation for λ on X with all widths equal to ϵ. A horocyclic spike-foliation is *generic* if each switch leaf contains exactly one frontier horocycle. Since $\mathcal{F}_\epsilon(\lambda)$ fails to be generic for only countably many ϵ, we will always assume ϵ is chosen so that $\mathcal{F}_\epsilon(\lambda)$ is generic. Two generic horocyclic spike-foliations $\mathcal{F}$ and $\mathcal{F}'$ are *equivalent* if there is a family of generic spike-foliations interpolating between them obtained by continuously varying the widths of $\mathcal{F}$.

Train tracks from spike-foliations. A horocyclic spike-foliation $\mathcal{F}$ for λ determines a train track $\lambda/\mathcal{F}$ carrying λ: specifically, $\mathcal{F}$ determines a differentiable map

[2] A priori, leaves of $\mathcal{F}_{\rm int}$ may be half-infinite or bi-infinite; we show each leaf of $\mathcal{F}$ (and hence $\mathcal{F}_{\rm int}$) has finite length in lemma 4.2.

$p\colon X \to X$, homotopic to the identity and nonsingular on the tangent spaces of leaves of λ, so that p collapses each leaf l of $\mathcal{F}$ to a point on X (see [**Th1**, §8.9]). The image $p(\lambda)$ is a train track τ; the image $p(b)$ of a transverse branch b of $\mathcal{F}$ is a branch of τ, and the image $p(l)$ of each switch leaf l of $\mathcal{F}$ is a switch v of τ. The train track τ minimally carries λ (see figure 7, which depicts a simple collapsing. Note that a leaf of $\mathcal{F}$ may intersect λ in a Cantor set).

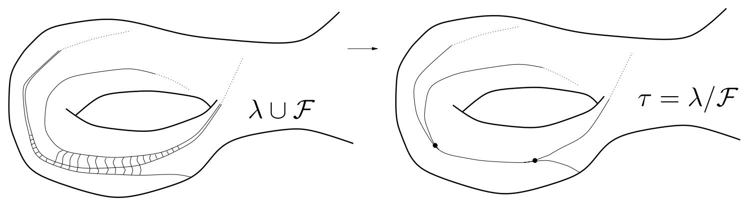

Figure 7. Collapsing horocyclic spike-foliations.

Here is a more direct construction the train track $\tau = \lambda/\mathcal{F}$: consider a collection of transverse branches so that each leaf of $\mathcal{F}$ that is not a switch leaf intersects precisely one transverse branch. Then, assuming $\mathcal{F}$ is generic, each switch leaf contains exactly three endpoints of transverse branches. Given a switch leaf l, choose a vertex v_l on l. Given a transverse branch b with an endpoint x in l, modify b by a smooth isotopy through arcs meeting l orthogonally that pulls x along l to the vertex v_l and fixes a neighborhood of the other endpoint of b. The resulting arc meets l at a right angle at v_l. Performing similar isotopies at each end of each transverse branch of $\mathcal{F}$ in such a way that the interiors of the branches remain pairwise disjoint, we obtain the branches of a train track τ on X with switches v_l for each switch leaf l of $\mathcal{F}$ (see figure 8). Equivalent spike-foliations

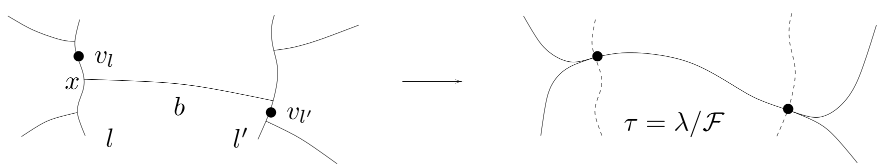

Figure 8. Building a train track out of transverse branches and switch leaves of $\mathcal{F}$.

for λ collapse to equivalent (indeed, smoothly ambient isotopic) train tracks.

We will now concern ourselves primarily with the case $\lambda = |\mu|$ where $\mu \in \mathcal{ML}(S)$ is a measured lamination with no compact leaves. Using information about $\mathcal{F}$, such as the lengths of its switch leaves and transverse branches, we can impose some geometric control on the resulting train-track $\tau = |\mu|/\mathcal{F}_\epsilon(\mu)$. As one expects, however, geometric quantities associated to $\mathcal{F}$ on X (and thus to $\tau = |\mu|/\mathcal{F}$) depend strongly on X. We will control this dependence by bounding from below the injectivity radius along the leaves of $|\mu|$.

Lemma 4.2. Short Leaves *Given $\epsilon_0 > 0$ and the surface S, there are constants $C_{\rm sw} > 1$ and $\epsilon_{\rm sw} > 0$ so that the following holds: if $\mu \in \mathcal{ML}(S)$ is a measured lamination with no compact leaves, $X \in \mathrm{Teich}(S)$ is a hyperbolic surface with $|\mu| \subset X_{\geq \epsilon_0}$, and $\mathcal{F}$ is a horocyclic spike-foliation for μ on X whose largest width is $\epsilon < \epsilon_{\rm sw}$, then any leaf l of $\mathcal{F}$ satisfies $\ell_X(l) < C_{\rm sw}\epsilon$.*

Proof: We begin by noting that since $|\mu|$ has 2-dimensional Lebesgue measure 0 on X [**Th2**, Prop. 5.3], the area of $|\mathcal{F}|$ is the sum of the areas of the spikes for $\mathcal{F}$. Since the area of a spike is equal to its width, we have:

$$\operatorname{area}_X(|\mathcal{F}|) < C_{\text{spike}}\epsilon.$$

Moreover, the length $\ell_X(l)$ is the sum $\sum_i \ell_X(l_i)$ over the (countably many) components l_i of the intersection $l \cap (X - |\mu|)$. Each component l_i cuts off a sub-spike σ_i of the spike s_{k_i} of $X - |\mu|$ to one side, $k_i = 1, \dots j$.

First of all, let $\epsilon_0' < 1$ be such that $\mathcal{N}_1(X_{\geq \epsilon_0})$ lies within $X_{\geq \epsilon_0'}$. Then we have $|\mathcal{F}| \subset X_{\geq \epsilon_0'}$ for $\epsilon < 1$. Lifting l to $\widetilde{l}$ in the universal cover $\mathbb{H}^2$, we consider any interval I of points in $\widetilde{l}$. At each point $y \in I \cap \widetilde{l}_i$ consider the geodesic arc g_y of length $\epsilon_0'/2$ orthogonal to $\widetilde{l}$ traveling from y into the lifted sub-spike $\widetilde{\sigma}_i$ cut off by $\widetilde{l}_i$. Let

$$A(I) = \overline{\{g_y | y \in I \cap (\cup_i \widetilde{l}_i)\}}$$

denote closure of the union of all these geodesic segments, and note that

$$\operatorname{area}_{\mathbb{H}^2} A(I) = \ell_{\mathbb{H}^2}(I)(1 - e^{-\epsilon_0'/2}).$$

Assume $\ell_X(l) > \epsilon_0'$ and let $\widetilde{x} \in \widetilde{l}$ be the center of an interval I_0 of length ϵ_0' in $\widetilde{l}$ (measured along $\widetilde{l}$). Each point in $A(I_0)$ lies within ϵ_0' of $\widetilde{x}$, so $A(I_0)$ embeds in the covering projection $\pi \colon \mathbb{H}^2 \to X$. Thus we have

$$\operatorname{area}_X(\pi(A(I_0))) = \ell_X(\pi(I_0))(1 - e^{-\epsilon_0'/2}) = \epsilon_0'(1 - e^{-\epsilon_0'/2}) < C_{\text{spike}}\epsilon.$$

Therefore, if

$$\epsilon < \frac{\epsilon_0'(1 - e^{-\epsilon_0'/2})}{C_{\text{spike}}} = \epsilon_{\text{sw}},$$

the above ensures that $\ell_X(l) \leq \epsilon_0'$.

Assume, then, that $\epsilon < \epsilon_{\text{sw}}$. If $x \in l$ is the midpoint of l (with respect to distance along l) then $A(\widetilde{l})$ again embeds under the covering projection to X so we have

$$\operatorname{area}_X(\pi(A(\widetilde{l}))) = \ell_X(l)(1 - e^{-\epsilon_0'/2}) < C_{\text{spike}}\epsilon.$$

Setting

$$C_{\text{sw}} = \frac{C_{\text{spike}}}{1 - e^{-\epsilon_0'/2}}$$

proves the lemma.

∎

While we have no guarantee that taking ϵ sufficiently small provides uniform separation of switch leaves of $\mathcal{F}_\epsilon(\mu)$, we will in practice be able to modify $\mathcal{F}_\epsilon(\mu)$ by uniformly bounded adjustments in spikes to obtain uniform separation. This is the import of the following lemma:

LEMMA 4.3. DEFINITE BRANCH LENGTHS *Given $\epsilon_0 > 0$ and the surface S, there are constants $\ell_0 > 0$, $C_{\text{adj}} > 0$, and $\epsilon_{\text{adj}} > 0$ so that the following holds: if $\mu \in \mathcal{ML}(S)$ is a lamination with no compact leaves, and $X \in \operatorname{Teich}(S)$ is a hyperbolic surface for which $|\mu|$ lies in $X_{\geq \epsilon_0}$, then for any positive $\epsilon < \epsilon_{\text{adj}}$ the horocyclic spike-foliation $\mathcal{F}_\epsilon(\mu)$ is equivalent by r-adjustments in spikes, $|r| < C_{\text{adj}}$, to a horocyclic spike-foliation $\mathcal{F}'$ for which:*

1. *each transverse branch has length at least ℓ_0, and*
2. *any leaf l of $\mathcal{F}'$ satisfies $\ell_X(l) < e^{C_{\text{adj}}} C_{\text{sw}}\epsilon$.*

Proof: We first argue that any train-path for $\tau = |\mu|/\mathcal{F}_\epsilon(\mu)$ that is homotopic to an embedding is homotopically essential in X.

To see this, add finitely many complete geodesics to the support of μ to obtain a *complete* geodesic lamination $|\mu|'$: the complementary regions of $|\mu|'$ are all ideal hyperbolic triangles embedded in X (see [**Th1**, Ch. 8]). The horocyclic spike-foliation $\widehat{\mathcal{F}} = \mathcal{F}_1(|\mu|')$ (each spike has width 1) restricts to a foliation of $\mathrm{int}(|\widehat{\mathcal{F}}|)$, which is all of X but closed triangular regions bounded by horocycles in each complementary ideal triangle of $X - |\mu|'$ (see figure 9). Each leaf of $\mathcal{F}_\epsilon(\mu)$ or any

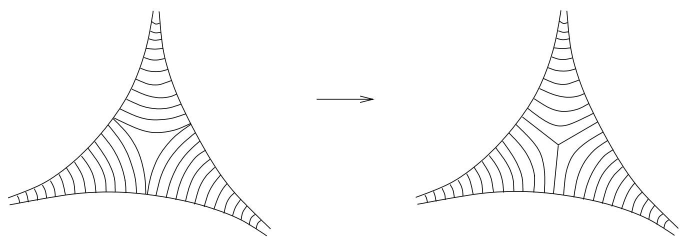

Figure 9. Collapsing a maximal horocyclic spike-foliation to a singular foliation.

allowable adjustment of $\mathcal{F}_\epsilon(\mu)$ is contained in a leaf of $\widehat{\mathcal{F}}$.

There is a natural collapsing map for the triangular components of $X - |\widehat{\mathcal{F}}|$ obtained by identifying each point x in such a region with the point(s) in the frontier horocycles at minimal distance from x. The image of $\widehat{\mathcal{F}}$ under this collapsing is a singular foliation $\widehat{\mathcal{F}}_{\mathrm{sing}}$ of X with one 3-pronged singularity at the center of each ideal triangle in $X - |\mu|'$ (see figure 9).

Any train-path on τ that is homotopic to an embedded simple closed curve in X is homotopic to an embedded loop $p\colon S^1 \to X$ that avoids the singularities of the foliation $\widehat{\mathcal{F}}_{\mathrm{sing}}$ and is transverse to $\widehat{\mathcal{F}}_{\mathrm{sing}}$. If p bounds an embedded disk D in X, then $\widehat{\mathcal{F}}_{\mathrm{sing}}$ restricts to a singular foliation transverse to ∂D. By the Poincaré index formula [**Th3**, Prop. 1.3.10], the Euler characteristic of D is given by the sum of the indices of the singularities of $\widehat{\mathcal{F}}_{\mathrm{sing}}$ inside of D. Since 3-prong singularities have negative index (-1/2), and the Euler characteristic of D is positive we have a contradiction. We conclude that p is homotopically essential.

Let C_{br} be the upper bound on the number of branches to τ (which depends only on S). As before, choose $\epsilon_0' < 1$ so that $\mathcal{N}_1(X_{\geq \epsilon_0})$ lies in $X_{\geq \epsilon_0'}$. Assume

$$\epsilon < \frac{\epsilon_0' \epsilon_{\mathrm{sw}}}{C_{\mathrm{sw}} 2 C_{\mathrm{br}} \exp(\epsilon_0'/4C_{\mathrm{br}})} = \epsilon_{\mathrm{adj}},$$

where C_{sw} and ϵ_{sw} are from lemma 4.2, and let

$$C_{\mathrm{adj}} = \epsilon_0'/4C_{\mathrm{br}}.$$

Then we have the following observations: if $\mathcal{F}'$ is obtained from $\mathcal{F}_\epsilon(\mu)$ by successive r-adjustments in distinct spikes with $|r| < C_{\mathrm{adj}}$ we have

- the support $|\mathcal{F}'|$ of $\mathcal{F}'$ lies in $X_{\geq \epsilon_0'}$,
- each switch leaf of $\mathcal{F}'$ has length less than $\epsilon_0'/2C_{\mathrm{br}}$ (see lemma 4.2), and
- each leaf of $\mathcal{F}'$ is contained in a leaf of $\widehat{\mathcal{F}}$.

We now argue that for ϵ so chosen, that we can perform r-adjustments in spikes to $\mathcal{F}_\epsilon(\mu)$, with $|r| < C_{\mathrm{adj}}$, to obtain a horocyclic spike-foliation $\mathcal{F}'$ equivalent to $\mathcal{F}_\epsilon(\mu)$ each transverse branch of which has length at least

$$\ell_0 = \frac{\epsilon_0'}{C_{\mathrm{br}} 4^{C_{\mathrm{br}}+1}}.$$

For the purposes of the argument we think of the train-track τ as providing a schematic for these adjustments to $\mathcal{F}_\epsilon(\mu)$: we label each branch $b \subset \tau$ with the corresponding length of the transverse branches that collapse to b. Then performing an r-adjustment in a spike corresponding to a switch v of τ corresponds to increasing (decreasing) the length assigned to each incoming branch to v by the same amount that we decrease (increase) the length of each outgoing branch (we assume $\mathcal{F}_\epsilon(\mu)$ is generic, so each switch v has exactly 3 branches incident on it). For the rest of the argument the "length" of a branch b of τ will refer to the length of the corresponding transverse branch.

Given k, $0 \le k \le C_{\mathrm{br}} - 1$, call a branch b "*k-short*" if b has length less than $\epsilon_0'/C_{\mathrm{br}}4^{k+1}$ and "*k-long*" otherwise. We say two branches b_1 and b_2 of τ are *adjacent at a switch* v if there is a train path through v that traverses b_1 and b_2 in succession.

We argue by induction on $k = 0, \ldots, C_{\mathrm{br}} - 1$. We claim that by performing r-adjustments, with $|r| < C_{\mathrm{adj}}$, in at most k distinct spikes of $\mathcal{F}_\epsilon(\mu)$ we can obtain an equivalent spike-foliation $\mathcal{F}_k$ with at least k branches that are k-long, and so that $\mathcal{F}_k$ has the following property:

(*) *if $\mathcal{F}_k$ has any k-short branch, we can find a k-short branch b_0 so that either b_0 is adjacent at one side to two k-long branches b_1 and b_2, or on one side b_0 is adjacent to a single branch b_1, and b_1 is k-long.*

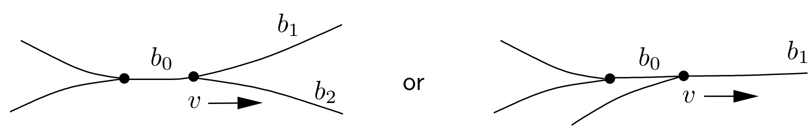

Figure 10. We may always lengthen a k-short branch b_0 by shortening only two adjacent k-long branches b_1 and b_2 or a single adjacent k-long branch b_1.

(See figure 10).

Let $\mathcal{F}_\epsilon(\mu) = \mathcal{F}_0$. We first argue that any spike-foliation $\mathcal{F}_k$ equivalent to $\mathcal{F}_0$ satisfies $(*)$. Otherwise, any k-short branch of $\mathcal{F}_k$ is adjacent to a k-short branch at each end. Then we can form a closed train-path in τ beginning at a k-short branch b_0 that always follows an adjacent k-short branch. By a pigeon-hole argument, after traversing at most $2C_{\mathrm{br}}$ k-short branches, this train path can be closed to form a single closed train path p_{short} that traverses only k-short branches. If p_{short} is not homotopic to an embedding, one may check that performing successive surgeries to eliminate crossings produces a train-path that is homotopic to an embedding and still traverses only short branches.

The requirement of monotonicity in the definition of a train-path prevents doubling back at a switch. This guarantees that p_{short} is homotopic to a closed path $p'_{\mathrm{short}} \colon S^1 \to |\mathcal{F}_k|$ that runs along segments of leaves of $|\mu|$ and then jumps between leaves of $|\mu|$ by running along switch leaves of $\mathcal{F}_k$ in such a way that p'_{short} enters and exits on opposite sides of each switch leaf of $\mathcal{F}_k$. Since each leaf of $\mathcal{F}_k$

is contained in a leaf of $\widehat{\mathcal{F}}$, the path p'_{short} can be perturbed in an arbitrarily small neighborhood of the switch leaves to an embedding $p''_{\text{short}}\colon S^1 \to |\widehat{\mathcal{F}}|$ transverse to $\widehat{\mathcal{F}}$ so that the length of p''_{short} differs from that of p'_{short} by an arbitrarily small amount.

But p''_{short}, being transverse to $\widehat{\mathcal{F}}$, is homotopically non-trivial in X and lies in $X_{\geq\epsilon'_0}$. Thus p''_{short} has length at least $2\epsilon'_0$, so p'_{short} also has length at least $2\epsilon'_0$. Since p'_{short} consists of an alternating sequence of at most $2C_{\text{br}}$ transverse branches and switch leaves, and the above choice of ϵ guarantees that each switch leaf has length at most $\epsilon'_0/2C_{\text{br}}$, the total length of p'_{short} along the transverse branches is at least ϵ'_0. Thus p_{short} must traverse a branch of length at least $\epsilon'_0/2C_{\text{br}}$, which is a contradiction, since p_{short} was assumed only to traverse k-short branches.

For $k = 0$, the contradiction implies that there is at least one branch with length at least $\epsilon'_0/4C_{\text{br}}$, so we take $\mathcal{F}_0 = \mathcal{F}_1$, and $\mathcal{F}_1$ satisfies $(*)$. Assume $\mathcal{F}_k$ is a spike foliation with at least k branches that are k-long, and $\mathcal{F}_k$ is equivalent to $\mathcal{F}_0$ by r-adjustments in spikes, with $|r| < C_{\text{adj}}$. Then $\mathcal{F}_k$ satisfies $(*)$, so if $\mathcal{F}_k$ has any k-short branch b_0, performing an $\big(\epsilon'_0/(C_{\text{br}}4^{k+2})\big)$-adjustment at v we may lengthen b_0 by $\epsilon'_0/(C_{\text{br}}4^{k+2})$ (figure 10) while shortening only k-long branches and shortening them by the same amount. The inequality

$$\frac{\epsilon'_0}{C_{\text{br}}4^{k+1}} - \frac{\epsilon'_0}{C_{\text{br}}4^{k+2}} > \frac{\epsilon'_0}{C_{\text{br}}4^{k+2}},$$

implies that the branches b_0, b_1, and b_2 all have length at least $\epsilon'_0/(C_{\text{br}}4^{k+2})$ after the adjustment, and the adjustment does not change the equivalence class of $\mathcal{F}_k$ (note that by induction we have not already adjusted v or each branch incident on v would be k-long). Thus, we have a new equivalent spike-foliation $\mathcal{F}_{k+1}$ (which also satisfies $(*)$) obtained by at most $k+1$ r-adjustments in distinct spikes, with $|r| < C_{\text{adj}}$, and $\mathcal{F}_{k+1}$ has at least $k+1$ branches that are $(k+1)$-long.

By induction, then, the spike-foliation $\mathcal{F}' = \mathcal{F}_{k+1}$ is equivalent to $\mathcal{F}_\epsilon(\mu)$ by r-adjustments in distinct spikes, with $|r| < C_{\text{adj}}$, and has the property that each transverse branch has length at least ℓ_0. Since each adjustment is uniformly bounded by C_{adj}, each switch leaf has length at most $\exp(C_{\text{adj}})C_{\text{sw}}\epsilon$. This proves the lemma.

■

Carrying laminations with compact leaves. A general lamination $\mu \in \mathcal{ML}(S)$ may have compact leaves: each such leaf is an isolated simple closed geodesic. Let $\mu = \mu_{\text{c}} \sqcup \mu_{\text{m}}$ be the decomposition of μ into its maximal sublaminations all of whose leaves are compact and non-compact, respectively.

Let ϵ_0 be so that $|\mu| \subset X_{\geq\epsilon_0}$. Choose ϵ sufficiently small to satisfy the hypotheses of lemmas 4.2 and 4.3 (ϵ less than ϵ_{adj} will do). Let $\mathcal{F}'$ be the horocyclic spike-foliation for μ_{m} on X guaranteed by lemma 4.3, and let τ_{m} be the train track $|\mu_{\text{m}}|/\mathcal{F}'$. Letting $|\mu_{\text{c}}| = \gamma_1 \sqcup \ldots \sqcup \gamma_p$, the disjoint union $\tau = \tau_{\text{m}} \sqcup \gamma_1 \sqcup \ldots \sqcup \gamma_p$ is a train track that carries μ. For future reference, equip each closed branch γ_q, $q = 1, \ldots, p$, with a single switch v_q. In lemma 4.5, we will argue that τ can be made nearly-straight in X. To avoid repetition, however, we first develop a technique to *enlarge* τ so that τ and its enlargement may be simultaneously made nearly-straight (cf. [**Min1**, §2]).

Enlarging train tracks. Given a weighted simple closed curve $tc \in \mathcal{ML}(S)$ that is close to μ, we might be fortunate enough that the above train track τ also carries

c; if τ were nearly-straight, this would tell us the length of tc is close to that of μ. In general, however, bits of c may lie far away from μ on X; c need not be carried by τ. To make a train-track carry more laminations we *enlarge* it:

Definition 4.4. *An* enlargement *of a train-track τ on X is a train-track $\check{\tau}$ obtained by adding branches to τ at switches of τ.*

The need to enlarge τ is related to the fact that the support of a measured lamination $|\mu| \in \mathcal{GL}(S)$ is not a continuous function of the measured lamination μ. A non-Hausdorff weakening of the topology on $\mathcal{GL}(S)$ called the *Thurston topology* is given by the following topology of convergence on a fixed surface $X \in \operatorname{Teich}(S)$ (see [**CEG**, 4.1.10]):

- $\lambda_i \to \lambda$ if for each $x \in \lambda$, there are $x_i \in \lambda_i$ so that $x_i \to x$.

Notice that the Thurston topology is not Hausdorff: a neighborhood of λ in the Thurston topology contains all geodesic laminations λ' such that $\lambda \subset \lambda'$. The Thurston topology is readily seen to give a notion of convergence on $\mathcal{GL}(S)$ independent of X (see [**CEG**, §4]), but we keep the metric X in the picture for the moment.

Given X and a measured lamination μ, geodesic laminations ν in a very close neighborhood of $|\mu|$ in the Thurston topology have leaves that make very small angles with leaves of $|\mu|$. Following Thurston, we define a projection map [**Th5**, §6] sending any such ν to the 'closest' geodesic lamination that contains $|\mu|$, as follows.

For ν sufficiently close to $|\mu|$ in the Thurston topology, define

$$\operatorname{cut}_{(X,|\mu|)}(\nu) \in \mathcal{GL}(S)$$

to be the geodesic lamination obtained by cutting X along $|\mu|$, and for each geodesic $\alpha \subset (\nu \cap \overline{(X - |\mu|)})$ either

1. discarding α if it lies in a close neighborhood of $|\mu|$ on X,
2. extending each remaining α to a bi-infinite piecewise geodesic α_∞ by adding the half-infinite segment of the leaf of $|\mu|$ at each endpoint of α that makes the smallest external angle with α, and
3. straightening each α_∞ *rel*-ideal endpoints to a geodesic α_∞^* asymptotic to $|\mu|$ in each direction.

Taking the union of the complete geodesics α_∞^* with $|\mu|$ (see figure 11)[3] yields a geodesic lamination $\operatorname{cut}_{(X,|\mu|)}(\nu) \in \mathcal{GL}(S)$ containing $|\mu|$.

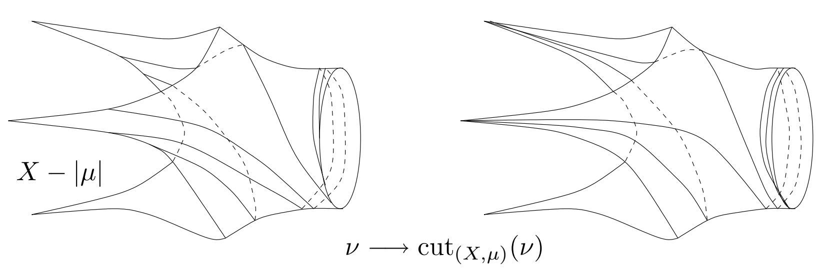

Figure 11. The projection map $\operatorname{cut}_{(X,\mu)}$.

[3]A similar figure appears in [**Th5**, pp. 24]

Given a convergent sequence $\mu_i \to \mu$ in $\mathcal{ML}(S)$, any limit λ of $|\mu_i|$ in $\mathcal{GL}(S)$ (with the Hausdorff topology) contains $|\mu|$ (see e.g. [**Otal**, A.3.2]). Hence for any sublamination $\lambda \subset |\mu|$, the supports $|\mu_i|$ converge to λ in the Thurston topology.

In particular, given a sequence $\mu_i \in \mathcal{ML}(S)$ of measured laminations tending to μ, $\mathrm{cut}_{(X,|\mu|)}(|\mu_i|)$ is well defined for all i sufficiently large, as is $\mathrm{cut}_{(X,\lambda)}(|\mu_i|)$ for any sublamination $\lambda \subset |\mu|$. For simplicity of notation, we will suppress the distinction between μ and its support $|\mu|$ in the definition of cut and use the notation $\mathrm{cut}_{(X,\mu)}(\mu_i) = \mathrm{cut}_{(X,|\mu|)}(|\mu_i|)$.

The map $\mathrm{cut}_{(X,\mu)}(.)$ depends strongly on the hyperbolic structure X. We will always fix X in our discussion. Given X, let $U_X(\mu) \subset \mathcal{GL}(S)$ be a sufficiently close neighborhood of $|\mu|$ in the Thurston topology so that $\mathrm{cut}_{(X,\mu)}(.)$ is well defined. Then $U_X(\mu)$ contains the supports of all measured laminations in a close neighborhood of any lamination containing μ as a sublamination.

The case of a simple closed curve c differs from the case of a general lamination ν in that $\mathrm{cut}_{(X,|\mu|)}(c)$ is obtained from $|\mu|$ by adding finitely many isolated leaves that either spiral towards closed leaves of $|\mu|$ or tend out a cusp of $X - |\mu|$. Given a simple closed geodesic $c \in U_X(\mu)$, we use the lamination $\lambda = \mathrm{cut}_{(X,\mu)}(c)$ to construct an enlargement $\check{\tau}$ of τ which will carry both μ and c.

To this end, let $|\mu_{\mathrm{c}}| = \gamma_1 \sqcup \ldots \sqcup \gamma_p$ as above, and assume each leaf γ_q has length bounded by L on X. Then there is an ϵ'_L depending only on L so that for $\epsilon < \epsilon'_L$ the annuli $\mathcal{N}_\epsilon(\gamma_1), \ldots, \mathcal{N}_\epsilon(\gamma_p)$ are pairwise disjoint, embedded on X, and disjoint from $\mathcal{F}_\epsilon(\mu_{\mathrm{m}})$. Let $\mathcal{F}_\epsilon(\mu_{\mathrm{c}})$ denote the foliation of these neighborhoods by geodesic arcs orthogonal to $|\mu_{\mathrm{c}}|$, and denote by $\mathcal{F}_\epsilon(\mu)$ the union of the horocyclic spike-foliation $\mathcal{F}_\epsilon(\mu_{\mathrm{m}})$ and the foliation $\mathcal{F}_\epsilon(\mu_{\mathrm{c}})$.

Let $\mathcal{F} = \mathcal{F}' \sqcup \mathcal{F}_\epsilon(\mu_{\mathrm{c}})$. Given λ and $\mathcal{F}$, we define the train track $\lambda/\mathcal{F}$ by enlarging τ as follows: given an arc α of $\lambda \cap (X - |\mathcal{F}|)$, each endpoint of α lies either in a switch leaf l of $\mathcal{F}'$ or in $\partial\mathcal{N}_\epsilon(\gamma_q)$ for some γ_q. As in the construction of τ_{m} above, we may modify α by a smooth isotopy to become an additional branch for τ: if an endpoint x of α lies in a switch leaf l, perturb α through arcs meeting l orthogonally until it meets l at the vertex v_l. If x lies in $\partial\mathcal{N}_\epsilon(\gamma_q)$, continue α geodesically until it meets the leaf l_q of $\mathcal{F}_\epsilon(\mu_{\mathrm{c}})$ containing v_q, and then smoothly perturb α through arcs that meet l_q at a definite angle to an arc that ends at v_q and meets l_q orthogonally. Performing these modifications for each arc of $\lambda \cap (X - |\mathcal{F}|)$ so that the resulting arcs are pairwise disjoint and disjoint from τ away from their endpoints, we obtain an enlargement $\check{\tau}$ of τ which we call $\lambda/\mathcal{F}$. The track $\check{\tau}$ is well defined up to equivalence and is equivalent to $\lambda/\mathcal{F}_\epsilon(\mu)$.

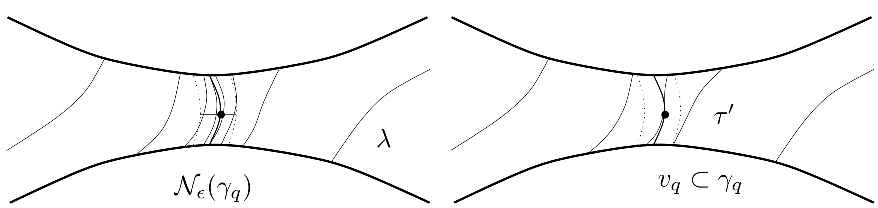

Figure 12. Enlarge τ at a closed branch γ_q by attaching at v_q.

Lemma 4.5. Nearly Straight *Given $\epsilon_0 > 0$, $L > 0$ and the surface S, there are constants $C_{\mathrm{ns}} > 1$ and $\epsilon_{\mathrm{ns}} > 0$ so that the following holds: let $\mu \in \mathcal{ML}(S)$ be* *a measured lamination realized by a pleated surface $(g, X) \in \mathcal{PS}(M)$, with $g(|\mu|) \subset M_{\geq \epsilon_0}$ and $\mathrm{length}_X(\gamma) < L$ for each compact leaf γ of μ. Let c be a simple closed curve in $U_X(\mu)$ and let $\lambda = \mathrm{cut}_{(X,\mu)}(c)$. If $(g_\lambda, X_\lambda) \in \mathcal{PS}(M)$ realizes λ and $\mathcal{F}_\epsilon(\mu)$ is taken on X_λ, then for positive $\epsilon < \epsilon_{\mathrm{ns}}$ the train track*

$$\tau = \lambda / \mathcal{F}_\epsilon(\mu)$$

has $C_{\mathrm{ns}}\epsilon$-nearly-straight realizations both in X_λ and in M.

Remark: Note that the above lemma contains the case when $\lambda = |\mu|$ by taking c sufficiently close to $|\mu|$ in the Hausdorff topology.

Proof: Assume we have chosen $\epsilon < \epsilon_{\mathrm{adj}} < \epsilon_{\mathrm{sw}}$; i.e. sufficiently small to satisfy the hypotheses of lemmas 4.2 and 4.3. We first consider the case $\lambda = |\mu|$, and take $X_\lambda = X$. The hypotheses imply that $|\mu|$ lies in $X_{\geq \epsilon_0}$.

Letting $\mu = \mu_{\mathrm{m}} \sqcup \mu_{\mathrm{c}}$ and letting $\tau_{\mathrm{m}} = |\mu_{\mathrm{m}}|/\mathcal{F}'$ be the train track described above (as guaranteed by lemma 4.3), we seek to determine the straightness of τ_{m} by applying the conclusions of lemma 4.3.

Let $\widehat{\tau}_{\mathrm{m}}$ denote the train track graph on X associated to τ_{m} (recall, $\widehat{\tau}_{\mathrm{m}}$ is obtained by straightening branches of τ_{m} to geodesics *rel*-endpoints). By lemma 4.3, the transverse branches of $\mathcal{F}'$ have length at least ℓ_0 and the switch leaves of $\mathcal{F}'$ have length at most $e^{C_{\mathrm{adj}}} C_{\mathrm{sw}} \epsilon$, so by lemma 4.3 the resulting edges of $\widehat{\tau}_{\mathrm{m}}$ have length at least

$$\ell_0 - 2e^{C_{\mathrm{adj}}} C_{\mathrm{sw}} \epsilon;$$

choosing

$$\epsilon < \frac{\ell_0}{4e^{C_{\mathrm{adj}}} C_{\mathrm{sw}}}$$

ensures each edge of $\widehat{\tau}_{\mathrm{m}}$ has length at least $\ell_0/2$.

If two transverse branches b_1 and b_2 are adjacent at a switch leaf l, an elementary hyperbolic trigonometry argument shows there is a constant C_{trig} depending on $\ell_0/2$ (which depends only on ϵ_0) so that if the switch leaves have length less than some $\epsilon_{\mathrm{trig}} > 0$ depending only on $\ell_0/2$, the train track graph $\widehat{\tau}_{\mathrm{m}}$ has external angles bounded by

$$C_{\mathrm{trig}} \cdot (\text{maximum length of switch leaves});$$

in other words, if $\epsilon < \epsilon_{\mathrm{trig}} / \left(e^{C_{\mathrm{adj}}} C_{\mathrm{sw}}\right)$ then the train track graph has external angles bounded by $C_{\mathrm{trig}} \left(e^{C_{\mathrm{adj}}} C_{\mathrm{sw}} \epsilon\right)$ on X. Likewise, since transverse branches of

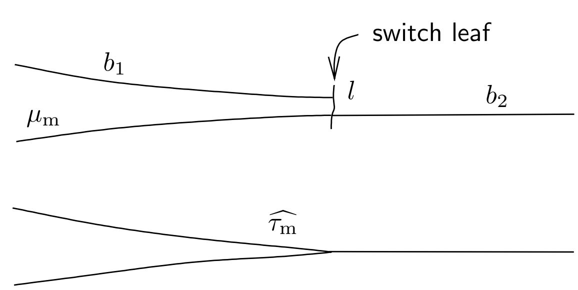

Figure 13. The train track graph $\widehat{\tau_{\mathrm{m}}}$ for τ_{m} has small external angles.

$\mathcal{F}'$ are segments of the pleating locus for (g, X), if we straighten the image of each edge of $\widehat{\tau}_{\mathrm{m}}$ under g to a geodesic in M *rel*-endpoints we obtain a train track graph $(\widehat{\tau_{\mathrm{m}}})^*$ with external angles bounded by $C_{\mathrm{trig}}\left(e^{C_{\mathrm{adj}}}C_{\mathrm{sw}}\epsilon\right)$ in M. The existence of C_{trig} follows from the fact that nearby points x and y on leaves ℓ_x and ℓ_y of a geodesic lamination λ realized on a pleated surface $(g, X) \in \mathcal{PS}(M)$, have lifts (x, ℓ_x) and (y, ℓ_y) to $\mathbf{P}X$ and lifts $(g(x), g(\ell_x))$ and $(g(y), g(\ell_y))$ to $\mathbf{P}M$ whose distance is bounded by a constant times the distance between x and y on X (which bounds the distance in M).

Thus for ϵ less than $\min\{\epsilon_{\mathrm{adj}}, \ell_0/(4e^{C_{\mathrm{adj}}}C_{\mathrm{sw}}), \epsilon_{\mathrm{trig}}/(e^{C_{\mathrm{adj}}}C_{\mathrm{sw}})\}$, the train track graph $\widehat{\tau}_{\mathrm{m}}$ is $\left(\ell_0/2, C_{\mathrm{trig}}e^{C_{\mathrm{adj}}}C_{\mathrm{sw}}\epsilon\right)$-nearly-straight in X, and straightening the image of each branch of $\widehat{\tau}_{\mathrm{m}}$ under g *rel*-endpoints to geodesics in M produces a $\left(\ell_0/2, C_{\mathrm{trig}}e^{C_{\mathrm{adj}}}C_{\mathrm{sw}}\epsilon\right)$-nearly-straight train track graph $(\tau_{\mathrm{m}})^*$ in M.

Add the closed geodesic leaves γ_q, $q = 1, \ldots, p$, each with its switch v_q, to obtain an enlarged train track τ and let $\widehat{\tau}$ denote its associated train track graph on X. Since we have only added geodesic branches, $\widehat{\tau}$ is still an $\left(\ell_0/2, C_{\mathrm{trig}}e^{C_{\mathrm{adj}}}C_{\mathrm{sw}}\epsilon\right)$-nearly-straight train track graph, and likewise for $(\tau_{\mathrm{m}})^* \cup_q g(\gamma_q)$ in M.

We now consider the case when $|\mu|$ is a proper sublamination of λ. If $(g_\lambda, X_\lambda) \in \mathcal{PS}(M)$ is a pleated surface realizing λ then the image $g_\lambda(|\mu|)$ is identical to $g(|\mu|)$. Thus, $g_\lambda(|\mu|)$ lies in $M_{\geq \epsilon_0}$, so $|\mu|$ lies in $(X_\lambda)_{\geq \epsilon_0}$. We may again, therefore, apply the conclusions of lemma 4.3; let $\mathcal{F}'$ be the spike-foliation equivalent to $\mathcal{F}_\epsilon(\mu_{\mathrm{m}})$ on X_λ that is guaranteed by lemma 4.3.

For each $\gamma \subset |\mu_{\mathrm{c}}|$, $\mathrm{length}_X(\gamma)$ lies in $[2\epsilon_0, L]$. Decrease ϵ if necessary so that $\epsilon < e^{-C_{\mathrm{adj}}}\epsilon'_L$. This ensures $\mathcal{F}'$ and $\mathcal{F}_\epsilon(\mu_{\mathrm{c}})$ are disjoint on X_λ, so we let $\mathcal{F} = \mathcal{F}' \sqcup \mathcal{F}_\epsilon(\mu_{\mathrm{c}})$ as above, but taken now on X_λ rather than X. Decreasing ϵ further if necessary so that $\epsilon < e^{(-\ell_0 - C_{\mathrm{adj}})}\epsilon'_L$, we ensure that arcs of $\lambda \cap (X_\lambda - |\mathcal{F}|)$ have length at least ℓ_0. If β is an arc of $\lambda \cap (X_\lambda - |\mathcal{F}|)$, β is a segment of a leaf of λ that is asymptotic to $|\mu| \subset \lambda$. Thus, it is close at its endpoints to leaves of $|\mu|$ in $\mathbf{P}X_\lambda$. Moreover, since (g_λ, X_λ) realizes λ, the image $g_\lambda(\beta)$ is a geodesic arc in M that is close at its endpoints to the realization $g_\lambda(|\mu|)$ in $\mathbf{P}M$.

Straightening branches of $\tau = \lambda/\mathcal{F}$ *rel*-endpoints, then, produces a train-track graph $\widehat{\tau}$ associated to $\lambda/\mathcal{F}$ that is again $\left(\ell_0/2, C_{\mathrm{trig}}e^{C_{\mathrm{adj}}}C_{\mathrm{sw}}\epsilon\right)$-nearly-straight in X_λ. Likewise, the graph $(\widehat{\tau})^*$ obtained by straightening the image of each edge under the mapping g_λ *rel*-endpoints in M is again a $\left(\ell_0/2, C_{\mathrm{trig}}e^{C_{\mathrm{adj}}}C_{\mathrm{sw}}\epsilon\right)$-nearly-straight train track graph associated to $\lambda/\mathcal{F}$ in M.

To promote the above nearly-straight train track graphs to nearly-straight train tracks, recall from the beginning of this section that given $\ell > 0$ there are constants $C_{\mathrm{curv}}(\ell) > 1$ and $\delta_\ell > 0$ so that if $\epsilon < \delta_\ell$ and $\widehat{\tau}$ is an (ℓ, ϵ)-nearly-straight train track graph associated to τ in X_λ or in M, then τ admits a $C_{\mathrm{curv}}(\ell)\epsilon$-nearly-straight realization in X_λ or in M.

Provided, then, that ϵ is bounded above by

$$\min\left\{\epsilon_{\mathrm{adj}}, \frac{\ell_0}{4e^{C_{\mathrm{adj}}}C_{\mathrm{sw}}}, \frac{\epsilon_{\mathrm{trig}}}{e^{C_{\mathrm{adj}}}C_{\mathrm{sw}}}, \frac{\epsilon'_L}{e^{(\ell_0 + C_{\mathrm{adj}})}}, \frac{\delta_{\ell_0/2}}{C_{\mathrm{trig}}e^{C_{\mathrm{adj}}}C_{\mathrm{sw}}}\right\} = \epsilon_{\mathrm{ns}},$$

which depends only on ϵ_0 and L, it follows that in each of the above cases τ admits $C_{\mathrm{curv}}(\ell_0/2)C_{\mathrm{trig}}e^{C_{\mathrm{adj}}}C_{\mathrm{sw}}\epsilon$-nearly-straight realizations in X_λ and in M. The lemma follows by setting $C_{\mathrm{ns}} = C_{\mathrm{curv}}(\ell_0/2)C_{\mathrm{trig}}e^{C_{\mathrm{adj}}}C_{\mathrm{sw}}$.

∎

5. Uniform estimates for train tracks

In this section, we harness theorem 2.2 and the results of section 3 to obtain uniform geometric estimates for train tracks in hyperbolic 3-manifolds. The discussion is technical at points, so we again encourage the reader to peruse sections 6 and 7 for motivation.

When enlarging train tracks in a 3-manifold $M \in AH(S)$, it will be convenient for us to remove the requirement that the added branches come from a train track on X:

DEFINITION 5.1. *Let τ be a train-track in $M \in AH(S)$. A* generalized enlargement *$\check{\tau}$ of τ is obtained by adding C^1 arcs to τ whose endpoints lie at switches of τ so that for any new arc b added at a switch v of τ, there is a branch b' of τ incident on v so that $b \cup b'$ forms a C^1 arc through v.*

A *train-path* on a generalized enlargement $\check{\tau}$ is a C^1 monotone immersion $r\colon \mathbb{R} \to M$ or $r\colon S^1 \to M$ with image in $\check{\tau}$ so that r is obtained by concatenating branches of τ and additional arcs used to form $\check{\tau}$. If $\check{\tau}$ is a generalized enlargement of τ in M, then $\check{\tau}$ *carries* λ if there is a smooth marking-preserving map $f\colon X \to M$ so that for each leaf ℓ of λ, $f|_\ell$ is homotopic (*rel*-ideal endpoints) through smooth immersions to a train-path on $\check{\tau}$. As with train tracks, given $\epsilon \in (0,1)$ a generalized enlargement $\check{\tau}$ of τ is *ϵ-nearly-straight* (cf. definition 4.1) if each train path r on $\check{\tau}$ is C^2 with geodesic curvature less than ϵ.

When a generalized enlargement $\check{\tau}$ of τ in M carries a simple closed curve c, the weight deposited by c on a branch b depends *a priori* on the train-path we choose to carry c; if $\check{c}\colon S^1 \to \check{\tau}$ is a train path on $\check{\tau}$ homotopic to c, we denote by $m_b(\check{c})$ the number of times the train path $\check{c}$ traverses the branch b. In what follows, there will always be an explicit train-path $\check{c}$ on which c is carried by $\check{\tau}$, so there will be no ambiguity in the weight $m_b(\check{c})$ deposited on b. If $tc \in \mathcal{ML}(S)$ is a weighted simple closed curve, let $m_b(t\check{c}) = tm_b(\check{c})$, and define the *track length* of $t\check{c}$ in M to be

$$\ell_{\check{\tau}}(t\check{c}) = \sum_{b \subset \check{\tau}} m_b(t\check{c})\ell_M(b).$$

Applying the results of the previous section, we now prove the lemma which serves as the central technical tool of the paper. This lemma generalizes a construction of Minsky [**Min1**, Thm. 2.4].

LEMMA 5.2. TRAIN TRACKS *Let $(M_i, t_i c_i) \to (M, \mu)$ in $AH(S) \times \mathcal{ML}(S)$ so that the realizable part $\mu^{\mathrm{r}} = \mathrm{R}_M(\mu)$ of μ lies in $\mathcal{ML}(S)_+$.*

Let $(g\colon X \to M) \in \mathcal{PS}(M)$ be a pleated surface realizing μ^{r} and let L_1 and ϵ_1 be positive constants so that $g(|\mu^{\mathrm{r}}|)$ lies in $M_{\geq \epsilon_1}$ and each compact leaf γ of $|\mu^{\mathrm{r}}|$ has length $\ell_X(\gamma) < L_1$. Then there are constants $C > 1$ and $\epsilon_{\mathrm{tt}} > 0$ depending at most on S, ϵ_1 and L_1 so that for each positive $\epsilon < \epsilon_{\mathrm{tt}}$ there exists:

1. *a train track τ minimally carrying μ^{r} with an ϵ-nearly-straight realization τ^* in M,*
2. *an integer $N_\epsilon > 0$ so that for all $i > N_\epsilon$, τ admits a realization τ_i in M_i with a $C\epsilon$-nearly-straight generalized enlargement $\check{\tau}_i$ so that c_i is carried by $\check{\tau}_i$ on a train-path $\check{c}_i$ for which if b is a branch of τ, we have $m_b(t_i\check{c}_i) \to m_b(\mu)$.*

Remark: In statement (2) of lemma 5.2 we abuse notation slightly and drop the distinction between a branch b of τ and its realization in M_i as a branch of the train track τ_i or its generalized enlargement $\check{\tau}_i$.

Proof: Give μ^{r} the decomposition

$$\mu^{\mathrm{r}} = \mu_{\mathrm{m}} \sqcup \mu_{\mathrm{c}},$$

into μ_{c}, its maximal sublamination all of whose leaves are compact (weighted simple closed curves), and μ_{m}, its maximal sublamination all of whose leaves are infinite. For convenience, we use the notation

$$\nu_{\mathrm{m}} = |\mu_{\mathrm{m}}| \quad \text{and} \quad \nu_{\mathrm{c}} = |\mu_{\mathrm{c}}|.$$

The proof occurs in three steps below. In Step I, we construct the train track τ with its nearly-straight realization τ^* in M and use algebraic convergence to provide realizations of τ in M_i for all $i > N_\epsilon$. In Step II, we treat the cases when $\mu^{\mathrm{r}} = \mu_{\mathrm{c}}$ and $\mu^{\mathrm{r}} = \mu_{\mathrm{m}}$ independently, constructing enlargements τ_i of τ with $C\epsilon$-nearly-straight realizations in M_i for all $i > N_\epsilon$; the first case employs the results of section 3, while the second case adapts the analogous result of Minsky [**Min1**, Thm. 2.4] to the setting of algebraic convergence using theorem 2.2. In Step III, we create a kind of convex combination or *melding* of the nearly-straight enlargements obtained for each case of Step II to form *generalized* enlargements of the nearly-straight realizations of τ in M_i; this final step makes direct use of Thurston's uniform injectivity theorem (theorem 2.1).

Constants: Let $\epsilon_0 > 0$ be such that for any hyperbolic 3-manifold we have

$$\mathcal{N}_1(M_{\geq \epsilon_1}) \subset M_{\geq \epsilon_0},$$

(see [**BM**]), and let $L = 2L_1$. Using ϵ_0, and L as the constants with the same names in the hypotheses of lemma 4.5, we assume $\epsilon < \min\{\epsilon_{\mathrm{ns}}, 1/2C_{\mathrm{tr}}\}$ from lemma 4.5 noting that ϵ_{ns} depends only on S, ϵ_0 and L, and that C_{tr} is universal. In particular, we may now apply the lemmas of the previous section to the realization of μ^{r} on the pleated surface $(g, X) \in \mathcal{PS}(M)$.

Step I: building $\boldsymbol{\tau}$. Letting $\epsilon' = \epsilon/C_{\mathrm{ns}}$, lemma 4.5 guarantees that taking $\mathcal{F}_{\epsilon'}(\nu_{\mathrm{m}})$ on X, the train track $\tau_{\mathrm{m}} = \nu_{\mathrm{m}}/\mathcal{F}_{\epsilon'}(\nu_{\mathrm{m}})$ can be realized as an ϵ-nearly-straight train track τ^*_{m} in M.

We write ν_{c} as the disjoint union $\nu_{\mathrm{c}} = \gamma_1 \sqcup \ldots \sqcup \gamma_p$ of simple closed geodesics γ_q on X. Introduce a single switch v_q on each γ_q, $q = 1, \ldots, p$ to obtain a train track τ_{c} carrying ν_c. The union $\tau = \tau_{\mathrm{m}} \sqcup \tau_{\mathrm{c}}$ is an ϵ-nearly-straight train track in X carrying μ^{r}. Let τ^*_{c} be the union of (geodesic) images of each γ_q equipped with its switch v_q in M under the pleated mapping g. Then the union

$$\tau^* = \tau^*_{\mathrm{m}} \cup \tau^*_{\mathrm{c}}$$

is an ϵ-nearly-straight realization of τ in M that carries μ^{r}.

By algebraic convergence (see section 2, [**Mc**, §3.1]), given a compact subset $K \subset M$ that contains τ^* and $g(X_{\geq \epsilon_1})$, there are smooth, marking-preserving homotopy equivalences $q_i \colon M \to M_i$ that tend C^∞ to a local isometry on K. Begin by taking N_ϵ sufficiently large so that for all $i > N_\epsilon$ we have

- $\|q_i\|_{C^1} < 2$ on K, and
- the image $q_i(\alpha)$ of any C^2 arc α in K with geodesic curvature less than ϵ has geodesic curvature less than 2ϵ.

Then the image $q_i(\tau^*)$ determines a 2ϵ-nearly-straight train track $(\tau^*)_i$ in M_i made up of 2ϵ-nearly-straight train tracks $(\tau^*_{\mathrm{c}})_i$ and $(\tau^*_{\mathrm{m}})_i$ determined by the images $q_i(\tau^*_{\mathrm{c}})$ and $q_i(\tau^*_{\mathrm{m}})$. Likewise, the image $q_i(\ell)$ of any leaf ℓ of $|\mu^{\mathrm{r}}|$ lifts to a bi-infinite arc in $\mathbb{H}^3$ that is homotopic to its geodesic representative *rel*-ideal endpoints by a homotopy that moves each point a distance less than $C_{\mathrm{tr}}2\epsilon$ for all $i > N_\epsilon$.

Since ϵ satisfies $\epsilon < 1/(2C_{\mathrm{tr}})$, the lamination μ^{r} is realized in $(M_i)_{\geq \epsilon_0}$ for all $i > N_\epsilon$. This implies, in particular, that for any pleated surface $(g_i \colon X_i \to M_i) \in \mathcal{PS}(M_i)$ that realizes μ^{r}, the support $|\mu^{\mathrm{r}}|$ lies in $(X_i)_{\geq \epsilon_0}$.

Finally, we note that since $L = 2L_1$, we have

$$\mathrm{length}_{M_i}(\gamma_q) < L$$

for each $\gamma_q \subset \nu_{\mathrm{c}}$ and all $i > N_\epsilon$. Thus we are in a position to apply the lemmas of the previous section to the realizations of μ^{r} on any sequence $(g_i \colon X_i \to M_i) \in \mathcal{PS}(M_i)$ of pleated surfaces realizing μ^{r} in M_i, once $i > N_\epsilon$.

In step II, we apply lemma 4.5 to build enlargements of $(\tau^*)_i$ that carry c_i. To do this, we must build pleated surfaces realizing the laminations $\mathrm{cut}_{(X,\mu^{\mathrm{r}})}(c_i)$ in M_i; we must first ensure these laminations are well-defined and realizable. We enlarge N_ϵ if necessary so that $c_i \in U_X(\mu^{\mathrm{r}})$ for all $i > N_\epsilon$ (so $\mathrm{cut}_{(X,\mu^{\mathrm{r}})}(c_i)$ is well defined), and since μ^{r} is realizable in M and in M_i for all $i > N_\epsilon$, we may apply theorem 2.3 to conclude that $\mathrm{cut}_{(X,\mu^{\mathrm{r}})}(c_i)$ is realizable in M and in M_i.

Step II: Assume $\boldsymbol{\mu}^{\mathrm{r}}$ has all leaves compact or all leaves infinite.

Case (i): each leaf of μ^{r} is compact. In this case, $|\mu^{\mathrm{r}}| = \nu_{\mathrm{c}} = \gamma_1 \sqcup \ldots \sqcup \gamma_p$. Let $i > N_\epsilon$, and assume $c_i \neq \gamma_q$ since otherwise there is nothing to prove. Let $\lambda_i = \mathrm{cut}_{(X,\nu_{\mathrm{c}})}(c_i)$, and consider the pleated surfaces $(g_i, X_i) \in \mathcal{PS}(M_i)$ realizing λ_i. By the above, for each $\gamma \subset \nu_{\mathrm{c}}$ it follows that $\mathrm{length}_{X_i}(\gamma) < L$ and γ lies in $(X_i)_{\geq \epsilon_0}$ for all $i > N_\epsilon$. Since we have assumed $\epsilon < \epsilon_{\mathrm{ns}}$, we may for such i consider the union of embedded annuli $\mathcal{N}_\epsilon(\nu_{\mathrm{c}})$ on X_i together with its foliation $\mathcal{F}_\epsilon(\nu_{\mathrm{c}})$. By lemma 4.5, the train track

$$(\check{\tau}_{\mathrm{c}})_i = \lambda_i / \mathcal{F}_\epsilon(\nu_{\mathrm{c}})$$

is an enlargement of τ_{c} with a $C_{\mathrm{ns}}\epsilon$-nearly-straight realization $(\check{\tau}^*_{\mathrm{c}})_i$ in M_i.

It remains to show that $(\check{\tau}_{\mathrm{c}})_i$ carries c_i and that given a branch b of $(\check{\tau}_{\mathrm{c}})_i$ that is also a branch of τ, the weight on $m_b(t_i c_i)$ deposited by $t_i c_i$ on b converges to $m_b(\mu)$. By corollary 3.10, if we pass to any subsequence (without loss of generality) so that c_i converges in the Hausdorff topology to $\lambda' \in \mathcal{GL}(S)$, given $\gamma \subset \nu_{\mathrm{c}}$ the lamination λ' is either to the right or left of γ (see figure 5) and we have

$$\mathcal{W}_{(\epsilon,X_i)}(c_i,\gamma) \to +\infty \quad \text{or} \quad -\infty$$

respectively. Moreover, for such a subsequence λ_i is either to the right or to the left of γ for all i sufficiently large.

By the definition of λ_i, we may modify c_i to a piecewise geodesic p_i on X_i so that $p_i \cap \overline{(X_i - \mathcal{N}_\epsilon(\nu_{\mathrm{c}}))}$ consists of arcs of $\lambda_i \cap \overline{(X_i - \mathcal{N}_\epsilon(\nu_{\mathrm{c}}))}$ and $A_i = p_i \cap \overline{\mathcal{N}_\epsilon(\nu_{\mathrm{c}})}$ is a collection pairwise disjoint geodesic arcs each of which crosses the annular component of $\mathcal{N}_\epsilon(\nu_{\mathrm{c}})$ in which it lies (figure 14 depicts the lifted picture on $\widetilde{X_i}$).

Since $\mathrm{length}_{X_i}(\gamma) > \epsilon_0$ for each i and each simple closed geodesic $\gamma \subset \nu_{\mathrm{c}}$, modifying c_i to p_i changes the winding about $\mathcal{N}_\epsilon(\gamma)$ a uniformly bounded amount: given $a \in A_i$ for which $a \subset \mathcal{N}_\epsilon(\gamma)$, one may check that the constant $K_1(\epsilon, \epsilon_0)$ from lemma 3.8 gives the uniform comparison

$$\mathcal{W}_{(\epsilon,X_i)}(a,\gamma) \asymp_{2K_1(\epsilon,\epsilon_0)} \mathcal{W}_{(\epsilon,X_i)}(c_i,\gamma).$$

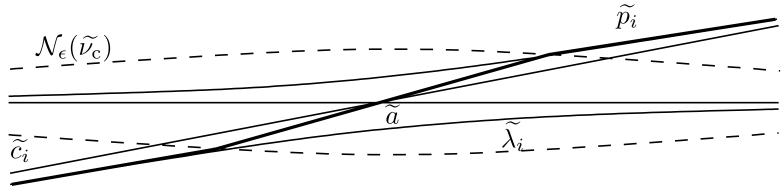

Figure 14. The lift $\widetilde{p_i}$ of the polygonal path p_i runs along $\widetilde{\lambda_i}$ outside of $\mathcal{N}_\epsilon(\widetilde{\nu_c})$.

It follows that we may enlarge N_ϵ so that for all $i > N_\epsilon$ any such arc a makes at least two full trips around the annulus $\mathcal{N}_\epsilon(\gamma)$; thus $(\check{\tau}_c)_i$ carries c_i. Each arc $a \in A_i$ that crosses $\mathcal{N}_\epsilon(\gamma)$ contributes to the weight $m_b(c_i)$ deposited on the branch $b \subset \gamma$. This contribution is uniformly comparable to the winding $\mathcal{W}_{(\epsilon,X_i)}(a,\gamma)$, so summing over arcs in A_i that cross $\mathcal{N}_\epsilon(\gamma)$ and applying the above, we have

$$i(c_i,\gamma)\mathcal{W}_{(\epsilon,X_i)}(c_i,\gamma) \asymp_{(i(c_i,\gamma)2K_1(\epsilon,\epsilon_0))} m_b(c_i)$$

for all $i > N_\epsilon$. Multiplying by the weight, then, we have that

$$t_i\, i(c_i,\gamma)\, \mathcal{W}_{(\epsilon,X_i)}(c_i,\gamma) \asymp_{(t_i\, i(c_i,\gamma)\, 2K_1(\epsilon,\epsilon_0))} m_b(t_i c_i).$$

Since $t_i c_i \to \mu$, we have $t_i i(c_i,\gamma) \to 0$, so

$$\lim_{i\to\infty} \frac{t_i\, i(c_i,\gamma)\, \mathcal{W}_{(\epsilon,X_i)}(c_i,\gamma)}{m_b(t_i c_i)} = 1.$$

Since the weight $m_b(\mu)$ deposited by μ on b is simply a real weight on the simple closed curve γ, on the fixed surface X we have

$$t_i\, i(c_i,\gamma)\, \mathcal{W}_{(\epsilon,X)}(c_i,\gamma) \to m_b(\mu).$$

Applying corollary 3.10, we have that the ratio of the winding of c_i about $\mathcal{N}_\epsilon(\gamma)$ on X_i and X satisfies

$$\lim_{i\to\infty} \frac{\mathcal{W}_{(\epsilon,X_i)}(c_i,\gamma)}{\mathcal{W}_{(\epsilon,X)}(c_i,\gamma)} = 1.$$

It follows that $\lim_{i\to\infty} m_b(t_i c_i) = m_b(\mu)$ for each branch b of τ.

Case (ii): each leaf of μ^{r} is infinite. This case follows from the analogous theorem of Minsky [**Min1**, Thm. 2.4] adapted to the setting of algebraic convergence; we reprise the argument.

In this case we have $|\mu^{\mathrm{r}}| = \nu_{\mathrm{m}}$. Let $\mathcal{F} = \mathcal{F}_{\epsilon'}(\nu_{\mathrm{m}})$, and let $\hat{c}_i$ be the geodesic representative of c_i on X. Then for i sufficiently large, each arc β of $\overline{\hat{c}_i \cap (X - |\mathcal{F}|)}$ either lies entirely within the ϵ'-neighborhood $\mathcal{N}_{\epsilon'}(\nu_{\mathrm{m}})$ on X or $\partial\beta$ lies in frontier horocycles of $\mathcal{F}$. Enlarge N_ϵ if necessary so this holds for all $i > N_\epsilon$. For such i, let B_i denote the subset of the collection of arcs $\overline{\hat{c}_i \cap (X - |\mathcal{F}|)}$ that do not lie entirely in $\mathcal{N}_{\epsilon'}(\nu_{\mathrm{m}})$.

Similarly to the above, let

$$\lambda_i = \mathrm{cut}_{(X,\nu_{\mathrm{m}})}(c_i).$$

Then for each $\beta \in B_i$, if α and α' are the frontier horocycles at its endpoints, β is homotopic with endpoints constrained to $(\alpha \sqcup \alpha')$ into a leaf ℓ_β of λ_i. Thus, the train track

$$(\check{\tau}_{\mathrm{m}})_i = \lambda_i/\mathcal{F}$$

is an enlargement of $\tau_{\mathrm{m}} = \nu_{\mathrm{m}}/\mathcal{F}$ (obtained by adding a branch b_β to τ_{m} for each $\beta \in B_i$) that carries ν_{m} and c_i.

Given a branch b of $(\check{\tau}_{\mathrm{m}})_i$ that is a branch of the sub-track τ_{m}, the weight $m_b(t_i c_i)$ comes from the transverse measure of $t_i c_i$ on a short arc transverse to ν_{m} that cuts across $\mathcal{N}_{\epsilon'}(\nu_{\mathrm{m}})$. Since the weighted simple closed curves $t_i c_i$ converge as measured laminations on X to μ, the transverse measures determined by $t_i c_i$ on this short arc converge weakly to the transverse measure determined by μ. Thus we have $m_b(t_i c_i) \to m_b(\mu)$.

Recall from Step I that the train track τ_{m} has an ϵ-nearly-straight realization τ^*_{m} in M with image $q_i(\tau^*_{\mathrm{m}})$ a 2ϵ-nearly-straight train track $(\tau^*_{\mathrm{m}})_i$ in M_i for all $i > N_\epsilon$. We now show that the enlargement $(\check{\tau}_{\mathrm{m}})_i$ can be realized as a nearly-straight enlargement $(\check{\tau}^*_{\mathrm{m}})_i$ of $(\tau^*_{\mathrm{m}})_i$ in M_i: we describe a procedure to add each additional branch b_β to $(\tau^*_{\mathrm{m}})_i$ in M_i in a nearly-straight manner.

As in the first case, let pleated surfaces $(g_i, X_i) \in \mathcal{PS}(M_i)$ realize λ_i. Then $(g_i \colon X_i \to M_i)$ realizes leaves of ν_{m} within $C_{\mathrm{tr}} 2\epsilon$ in $\mathbf{P}M_i$ of their associated train-paths on $(\tau^*_{\mathrm{m}})_i$.

Let $\partial\alpha = x \sqcup y$, and let ℓ_x and ℓ_y be the asymptotic leaves of ν_{m} corresponding to the spike α bounds. Each leaf ℓ of the realization of ν_{m} in M has image $q_i(\ell)$ with geodesic curvature bounded by 2ϵ. The frontier horocycle α, then, has image $q_i(\alpha)$ with endpoints within $C_{\mathrm{tr}} 2\epsilon$ of the corresponding leaves of ν_{m} on X_i. Let x_i and y_i denote images of the endpoints of $q_i(\alpha)$ under the natural projections from $q_i(\ell_x)$ and $q_i(\ell_y)$ to the representatives of ℓ_x and ℓ_y on X_i. Then the concatenation of the orthogonal projection from $q_i(x)$ to x_i, the arc $q_i(\alpha)$, and the orthogonal projection from $q_i(y)$ to y_i gives a homotopy class of paths from x_i to y_i in M_i; let ξ be its geodesic representative *rel*-endpoints (see figure 15).

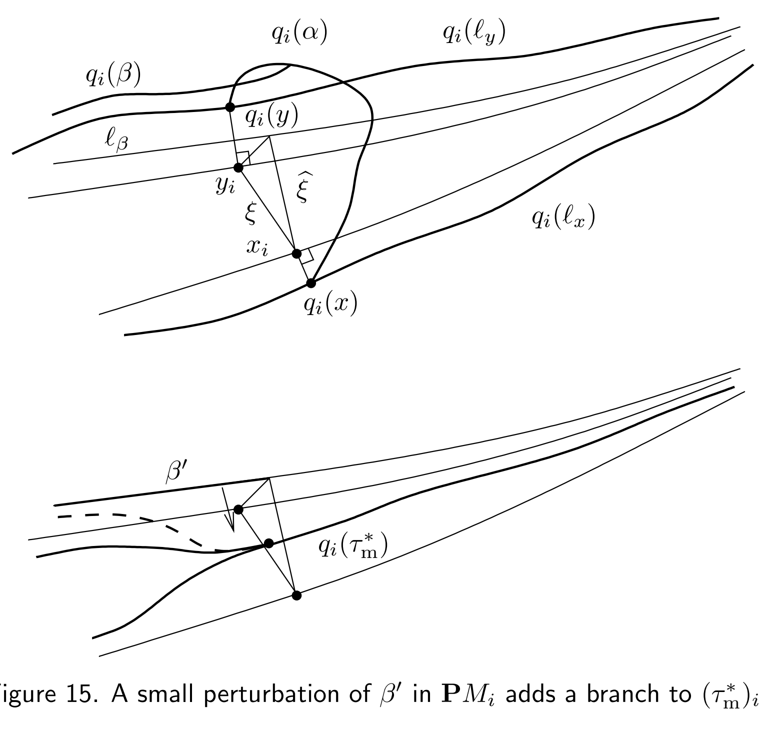

Figure 15. A small perturbation of β' in $\mathbf{P}M_i$ adds a branch to $(\tau^*_{\mathrm{m}})_i$.

Since each projection has length less than $C_{\mathrm{tr}}2\epsilon$ and $q_i(\alpha)$ has length less than 2ϵ, we have

$$\ell_{M_i}(\xi) < (2C_{\mathrm{tr}} + 2)\epsilon.$$

By theorem 2.2, there is a $C_{\mathrm{inj}} > 1$ and a $\delta_{\mathrm{inj}} > 0$ depending only on S so that if

$$\epsilon < \frac{\delta_{\mathrm{inj}}}{2C_{\mathrm{tr}} + 2}, \tag{5.6}$$

then ξ is homotopic *rel*-endpoints in M_i into X_i to an arc $\widehat{\xi}$ bounding the spike determined by the realizations of ℓ_x and ℓ_y on X_i for which

$$\ell_{X_i}(\widehat{\xi}) < C_{\mathrm{inj}}(2C_{\mathrm{tr}} + 2)\epsilon.$$

For each b_β that we wish to add to $(\tau^*_{\mathrm{m}})_i$, the leaf ℓ_β of λ_i crosses two such arcs $\widehat{\xi}$ and $\widehat{\xi}'$ determined by the original frontier horocycles α and α' as it enters spikes at each end. Cutting ℓ_β at its intersections with $\widehat{\xi}$ and $\widehat{\xi}'$ and taking β' to be the finite arc left over, if we let $C_2' = (2C_{\mathrm{tr}} + C_{\mathrm{trig}}C_{\mathrm{inj}}(2C_{\mathrm{tr}} + 2))$ then by a $C_2'\epsilon$ perturbation of the endpoints of β' in $\mathbf{P}M_i$ we may attach β' to $(\tau^*_{\mathrm{m}})_i$ as a new branch (figure 15). As in the proof of lemma 4.5, if $\epsilon < \epsilon_L'/(C_{\mathrm{inj}}2(C_{\mathrm{tr}} + 1)e^{\ell_0})$ any such arc β' has length at least $\ell_0/2$. Letting

$$C_2 = C_{\mathrm{curv}}(\ell_0/2)C_{\mathrm{trig}}C_2', \tag{5.7}$$

if $\epsilon < \epsilon_{\mathrm{trig}}/C_2'$ we can make the resulting enlargement $C_2\epsilon$-nearly-straight in M_i. Since this perturbation moves the endpoints of β' within the thick part, the added arc is homotopic to $q_i(b_\beta)$ *rel*-endpoints.

Adding the branches b_β for each $\beta \in B_i$ we obtain a realization $(\check{\tau}^*_{\mathrm{m}})_i$ of $(\check{\tau}_{\mathrm{m}})_i$ in M_i as an enlargement of $(\tau^*_{\mathrm{m}})_i$. Thus, we have that each $(\check{\tau}^*_{\mathrm{m}})_i$ is $C_2\epsilon$-nearly-straight in M_i where C_2 depends only on ϵ_0 and S, provided ϵ is chosen sufficiently (uniformly) small to satisfy the above constraints, and the index N_ϵ is chosen accordingly.

Step III: melding enlargements of train tracks. We now bootstrap our way to the case when μ^{r} has compact leaves and infinite leaves by melding the realizations $(\check{\tau}^*_{\mathrm{c}})_i$ and $(\check{\tau}^*_{\mathrm{m}})_i$ in M_i that we obtained in the previous step into a single (possibly slightly less) nearly-straight *generalized* enlargement of $(\tau^*)_i$. Our aim is to show that we lose a uniformly bounded multiplicative amount of straightness in the process.

To avoid excessive sub-scripts, let $\mathbf{T}_{\mathrm{c}}$ and $\mathbf{T}_{\mathrm{m}}$ be train tracks minimally carrying ν_{c} and ν_{m} respectively that are ϵ-nearly-straight in M_i. For a given simple closed curve c, assume we are given enlargements $\check{\mathbf{T}}_{\mathrm{c}}$ and $\check{\mathbf{T}}_{\mathrm{m}}$ of $\mathbf{T}_{\mathrm{c}}$ and $\mathbf{T}_{\mathrm{m}}$, so that each enlargement minimally carries c and is $C_3\epsilon$-nearly-straight in M_i, with $C_3 > 1$. Then we claim there is a $C_{\mathrm{meld}} > 1$ so that these train tracks may be melded into a single $C_{\mathrm{meld}}C_3\epsilon$-nearly-straight *generalized* enlargement $\mathbf{T}$ of $\mathbf{T}_{\mathrm{c}} \sqcup \mathbf{T}_{\mathrm{m}}$ obtained by adding branches to $\mathbf{T}_{\mathrm{c}} \sqcup \mathbf{T}_{\mathrm{m}}$ so that $\mathbf{T}$ carries c.

To see this, let c^* be the geodesic representative of c in M_i, and lift the picture to $\mathbf{P}M_i$. To represent lifting to $\mathbf{P}M_i$, we prepend by "$\mathbf{P}$": i.e. $\mathbf{P}c^*$ denotes the lift of c^* to $\mathbf{P}M_i$. The train tracks $\mathbf{T}_{\mathrm{c}}$ and $\mathbf{T}_{\mathrm{m}}$ and their enlargements and train-paths all admit natural lifts to $\mathbf{P}M_i$.

Since c is carried by both $\check{\mathbf{T}}_{\mathrm{c}}$ and $\check{\mathbf{T}}_{\mathrm{m}}$, there are train-paths $r_{\mathrm{c}} \colon S^1 \to \check{\mathbf{T}}_{\mathrm{c}}$ and $r_{\mathrm{m}} \colon S^1 \to \check{\mathbf{T}}_{\mathrm{m}}$ homotopic to c. As $\check{\mathbf{T}}_{\mathrm{c}}$ and $\check{\mathbf{T}}_{\mathrm{m}}$ are $C_3\epsilon$-nearly-straight, there are smooth homotopies $h_{\mathrm{c}}(x,t)$ and $h_{\mathrm{m}}(x,t)$ to c^*, so that we have

$$h_{\mathrm{c}}(x,0) = r_{\mathrm{c}}(x) \quad \text{and} \quad h_{\mathrm{m}}(x,0) = r_{\mathrm{m}}(x),$$

so that $h_{\mathrm{c}}(x,1)$ and $h_{\mathrm{m}}(x,1)$ are each smooth parameterizations of c^*, and so that h_{c} and h_{m} lift to homotopies $\mathbf{P}h_{\mathrm{c}}$ and $\mathbf{P}h_{\mathrm{m}}$ in $\mathbf{P}M_i$ from $\mathbf{P}r_{\mathrm{c}}$ and $\mathbf{P}r_{\mathrm{m}}$ to $\mathbf{P}c^*$ so that tracks of $\mathbf{P}h_{\mathrm{c}}$ and $\mathbf{P}h_{\mathrm{m}}$ have length at most $C_{\mathrm{tr}}C_3\epsilon$ (the homotopy h_{c} is lifted by lifting each smooth closed curve $c_{t_0}(x) = h_{\mathrm{c}}(x,t_0)$ to $\mathbf{P}M_i$ for each $t_0 \in [0,1]$, and likewise for h_{m})

Given a branch b of $\mathbf{T}_{\mathrm{c}}$, we let $I_{\mathrm{c}}(b) \subset \mathbf{P}M_i$ denote the union of intervals b^* on $\mathbf{P}c^*$ for which $\mathbf{P}r_{\mathrm{c}}(x) \in \mathbf{P}b$ if and only if $\mathbf{P}h_{\mathrm{c}}(x,1) \in b^*$: i.e. $I_{\mathrm{c}}(b)$ is the set of arcs on $\mathbf{P}c^*$ sent to $\mathbf{P}b$ under the homotopy $\mathbf{P}h_{\mathrm{c}}$. Let $I_{\mathrm{m}}(b)$ be defined similarly for each branch b of $\mathbf{T}_{\mathrm{m}}$.

Conceivably, branches $b_{\mathrm{c}} \subset \mathbf{T}_{\mathrm{c}}$ and $b_{\mathrm{m}} \subset \mathbf{T}_{\mathrm{m}}$ could determine overlapping sets $I_{\mathrm{c}}(b_{\mathrm{c}})$ and $I_{\mathrm{m}}(b_{\mathrm{m}})$ on $\mathbf{P}c^*$. We now use Thurston's uniform injectivity theorem to show the following claim: for ϵ sufficiently small, any pair of intervals $a_{\mathrm{c}} \subset I_{\mathrm{c}}(b_{\mathrm{c}})$ and $a_{\mathrm{m}} \subset I_{\mathrm{m}}(b_{\mathrm{m}})$ are disjoint and separated by distance 2 along $\mathbf{P}c^*$.

Let $\lambda \supset |\mu^{\mathrm{r}}|$ be a maximal lamination realized by (g,X). Let $(X_\lambda)_i \in \mathcal{PS}(M_i)$ be the unique pleated surfaces realizing λ in M_i. Then the (marked) surfaces $\{(X_\lambda)_i\}$ converge to X in Teich(S) (see [**Bon3**, Thm. D]). We have chosen ϵ_{ns} (see lemma 4.5) so that the neighborhoods $\mathcal{N}_{\epsilon_{\mathrm{ns}}}(\nu_{\mathrm{c}})$ and $\mathcal{N}_{\epsilon_{\mathrm{ns}}}(\nu_{\mathrm{m}})$ of ν_{c} and ν_{m} on X are embedded and disjoint, so we may enlarge N_ϵ if necessary so that so that for all $i > N_\epsilon$ the neighborhoods $\mathcal{N}_{\epsilon_{\mathrm{ns}}/2}(\nu_{\mathrm{c}})$ and $\mathcal{N}_{\epsilon_{\mathrm{ns}}/2}(\nu_{\mathrm{m}})$ on $(X_\lambda)_i$ are embedded and disjoint as well.

Since the realization of μ^{r} on $(X_\lambda)_i$ lies in $(M_i)_{\geq \epsilon_0}$ for all $i > N_\epsilon$, any pair of points $x \in \nu_{\mathrm{c}}$ and $y \in \nu_{\mathrm{m}}$ as they sit on $(X_\lambda)_i$ lie in $((X_\lambda)_i)_{\geq \epsilon_0}$. By Thurston's uniform injectivity theorem (theorem 2.1) we have that corresponding to $\epsilon_{\mathrm{ns}}/2$ (and ϵ_0), there is a δ so that if $\ell_x \subset \nu_{\mathrm{c}}$ and $\ell_y \subset \nu_{\mathrm{m}}$ are the geodesic leaves through x and y, the distance in the projective tangent bundle satisfies

$$d_{\mathbf{P}M_i}((x,\ell_x),(y,\ell_y)) > \delta \tag{5.8}$$

for all $i > N_\epsilon$.

Thus, the subsets

$$Z_{\mathrm{c}} = \{(z,c^*) \in \mathbf{P}c^* \mid \text{for some } x \in \nu_{\mathrm{c}},\ d_{\mathbf{P}M_i}((z,c^*),(x,\ell_x)) < \delta/2\}$$

and

$$Z_{\mathrm{m}} = \{(z,c^*) \in \mathbf{P}c^* \mid \text{for some } y \in \nu_{\mathrm{m}},\ d_{\mathbf{P}M_i}((z,c^*),(y,\ell_y)) < \delta/2\}$$

where $\mathbf{P}c^*$ runs within $\delta/2$ of $\mathbf{P}\nu_{\mathrm{c}}$ and $\mathbf{P}\nu_{\mathrm{m}}$ in $\mathbf{P}M_i$ respectively, are disjoint.

Moreover, by a hyperbolic trigonometry argument, there is a constant $C_4 > 1$ so that if $\delta > 0$ is sufficiently small and we have

$$d_{\mathbf{P}M_i}((x,\ell_x),(z,c^*)) < \delta$$

then for any point (z',c^*) within distance 1 along $\mathbf{P}c^*$ of (z,c^*), there is a point x' on ℓ_x so that

$$d_{\mathbf{P}M_i}((x',\ell_x),(z',c^*)) < C_4\delta.$$

Assuming $\delta < 1$ is sufficiently small for this relation to hold, we collect our constraints on ϵ (see inequality 5.6 and equation 5.7): let

$$\epsilon_{\mathrm{tt}} = \min\left\{\epsilon_{\mathrm{ns}}, \frac{\delta}{4C_4C_3C_{\mathrm{tr}}}, \frac{\delta_{\mathrm{inj}}}{2(C_{\mathrm{tr}}+1)}, \frac{\epsilon'_L}{C_{\mathrm{inj}}2(C_{\mathrm{tr}}+1)e^{\ell_0}}, \frac{\epsilon_{\mathrm{trig}}}{C'_2C_3C_{\mathrm{tr}}}\right\}.$$

Given a branch $b_{\rm c}$ of $\mathbf{T}_{\rm c}$, let (z,c^*) be a point lying in $I_{\rm c}(b_{\rm c})$. Then for $\epsilon < \epsilon_{\rm tt}$ there is a point x on ℓ_x, a leaf of $\nu_{\rm c}$, so that

$$d_{\mathbf{P}M_i}((z,c^*),(x,\ell_x)) < 2C_{\rm tr}C_3\epsilon.$$

It follows that if $(z',c^*) \in \mathbf{P}c^*$ is any point within distance 1 of (z,c^*) along $\mathbf{P}c^*$, then (z',c^*) lies within $C_4(2C_{\rm tr}C_3\epsilon) < \delta/2$ of (x',ℓ_x) in $\mathbf{P}M_i$ for some point x' on ℓ_x. Thus, the radius 1 neighborhood of (z,c^*) along $\mathbf{P}c^*$ lies in $Z_{\rm c}$. We conclude that for each branch $b_{\rm c}$ of $\mathbf{T}_{\rm c}$, the radius 1 neighborhood along $\mathbf{P}c^*$ of the collection of arcs $I_{\rm c}(b_{\rm c})$ lies in $Z_{\rm c}$; similarly, for each branch $b_{\rm m}$ of $\mathbf{T}_{\rm m}$ the radius 1 neighborhood of $I_{\rm m}(b_{\rm m})$ along $\mathbf{P}c^*$ lies in $Z_{\rm m}$. Since $Z_{\rm c} \cap Z_{\rm m} = \varnothing$, the claim is proven.

If e is any arc in the complement

$$\mathbf{P}c^* - \{\cup_{b_{\rm c}\in\mathbf{T}_{\rm c}} I_{\rm c}(b_{\rm c})\} \bigcup \{\cup_{b_{\rm m}\in\mathbf{T}_{\rm m}} I_{\rm m}(b_{\rm m})\}$$

then either

1. there is a branch $\check{b}_{\rm c} \in \check{\mathbf{T}}_{\rm c}$ so that e is an arc of $I_{\rm c}(\check{b}_{\rm c})$,
2. there is a branch $\check{b}_{\rm m} \in \check{\mathbf{T}}_{\rm m}$ so that e is an arc of $I_{\rm m}(\check{b}_{\rm m})$, or
3. $\partial e = z_{\rm c} \sqcup z_{\rm m}$ where $z_{\rm c}$ lies in $\partial I_{\rm c}(b_{\rm c})$ and $z_{\rm m}$ lies in $\partial I_{\rm m}(b_{\rm m})$.

If e satisfies (1) let $b_e = \check{b}_{\rm c}$, and if e satisfies (2) let $b_e = \check{b}_{\rm m}$. If e satisfies

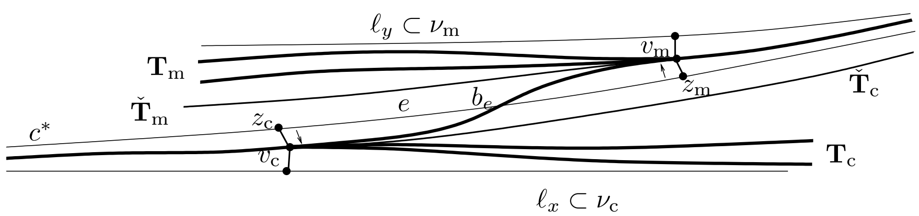

Figure 16. Melding train tracks.

(3), a perturbation of e moving each point at most $C_3C_{\rm tr}\epsilon$ in $\mathbf{P}M_i$ yields the lift to $\mathbf{P}M_i$ of a branch b_e with endpoints in switches $v_{\rm c} \in \partial b_{\rm c}$ and $v_{\rm m} \in \partial b_{\rm m}$ with which we may enlarge $\mathbf{T}_{\rm c} \sqcup \mathbf{T}_{\rm m}$. (See figure 16. The perturbation may be performed by taking a smooth convex combination along e of the homotopies $\mathbf{P}h_{\rm c}$ at $(z_{\rm c},c^*)$ and $\mathbf{P}h_{\rm m}$ at $(z_{\rm m},c^*)$).

For each e, then, add the branch b_e to $\mathbf{T}_{\rm c} \sqcup \mathbf{T}_{\rm m}$, collapsing any two branches b_e and b'_e from case (3) that are homotopic *rel*-endpoints into a single branch. The result is a generalized enlargement of $\mathbf{T}_{\rm c} \sqcup \mathbf{T}_{\rm m}$. By construction, the generalized enlargement $\mathbf{T}$ is the image of c^* under an explicit homotopy of c^* to a train path $\check{c}\colon S^1 \to M_i$ on $\mathbf{T}$.

Furthermore, since each e satisfying (3) has length at least $2 > \ell_0/2$ and we have $\epsilon < \epsilon_{\rm trig}/C_3C_{\rm tr}$, by repeating the argument of lemma 4.5 and letting

$$C_{\rm meld} = C_{\rm curv}(\ell_0/2)C_{\rm trig}C_{\rm tr},$$

the branches b_e from case (3) may be adjusted so that the generalized enlargement $\mathbf{T}$ is $C_{\rm meld}C_3\epsilon$-nearly-straight in M_i, provided $i > N_\epsilon$.

The enlargements $\check{\mathbf{T}}_{\rm c}$ and $\check{\mathbf{T}}_{\rm m}$ each carry c. Given a branch b of the sub-track $\mathbf{T}_{\rm c}$ or $\mathbf{T}_{\rm m}$, let $m_b(c)$ denote the weight deposited by c on b. The weight $m_b(c)$ is precisely the number of arcs on $\mathbf{P}c^*$ that are homotopic into $\mathbf{P}b$ under the homotopy of $\mathbf{P}c^*$ to $\mathbf{P}\check{c}$, which is precisely the number of arcs in $I_{\mathrm{c}}(b)$ if b is a branch of $\mathbf{T}_{\mathrm{c}}$ or $I_{\mathrm{m}}(b)$ if b is a branch of $\mathbf{T}_{\mathrm{m}}$. Thus we have

$$m_b(\check{c}) = m_b(c) \tag{5.9}$$

for any branch b of $\mathbf{T}_{\mathrm{c}}$ or $\mathbf{T}_{\mathrm{m}}$.

Conclusion. Provided $\epsilon < \epsilon_{\mathrm{tt}}$ and N_ϵ is chosen accordingly, then, for each $i > N_\epsilon$, letting $C_3 = \max\{C_{\mathrm{ns}}, C_2\}$, setting

$$(\tau^*_{\mathrm{m}})_i = \mathbf{T}_{\mathrm{m}} \quad \text{and} \quad (\tau^*_{\mathrm{c}})_i = \mathbf{T}_{\mathrm{c}}$$

and setting

$$(\check{\tau}^*_{\mathrm{m}})_i = \check{\mathbf{T}}_{\mathrm{m}} \quad \text{and} \quad (\check{\tau}^*_{\mathrm{c}})_i = \check{\mathbf{T}}_{\mathrm{c}}$$

in the above construction, we denote by $\check{\tau}_i = \mathbf{T}$ the resulting generalized enlargement of $(\tau^*_{\mathrm{m}})_i \sqcup (\tau^*_{\mathrm{c}})_i$ in M_i. The train track $\check{\tau}_i$ is a $C_{\mathrm{meld}}C_3\epsilon$-nearly-straight generalized enlargement of $(\tau^*)_i = (\tau^*_{\mathrm{m}})_i \sqcup (\tau^*_{\mathrm{c}})_i$ in M_i that carries μ^{r} and carries c_i by the train-path $\check{c}_i$.

Referring back to the remark before the proof, if $(b)_i$ is the image of the branch b of τ in its realization $(\tau^*)_i$ and generalized enlargement $\check{\tau}_i$ in M_i, we denote by $m_b(\check{c}_i)$ the weight $m_{(b)_i}(\check{c}_i)$ deposited by $\check{c}_i$ on $(b)_i$. Then by equation 5.9, we have that the weight $m_b(\check{c}_i)$ satisfies

$$m_b(t_i\check{c}_i) \to m_b(\mu)$$

for each branch b of τ. Thus, letting $C = C_{\mathrm{meld}}C_3$ proves the lemma.

∎

Although our primary aim here is to control the lengths and positions of *measured* laminations, we remark that the geometric techniques above did not use the convergence of the measured laminations $t_ic_i \to \mu$ but merely convergence in the Thurston topology of c_i to sublaminations $\nu \subset |\mu|$.

Corollary 5.3. *Let simple closed curves c_i converge in the Thurston topology to $\lambda = |\nu|$, for $\nu \in \mathcal{R}(M)$. Then given $\epsilon_1 > 0$ and $L_1 > 0$ so that λ is realized in $M_{\geq \epsilon_1}$ and each compact leaf of ν has length less than L_1 in M, there is a constant C depending only on S and ϵ_1 so that for positive $\epsilon < \epsilon_{\mathrm{tt}}$ there exists:*

1. *a train track τ minimally carrying λ with an ϵ-nearly-straight realization τ^* in M,*
2. *an integer $N_\epsilon > 0$ so that for all $i > N_\epsilon$, τ admits a realization in M_i with a $C\epsilon$-nearly-straight generalized enlargement $\check{\tau}_i$ so that c_i is carried by $\check{\tau}_i$ on a train-path $\check{c}_i$ for which if b is a branch of τ, we have $m_b(\check{c}_i) \to \infty$.*

∎

6. Continuity on realizable pairs

In this section we bring the nearly-straight train tracks of the previous section to bear on the length function.

Theorem 6.1. Length Continuous on Realizables *Let $(M_i, \mu_i) \in \mathfrak{R}$ converge to $(M, \mu) \in \mathfrak{R}$ in the product topology. Then*

$$\lim_{i\to\infty} \mathrm{length}_{M_i}(\mu_i) = \mathrm{length}_M(\mu).$$

Proof: Let μ lie in $\mathcal{R}(M)$. Since weighted isotopy classes of simple closed curves are dense in $\mathcal{ML}(S)$, it suffices to show that if the weighted isotopy classes $t_i c_i \in \mathcal{R}(M_i)$ converge to μ then the sequence $\{\mathrm{length}_{M_i}(t_i c_i)\}_{i=1}^{\infty}$ converges to $\mathrm{length}_M(\mu)$ and apply a diagonal argument.

Upper semi-continuity: Let $g\colon X \to M$ be a pleated surface realizing a maximal lamination λ containing μ. Then for i sufficiently large, λ is realizable in M_i with unique realizing pleated surface $(X_\lambda)_i \in \mathcal{PS}(M_i)$. The hyperbolic structures $(X_\lambda)_i$ converge to X in $\mathrm{Teich}(S)$ (again, by Bonahon's theorem [**Bon3**, Thm. D]).

By continuity of $\mathrm{length}\colon \mathrm{Teich}(S) \times \mathcal{ML}(S) \to \mathbb{R}$, it follows that

$$\lim_{i\to\infty} \mathrm{length}_{(X_\lambda)_i}(t_i c_i) = \mathrm{length}_X(\mu) = \mathrm{length}_M(\mu)$$

so, since $\mathrm{length}_{(X_\lambda)_i}(t_i c_i)$ bounds $\mathrm{length}_{M_i}(t_i c_i)$ for all i, we have

$$\limsup_{i\to\infty} \mathrm{length}_{M_i}(t_i c_i) \leq \mathrm{length}_M(\mu). \tag{6.10}$$

This proves upper semi-continuity.

We make the additional observation that since $\mathrm{length}_{M_i}(\mu) = \mathrm{length}_{(X_\lambda)_i}(\mu)$ for all i, we have

$$\lim_{i\to\infty} \mathrm{length}_{M_i}(\mu) = \mathrm{length}_M(\mu) \tag{6.11}$$

for later reference.

Lower semi-continuity: Assume μ is positive (otherwise there is nothing to prove). Fix $\epsilon > 0$ sufficiently small to satisfy the hypotheses of lemma 5.2, and let N_ϵ be the corresponding integer so that for $i > N_\epsilon$ we have

1. a train track τ minimally carrying μ with an ϵ-nearly-straight realization τ^* in M
2. $C\epsilon$-nearly-straight generalized enlargements $\check{\tau}_i$ of realizations τ_i of τ in M_i and train-paths $\check{c}_i\colon S^1 \to \check{\tau}_i$ carrying c_i so that for each branch b of τ, we have $m_b(t_i \check{c}_i) \to m_b(\mu)$.

Since $m_b(t_i \check{c}_i)$ converges to $m_b(\mu)$ for each branch b of τ, there is a sequence $\delta_i > 1$, tending to 1, so that we have

$$\delta_i m_b(\mu) \geq m_b(t_i \check{c}_i) \geq \frac{1}{\delta_i} m_b(\mu). \tag{6.12}$$

Since $\check{\tau}_i$ is $C\epsilon$-nearly-straight in M_i, we have

$$\mathrm{length}_{M_i}(t_i c_i) \geq \frac{1}{K(C\epsilon)} \ell_{\check{\tau}_i}(t_i \check{c}_i). \tag{6.13}$$

By equation 6.12 we have the inequality

$$\ell_{\check{\tau}_i}(t_i \check{c}_i) \geq \frac{1}{\delta_i} \ell_{\tau_i}(\mu).$$

Hence for all $i > N_\epsilon$, we have

$$\mathrm{length}_{M_i}(t_i c_i) \geq \frac{1}{K(C\epsilon)\delta_i} \ell_{\tau_i}(\mu) \geq \frac{1}{K(C\epsilon)\delta_i} \mathrm{length}_{M_i}(\mu).$$

Since $\mathrm{length}_{M_i}(\mu) \to \mathrm{length}_M(\mu)$ (6.11), it follows that

$$\liminf_{i\to\infty} \mathrm{length}_{M_i}(t_i c_i) \geq \frac{1}{K(C\epsilon)} \mathrm{length}_M(\mu).$$

Letting ϵ approach 0, $K(C\epsilon)$ tends to 1, so we may conclude the lower-semi-continuity

$$\liminf_{i\to\infty} \operatorname{length}_{M_i}(t_i c_i) \geq \operatorname{length}_M(\mu). \tag{6.14}$$

The theorem follows.

∎

Corollary 6.2. *Let $M_i \to M$ in $AH(S)$ and let $\mu \in \mathcal{R}(M)$. Then for i sufficiently large there are realizations τ_i in M_i of (finer and finer) train tracks minimally carrying μ so that*

$$\lim_{i\to\infty} \ell_{\tau_i}(\mu) = \operatorname{length}_M(\mu).$$

Proof: This follows easily from a diagonal argument using the observation that by analogy with inequality 6.13, τ_i satisfies the double inequality

$$\ell_{\tau_i}(\mu) \geq \operatorname{length}_{M_i}(\mu) \geq \frac{1}{K(C\epsilon)}\ell_{\tau_i}(\mu). \tag{6.15}$$

∎

7. Extending to non-realizable pairs

Given a pair $(M,\mu) \in AH(S) \times \mathcal{ML}(S)$, the function

$$(M,\mu) \to \operatorname{length}_M(\mathrm{R}_M(\mu))$$

is a natural extension of length: $\mathfrak{R} \to \mathbb{R}$ to all of $AH(S) \times \mathcal{ML}(S)$. When μ is connected and non-realizable, it agrees with the function $\underline{\operatorname{length}}_M(\mu)$ given by taking lim inf of the infima of $\operatorname{length}_M(\nu)$ over realizable laminations $\nu \in \mathcal{R}(M)$ in smaller and smaller neighborhoods of μ (see [**Th1**, Ch. 9] [**Th4**, §3], [**Bon1**, Lem. 5.1]). In this section we show the these functions are the same continuous function.

Theorem 7.1. Length Extends Continuously *The function*

$$\underline{\operatorname{length}}\colon AH(S) \times \mathcal{ML}(S) \to \mathbb{R}$$

is continuous.

By lemma 5.2, the lower-semi-continuity argument of theorem 6.1 applies with $\mathrm{R}_M(\mu) = \mu^{\mathrm{r}}$ in place of μ to establish the following analog of equation 6.14 from its proof.

Lemma 7.2. *Let $\{(M_i, t_i c_i)\} \subset \mathfrak{R}$ converge to $(M,\mu) \in AH(S) \times \mathcal{ML}(S)$. Then*

$$\liminf_{i\to\infty} \operatorname{length}_{M_i}(t_i c_i) \geq \operatorname{length}_M(\mu^{\mathrm{r}}).$$

∎

Proof: (of theorem 7.1). Let $\{(M_i, t_i c_i)\} \subset \mathfrak{R}$ converge to $(M, \mu) \in AH(S) \times \mathcal{ML}(S)$ as above. Let μ admit the decomposition $\mu = \mu^{\mathrm{r}} \cup \mu^{\mathrm{nr}}$, where $\mu^{\mathrm{r}} = \mathrm{R}_M(\mu)$, and $\mu^{\mathrm{nr}} = \mu - \mathrm{R}_M(\mu)$ is the maximal non-realizable sublamination of μ in M.

Shortening non-realizables. By the shortening process of [**Bon1**, Lem. 5.1], there is a sequence of train tracks τ_n^{nr} in M, so that

1. each τ_n^{nr} minimally carries μ^{nr}, and
2. we have the limit
$$\lim_{n \to \infty} \ell_{\tau_n^{\mathrm{nr}}}(\mu^{\mathrm{nr}}) = 0.$$

Straightening realizables. Using corollary 6.2 of lemma 5.2, we construct a sequence of train tracks τ_n^{r} in M so that

1. each τ_n^{r} minimally carries μ^{r}, and
2. we have the limit
$$\lim_{n \to \infty} \ell_{\tau_n^{\mathrm{r}}}(\mu^{\mathrm{r}}) = \mathrm{length}_M(\mu^{\mathrm{r}}).$$

We let τ_n be the train track
$$\tau_n = \tau_n^{\mathrm{r}} \cup \tau_n^{\mathrm{nr}}$$
in M, and note that
$$\lim_{n \to \infty} \ell_{\tau_n}(\mu) = \mathrm{length}_M(\mu^{\mathrm{r}}).$$

Pass to a subsequence of $\{c_i\}$ that converges in the Hausdorff topology. Then there is a sequence of enlargements of τ_n to a train track τ_n' in M such that τ_n' carries c_i for all $i > N_n$.

The track-length $\ell_{\tau_n'}(t_i c_i)$ satisfies
$$\mathrm{length}_M(t_i c_i) \le \ell_{\tau_n'}(t_i c_i)$$
for each $i > N_n$. Since for any branch b of $\tau_n' - \tau_n$ the masses $\{m_{t_i c_i}(b)\}$ tend to 0 as i tends to ∞, it follows that
$$\lim_{i \to \infty} \mathrm{length}_M(t_i c_i) \le \ell_{\tau_n}(\mu)$$
for each n.

Algebraic convergence. For each n, the train track τ_n' is contained in a compact set $K_n \subset M$. Let real numbers $L_n > 1$ tend to 1. For each n, let $N_n' > N_n$ be a positive integer such that there are smooth, marking-preserving homotopy equivalences
$$q_i \colon M \to M_i$$
so that $\|q_i\|_{C^1} < L_n$ on K_n for all $i > N_n'$. Then we have
$$\ell_{q_i(\tau_n')}(t_i c_i) \le L_n \ell_{\tau_n'}(t_i c_i)$$
for all $i > N_n'$. Diagonalizing, we have
$$\limsup_{i \to \infty} \mathrm{length}_{M_i}(t_i c_i) \le \mathrm{length}_M(\mu^{\mathrm{r}}),$$
and thus by lemma 7.2
$$\lim_{i \to \infty} \mathrm{length}_{M_i}(t_i c_i) = \mathrm{length}_M(\mu^{\mathrm{r}}).$$

It follows from an additional diagonal argument that the function $\mathrm{length}_M(\mu)$ on $\mathfrak{R}$ extends continuously to the function
$$(M, \mu) \to \mathrm{length}_M(\mathrm{R}_M(\mu))$$

on $\overline{\mathfrak{R}} = AH(S) \times \mathcal{ML}(S)$ which is therefore equal to $\underline{\text{length}}$ on $AH(S) \times \mathcal{ML}(S)$. ∎

The following corollary is an immediate consequence:

Corollary 7.3. Zero-Locus *Let* $(M_i, \mu_i) \to (M, \mu)$ *in* $AH(S) \times \mathcal{ML}(S)$ *so that the sequence*

$$\left\{ \underline{\text{length}}_{M_i}(\mu_i) \right\}_{i=1}^{\infty}$$

converges to 0*. Then* $\mathrm{R}_M(\mu) = 0$.

∎

References

[Ah] L. Ahlfors. An extension of Schwarz's lemma. *Trans. Amer. Math. Soc.* **43**(1938), 359–364.

[BP] R. Benedetti and C. Petronio. *Lectures on Hyperbolic Geometry.* Springer-Verlag, 1992.

[Bon1] F. Bonahon. Bouts des variétés hyperboliques de dimension 3. *Annals of Math.* **124**(1986), 71–158.

[Bon2] F. Bonahon. The geometry of Teichmüller space via geodesic currents. *Invent. math.* **92**(1988), 139–162.

[Bon3] F. Bonahon. Shearing hyperbolic surfaces, bending pleated surfaces, and Thurston's symplectic form. *Ann. Fac. Sci. Toulouse Math.* **5**(1996), 233–297.

[Bon4] F. Bonahon. Geodesic laminations with transverse Hölder distributions. *Ann. scient. Éc. Norm. Sup.* **30**(1997), 205–240.

[BO] F. Bonahon and J. P. Otal. Variétés hyperboliques à géodésiques arbitrairement courtes. *Bull. London Math. Soc.* **20**(1988), 255–261.

[Br1] J. Brock. Boundaries of Teichmüller spaces and geodesic laminations. *To appear, Duke Math. J.*

[Br2] J. Brock. Iteration of mapping classes and limits of hyperbolic 3-manifolds. *Submitted to Invent. Math.*

[BM] R. Brooks and J. P. Matelski. Collars for Kleinian Groups. *Duke Math. J.* **49**(1982), 163–182.

[Bus] P. Buser. *Geometry and Spectra of Compact Riemann Surfaces.* Birkhauser Boston, 1992.

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

[CEG] R. D. Canary, D. B. A. Epstein, and P. Green. Notes on notes of Thurston. In *Analytical and Geometric Aspects of Hyperbolic Space*, pages 3–92. Cambridge University Press, 1987.

[Fl] W. Floyd. Group completions and limit sets of Kleinian groups. *Invent. Math.* **57**(1980), 205–218.

[Hat] A. Hatcher. On triangulations of surfaces. *Topology Appl.* **40**(1991), 189–194.

[Mc] C. McMullen. *Renormalization and 3-Manifolds Which Fiber Over the Circle.* Annals of Math. Studies 142, Princeton University Press, 1996.

[Min1] Y. Minsky. Harmonic maps into hyperbolic 3-manifolds. *Trans. AMS* **332**(1992), 539–588.

[Min2] Y. Minsky. Harmonic maps, length, and energy in Teichmüller space. *J. Diff. Geom.* **35**(1992), 151–217.

[Min3] Y. Minsky. Dehn twists have linear growth. *Preprint.*

[Min4] Y. Minsky. Kleinian groups and the complex of curves. *Preprint.*

[Ohs] K. Ohshika. Divergent sequences of Kleinian groups. *Geometry and Topology Monographs Volume 1: The Epstein Birthday Schrift, paper no. 21* **1**(1998), 419–450.

[Otal] J. P. Otal. *Le théorème d'hyperbolisation pour les variétés fibrées de dimension trois.* Astérisque, 1996.

[PH] R.C. Penner and J.L. Harer. *Combinatorics of Train Tracks.* Annals of Math. Studies 125, Princeton University Press, 1992.

[Th1] W. P. Thurston. *Geometry and Topology of Three-Manifolds*. Princeton lecture notes, 1979.

[Th2] W. P. Thurston. Hyperbolic structures on 3-manifolds I: Deformations of acylindrical manifolds. *Annals of Math.* **124**(1986), 203–246.

[Th3] W. P. Thurston. *Three-Dimensional Geometry and Topology*. Princeton University Press, Princeton, NJ, 1997.

[Th4] W. P. Thurston. Hyperbolic structures on 3-manifolds II: Surface groups and 3-manifolds which fiber over the circle. *Preprint, 1986*.

[Th5] W. P. Thurston. Minimal stretch maps between hyperbolic surfaces. *Preprint, 1986*.

Department of Mathematics, Stanford University, Stanford, CA 94305
E-mail address: brock@math.stanford.edu