

THE INFINITESIMAL AND GLOBAL THURSTON GEOMETRY OF TEICHMÜLLER SPACE

YI HUANG, KEN'ICHI OHSHIKA AND ATHANASE PAPADOPOULOS

ABSTRACT. We undertake a systematic study of the infinitesimal geometry of the Thurston metric, showing that the topology, convex geometry and metric geometry of the tangent and cotangent spheres based at any marked hyperbolic surface representing a point in Teichmüller space can recover the marking and geometry of this marked surface. We then translate the results concerning the infinitesimal structures to global geometric statements for the Thurston metric, most notably deriving rigidity statements for the Thurston metric analogous to the celebrated Royden theorem.

Keywords.— Teichmüller space, Thurston metric, rigidity, hyperbolic surfaces, stretch map, convex geometry.

AMS classification.— 32G15, 30F60, 30F10, 52A21.

1. INTRODUCTION

An important and beautiful legacy of S. S. Chern's œuvre in differential geometry is the idea that many *global* geometric and topological properties of a manifold may be inferred from its *infinitesimal* structure. Our goal is to investigate Thurston's Finsler metric geometry on Teichmüller space along this theme. Throughout this paper,

- $S = S_{g,n}$ denotes an orientable hyperbolic surface of genus $g \geq 0$, and with $n \geq 0$ cusps (but no other boundaries),
- and $\mathcal{T}(S)$ denotes the Teichmüller space of S , i.e., the space of isotopy classes of complete finite-area hyperbolic metrics on S .

Here are some highlights of our work.

STRUCTURED STRATIFICATION: for any x in $\mathcal{T}(S)$, we establish and study stratifications on the Thurston Finsler unit norm ball in $T_x\mathcal{T}(S)$ and on the unit co-norm ball in $T_x^*\mathcal{T}(S)$ induced by their convex geometries. These stratifications are highly heterogeneous, with polytope-like strata, many of which are naturally labelled by geodesic laminations in a functorially structured manner.

LENGTH EXTRACTION: we obtain a limiting procedure for recovering any hyperbolic metric $x \in \mathcal{T}(S)$ from the Thurston Finsler norm on the tangent space $T_x\mathcal{T}(S)$ at x .

INFINITESIMAL RIGIDITY: we prove the Thurston metric analogue of the (stronger and) infinitesimal form of the celebrated Royden theorem. We show that isometries of the Thurston Finsler norm $\|\cdot\|_{\text{Th}}$ on $T_x\mathcal{T}(S)$ are precisely given by extended mapping classes on S . Stated precisely,

Theorem 1.1 (Infinitesimal rigidity). *Consider two arbitrary points $x, y \in \mathcal{T}(S)$. The normed vector spaces*

$$(T_x\mathcal{T}(S), \|\cdot\|_{\text{Th}}) \text{ and } (T_y\mathcal{T}(S), \|\cdot\|_{\text{Th}}) \text{ are isometric}$$

if and only if the hyperbolic surfaces (S, x) and (S, y) are isometric.

This project grew out of our attempt to prove Theorem 1.1, and the second author presented this earlier work in a series of lectures given at a CIMPA school in Varanasi in December 2019 [21]. In the course of preparing the present paper, we learned of Pan's independent work [22] establishing Theorem 1.1. Our respective proof strategies and the flavour of our work differ substantially: Pan's approach makes beautiful use of [30, Theorem 3.1 and Theorem 10.1] to efficiently establish infinitesimal rigidity. In contrast, in part informed by conversations with David Dumas and Kasra Rafi, we opt to follow Thurston's philosophy [30]:

In the course of this paper, we will develop some other ideas which are of interest in their own right. The intent is not to give the slickest proof of the main theorem, but to develop a good picture.

We now present an overview of the picture we have learned.

1.1. The Thurston metric. Let $\mathcal{S} = \mathcal{S}(S)$ denote the set of free homotopy classes of essential simple closed curves on S , that is, simple closed curves which are neither homotopic to a point nor to a cusp. In [30], Thurston introduces the following quantity, defined for any *ordered* pair $x, y \in \mathcal{T}(S)$

$$(1) \quad d_{\text{Th}}(x, y) := \sup_{\gamma \in \mathcal{S}} \log \frac{\ell_y(\gamma)}{\ell_x(\gamma)},$$

where $\ell_z(\gamma)$ denotes the hyperbolic length of the unique geodesic on (S, z) freely homotopic to γ . He shows that $e^{d_{\text{Th}}(x, y)}$ is the optimal Lipschitz constant for self-homeomorphisms of S homotopic to $\text{id}_S : (S, x) \rightarrow (S, y)$ (see Eq. (9) in §2.1), and that d_{Th} defines an asymmetric Finsler metric on $\mathcal{T}(S)$. In particular, Thurston's theory furnishes the following description of the asymmetric Finsler norm of an arbitrary tangent vector $v \in T_x\mathcal{T}(S)$:

$$(2) \quad \|v\|_{\text{Th}} = \sup_{\gamma \in \mathcal{S}} (d \log \ell_{(\cdot)}(\gamma))_x(v) = \sup_{\gamma \in \mathcal{S}} \left(\frac{v(\ell_{(\cdot)}(\gamma))}{\ell_x(\gamma)} \right),$$

where $\ell_{(\cdot)}(\gamma) : \mathcal{T}(S) \rightarrow \mathbb{R}$ measures the length of the geodesic representative of γ on x for each $\gamma \in \mathcal{S}$.

1.2. Infinitesimal structures on tangent and cotangent space. The derivative $d \log \ell(\lambda)$ of $\log \ell_{(\cdot)}(\lambda) : \mathcal{T}(S) \rightarrow \mathbb{R}$ at $x \in \mathcal{T}(S)$ defines a linear function from the tangent space $T_x \mathcal{T}(S)$ to \mathbb{R} , and hence defines an element of the cotangent space $T_x^* \mathcal{T}(S)$. Moreover, $d \log \ell(\lambda)$ depends only on the projective class of λ and defines a map

$$(3) \quad \iota_x : \mathcal{PM}\mathcal{L}(S) \rightarrow T_x^* \mathcal{T}(S), \quad [\lambda] \mapsto d_x \log \ell_{(\cdot)}(\lambda)$$

Theorem 1.2 (Thurston, [30]). *The map ι_x is an embedding of $\mathcal{PM}\mathcal{L}(S)$ into $T_x^* \mathcal{T}(S)$. The image set*

$$\mathbf{S}_x^* := \iota_x(\mathcal{PM}\mathcal{L}(S)) \subsetneq T_x^* \mathcal{T}(S)$$

bounds a convex ball of full dimension in $T_x^ \mathcal{T}(S)$ containing the origin.*

Setting $\mathbf{S}_x^* \subsetneq T_x^* \mathcal{T}(S)$ to be a unit co-norm sphere and extending by positive homothety defines an asymmetric norm $\|\cdot\|_{\text{Th}}^*$ on each $T_x^* \mathcal{T}(S)$. Its dual norm $\|v\|_{\text{Th}}^{**}$ on the tangent space $T_x \mathcal{T}(S)$ coincides with the original norm $\|\cdot\|_{\text{Th}}$. To see this, observe that

$$(4) \quad \begin{aligned} \|v\|_{\text{Th}}^{**}(v) &:= \sup \{w^*(v) \mid w^* \in \mathbf{S}_x^*\} \\ &= \sup \{\iota_x([\lambda])(v) \in \mathbb{R} \mid [\lambda] \in \mathcal{PM}\mathcal{L}(S)\} = \|v\|_{\text{Th}}(v). \end{aligned}$$

The right-most equality in Eq. (4) follows from the denseness of the set of projectivised simple closed closed geodesics in $\mathcal{PM}\mathcal{L}(S)$. We denote the unit sphere for the norm $\|v\|_{\text{Th}}$ by $\mathbf{S}_x \subsetneq T_x \mathcal{T}(S)$.

1.3. The unit norm sphere \mathbf{S}_x . In Section 5, we develop the general theory of convex bodies, show that any convex body has a *convex stratification* Definition 5.11 which decomposes it into (generically uncountably many) convex strata given by the interior of its *faces* (Definition 5.6). Thurston shows, in a *deep* result [30, Theorem 10.1], that the top-dimensional faces of are naturally indexed by simple closed curves and cover almost almost all of the unit Thurston norm sphere $\mathbf{S}_x \subsetneq T_x \mathcal{T}(S)$.

The aforementioned correspondence between the top-dimensional faces and simple closed curves is explicit: for any $\gamma \in \mathcal{S}$, its associated top-dimensional face is given by

$$\{v \in \mathbf{S}_x \mid \iota_x([\gamma])(v) := (d \log \ell_{(\cdot)}(\gamma))_x(v) = \|v\|_{\text{Th}}\}.$$

The above expression depends only on the projective measure class of γ , and we generalise it to general projective measured laminations $[\lambda] \in \mathcal{PM}\mathcal{L}(S)$ via

$$(5) \quad N_x([\lambda]) := \{v \in \mathbf{S}_x \mid \iota_x([\lambda])(v) = \|v\|_{\text{Th}}\}.$$

The set $N_x([\lambda])$ is necessarily an exposed face of \mathbf{S}_x .

Corollary 1.3 (contravariant labelling, Corollary 6.15). *For two arbitrary projective multicurves $[\Gamma_1], [\Gamma_2] \in \mathcal{PML}(S)$, denote their respective supports by $|\Gamma_1|$ and $|\Gamma_2|$. Then,*

$$(6) \quad |\Gamma_1| \subseteq |\Gamma_2| \text{ if and only if } N_x([\Gamma_2]) \subseteq N_x([\Gamma_1]),$$

with equality on the left if and only if we have equality on the right.

The above proposition induces the embedding of a naïve “dual” to the curve complex as a subcomplex of \mathbf{S}_x .

Corollary 1.4 (embedded dual curve complex, Corollary 6.32). *For every $x \in \mathcal{T}(S)$, the convex stratification on \mathbf{S}_x contains an ideal cell-complex \mathbf{C}_x which is dual to the curve complex in the sense that*

- *the support $|\Gamma|$ for any projective multicurve $[\Gamma] \in \mathcal{PML}$ is assigned to the face $N_x([\Gamma]) \subsetneq \mathbf{S}_x$;*
- *the subset-relation for cells in the curve complex inverts to the superset-relation for faces in \mathbf{C}_x ;*
- *the dimension of each cell in the curve complex is equal to the codimension of its corresponding face in \mathbf{C}_x as a subset of \mathbf{S}_x .*

This polytope-like convex geometry of \mathbf{S}_x^* tells us that the Finsler structure of Thurston’s metric is highly heterogeneous, lending further support for the infinitesimal rigidity of the Thurston metric.

1.4. The unit co-norm sphere \mathbf{S}_x^* . The convex stratification of the dual unit co-norm sphere \mathbf{S}_x^* enjoys a cleaner and more natural “dual” statement:

Corollary 1.5 (embedded curve complex, Corollary 6.13). *The convex stratification of \mathbf{S}_x^* contains the curve complex for S as a subcomplex \mathbf{C}_x^* .*

The convex stratification of \mathbf{S}_x^* is comprised of uncountable many convex faces of dimensions varying from 0 to $3g - 4 + n$ (Corollary 6.11), with many of the faces “covariantly” labelled by geodesic laminations which support measured laminations. Let $|\mathcal{ML}|$ denote the set of supporting laminations for (projective) measured laminations on S , and let \mathcal{F} denote the set of faces of \mathbf{S}_x^* . We define an injective map (see Lemma 6.4)

$$F_{(\cdot)}: |\mathcal{ML}| \rightarrow \mathcal{F}, \quad |\lambda| \mapsto F_{|\lambda|}$$

which assigns to the support $|\lambda| \in |\mathcal{ML}|$ of an arbitrary measured lamination $\lambda \in \mathcal{ML}$ the minimal face $F_{|\lambda|}$ containing the ι_x -image of the set of (projective) measured laminations supported on $|\lambda|$. We show in Theorem 6.7 that $F_{|\lambda|}$ consists precisely of the ι_x -image of projective measured laminations supported on $|\lambda|$, thereby deriving the following

Corollary 1.6 (covariant labelling, Corollary 6.9). *For two arbitrary projective laminations $[\lambda], [\mu] \in \mathcal{PML}(S)$,*

$$|\mu| \subsetneq |\lambda| \text{ if and only if } F_{|\mu|} \subsetneq F_{|\lambda|}.$$

1.5. Closure correspondence. There is deeper structure attached to the face-assigning map $F_{(\cdot)}$. We introduce two notions of closure,

- *support closure* (Definition 6.3) for laminations: given an arbitrary measured lamination λ , we add to its support lamination $|\lambda|$ the boundary of the smallest (up to isotopy) incompressible subsurface of S containing λ and denote the new geodesic lamination by $\widehat{|\lambda|}$;
- *adherence closure* (Definition 5.15) for faces of convex bodies: this assigns to each face F a specific superface \widehat{F} in a manner purely determined by the convex geometry of \mathbf{S}_x^* .

We show, as a corollary to Theorem 6.16 and Proposition 6.25, that

Corollary 1.7 (Corollary 6.26). *For any projective measured lamination $[\lambda]$,*

$$\widehat{F_{[\lambda]}} = F_{\widehat{[\lambda]}}.$$

The above result asserts that the closure operation $\widehat{}$ defines a natural transformation on $F_{(\cdot)}$ — regarded as a functor between two subcategories of the category of sets (see Question 7.13).

1.6. Topological rigidity. Any homeomorphism $f^* : \mathbf{S}_y^* \rightarrow \mathbf{S}_x^*$ between two unit co-norm spheres induces a homeomorphism

$$f_{y,x} := \iota_x^{-1} \circ f^* \circ \iota_y : \mathcal{PML}(S) \rightarrow \mathcal{PML}(S).$$

Corollary 6.32, combined with the rigidity of the curve complex [12, 15, 18] and the denseness of projectivised multicurves in \mathcal{PML} , tells us that if f^*

- preserves the convex stratification, and
- takes the embedded curve complex $\mathbf{C}_y^* \subsetneq \mathbf{S}_y^*$ to the curve complex in $\mathbf{C}_x^* \subsetneq \mathbf{S}_x^*$,

then $\iota_x \circ f_{y,x} \circ \iota_y^{-1}$ is induced by an extended mapping class.

In particular, linear isometries from $(T_y^* \mathcal{T}(S), \|\cdot\|_{\text{Th}})$ to $(T_x^* \mathcal{T}(S), \|\cdot\|_{\text{Th}})$ restrict to homeomorphisms from \mathbf{S}_y^* to \mathbf{S}_x^* which preserve the convex stratification (Theorem 5.34) and take \mathbf{C}_y^* to \mathbf{C}_x^* — we prove the latter claim using various adherence-related properties to distinguish the faces in \mathbf{C}_x^* from the other faces in \mathbf{S}_x^* . This yields:

Theorem 1.8 (Topological rigidity). *Let $f^* : T_y^* \mathcal{T}(S) \rightarrow T_x^* \mathcal{T}(S)$ denote the linear isometry dual to a linear Thurston-norm isometry $f : T_x \mathcal{T}(S) \rightarrow T_y \mathcal{T}(S)$. Then the map $f_{y,x} = \iota_x^{-1} \circ f^* \circ \iota_y : \mathcal{PML}(S) \rightarrow \mathcal{PML}(S)$ is necessarily induced by some diffeomorphism $h : (S, x) \rightarrow (S, y)$.*

Remark 1.9. Theorem 1.8 is key to our strategy for proving infinitesimal rigidity of the Thurston metric, and relies upon Ivanov’s theorem [12, 15, 18]. This argument fails when $(g, n) = (1, 1)$ or $(0, 4)$, which are precisely the two cases are (effectively) covered by [6].

1.7. Thurston's stretch maps. We have hitherto mainly considered the infinitesimal theory of the Thurston metric, and Thurston bridges this to the global perspective via objects called *stretch paths* (Section 2.3). He (implicitly) proves that the Finsler metric defined by Eq. (2) is equal to the asymmetric metric defined by Eq. (1) by showing that they are both geodesic metrics with common geodesics obtained from concatenating specific sequences of stretch paths.

Stretch paths are parametrised by three inputs: an initial point $x \in \mathcal{T}(S)$, a complete geodesic lamination Λ on (S, x) , and a time period over which to stretch. A time- t stretch uniformly increases the length of every geodesic segment in Λ by a factor of e^t whilst non-uniformly shrinking the “dual” perpendicular measured foliation μ , which is determined by the initial point $x \in \mathcal{T}(S)$. The third author shows in [24] that stretch paths always converge in forward time to the projective measured lamination representing μ in the Thurston compactification. The forward-time flow induced by stretch paths is highly dependent on the starting point and provides a codimension-1 projection map from $\mathcal{T}(S)$ to $\mathcal{PML}(S)$ when S has no cusps [25].

In Section 3, we explicitly express stretch paths for all complete finite-leaf laminations as real analytic functions of certain Fenchel-Nielsen length and twist parameters, and Thurston's shearing parameters on Teichmüller space (see [9] and [5]), and give a new proof to a result due to Guillaume Th  ret [29, Theorem 4.1]:

Theorem 1.10 (back-time convergence Theorem 3.8). *Consider a finite-leaf lamination Λ , and denote its closed leaves by $\gamma_1, \dots, \gamma_k$. As $t \rightarrow \infty$, for any $x \in \mathcal{T}(S)$, the e^{-t} -stretch path with respect to Λ $x_{-t} := \text{stretch}(x, \Lambda, -t)$ converges to the uniform weighted projectivised multicurve*

$$[\gamma_1 + \gamma_2 + \dots + \gamma_k] \in \mathcal{PML}(S) = \partial_\infty \mathcal{T}(S),$$

where $\mathcal{PML}(S) = \partial_\infty \mathcal{T}(S)$ denotes the boundary of the Thurston compactification of $\mathcal{T}(S)$. Note that this is independent of the starting point x .

Remark 1.11. There is a proto-dictionary between the structure of various \mathcal{PML} -related flows on $\mathcal{T}(S)$ and the Anosov nature of geodesic flow on the unit tangent bundle $T^1\mathbb{H}$ of hyperbolic space \mathbb{H} :

- the projectivisation $\mathcal{PML}(S)$ of the measured lamination space $\mathcal{ML}(S)$ defines a boundary at infinity of $\mathcal{T}(S)$, just as the projectivisation $\mathbb{P}(L^+) = \mathbb{S}^1$ of the positive light-cone L^+ in Minkowski space $\mathbb{R}^{2,1}$ (see [26, pg. 302]) defines the ideal boundary of (the hyperboloid model of) $\mathbb{H} \subsetneq \mathbb{R}^{2,1}$;
- the Thurston Finsler cometric on the cotangent bundle of $\mathcal{T}(S)$ has conorm given by taking $\iota_x(\mathcal{PML}(S)) \subsetneq T_x^*\mathcal{T}(S)$ as the unit sphere [30, Theorem 5.1]. Analogously, the standard hyperbolic cometric on the cotangent bundle of \mathbb{H} has conorm given by embedding $\mathbb{P}L^+$ as

the unit circle in $T^*\mathbb{H}$, where the embedding is defined (cf. Eq. (3)) as follows: for $u \in \mathbb{H} \subsetneq \mathbb{R}^{2,1}$ and $[v] \in \mathbb{P}L^+$, define

$$(7) \quad \iota_u: \mathbb{P}L^+ \rightarrow T_u^*\mathbb{H}, \quad [v] \mapsto d_u \log \langle \cdot, v \rangle,$$

where $\langle \cdot, \cdot \rangle$ here denotes the Minkowski metric on $\mathbb{R}^{2,1}$;

- consequentially, the Thurston Finsler metric on $\mathcal{T}(S)$ parallels the hyperbolic metric on \mathbb{H} , and hence the unit tangent bundle $T^1\mathcal{T}(S)$ is akin to the unit tangent bundle $T^1\mathbb{H}$;
- the flow induced by Thurston stretch maps is akin to the geodesic flow — it is expanding in positive time and contracting in negative time;
- the earthquake flow with respect uniformly weighted projectivised multicurves (which constitute the closed leaves of a finite-leaf lamination) is akin to the unstable horocyclic flow;
- the earthquake flow with respect to the dual perpendicular measured foliation is analogous to the stable horocyclic flow.

1.8. Stretch vectors. Thurston shows that his stretch path are analytic with respect to its time parameter [30, Corollary 4.2], and we refer to their tangent vectors as *stretch vectors* Definition 4.1. Thurston's concatenated stretch path construction for Thurston metric geodesics [30, Theorem 8.5] between arbitrary points $x \in \mathcal{T}(S)$ and $y \in \mathcal{T}(S)$ is analogous Dantzig's simplex method in linear programming:

- the role of the finitely-constrained feasible region is served by the infinitely-constrained *envelope* $\text{Env}(x, y) \subsetneq \mathcal{T}(S)$ (see [6]) comprised of points on Thurston geodesics from x to y ;
- just as the feasible region is convex, one naturally expects $\text{Env}(x, y)$ to be geodesically convex with respect to the Thurston metric;
- the point x is the chosen initial basic feasible solution (i.e., extreme point), and the goal is to maximise the Thurston metric distance from x — the maximum distance is uniquely attained by y ;
- one expects the edges of $\text{Env}(x, y)$ to be given by stretch paths with respect to maximal chain-recurrent laminations, and so we start at x and traverse along arbitrary stretch paths whilst increasing the distance from x , and occasionally pivoting to new basic feasible solutions until we reach y .

Consider the cross-sections of $\text{Env}(x, y)$ given distance from x . As the distance tends to 0, we expect to get the Thurston-norm unit sphere \mathbf{S}_x , and so the above conjectural description suggests the following conjecture

Conjecture 1.12. The set of stretch vectors in \mathbf{S}_x of stretch paths with respect to maximal chain-recurrent laminations characterise the set of extreme points of \mathbf{S}_x .

Conjecture 1.12 is true when S is either the once-punctured torus or the 4-punctured sphere [6], but is unknown for surfaces of general type.

In fact, neither direction of the characterisation is known. Nevertheless, the conjecture being true would imply the set of stretch vectors (with respect to maximal chain-recurrent laminations) would be preserved under linear Thurston-norm isometries. We provide support for Conjecture 1.12 by showing something stronger:

Corollary 1.13 (equivariant stretch vectors, Corollary 7.7). *Given an arbitrary Thurston-norm isometry $f: T_x\mathcal{T}(S) \rightarrow T_y\mathcal{T}(S)$, let $[h] \in \text{Mod}^*(S)$ denote its inducing extended mapping class (topological rigidity, Theorem 1.8). For an arbitrary maximal chain-recurrent lamination Λ of S , the stretch vector $v_\Lambda \in T_x\mathcal{T}(S)$ satisfies the following equivariance property:*

$$f(v_\Lambda) = v_{h(\Lambda)} \in T_y\mathcal{T}(S).$$

Remark 1.14. It is tempting to wonder if Corollary 7.7 (and even Conjecture 1.12) might be derived from Thurston's [30, Theorem 8.4]. The difficulty lies in showing that faces $N_x([\lambda])$ vary Hausdorff continuously with respect to varying $x \in \mathcal{T}(S)$. We prove Corollary 7.7 by showing that stretch vectors with respect to maximal chain-recurrent laminations necessarily arise as the Hausdorff limit of a sequence of shrinking faces of \mathbf{S}_x (Theorem 7.6) and then invoking topological rigidity.

1.9. Geometric rigidity. We show in Proposition 7.8 that the simple length spectrum for any hyperbolic metric $x \in \mathcal{T}(S)$ may be extracted from the asymptotic behaviour of Thurston norms of differences between certain stretch vector pairs.

For any simple closed geodesic $\gamma \in \mathcal{S}$, there are sequences $\{v_{+,m}, v_{-,m}\}$ of stretch vector pairs such that

$$\ell_x(\gamma) = \lim_{m \rightarrow \pm\infty} \frac{-\log \|v_{+,m} - v_{-,m}\|_{\text{Th}}}{i(\gamma, \alpha_0) \cdot |m|},$$

where α_0 is a transverse simple closed geodesic, whose Dehn twists are used to generate $\{v_{+,i}\}$ and $\{v_{-,i}\}$. This observation, coupled with the equivariance of stretch vectors under Thurston-norm isometries (Corollary 7.7) promotes topological rigidity to *geometric rigidity*:

Theorem 1.15 (Geometric rigidity). *Let $f: T_x\mathcal{T}(S) \rightarrow T_y\mathcal{T}(S)$ be a linear Thurston-norm isometry, and let h be a diffeomorphism representing the mapping class inducing f as given by topological rigidity (Theorem 1.8). Then, for every simple closed curve $\gamma \in \mathcal{S}$ on S , we have*

$$\ell_x(\gamma) = \ell_y(h(\gamma)),$$

and hence $(S, y) = (S, h(x))$.

1.10. Infinitesimal rigidity. We begin with a remark:

Remark 1.16. The statement of Theorem 1.1 *appears to be* slightly stronger than the main result in Pan's paper [22], as the latter assumes linearity of the isometry. However, this linearity condition is unnecessary, as isometries of

the Thurston norm are also isometries of its (additive) symmetrisation, and the Mazur–Ulam theorem [20] ensures that any isometry of the Thurston-norm must be affine, and hence there exists an isometric translate of the affine isometry which is a linear map.

When combined, geometric rigidity (Theorem 1.15) and the Mazur–Ulam theorem (Remark 1.16) suffice to prove infinitesimal rigidity (Theorem 1.1). Further combining the infinitesimal rigidity with [6, Theorem 6.1] then yields:

Corollary 1.17 (Local rigidity). *Consider a connected open set $U \subset \mathcal{T}(S)$. Then any isometric embedding $(U, d_{\text{Th}}) \rightarrow (\mathcal{T}(S), d_{\text{Th}})$ is the restriction to U of an element of the extended mapping class group of S .*

And this in turn implies the following result:

Corollary 1.18 (Global rigidity). *Isometries of $(\mathcal{T}(S), d_{\text{Th}})$ are precisely those induced by extended mapping classes of S .*

As previously noted, Corollary 1.18, which was first proven by Walsh [34], is an analogue of the celebrated (global form of the) Royden Theorem [28], with the role of the Teichmüller metric supplanted by the Thurston metric. Walsh’s proof is based on the action of the mapping class group on the horocyclic boundary of the Teichmüller space of the surface equipped with the Thurston metric. Dumas, Lenzhen, Rafi and Tao [6] utilise a different strategy to derive infinitesimal rigidity in the two remaining cases, thereby completing the Thurston metric analogue of Royden’s theorem. Our method of proof shares many similar ingredients to theirs. In particular, the 1-cusped tori and 4-cusped sphere cases of our main results are due to [6].

1.11. Paper outline. This paper is structured as follows:

- In Section 2 we provide a sketch of the preliminaries underpinning the theory of the Thurston metric on hyperbolic surfaces. The basic construction is that of a K -Lipschitz homeomorphism of an ideal triangle. Such maps glue to constitute stretch maps between hyperbolic surfaces, which define stretch lines in Teichmüller space. We highlight the main results from Thurston’s paper [30] needed in the constructions for this paper.
- In Section 3 we describe all K -stretch maps on crowned annuli, i.e., annuli with punctures on their boundaries, in terms of Fenchel–Nielsen coordinates and shearing parameters. This also paves the way for describing stretch maps with respect to finite-leaf laminations on (finite type) hyperbolic surfaces.
- In Section 4 we take the stretch maps described in Section 3 and we determine the behaviour of stretch vectors, i.e., tangent vectors to stretch maps. These estimates are used to isolate the lengths of simple closed geodesics in the limiting behaviour of the difference between certain sequences of pairs of stretch vectors, and hence play

a major role in establishing geodesic rigidity. Moreover, they help us to determine the codimension of ι_x -images of multicurves in \mathbf{S}_x^*

- In Section 5 we develop the general theory of convex bodies D in finite-dimensional vector spaces \mathbb{V} , using faces on the boundary ∂D to stratify ∂D . In this section we introduce various linear invariants associated with points and faces on D , including dimension, face-dimension, adherence-dimension and codimension.
- In Section 6 we explore the convex geometry of the unit tangent and contangent spheres in $T_x\mathcal{T}(S)$ and $T_x^*\mathcal{T}(S)$, showing that much of the inclusion structure for many faces of \mathbf{S}_x^* is encoded in terms of geodesic laminations associated with faces. In this section we also produce bounds and derivations for various notions of dimension associated with points in \mathbf{S}_x^* .
- In Section 7 we give two distinct proofs of topological rigidity before promoting topological to geometric rigidity, thereby establishing the infinitesimal rigidity of the Thurston metric. We also put forward consequences of infinitesimal rigidity result and raise further questions and lines of investigation.

We end this introduction with a word on references. The bases of Thurston's theory of surfaces including hyperbolic structures, geodesic lamination, Teichmüller spaces and mapping class groups are contained in Chapter 8 of his lecture notes [33] and in his article [31], with more details and developments in [7], [8] and CB. In the next section, we shall give references for Thurston's theory of best Lipschitz maps and Thurston's metric on Teichmüller space.

2. PRELIMINARIES ON THURSTON'S METRIC

We shall use the fact that the length function defined on the set \mathcal{S} of homotopy classes of simple closed curves has a continuous and positively homogeneous extension to the space $\mathcal{ML} = \mathcal{ML}(S)$ of compactly supported measured laminations on S (see Thurston [32] and Kerckhoff [14]) and that the supremum in Eq. (1) is realised over the projective space $\mathcal{PML} = \mathcal{PML}(S)$, as it is compact. Thus, the distance function d_{Th} defined by Eq. (1) is also equal to the function

$$(8) \quad R(x, y) := \max_{[\lambda] \in \mathcal{PML}} \log \frac{\ell_y(\lambda)}{\ell_x(\lambda)}$$

(the letter R here stands for “Ratio”), where $\lambda \in \mathcal{ML}(S)$ is an arbitrary measured lamination representative of $[\lambda] \in \mathcal{PML}(S)$.

2.1. Lipschitz constants formulation of the Thurston metric. In [30], Thurston gives another formula for the metric defined in Eq. (1). This alternative formulation is based on the comparison of the shapes of two marked hyperbolic surfaces by examining the smallest Lipschitz constant of a homeomorphism between them in the right homotopy class.

Given any two hyperbolic structures $x, y \in \mathcal{T}(S)$ on S and any homeomorphism $\varphi : (S, x) \rightarrow (S, y)$, the Lipschitz constant $\text{Lip}(\varphi)$ of φ is defined as the quantity

$$\text{Lip}(\varphi) := \sup_{p \neq q \in S} \frac{d_x(\varphi(p), \varphi(q))}{d_y(p, q)} \in \mathbb{R} \cup \{\infty\},$$

where d_x denotes the distance function on the hyperbolic surface (S, x) . Consider the expression

$$(9) \quad L(x, y) := \inf_{\varphi \sim \text{Id}_S} \log \text{Lip}(\varphi) \in \mathbb{R}.$$

The quantity $L(x, y)$ is invariant under isotopic deformations of the hyperbolic metrics x and y on S , and hence yields a well-defined function $L : \mathcal{T}(S) \times \mathcal{T}(S) \rightarrow \mathbb{R}$. Thurston shows that this function satisfies all the properties of an asymmetric metric [30, §2] and he proves the following:

Theorem 2.1 (Thurston [30, Theorem 8.5]). *The asymmetric metric defined in Eq. (9) coincides with the asymmetric metric d_{Th} defined in Eq. (1), i.e.,*

$$(10) \quad L \equiv R$$

2.2. Stretch maps on ideal triangles. The key advantage of Eq. (8) over Eq. (1) is that, unlike taking the supremum over $\mathcal{S}(S)$ for L , the $[\lambda] \in \mathcal{PML}(S)$ maximising R is realised. The proof that Thurston gives of Theorem 2.1 encodes the fact the best Lipschitz constant of a homeomorphism between the two hyperbolic surfaces (S, x) and (S, y) is attained on some compactly supported measured lamination λ , and hence any optimal Lipschitz constant-realising homeomorphism must be stretching “maximally” at a constant rate on (the unit tangent bundle of) the support of λ . In particular, when some projective measured lamination $[\lambda]$ that realises the maximum of R between (S, x) and (S, y) is supported on a complete geodesic lamination $|\lambda|$, the complementary regions on $S \setminus |\lambda|$ are composed of ideal triangles. In this case, the fact that a uniquely determined Lipschitz map on $|\lambda|$ homeomorphically extends to all of S , which means that the theory of Lipschitz maps on ideal triangles is both natural and unavoidable in Thurston’s theory of Lipschitz homeomorphisms between hyperbolic surfaces.

In [30, §4], Thurston concretely describes the construction of a canonical K -Lipschitz self-map of an ideal triangle, for every $K \geq 1$. We recall now this construction. It is illustrated in Fig. 1.

The ideal triangle is foliated by pieces of horocycles that are perpendicular to the boundary except for a central non-foliated region bordered by 3 pieces of such horocycles.

Definition 2.2 (central stable triangle and anchors). We refer to the central nonfoliated region of an ideal triangle as the *central stable triangle*, and refer

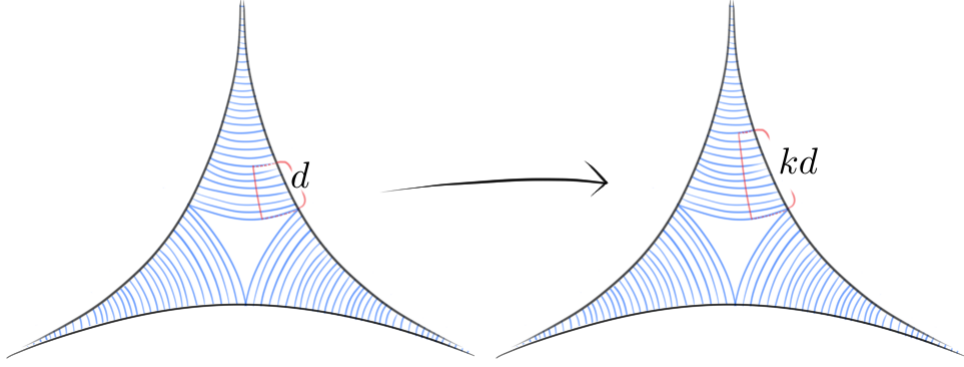


FIGURE 1. A stretch map between two ideal triangles.

to the three vertices of the stable triangle, i.e. the three points where two distinct horocycle leaves meet, as the *anchors*.

The ideal triangle equipped with its horocyclic foliation is unique up to isometry. The K -Lipschitz self-map of the triangle is taken to be the identity on the central stable triangle and to send every piece of horocycle at distance d from this region to a piece of horocycle at distance Kd . Furthermore, on each piece of horocycle at distance d from the central stable triangle, the map uniformly contracts horocyclic arc-length by a factor of e^{-dK} . These properties uniquely determine the map. Restricted to the boundary geodesics, this map expands distance by the factor K . One can show that this map is K -Lipschitz. It is called the K -stretch map of the ideal triangle.

Definition 2.3 (stretch-invariant foliations). There are two foliations that are invariant under stretching and supported on the complement of the central stable triangle: the horocyclic foliation illustrated in Fig. 1 and the geodesic foliation orthogonal to it. We refer to these respectively as the *stretch-invariant horocyclic foliation* and the *stretch-invariant geodesic foliation*.

2.3. Stretch maps on surfaces. Thurston uses K -stretch maps on ideal triangles to generate K -stretch maps between hyperbolic surfaces [30]. Given a complete geodesic lamination ν on the surface S equipped with a hyperbolic structure $x \in \mathcal{T}(S)$. We assume that not all leaves of ν converge to cusps at both ends, i.e. ν is not an ideal geodesic triangulation, but impose no other condition; ν need not carry a transverse measure of full support. Thurston constructs a family $\{x_t\}_{t \in [0, \infty)}$ of hyperbolic structures on S such that for every $t \geq 0$,

- $L(x, x_t) = R(x, x_t) = t$,
- there exists a homeomorphism $\phi_t: (S, x) \rightarrow (S, x_t)$ with $\text{Lip}(\phi) = e^t$,

- ϕ_t sends the geodesic lamination ν on (S, x) to the (unique geodesic) lamination (isotopic to) ν on (S, x_t) ,
- the restriction of ϕ_t to the lamination ν on (S, x) is an e^t -Lipschitz map which uniformly expands arc-length on ν by e^t .

Thurston's construction is fairly straightforward when the lamination ν consists of finitely many geodesic leaves: Thurston's e^t -stretch maps on each of the $4g - 4 + 2n$ ideal triangles complementary to ν stretches its boundary geodesics uniformly by a factor of e^t , and hence glue together continuously. In particular, all of the leaves of ν are boundary geodesics and hence this defines an e^t -stretch map from x on S to a different hyperbolic structure x_t on S satisfying the properties listed above.

However, when the lamination ν has (uncountably) infinitely many leaves, the situation is much subtler as only finitely many leaves in ν bound the complementary ideal triangles, and it is far from obvious as to whether Thurston's e^t -stretch maps on these complementary ideal triangles continuously extend to all the leaves in ν . Instead, Thurston studies the space of all possible hyperbolic metrics on small ϵ -neighbourhoods of ν , and parametrises this space in terms of functions describing transverse measured foliations (with mainly horocyclic leaves) transverse to ν on the aforementioned ϵ -neighbourhood. To get a sense of how this works, observe that the stretch-invariant horocyclic foliation on an ideal triangle (see Fig. 1) is endowed with a natural transverse measure which measures the arc-length of subsegments of the leaves of ν (this determines the transverse measure).

Thurston shows [30, Proposition 4.1] that any such transverse measured foliation (with closed leaves around cusps and with infinite transverse measure in the neighbourhoods of cusps) is realised by a hyperbolic metric on the ϵ -neighbourhood of ν in such a way that the transverse measure of the foliation measures arc-length on subsegments of the leaves of ν . In so doing, he shows that the initial transverse measured foliation describing the neighbourhood of ν in $x \in \mathcal{T}(S)$, when multiplied by e^t , $t \in [0, \infty)$, corresponds to a ray of hyperbolic metrics in the ϵ -neighbourhood of ν . The requisite ray $\{x_t\}_{t \in [0, \infty)}$ of metrics and the associated Lipschitz maps $\{\phi_t\}$ are obtained by filling in the rest of the surface using e^t -stretch maps on ideal triangles. The family of hyperbolic metrics defined by this ray satisfy the previously listed properties.

Definition 2.4 (stretch lines and stretch rays). The ray $\{x_t\}_{t \geq 0}$ in Teichmüller space is called a *stretch ray* starting at x_0 , and a line $\{x_t\}_{t \in \mathbb{R}}$ obtained in this way is called a *stretch line*. we use the term *stretch path* to refer to either a stretch line or a stretch ray.

Remark 2.5. We see by construction that a $K = e^t$ -stretch map supported by a geodesic lamination ν on a hyperbolic surface (S, x) stretches the stretch-invariant geodesic foliation on each complementary ideal triangle by a factor of $K > 1$ and contracts in the direction of the (orthogonal)

stretch-invariant horocyclic foliation. Therefore, the only geodesic segments stretched by K are those that lie on ν and the stretch-invariant geodesic foliation. In particular, since the leaves of the stretch-invariant geodesic foliation which do not lie on the boundary of the complementary ideal triangles (and hence ν) necessarily meet one central stable triangle, such leaves can never be complete geodesics. Hence, any complete geodesic which is stretched by K is necessarily a leaf of ν . In particular, this means that (compactly supported) measured laminations which transversely intersect ν cannot have their length increase by a factor of K . This fact is used by Thurston, for example, in the proof of [30, Theorem 5.1], and we highlight it for future reference. We shall give a precise infinitesimal version of this fact in Lemma 4.2.

2.4. Laminations. We specify several special classes of geodesic laminations which take a crucial role in the Lipschitz theory of hyperbolic surfaces.

Definition 2.6 (chain-recurrent laminations). A geodesic lamination λ on a hyperbolic surface is said to be *chain-recurrent* if for any point x on the support of λ and for any $\epsilon > 0$ there exists a parametrised C^1 path δ on S containing x such that for any unit length path I_1 contained in δ there exists a unit length path I_2 contained in the support of λ such that the two paths I_1 and I_2 are ϵ -close in the C^1 topology.

We advise readers to see [30, p. 24-25] and elsewhere in Thurston's paper for some basic properties of chain-recurrent laminations. Crucially, the support of any measured lamination is chain-recurrent. We note, however, that there is a small error in the statement of [30, Lemma 8.3]: it is *not* true that any chain-recurrent lamination is approximated arbitrarily closely by simple closed curves. The correct statement is for chain-recurrent laminations which are *connected*. This correction is supported by the second sentence on [30, Pg. 38], where Thurston invokes the "*hypothesis that λ is connected*". The following is an immediate corollary to Thurston's Lemma 8.3 [30]:

Lemma 2.7. *For any chain-recurrent lamination λ there is a (simple closed) multicurve which approximates λ arbitrarily closely in the Hausdorff topology.*

Definition 2.8 (maximal chain-recurrent laminations). A chain-recurrent lamination λ on S is *maximal* if there are no chain-recurrent laminations on S properly containing λ .

Remark 2.9. Thurston states in the paragraph before Theorem 10.2 of [30] that the complement of a maximal chain-recurrent lamination λ consists of

- ideal triangles and/or
- once-punctured monogons (the latter arises only in the setting when S has cusps), as well as

- a single once-punctured bigon in the special case when S is the 1-cusped torus¹.

A careful reading of the proof of [30, Theorem 8.5] tells us that every complete geodesic lamination extending a specified maximal chain-recurrent lamination induces the same stretch path. This observation is highlighted in [6, Corollary 2.3] and utilised in [11], and allows us to adopt the notation

$$\text{stretch}(x, \lambda, t)$$

even when a maximal chain-recurrent lamination λ is not complete (this arises precisely when S has cusps).

The set of compactly supported geodesic laminations on S is a subset of the set of compact subsets of S and hence naturally inherits the Hausdorff metric; indeed, it forms a compact metric space (see Chap. 4 of [3]). Thurston shows that the set of chain-recurrent laminations forms a closed (and hence also compact) subset, and we use this fact to establish the following lemma:

Lemma 2.10. *Given a sequence of chain-recurrent laminations $\{\lambda_i\}$ whose Hausdorff limit Λ is a maximal chain-recurrent lamination, let $\{\Lambda_i\}$ be a sequence of chain-recurrent laminations such that $\lambda_i \subseteq \Lambda_i$. Then, the sequence $\{\Lambda_i\}$ also converges to Λ in the Hausdorff topology.*

Proof. [30, Proposition 6.2] tells us that the sequence $\{\Lambda_i\}$ has limit points, and our goal is show that every convergent sequence tends to Λ . By reducing to a subsequence, we may assume without loss of generality that $\{\Lambda_i\}$ converges to some geodesic lamination Λ_∞ . By the definition of the Hausdorff topology, the limit has the form

$$\Lambda_\infty = \{ p \in (S, x) \mid \text{there is a sequence } \{p_m \in \Lambda_i\} \text{ converging to } p \},$$

and by the closedness of the set of chain-recurrent laminations, Λ_∞ is chain-recurrent. This necessarily contains

$$\Lambda = \{ p \in (S, x) \mid \text{there is a sequence } \{p_m \in \lambda_i\} \text{ converging to } p \}.$$

The maximality of Λ amongst chain-recurrent laminations means that $\Lambda_\infty = \Lambda$ and hence the Hausdorff limit of $\{\Lambda_i\}$ is indeed Λ . □

Definition 2.11 (ratio-maximising lamination). Given an ordered pair of distinct points $x, y \in \mathcal{T}(S)$, we say that a geodesic lamination λ is *ratio-maximising* from x to y if

- there exists a Lipschitz homeomorphism, in the correct homotopy class, with optimal (i.e., minimal) Lipschitz constant $e^{\text{d}_{\text{Th}}(x, y)}$ mapping from a neighbourhood of λ in (S, x) to a neighbourhood of λ in (S, y) , and

¹Thurston does not specify this last case, and it is a consequence of the hyperelliptic involution on the once-punctured tori. Interested readers may consult [6] or [11, Remark 3.4].

- every such optimal Lipschitz constant homeomorphism takes the leaves of λ in (S, x) to the leaves of λ in (S, y) by multiplying ar-length on λ by $e^{\text{dTh}(x,y)}$.

Thurston unfortunately gives the incorrect definition for ratio-maximising laminations in the paragraph prior to [30, Lemma 8.2]. This is corrected in both [6, Section 2.6] and [11, Definition 4.2]. In any case, the main motivation for defining ratio-maximising laminations is their central role in Thurston's concatenated stretch path construction for Thurston metric geodesics [30, Theorem 8.5].

Definition 2.12 (maximal ratio-maximising chain-recurrent lamination, [30, Theorem 8.2]). For any ordered pair of distinct points $x, y \in \mathcal{T}(S)$, there is a (unique) ratio-maximising chain-recurrent lamination which contains all other ratio-maximising chain-recurrent laminations from x to y . This is called the *maximal ratio-maximising chain-recurrent lamination*, and we denote it by $\mu(x, y)$.

The maximal ratio-maximising chain-recurrent lamination is central to Thurston's construction of concatenated stretch path geodesics [30, Theorem 8.5], the proof of which in turn may be used to show that $\mu(x, y)$ is the unique maximal ratio-maximising lamination among all ratio-maximising laminations, not just among the chain-recurrent ones.

Theorem 2.13. *For arbitrary points $x, y \in \mathcal{T}(S)$, the maximal ratio-maximising chain-recurrent lamination $\mu(x, y)$ contains every ratio-maximising lamination from x to y .*

Proof. Let λ be an arbitray ratio-maximising lamination from x to y . A K -stretch map with respect to a complete geodesic lamination Λ gives K -Lipschitz maps between x and $\text{stretch}(x, \Lambda, \log(K))$ which attains the Lipschitz constant K precisely along geodesic subsegments of Λ . Hence, the composition of K_i -stretch maps with respect to complete laminations Λ_i gives a Lipschitz map which attains the Lipschitz constant $\prod K_i$ precisely along geodesic subsegments contained in $\bigcap \Lambda_i$ (which may be empty).

The proof of Theorem 8.5 of [30] not only constructs a Thurston metric geodesic between (S, x) and (S, y) comprised of stretch paths segments, but also produces a $K = e^{\text{dTh}(x,y)}$ -Lipschitz map from (S, x) to (S, y) which is obtained by composing the stretch maps corresponding to the aforementioned stretch path segments. In particular, if the first stretch path segment is taken with respect to some complete geodesic lamination Λ , then the eventual K -Lipschitz between (S, x) and (S, y) can only attain the Lipschitz constant K on some sublamination of Λ . Being ratio-maximising laminations, both $\mu(x, y)$ and λ must therefore be contained in Λ . However, Thurston's construction allows us to take Λ to be *any* complete geodesic lamination containing $\mu(x, y)$, which means that λ must be contained in the intersection of all complete geodesic laminations containing $\mu(x, y)$. The only components of $S \setminus \mu(x, y)$ that can support leaves of λ are punctured monogons; for

- ideal triangles contain no leaves, and
- every other surface admits too many intersecting simple geodesics.

Therefore, $\lambda \setminus \mu(x, y)$, if non-empty, is comprised of bi-infinite geodesics contained on punctured monogons in the complement of $\mu(x, y)$. We shall show that even this is impossible because there are K -Lipschitz maps which K -stretch subsegments of the boundary of a punctured monogon, but not the interior bi-infinite geodesic.

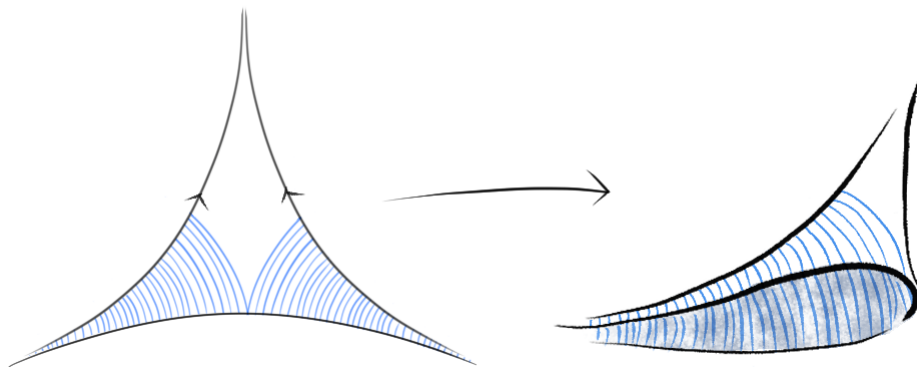


FIGURE 2. A partial horocyclic foliation on an ideal triangle (left) glue to form a partial foliation of a punctured monogon (right).

Consider the partial horocyclic foliation of the ideal triangle depicted in Fig. 2, which contains two of the horocyclic foliation regions of Fig. 1. We can similarly construct a partial K -stretch map of this ideal triangle and observe that whilst every subsegment of the bottom boundary edge is stretched by a factor of K , the other two sides are not. By symmetry, the left and right boundary edges glue to form a punctured monogon in such a way that the partial horocyclic foliation matches up, and hence defines a partial K -stretch map on the punctured monogon which K -stretches the boundary geodesic but not the interior bi-infinite geodesic. Replacing the usual stretch map on punctured monogon components in the complement of $\mu(x, y)$ by this partial K -stretch map which produces the same metric as $\text{stretch}(x, \Lambda, \log(K))$. This shows that $\lambda \setminus \mu(x, y)$ must be empty. \square

Remark 2.14 (maximal ratio-maximising laminations). The upshot of Theorem 2.13 is that we may simply refer to $\mu(x, y)$ as the (unique) *maximal ratio-maximising lamination* from x to y . Note that Thurston also refers to maximal ratio-maximising laminations as *maximal maximally-stretched laminations* on two occasions in [30]:

- in the last paragraph of the proof of Theorem 8.5, and
- in the statement of Theorem 10.7,

and similar nomenclature is adopted in [6, Section 2.6].

3. STRETCH MAPS AND STRETCH VECTORS

The goal of this section is to better understand the trajectory and behaviour of Thurston's stretch paths. We focus on the specific instance where S is a complete hyperbolic surface which is topologically an annulus with punctures on both of its boundary components². We refer to such surfaces as (n_L, n_R) -crowned annuli [4, 9], where $n_L, n_R \geq 1$ denote the number of punctures on the “left”, respectively “right” boundary component. Up to the action of the mapping class group relative to the boundary, crowned annuli only admit finitely many complete geodesic laminations, and such laminations can only contain finitely many leaves. As before, we assume that the leaves of such a lamination do not all go from cusp to cusp. The relative restrictiveness of the context will allow us to describe all stretch paths on the Teichmüller spaces for crowned annuli. At the same time, crowned annuli constitute common examples of subsurfaces and it will aid us also in the understanding of the behaviour of stretch maps on surfaces of greater topological complexity.

3.1. Stretch maps for $(1, 1)$ -crowned annuli. We begin with the case where S is a $(1, 1)$ -crowned annulus, see Figure 3. Since S is an annulus, it supports a unique essential simple closed curve γ_S up to isotopy, which is a retract of S .

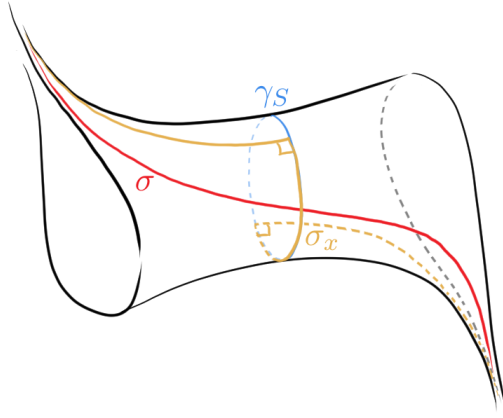


FIGURE 3. A depiction of σ , σ_x and γ_S .

²Taking the double of such a surface, one gets a cusped hyperbolic surface in the sense we considered before. Thurston's stretch maps are still well-defined on these surfaces since they can be decomposed into ideal triangles.

The Teichmüller space for $(1,1)$ -crowned annuli is $\mathcal{T}(S) \cong \mathbb{R}_{>0} \times \mathbb{R}$ (see [9]), where, for a given marked hyperbolic surface $x \in \mathcal{T}(S)$, the first coordinate parametrises the length of the unique simple closed geodesic representative of γ_S on x . For expositional simplicity, we henceforth conflate γ_S with its unique geodesic representative for any $x \in \mathcal{T}(S)$. The second coordinate parametrises the twist parameter for γ_S . There is a one-parameter family of choices for how one might define the twist parameter (see, e.g. [32, p. 272]), and we adopt the following convention:

- fix an (arbitrary simple) complete geodesic σ on (S, x) with endpoints at the two punctures of this surface;
- for a given hyperbolic metric $x \in \mathcal{T}(S)$, homotope σ so that it is the unique piecewise geodesic σ_x traveling from cusp 1 until it orthogonally meets γ_S , then travels along γ_S without back-tracking, but possibly lapping γ_S several times, before orthogonally leaving this curve to approach cusp 2 (see Fig. 3);
- the twist parameter is then the signed length of the subsegment of σ_x traversing over γ_S , with each additional loop along γ_S contributing to the twist parameter by $\ell_x(\gamma_S)$;
- we denote the length of γ_S and the twist parameter in the form (ℓ, τ) and call this pair of numbers the *Fenchel–Nielsen coordinates* of (S, x) relative to our choice of geodesic σ .

The main goal of the rest of this subsection is to explicitly write down all K -stretch maps for the $(1,1)$ -annulus S . To do so, we first determine the collection of all complete geodesic laminations on S where not all leaves have both ends in the boundary.

The limit set of the non-boundary end of a leaf which does not have both ends in the boundary is a compactly supported geodesic lamination on S , and hence must be γ_S . Therefore, any complete geodesic lamination on S necessarily contains γ_S . (Recall that, by definition, the support of a geodesic lamination must be a closed subset of S .) Now, we can complete γ_S to a (non-compactly supported) geodesic lamination of (S, x) by adding two geodesic leaves α_1, α_2 , where each α_i shoots out from cusp i to spiral around γ_S . There are four possible configurations of the directions in which α_1 and α_2 approach γ_S , they are represented in Fig. 4. Cases 1 and 3 (from left to right) in the figure are those where the complete geodesic lamination is chain-recurrent in the sense of Thurston (Definition 2.6).

Note that we intentionally conflate the two laminations where α_1 and α_2 spiral toward γ_S in the same direction and denote either of these two laminations as λ_0 . Our notation abuse is motivated by the fact that the Thurston stretch ray on $\mathcal{T}(S)$ with respect to either choice of λ_0 produces the same parametrised Thurston geodesic ray on $\mathcal{T}(S)$, as we shall see in the next theorem, despite seemingly defining distinct operations on the actual

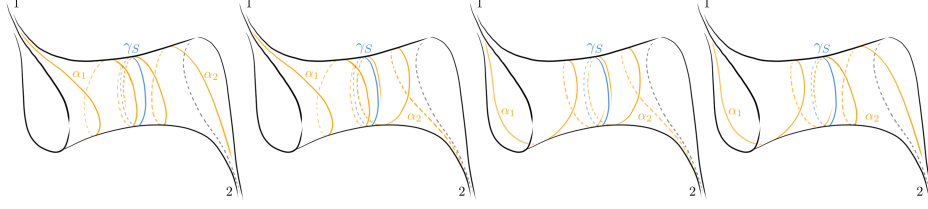


FIGURE 4. The four possible configurations of a complete lamination composed of γ_S , α_1 and α_2 . We respectively denote the maximal laminations for cases 1 and 3 (start counting from the left) as λ_- and λ_+ . Cases 2 and 4 are those where the two geodesics α_1 and α_2 spiral around γ_S in the same direction, and where we denote the resulting maximal lamination by λ_0 .

hyperbolic surface. We denote the other two laminations by λ_+ and λ_- (see Fig. 5).

Theorem 3.1 (stretching $(1,1)$ -crowned annuli). *Consider a hyperbolic metric $x \in \mathcal{T}(S)$ on S given by the Fenchel–Nielsen coordinates (ℓ_0, τ_0) and let $K \geq 1$ be an arbitrary real number ≥ 1 . The K -stretch map with respect to λ_0 deforms the Fenchel–Nielsen coordinates (ℓ_0, τ_0) to $(K\ell_0, K\tau_0)$, whereas the K -stretch map with respect to λ_{\pm} takes (ℓ_0, τ_0) to*

$$\left(K\ell_0, K\tau_0 \pm 2 \left(\log(1 - e^{-K\ell_0}) - K \log(1 - e^{-\ell_0}) \right) \right).$$

Proof. We first note that since γ_S is in the closure of the leaves of the α_i , and since the latter are stretched by a factor of K , γ_S is also stretched by K . Thus the length coordinate ℓ_0 increases to $K\ell_0$ in all three cases.

The λ_0 case: We now turn to the twist parameter in the λ_0 case. The fact that the K -stretch map takes (ℓ_0, τ_0) to $(K\ell_0, K\tau_0)$ is heuristically reasonable: both α_1 and α_2 spiral towards γ_S in the same direction, and so any twisting applied on the left of γ_S is matched by the twist on the right, thereby producing 0 net change in the ℓ -normalised twist parameter $\frac{\tau_0}{\ell_0} = \frac{K\tau_0}{K\ell_0}$. It is nevertheless helpful and reassuring to do the full computation, and we shall see that essentially the same computations are used in the proof of the λ_{\pm} case formulae.

The universal cover of the hyperbolic surface (S, x) is a contractible bordered domain embedded in the hyperbolic plane \mathbb{H} . In Fig. 5, we depict \mathbb{H} as the unit disc model, but for computational purposes, we shall regard this as the upper half-space model. Consider a lift $\tilde{\gamma}_S$ of γ_S . On each side of $\tilde{\gamma}_S$, we see a \mathbb{Z} family of lifts of α_1 and α_2 , each spiralling towards an endpoint of $\tilde{\gamma}_S$ which we shall assume, without loss of generality, to be the attracting fixed point of an automorphism of the universal cover corresponding to the closed geodesic γ_S . This convergence in the universal cover means that in

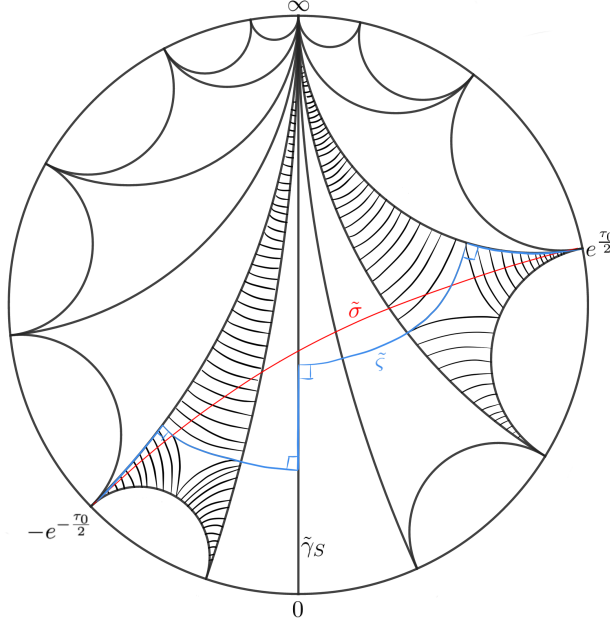


FIGURE 5. A lift of λ_0 to the universal cover with specifications of $\tilde{\gamma}_S$, $\tilde{\sigma}$ and $\tilde{\zeta}$. The infinite-sided polygon represents the universal cover of the surface S . the coordinates on the circle at infinity are marked after identifying this circle with the upper half-plane

the surface S , the images of the corresponding geodesics spiral around γ_S . We mark each ideal triangle in the complement of the lifts of the α_i with its central stable triangle and note that the distance, along a lift $\tilde{\alpha}_i$, of the anchors for any two adjacent ideal triangles is ℓ_0 : one can see this geometrically by extending the horocyclic segment running through the given anchors all the way to $\tilde{\gamma}_S$ and observing that their endpoints on $\tilde{\gamma}_S$ are translates of each other with respect to the action of γ_S . Fix an arbitrary lift $\tilde{\sigma}$ of σ . We may assume without loss of generality that $\tilde{\gamma}_S$ has ∞ and 0 as its ideal points. Since the initial twist parameter is τ_0 , the ideal points for $\tilde{\sigma}$ are $-e^{-\frac{\tau_0}{2}}$ and $e^{\frac{\tau_0}{2}}$ for some convenient parametrisation of the boundary at infinity by $\mathbb{R} \cup \{\infty\}$. The Thurston stretch map preserves γ_S and hence fixes 0 and ∞ in its lift to the universal cover of x .

To analyse the action of the stretch map on σ , we use a homotopy representative of σ whose behaviour under this map is more transparent. Namely, consider the piecewise-smooth path ς homotopic to σ which:

- (1) travels along α_1 until it reaches the central stable triangle (for the first time);

- (2) then turns onto the stretch-invariant horocyclic foliation and travels along this foliation until γ_S ;
- (3) we similarly produce such a piecewise (non-geodesic) ray emanating along α_2 from cusp 2 which eventually reaches γ_S , and join these two rays with a geodesic segment running over γ_S without back-tracking, but possibly traversing γ_S multiple times.

See Fig. 5 for a depiction of the lift $\tilde{\zeta}$ of ζ which is homotopic to $\tilde{\sigma}$.

The Thurston stretch map acts on ζ by uniformly stretching the geodesic segments shown in the picture by a factor of K and non-uniformly shrinking the two horocyclic segments, each joining α_i and γ_S . The reason for the non-uniformity of the latter is because each of the two horocyclic segments is comprised of (countably infinitely many) horocyclic segments taken from γ_S -translates of lifts of the two triangles in $S \setminus (\alpha_1 \cup \alpha_2 \cup \gamma_S)$. A little elementary hyperbolic geometry (e.g. [30, p. 16] or [26, Proposition 2.8]) shows that the lengths of these composite horocyclic segments on either side of $\tilde{\gamma}_S$ are both given by the following geometric series:

$$(11) \quad 1 + e^{-\ell_0} + e^{-2\ell_0} + e^{-3\ell_0} + \dots = (1 - e^{-\ell_0})^{-1}.$$

Take special note that this horocyclic length is only dependent on ℓ_0 — the length of γ_S (see the statement of Theorem 3.1).

Eq. (11) implies that, when depicted in the upper half-plane model, as it is here, the first (left) horocyclic segment is positioned at height $e^{-\frac{\tau_0}{2}}(1 - e^{-\ell_0})$, and the second is positioned at height $e^{-\frac{\tau_0}{2}}(1 - e^{-\ell_0})$. Note in particular that this means that the midpoint of the vertical geodesic segment in the middle of $\tilde{\zeta}$ is at height $(1 - e^{-\ell_0})$.

After a K -stretch, the new metric deforms the surface in such a manner that the geodesic paths α_1 and α_2 remain as geodesics which spiral to γ_S , and the transverse horocyclic foliation (depicted in Fig. 5) is preserved. We develop the universal cover of the K -stretched surface subject to the normalisation condition that the oriented geodesic $\tilde{\gamma}_S$ is preserved, i.e., 0 and ∞ are respectively fixed, and that the midpoint $i(1 - e^{-\ell_0}) \in \mathbb{H}$ of the geodesic subsegment of $\tilde{\zeta}$ along $\tilde{\gamma}_S$ is fixed. Then, the configuration depicted in Fig. 5 still holds, albeit with new positions $u_1, u_2 \in \mathbb{R}$ instead of $e^{\frac{\tau_0}{2}}$ and $e^{-\frac{\tau_0}{2}}$. To determine the new u_1, u_2 , we first observe that the metric along $\tilde{\gamma}_S$ is stretched by a factor of K , the length of the vertical geodesic segment in the middle of $\tilde{\zeta}$ increases from τ_0 to $K\tau_0$ and hence the new heights of the two ends of this vertical segment are given by

$$e^{-\frac{K\tau_0}{2}}(1 - e^{-\ell_0}) \quad \text{and} \quad e^{\frac{K\tau_0}{2}}(1 - e^{-\ell_0}).$$

We next consider what happens to the two composite horocyclic segments on $\tilde{\zeta}$. Recall from above that each composite horocyclic segment is made up of infinitely many horocyclic segments, each contained in a distinct lift of

one of the two ideal triangles in $S - \lambda_0$. The total length of each of the new composite horocyclic segments is given by

$$(12) \quad 1 + e^{-K\ell_0} + e^{-2K\ell_0} + e^{-3K\ell_0} + \dots = (1 - e^{-K\ell_0})^{-1}.$$

We now have enough information to determine the transformed ideal endpoints $u_1 < 0 < u_2$ of the transformed $\tilde{\zeta}$. To find u_1 , we note that a horocyclic segment of length $(1 - e^{-K\ell_0})^{-1}$ at height $e^{\frac{-K\tau_0}{2}}(1 - e^{-\ell_0})$ has Euclidean width (with respect to the upper half-plane model) given by

$$\frac{e^{\frac{-K\tau_0}{2}}(1 - e^{-\ell_0})}{(1 - e^{-K\ell_0})},$$

hence

$$u_1 = -\frac{e^{\frac{-K\tau_0}{2}}(1 - e^{-\ell_0})}{(1 - e^{-K\ell_0})}.$$

Similarly, the Euclidean width of a horocyclic segment of length $(1 - e^{-K\ell_0})^{-1}$ at height $e^{\frac{K\tau_0}{2}}(1 - e^{-\ell_0})$ is

$$\frac{e^{\frac{K\tau_0}{2}}(1 - e^{-\ell_0})}{(1 - e^{-K\ell_0})},$$

hence

$$u_2 = \frac{e^{\frac{K\tau_0}{2}}(1 - e^{-\ell_0})}{(1 - e^{-K\ell_0})}.$$

Having obtained u_1 and u_2 , we can compute the K -stretched image of $\tilde{\sigma}$ and hence derive that the new twist coordinate is $\log \left| \frac{u_2}{u_1} \right| = K\tau_0$, as desired. There are various methods for showing that the new twist coordinate is indeed given by this expression: one can drop perpendiculars from u_1 and u_2 to $\tilde{\gamma}_S$ and measure the difference between the respective perpendicular points of these two rays, or one can observe that the twist parameter is the logarithm of (one of the permutations of) the cross ratio between $0, \infty$ and the two end points of $\tilde{\sigma}$.

The λ_+ case: we repeat the same strategy as for λ_0 and consider lifts $\tilde{\gamma}_S, \tilde{\zeta}$ respectively of γ_S and ς as before. In particular, we again set $\tilde{\gamma}_S$ as the geodesic joining 0 and ∞ and set $\tilde{\sigma}$ as the geodesic joining $-e^{-\frac{\tau_0}{2}}$ and $e^{\frac{\tau_0}{2}}$. We know immediately from the computation in the λ_0 case that the horocyclic subsegment of $\tilde{\zeta}$ to the left of $\tilde{\gamma}_S$ is placed at height $e^{-\frac{\tau_0}{2}}(1 - e^{-\ell_0})$. Observe that the hyperbolic involution $z \mapsto -z^{-1}$ preserves $\tilde{\gamma}_S, \tilde{\zeta}$ and permutes the two horocyclic subsegments of $\tilde{\zeta}$. This immediately tells us that the horocyclic segment to the right of $\tilde{\gamma}_S$ lies on the Euclidean circle passing through 0 and $ie^{\frac{\tau_0}{2}}(1 - e^{-\ell_0})^{-1}$ and tangent to 0 . Further observe that the hyperbolic length of the left composite horocyclic segment is still given by $(1 - e^{-\ell_0})^{-1}$, and hence, by the $z \mapsto -z^{-1}$ involution, this is also the length of the right composite horocyclic segment. Furthermore, these

two composite horocyclic segments both K -stretched to horocyclic segments of hyperbolic length $(1 - e^{-K\ell_0})^{-1}$.

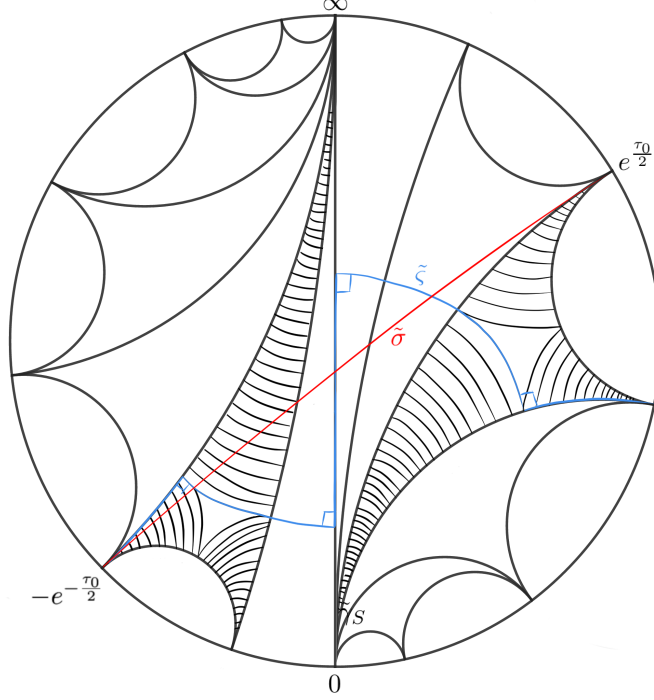


FIGURE 6. A lift of λ_+ to the universal cover with specifications of $\tilde{\gamma}_S$, $\tilde{\sigma}$ and $\tilde{\zeta}$.

We now derive the K -stretched twist parameter for λ_+ (see Fig. 6). We develop the K -stretched image of $\tilde{\zeta}$ subject to the normalisation condition that preserves $\tilde{\gamma}_S$ and fixes the midpoint $i \in \mathbb{H}$ of the geodesic subsegment of $\tilde{\zeta}$ along $\tilde{\gamma}_S$. One reason for choosing this normalisation comes from the observation that it is preserved under the involution $z \mapsto -z^{-1}$, and since we are K -stretching with respect to a lamination which is preserved by $z \mapsto -z^{-1}$, we ensure that this involution remains an involution of the K -stretched universal cover of S we develop (subject to this normalisation condition). Note in particular that the mid-point of $\tilde{\zeta}$ is necessarily given by $i \in \mathbb{H}$ (both before and after the K -stretch) as this is the unique fixed point of the involution $z \mapsto -z^{-1}$.

We now determine the pre-stretch distance between the mid-point i and the endpoints of the central geodesic subsegment of $\tilde{\zeta}$. Our computations for λ_0 show that the height of the lower endpoint is $e^{-\frac{\tau_0}{2}}(1 - e^{-\ell_0})$, hence the distance between this end and the mid-point i is $\log(e^{\frac{\tau_0}{2}}(1 - e^{-\ell_0})^{-1}) = \frac{\tau_0}{2} - \log(1 - e^{-\ell_0})$. This is the same as the distance between i and the upper

endpoint of the central geodesic subsegment of $\tilde{\zeta}$, hence the height of the higher endpoint of this segment is $e^{\frac{\tau_0}{2}}(1 - e^{-\ell_0})^{-1}$. The K -stretch distorts distances along $\tilde{\gamma}_S$ by a factor of K , hence the distance between the new mid-point (still at $i \in \mathbb{H}$ by normalisation) and the new endpoints of the central geodesic subsegment of $\tilde{\zeta}$ is $\frac{K\tau_0}{2} - K \log(1 - e^{-\ell_0})$. Therefore, the new endpoints of the K -stretched central (vertical) geodesic segment of the K -stretched $\tilde{\zeta}$ are at

$$ie^{\frac{-K\tau_0}{2}}(1 - e^{-\ell_0})^K \text{ and } ie^{\frac{K\tau_0}{2}}(1 - e^{-\ell_0})^{-K} \in \mathbb{H}.$$

We again denote the ideal points of the K -stretched $\tilde{\zeta}$ by $u_1 < 0 < u_2$. Then, Eq. (12) again tells us that the length of the K -stretched composite horocyclic segment to the left of $\tilde{\gamma}_S$ has length $(1 - e^{-K\ell_0})^{-1}$, hence

$$(13) \quad u_1 = -\frac{e^{\frac{-K\tau_0}{2}}(1 - e^{-\ell_0})^K}{(1 - e^{-K\ell_0})}.$$

To determine the position of u_2 , we again apply the $z \mapsto -z^{-1}$ involution and see that

$$(14) \quad u_2 = -(u_1)^{-1} = \frac{e^{\frac{K\tau_0}{2}}(1 - e^{-K\ell_0})}{(1 - e^{-\ell_0})^K}.$$

Having determined u_1 and u_2 , we now know the K -stretched image of $\tilde{\sigma}$ and hence derive that the new twist coordinate is

$$\begin{aligned} \log \left| \frac{u_2}{u_1} \right| &= 2 \log \left(\frac{e^{\frac{K\tau_0}{2}}(1 - e^{-K\ell_0})}{(1 - e^{-\ell_0})^K} \right) \\ &= K\tau_0 + 2 \left(\log(1 - e^{-K\ell_0}) - K \log(1 - e^{-\ell_0}) \right), \text{ as claimed.} \end{aligned}$$

This computation again makes use of the fact that the twist parameter is the logarithm of a cross-ratio between $0, \infty$ and the two endpoints of $\tilde{\sigma}$.

The λ_- case: this follows by symmetry from the λ_+ case. Reversing the orientation on S changes the lamination λ_- to λ_+ , whilst affecting the Fenchel–Nielsen twist parameter by a sign change (and leaving the length parameter unchanged). This means that the new twist coordinate after a K -stretch map is

$$\begin{aligned} &- \left(K \times (-\tau_0) + 2 \left(\log(1 - e^{-K\ell_0}) - K \log(1 - e^{-\ell_0}) \right) \right) \\ &= K\tau_0 - 2 \left(\log(1 - e^{-K\ell_0}) - K \log(1 - e^{-\ell_0}) \right), \text{ as claimed.} \end{aligned}$$

□

3.2. Stretch maps for general crowned annuli. We now seek to understand stretch maps on an (n_L, n_R) -crowned annulus S , i.e., annuli with $n_L, n_R > 0$ punctures on their boundaries and endowed with a complete finite-area geodesic-bordered hyperbolic metric [4, 9]. We (again) denote the unique simple closed curve on S by γ_S . The geodesic γ_S cuts S into

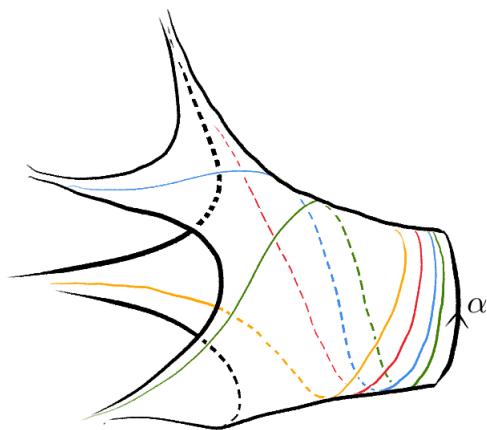
two annuli, and we refer to the annulus containing the n_L boundary cusps as the left component, and the one with the n_R boundary cusps as the right component.

We first note that any complete lamination Λ of S that contains at least one leaf which does not have both ends going into the boundary of S necessarily contains γ_S , hence no leaves of Λ pass from the left component of $S \setminus \gamma_S$ to the right. Topological constraints tell us that the left component of $S \setminus \gamma_S$ contains precisely n_L leaves (not including γ_S , and the right contains n_R leaves. This means that we can parametrise the Teichmüller space $\mathcal{T}(S)$ of S with

- n_L shearing parameters, one for each of the n_L leaves on the left component of $S \setminus \gamma_S$ with the linear constraint that the shearing parameters which spiral around γ_S sum to ℓ_{γ_S} ,
- n_R shearing parameters, one for each of the n_R leaves on the right component of $S \setminus \gamma_S$ with the linear constraint that the shearing parameters which spiral around γ_S sum to ℓ_{γ_S} ,
- one Fenchel–Nielsen length parameter ℓ_{γ_S} and one twist parameter τ_{γ_S} for γ_S .

The $K(> 1)$ -stretch map along Λ multiplies all of the shearing parameters and the Fenchel–Nielsen length parameter ℓ_{γ_S} by a factor of K (the two linear constraints are still satisfied). Thus, we need only determine what happens to the twist parameter to determine the stretch path given by e^t -stretch maps along Λ , for $t \geq 0$. To this end, we further observe that only leaves spiralling into γ_S can actually affect τ_{γ_S} , and hence we can determine τ_{γ_S} completely from the subsurface of S which consists of the convex hull of the leaves of Λ spiralling into γ_S . This convex hull is a (n'_L, n'_R) -crowned annulus with $n'_L \leq n_L$ and $n'_R \leq n_R$, and, by possibly reducing to a smaller crowned annulus, we need only determine τ_{γ_S} for stretch maps with respect to Λ which consist of γ_S and $n_L + n_R$ bi-infinite leaves spiralling into γ_S . To avoid intersection, the leaves on the left component of $S \setminus \gamma_S$ must all spiral toward γ_S in the same direction, and likewise for the right component. This means that we need only consider four possible cases for Λ : there are two cases where the left and right leaves spiral in the same direction toward γ_S (like for λ_0 in Fig. 4), and there are two Λ where the left and right leaves spiral in opposite directions toward γ_S (like for λ_{\pm} in Fig. 4). We shall handle these two situations separately, but first, we establish a useful lemma for a single component of $S \setminus \gamma_S$, i.e., a $(n, 0)$ -crowned annulus (see Fig. 7).

Consider an $(n, 0)$ -crowned annulus A , that is: a complete finite-area geodesic bordered hyperbolic surface which is topologically an annulus with $n \geq 1$ punctures on one of its boundaries. Denote the closed geodesic boundary of A by α (endowed with an auxiliary orientation), and a complete geodesic lamination μ on A comprised of α and n bi-infinite geodesic leaves


 FIGURE 7. A $(4,0)$ -crowned annulus.

which spiral towards α in the same direction as α . See Fig. 8 for a depiction of the universal cover of A .

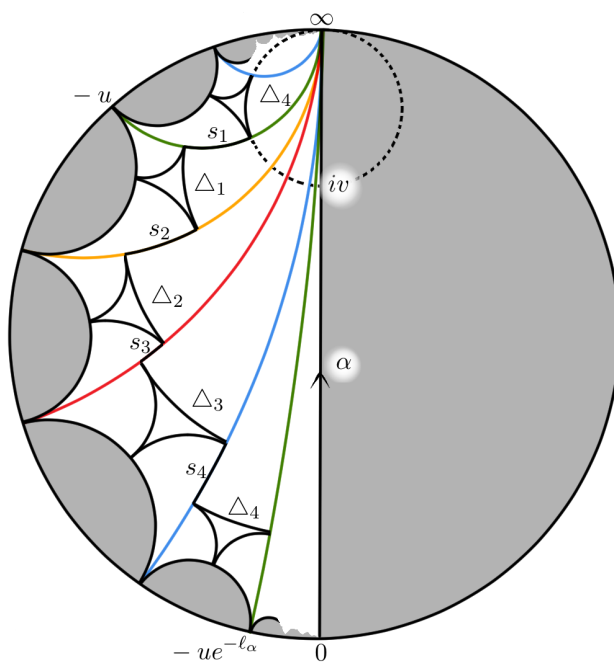


FIGURE 8. The universal cover for the annulus in Fig. 7.

We sequentially label the ideal triangles on A by $\Delta_1, \dots, \Delta_n$. Fig. 8 represents the situation in the universal cover. Let $s_i \in \mathbb{R}$ denote the shearing

parameter for the edge shared by Δ_i and Δ_{i+1} (cyclically indexed so that $\Delta_1 = \Delta_{n+1}$). In other words, s_i denotes the signed hyperbolic distance travelled, along the ideal geodesic shared by Δ_i and Δ_{i+1} , from the anchor for Δ_{i+1} to the anchor for Δ_i . This is signed to be positive if one moves towards closer to α and negative if one moves farther from α . In the example depicted in Fig. 8, all of the s_i are positive. We shall assume without loss of generality that the universal cover of A is positioned so that the (unique) lift of α joins 0 and ∞ , with the cover of A positioned to the left of the imaginary axis. Take an arbitrary lift σ of the ideal edge shared by Δ_n and Δ_1 and let $-u$ denote the non- ∞ end of σ . Further consider the lift $\hat{\Delta}_n$ of Δ_n bordered by σ and the horocyclic arc on $\hat{\Delta}_n$ which extends to a horocycle η tangent to ∞ . Since σ is a geodesic that joins $-u$ and ∞ , the (unique) intersection point of σ and η necessarily takes the form $-u + iv$. Then, the (hyperbolic) length of the horocyclic subsegment of η that traverses from $-u + iv$ to iv is $\frac{u}{v}$. We have the following:

Lemma 3.2. *With the above notation, we have the following relation:*

$$(15) \quad \frac{u}{v} = \frac{1 + e^{-s_1} + e^{-(s_1+s_2)} + \dots + e^{-(s_1+s_2+\dots+s_{n-1})}}{1 - e^{-\ell_\alpha}}.$$

Proof. The horocyclic subsegment of η that traverses from $-u + iv$ to iv and whose length is $\frac{u}{v}$ is partitioned by the leaves of the complete geodesic lamination μ into (countably) infinitely many horocyclic segments. The first of these segments is a boundary of the central stable triangle of Δ_1 and hence has length 1. The second horocyclic segment, which is on Δ_2 , is at distance s_1 away from a length 1 horocyclic segment which constitutes the boundary of the central stable triangle of Δ_2 and hence has length e^{-s_1} . We continue onwards, noting that for Δ_3 , the horocyclic segment in question is at distance $s_1 + s_2$ away from the relevant boundary of the central stable triangle on Δ_3 and hence has length $e^{-(s_1+s_2)}$. We continue in this manner until we get through all n ideal triangles and return to Δ_1 , at which point we now know that the $(n+1)$ -th horocyclic segment is at distance ℓ_α away from the original length 1 horocyclic segment, and hence is of length $e^{-\ell_\alpha}$. Similarly, the $(n+2)$ -th horocyclic segment, which lies on Δ_2 , is of length $e^{-(\ell_\alpha+s_1)}$, and so forth. Therefore:

$$\begin{aligned} \frac{u}{v} &= 1 + e^{-s_1} + e^{-(s_1+s_2)} + \dots + e^{-(s_1+s_2+\dots+s_{n-1})} \\ &\quad + e^{-\ell_\alpha} + e^{-(\ell_\alpha+s_1)} + e^{-(\ell_\alpha+s_1+2)} + e^{-(\ell_\alpha+s_1+s_2+\dots+s_{n-1})} \\ &\quad + e^{-2\ell_\alpha} + e^{-(2\ell_\alpha+s_1)} + e^{-(2\ell_\alpha+s_1+2)} + e^{-(2\ell_\alpha+s_1+s_2+\dots+s_{n-1})} + \dots \\ &= \left(1 + e^{-s_1} + \dots + e^{-(s_1+s_2+\dots+s_{n-1})}\right) \left(1 + e^{-\ell_\alpha} + e^{-2\ell_\alpha} + \dots\right) \\ &= \frac{1 + e^{-s_1} + e^{-(s_1+s_2)} + \dots + e^{-(s_1+s_2+\dots+s_{n-1})}}{1 - e^{-\ell_\alpha}}, \text{ as claimed.} \end{aligned}$$

□

3.2.1. *Parallel spiralling case.* We now move onto working with stretch maps of S , and we first consider the case where all bi-infinite leaves of the maximally stretched lamination spiral in the same direction. Let μ_+ denote the complete lamination on S containing, as its leaves (see Fig. 9)

- the closed geodesic γ_S ;
- $n_L + n_R$ geodesics going from the boundary cusps of S and spiralling towards γ_S in the direction of γ_S ;
- the boundary geodesics on S .

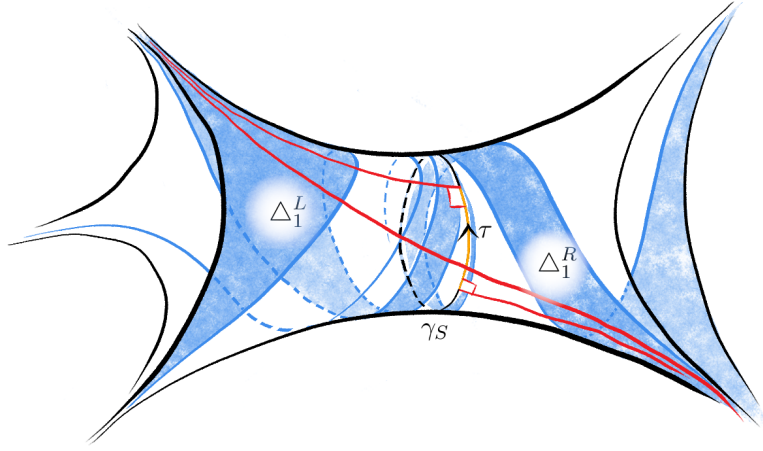


FIGURE 9. A $(3, 2)$ -crowned annulus with τ coordinate given by the signed length of the orange geodesic arc.

We label the ideal triangles on the left component by $\Delta_1^L, \dots, \Delta_{n_L}^L$ and those on the right by $\Delta_1^R, \dots, \Delta_{n_R}^R$ ordered as in Fig. 10. Let $s_i^L \in \mathbb{R}$ denote the shearing parameter for the edge shared by Δ_i^L and Δ_{i+1}^L (cyclically indexed), and likewise let $s_i^R \in \mathbb{R}$ denote the shearing parameter for the edge shared by Δ_i^R and Δ_{i+1}^R . Recall that these shearing parameters may differ in sign from the usual convention, and s_i^X (for $X = L, R$) is defined as the signed hyperbolic distance travelled, along the ideal geodesic shared by Δ_i^X and Δ_{i+1}^X , from the anchor for Δ_{i+1}^X to the anchor for Δ_i^X . This is signed to be positive if one moves towards closer to γ_S and negative if one moves farther from γ_S . As an example, every shearing parameter in Fig. 10 is positive except for s_2^L .

Lemma 3.3. *Given an (n_L, n_R) -crowned hyperbolic annulus structure x on S , set τ_{γ_S} to be the twist parameter for γ_S obtained by homotoping the simple ideal geodesic going from the “leftmost” cusp on Δ_1^L to the “rightmost” cusp*


$$\begin{aligned} & \tau_{\gamma_S}(\text{stretch}(x, \mu_+, t)) \\ &= e^t \tau_{\gamma_S}(x) \\ &+ e^t \log \left(1 + e^{-s_1^L(x)} + e^{-(s_1^L(x) + s_2^L(x))} + \dots + e^{-(s_1^L(x) + \dots + s_{n_L-1}^L(x))} \right) \\ &- \log \left(1 + e^{-e^t s_1^L(x)} + e^{-e^t(s_1^L(x) + s_2^L(x))} + \dots + e^{-e^t(s_1^L(x) + \dots + s_{n_L-1}^L(x))} \right) \\ &- e^t \log \left(1 + e^{-s_1^R(x)} + e^{-(s_1^R(x) + s_2^R(x))} + \dots + e^{-(s_1^R(x) + \dots + s_{n_R-1}^R(x))} \right) \\ &+ \log \left(1 + e^{-e^t s_1^R(x)} + e^{-e^t(s_1^R(x) + s_2^R(x))} + \dots + e^{-e^t(s_1^R(x) + \dots + s_{n_R-1}^R(x))} \right), \end{aligned}$$

Proof. Let τ_0 denote the twist parameter $\tau_{\gamma_S}(x)$ for γ_S before the e^t -stretching and let τ_t denote the twist parameter $\tau_{\gamma_S}(\text{stretch}(x, \mu_+, t))$ after the e^t -stretching. Then, we know that the correspondingly labelled u_0^L, u_0^R and u_t^L, u_t^R (see Fig. 10) satisfy

$$(16) \quad \tau_0 = \log \left(\frac{u_0^R}{u_0^L} \right) \quad \text{and} \quad \tau_t = \log \left(\frac{u_t^R}{u_t^L} \right).$$

Our goal is to express τ_t in terms of τ_0 and other $\mathcal{T}(S)$ coordinates. From Lemma 3.2, we obtain that

$$\begin{aligned}
 \tau_0 &= \log \left(\frac{u_0^R}{v_0^R} \cdot \frac{v_0^L}{u_0^L} \cdot \frac{v_0^R}{v_0^L} \right) = \log \left(\frac{u_0^R}{v_0^R} \right) - \log \left(\frac{u_0^L}{v_0^L} \right) + \log \left(\frac{v_0^R}{v_0^L} \right) \\
 &= \log \left(\frac{v_0^R}{v_0^L} \right) \\
 &\quad - \log \left(\frac{1 + e^{-s_1^L(x)} + e^{-(s_1^L(x)+s_2^L(x))} + \dots + e^{-(s_1^L(x)+s_2^L(x)+\dots+s_{n_L-1}^L(x))}}{1 - e^{-\ell_{\gamma_S}(x)}} \right) \\
 (17) \quad &+ \log \left(\frac{1 + e^{-s_1^R(x)} + e^{-(s_1^R(x)+s_2^R(x))} + \dots + e^{-(s_1^R(x)+s_2^R(x)+\dots+s_{n_R-1}^R(x))}}{1 - e^{-\ell_{\gamma_S}(x)}} \right).
 \end{aligned}$$

The fact that the last term of Eq. (17) is equal to $\log \left(\frac{v_0^R}{v_0^L} \right)$ follows from Lemma 3.2 by applying the reflection isometry with respect to the imaginary axis in \mathbb{H} .

We draw particular attention to the term $\log \left(\frac{v_0^R}{v_0^L} \right)$, which measures the length of a geodesic segment that traverses over γ_S . After e^t -stretching, the length of this segment must therefore be $\log \left(\frac{v_t^R}{v_t^L} \right) = e^t \log \left(\frac{v_0^R}{v_0^L} \right)$. Substituting this into the e^t -stretched version of Eq. (17) yields

$$\begin{aligned}
 \tau_t &= \log \left(\frac{u_t^R}{v_t^R} \cdot \frac{v_t^L}{u_t^L} \cdot \frac{v_t^R}{v_t^L} \right) = \log \left(\frac{u_t^R}{v_t^R} \right) - \log \left(\frac{u_t^L}{v_t^L} \right) + \log \left(\frac{v_t^R}{v_t^L} \right) \\
 &= e^t \log \left(\frac{v_0^R}{v_0^L} \right) \\
 &\quad - \log \left(\frac{1 + e^{-e^t s_1^L(x)} + e^{-e^t (s_1^L(x)+s_2^L(x))} + \dots + e^{-e^t (s_1^L(x)+s_2^L(x)+\dots+s_{n_L-1}^L(x))}}{1 - e^{-e^t \ell_{\gamma_S}(x)}} \right) \\
 &\quad + \log \left(\frac{1 + e^{-e^t s_1^R(x)} + e^{-e^t (s_1^R(x)+s_2^R(x))} + \dots + e^{-e^t (s_1^R(x)+s_2^R(x)+\dots+s_{n_R-1}^R(x))}}{1 - e^{-e^t \ell_{\gamma_S}(x)}} \right).
 \end{aligned}$$

Using Eq. (17) to replace the $\log \left(\frac{v_0^R}{v_0^L} \right)$ term here with expressions in τ_0 , $\ell_{\gamma_S}(x)$ and shearing parameters $s_i^L(x)$, $s_j^R(x)$ then gives the desired result. \square

Remark 3.4. Similarly define μ_- but with the $n_L + n_R$ bi-infinite geodesic spiralling towards γ_S in the direction of γ_S^{-1} . By symmetry, the relevant

Dehn-twist parameter τ_{γ_S} satisfies:

$$\begin{aligned}
& \tau_{\gamma_S}(\text{stretch}(x, \mu_-, t)) \\
&= e^t \tau_{\gamma_S}(x) \\
&- e^t \log \left(1 + e^{-s_1^L(x)} + e^{-(s_1^L(x) + s_2^L(x))} + \dots + e^{-(s_1^L(x) + \dots + s_{n_L-1}^L(x))} \right) \\
&+ \log \left(1 + e^{-e^t s_1^L(x)} + e^{-e^t(s_1^L(x) + s_2^L(x))} + \dots + e^{-e^t(s_1^L(x) + \dots + s_{n_L-1}^L(x))} \right) \\
&+ e^t \log \left(1 + e^{-s_1^R(x)} + e^{-(s_1^R(x) + s_2^R(x))} + \dots + e^{-(s_1^R(x) + \dots + s_{n_R-1}^R(x))} \right) \\
&- \log \left(1 + e^{-e^t s_1^R(x)} + e^{-e^t(s_1^R(x) + s_2^R(x))} + \dots + e^{-e^t(s_1^R(x) + \dots + s_{n_R-1}^R(x))} \right).
\end{aligned}$$

However, we should emphasise that in changing the lamination from μ_+ to μ_- , the shearing parameters s_i^L and s_j^R are different to the identically denoted variables for μ_+ . In particular, it is *generically false* that

$$\tau_{\gamma_S}(\text{stretch}(x, \mu_+, t)) + \tau_{\gamma_S}(\text{stretch}(x, \mu_-, t)) = 2e^t \tau_{\gamma_S}(x).$$

3.2.2. Opposite spiralling case. We now consider the case where we stretch with respect to complete geodesic laminations where the leaves on the left and right components of $S \setminus \gamma_S$ spiral toward γ_S in opposite directions. Let λ_+ denote the complete lamination on S having, as its leaves (see Fig. 11)

- the closed geodesic γ_S ;
- $n_L + n_R$ geodesics going from the boundary cusps of S and spiralling towards γ_S in such a way that those on the left component spiral in the direction of γ_S , and those on the right component spiral in the direction of γ_S^{-1} ;
- the boundary geodesics on S .

We label the ideal triangles on the left component by $\Delta_1^L, \dots, \Delta_{n_L}^L$ and those on the right by $\Delta_1^R, \dots, \Delta_{n_R}^R$ ordered as in Fig. 12. Let $s_i^L \in \mathbb{R}$ denote the shearing parameter for the edge shared by Δ_i^L and Δ_{i+1}^L (cyclically indexed) and likewise let $s_i^R \in \mathbb{R}$ denote the shearing parameter for the edge shared by Δ_i^R and Δ_{i+1}^R . We once again emphasise that these shearing parameters may differ in sign from the usual convention, and s_i^X (for $X = L, R$) is defined as the signed hyperbolic distance travelled, along the ideal geodesic shared by Δ_i^X and Δ_{i+1}^X , from the anchor for Δ_{i+1}^X to the anchor for Δ_i^X . This is signed to be positive if one moves towards closer to γ_S , and negative if one moves farther from γ_S . As an example, all of the shearing parameters in Fig. 12 are positive.

Lemma 3.5. *Given an (n_L, n_R) -crowned hyperbolic annulus structure on S , set τ_{γ_S} to be the twist parameter for γ_S obtained by homotoping the ideal geodesic going from the “leftmost” cusp on Δ_1^L to the “leftmost” cusp on*

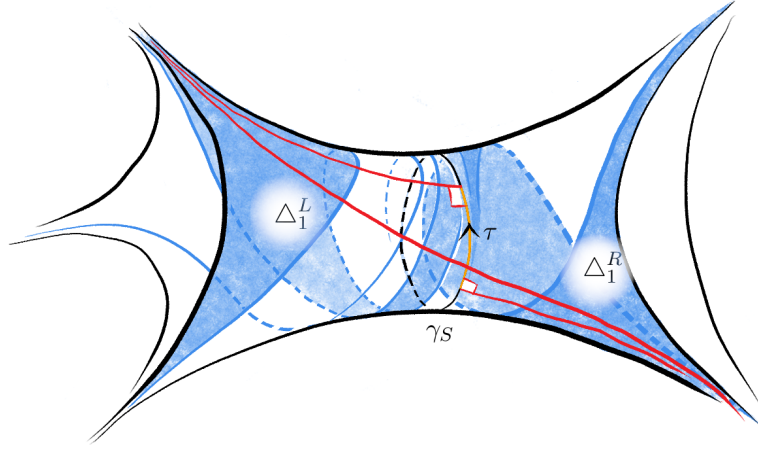


FIGURE 11. A $(3,2)$ -crowned annulus with τ coordinate given by the signed length of the orange geodesic arc.

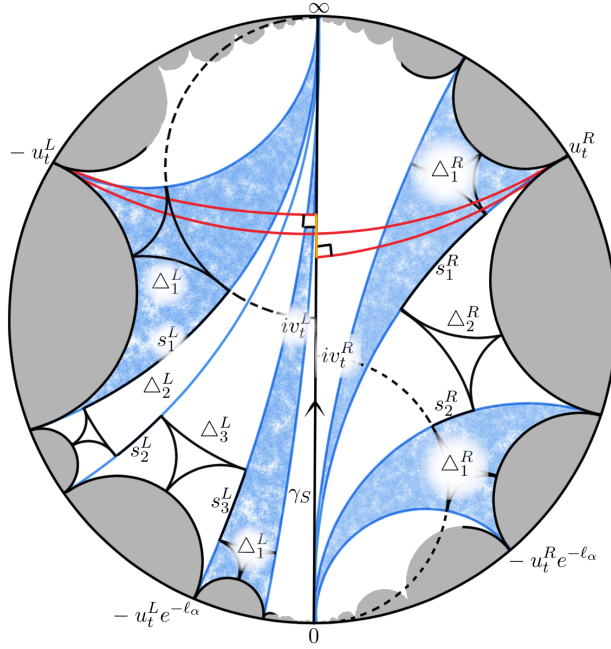


FIGURE 12. The universal cover of the $(3,2)$ -annulus in Fig. 11.

Δ_1^R (see Fig. 12). Then,

$$\begin{aligned}
 & \tau_{\gamma_S}(\text{stretch}(x, \lambda_+, t)) \\
 &= e^t \tau_{\gamma_S}(x) + 2 \left(\log(1 - e^{-e^t \ell_{\gamma_S}(x)}) - e^t \log(1 - e^{-\ell_{\gamma_S}(x)}) \right) \\
 &+ e^t \log \left(1 + e^{-s_1^L(x)} + e^{-(s_1^L(x) + s_2^L(x))} + \dots + e^{-(s_1^L(x) + \dots + s_{n_L-1}^L(x))} \right) \\
 &- \log \left(1 + e^{-e^t s_1^L(x)} + e^{-e^t (s_1^L(x) + s_2^L(x))} + \dots + e^{-e^t (s_1^L(x) + \dots + s_{n_L-1}^L(x))} \right) \\
 &+ e^t \log \left(1 + e^{-s_1^R(x)} + e^{-(s_1^R(x) + s_2^R(x))} + \dots + e^{-(s_1^R(x) + \dots + s_{n_R-1}^R(x))} \right) \\
 &- \log \left(1 + e^{-e^t s_1^R(x)} + e^{-e^t (s_1^R(x) + s_2^R(x))} + \dots + e^{-e^t (s_1^R(x) + \dots + s_{n_R-1}^R(x))} \right).
 \end{aligned}$$

Proof. We employ the same strategy of proof as for Lemma 3.3. Let τ_0 and τ_t respectively denote the twist parameter τ_{γ_S} evaluated at x and $\text{stretch}(x, \lambda_+, t)$. As before, this tells us that the correspondingly labelled u_0^L, u_0^R and u_t^L, u_t^R (see Fig. 12) satisfy

$$(18) \quad \tau_0 = \log \left(\frac{u_0^R}{u_0^L} \right) \quad \text{and} \quad \tau_t = \log \left(\frac{u_t^R}{u_t^L} \right).$$

From Lemma 3.2, we obtain

$$(19) \quad \begin{aligned} \tau_0 &= \log \left(\frac{u_0^R}{v_0^R} \cdot \frac{v_0^L}{u_0^L} \cdot \frac{v_0^R}{v_0^L} \right) = \log \left(\frac{v_0^R}{v_0^L} \right) - \log \left(\frac{v_0^R}{u_0^R} \right) - \log \left(\frac{u_0^L}{v_0^L} \right) \\ &= \log \left(\frac{v_0^R}{v_0^L} \right) \\ &\quad - \log \left(\frac{1 + e^{-s_1^L(x)} + e^{-(s_1^L(x)+s_2^L(x))} + \dots + e^{-(s_1^L(x)+s_2^L(x)+\dots+s_{n_L-1}^L(x))}}{1 - e^{-\ell_{\gamma_S}(x)}} \right) \\ &\quad - \log \left(\frac{1 + e^{-s_1^R(x)} + e^{-(s_1^R(x)+s_2^R(x))} + \dots + e^{-(s_1^R(x)+s_2^R(x)+\dots+s_{n_R-1}^R(x))}}{1 - e^{-\ell_{\gamma_S}(x)}} \right). \end{aligned}$$

To see how the last term in Eq. (19) follows from Lemma 3.2, apply the Möbius transformation $z \mapsto -z^{-1}$ to the universal cover of the right component of $S \setminus \gamma_S$ and observe that its image precisely satisfies the configuration of Lemma 3.2, and that our particular definition of shearing parameter remains unchanged under such a transformation, and the only changes are for u_0^R and $v_0^R i$, which respectively get sent to $-(u_0^R)^{-1}$ and $(v_0^R)^{-1} i$. Hence

$$\frac{v_0^R}{u_0^R} = \frac{(u_0^R)^{-1}}{(v_0^R)^{-1}} = \frac{1 + e^{-s_1^R(x)} + e^{-(s_1^R(x)+s_2^R(x))} + \dots + e^{-(s_1^R(x)+s_2^R(x)+\dots+s_{n_R-1}^R(x))}}{1 - e^{-\ell_{\gamma_S}(x)}}.$$

Proceeding as in the proof of Lemma 3.3, we observe that the segment between v_0^L and v_0^R is e^t -stretched to a segment of length $\log \left(\frac{v_0^R}{v_0^L} \right) = e^t \log \left(\frac{v_0^R}{v_0^L} \right)$. Substituting this into the e^t -stretched version of Eq. (19) yields

$$\begin{aligned} \tau_t &= \log \left(\frac{u_t^R}{v_t^R} \cdot \frac{v_t^L}{u_t^L} \cdot \frac{v_t^R}{v_t^L} \right) = \log \left(\frac{v_t^R}{v_t^L} \right) - \log \left(\frac{v_t^R}{u_t^R} \right) - \log \left(\frac{u_t^L}{v_t^L} \right) \\ &= e^t \log \left(\frac{v_0^R}{v_0^L} \right) \\ &\quad - \log \left(\frac{1 + e^{-e^t s_1^L(x)} + e^{-e^t(s_1^L(x)+s_2^L(x))} + \dots + e^{-e^t(s_1^L(x)+s_2^L(x)+\dots+s_{n_L-1}^L(x))}}{1 - e^{-e^t \ell_{\gamma_S}}} \right) \\ &\quad - \log \left(\frac{1 + e^{-e^t s_1^R(x)} + e^{-e^t(s_1^R(x)+s_2^R(x))} + \dots + e^{-e^t(s_1^R(x)+s_2^R(x)+\dots+s_{n_R-1}^R(x))}}{1 - e^{-e^t \ell_{\gamma_S}(x)}} \right). \end{aligned}$$

Replacing the $\log\left(\frac{v_0^R}{v_0^L}\right)$ term here, using Eq. (19), with expressions in τ_0 , $\ell_{\gamma_S}(x)$ and shearing parameters then yields the desired result. \square

Remark 3.6. We can analogously define λ_- but with $n_L + n_R$ geodesics going from the boundary cusps of S and spiralling towards γ_S in such a way that those on the left component spiral in the direction of γ_S^{-1} , and those on the right component spiral in the direction of γ_S . By symmetry, the relevant Dehn-twist parameter τ_{γ_S} satisfies:

$$\begin{aligned} & \tau_{\gamma_S}(\text{stretch}(x, \lambda_-, t)) \\ &= e^t \tau_{\gamma_S}(x) - 2 \left(\log(1 - e^{-e^t \ell_{\gamma_S}(x)}) - e^t \log(1 - e^{-\ell_{\gamma_S}(x)}) \right) \\ & - e^t \log \left(1 + e^{-s_1^L(x)} + e^{-(s_1^L(x) + s_2^L(x))} + \dots + e^{-(s_1^L(x) + \dots + s_{n_L-1}^L(x))} \right) \\ & + \log \left(1 + e^{-e^t s_1^L(x)} + e^{-e^t (s_1^L(x) + s_2^L(x))} + \dots + e^{-e^t (s_1^L(x) + \dots + s_{n_L-1}^L(x))} \right) \\ & - e^t \log \left(1 + e^{-s_1^R(x)} + e^{-(s_1^R(x) + s_2^R(x))} + \dots + e^{-(s_1^R(x) + \dots + s_{n_R-1}^R(x))} \right) \\ & + \log \left(1 + e^{-e^t s_1^R(x)} + e^{-e^t (s_1^R(x) + s_2^R(x))} + \dots + e^{-e^t (s_1^R(x) + \dots + s_{n_R-1}^R(x))} \right). \end{aligned}$$

We again emphasise that in changing the lamination from λ_+ to λ_- , the new shearing parameters s_i^L and s_j^R are (a priori) different to the corresponding variables for λ_+ , and it is *generically false* that

$$\tau_{\gamma_S}(\text{stretch}(x, \lambda_+, t)) + \tau_{\gamma_S}(\text{stretch}(x, \lambda_-, t)) = 2e^t \tau_{\gamma_S}(x).$$

3.3. Stretch maps for complete finite-area hyperbolic surfaces. We now explain how we can utilise the expressions for stretch maps on crowned annuli to describe stretch maps for finite-leaf laminations of complete finite-area hyperbolic surfaces, i.e., general crowned hyperbolic surfaces with finitely many boundaries, each with finitely many punctures.

A finite-leaf complete lamination Λ on a complete finite-area hyperbolic surface (S, x) consists of simple closed geodesics and simple bi-infinite geodesics which either spiral to (and from) the aforementioned simple closed geodesics or to cusps. To each closed geodesic γ , we can associate its length and Fenchel–Nielsen twist functions $\ell_\gamma: \mathcal{T}(S) \rightarrow \mathbb{R}_{>0}$ and $\tau_\gamma: \mathcal{T}(S) \rightarrow \mathbb{R}$. We observe that each bi-infinite geodesic in Λ is the shared edge between two (possibly non-distinct) ideal triangles and hence has a well-defined shearing parameter s_α . The combination of the Fenchel–Nielsen twist parameters and the shearing parameters globally parameterises $\mathcal{T}(S)$: the shearing parameters tell us how to glue ideal triangles to constitute the complement of the closed geodesics on S , and the twist parameter tells us how to glue together those surfaces to form (S, x) (also see [5, Section 7] and [9, Section 2.4.1]). We further note that the length parameter ℓ_γ is a linear combinations of the shearing parameters of bi-infinite geodesics spiralling to α (see, e.g. [2, Section 2.3.1]). Specifically, if the geodesic *rays* (i.e., rays consisting of one end

of a bi-infinite geodesic) spiralling to a given side of γ lie on the bi-infinite geodesics $\alpha_1, \dots, \alpha_k$, then there is a choice of signs $\epsilon_i \in \{+, -\}$ such that

$$(20) \quad \ell_\gamma = \epsilon_1 s_{\alpha_1} + \dots + \epsilon_k s_{\alpha_k}.$$

The sign choices here depend on the initial choice of sign for the shearing parameters, and can be determined without mysterious numerology. The upshot here is that the expressions for Thurston's stretch map in Theorem 3.1 and Lemmas 3.3 and 3.5 are given in terms of a natural coordinate system on $\mathcal{T}(S)$ which is a mixture of Fenchel–Nielsen twist coordinates and Thurston's shearing parameters.

Thurston's shearing parameter measures the signed length of the geodesic segment on a bi-infinite geodesic leaf α in Λ bounded by the anchor points of adjoining ideal triangles to α . Under the action of a K -stretch map, these parameters are multiplied by a factor of K . Likewise, under the action of a K -stretch map, the length ℓ_γ of a simple closed geodesic γ in Λ is also multiplied by a factor of K (note that this agrees with Eq. (20)).

The action of a K -stretch map on Fenchel–Nielsen twist parameters is more complex, and is precisely spelled out by Theorem 3.1 and Lemmas 3.3 and 3.5. To be more concrete: given a simple closed geodesic $\gamma \subset \Lambda$, consider an arbitrary lift $\tilde{\gamma}$ of γ to the universal cover \tilde{S} of S , and consider the subset of \tilde{S} comprised of $\tilde{\gamma}$ and all of ideal triangles in the complement of the lift of Λ in \tilde{S} which spiral to $\tilde{\gamma}$. The metric completion \tilde{A} of this collection of ideal triangles is invariant under translation along $\tilde{\gamma}$ by the deck transformation in $\pi_1(S)$ corresponding to $\tilde{\gamma}$. The quotient A of \tilde{A} by the aforementioned deck transformation is a crowned annulus around a copy of γ , and the restriction of Λ on A simultaneously takes the form of the laminations described in Theorem 3.1 and Lemmas 3.3 and 3.5 whilst having the same shearing parameters and twist coordinates as corresponding geodesic leaves on Λ on (S, x) , which allows us to invoke Theorem 3.1 and Lemmas 3.3 and 3.5 to determine the deformation of the Fenchel–Nielsen twist coordinate τ_γ for γ .

3.4. Antistretch paths.

Definition 3.7 (antistretch maps, [29]). Although we have hitherto defined e^t -stretch maps and stretch paths exclusively for $t \geq 0$, the notion of stretch maps and stretch path are perfectly well-defined for $t < 0$, and are respectively referred to as the *antistretch map* and *antistretch path*.

Theorem 3.8 (back-time convergence [29, Theorem 4.1]). *Consider a finite-leaf lamination Λ , and denote its closed leaves by $\gamma_1, \dots, \gamma_k$. As $t \rightarrow \infty$, for any $x \in \mathcal{T}(S)$, the e^{-t} -stretch path with respect to Λ $x_{-t} := \text{stretch}(x, \Lambda, -t)$ converges to the uniform weighted projectivised multicurve*

$$[\gamma_1 + \gamma_2 + \dots + \gamma_k] \in \text{PML}(S) = \partial_\infty \mathcal{T}(S),$$

where $\text{PML}(S) = \partial_\infty \mathcal{T}(S)$ denotes the boundary of the Thurston compactification of $\mathcal{T}(S)$. Note that this is independent of the starting point x .

Proof. We only sketch the proof of the above result. To begin with, we make the observation that the shearing lengths assigned to the bi-infinite leaves in Λ , and the Fenchel-Nielsen parameters attached to $\gamma_1, \dots, \gamma_k$ suffice to parametrise the Teichmüller space $\mathcal{T}(S)$. As $s \rightarrow \infty$ the e^{-s} -stretch map (i.e., the antistretch map) with respect to Λ ,

- (1) all of the shearing lengths for the bi-infinite leaves in Λ exponentially shrink to 0,
- (2) the Fenchel-Nielsen length parameter for γ_j also exponentially shrinks to 0, and
- (3) the Fenchel-Nielsen twist parameter for γ_j behaves as $\pm 2s + o(1)$.

Point 3 comes from analysing the behaviour of Lemma 3.5 and Remark 3.6 as $t = -s \rightarrow -\infty$ there: every term tends either to 0 or a constant except for

$$\pm 2 \log \left(1 - e^{-e^s \ell_j(x)} \right) = \mp 2s + O(1).$$

First two points tell us that the geometry of $S \setminus (\gamma_1 \cup \dots \cup \gamma_k)$ converges to a cusped surface without open ends, possibly with multiple connected components, in which each of $\gamma_1, \dots, \gamma_k$ corresponds to a cusp. In particular, each side of the standard collar neighbourhood around γ_j is an annulus of width

$$\operatorname{arcsinh} \left(\frac{1}{\sinh(\frac{1}{2} e^{-s} \ell_{\gamma_j})} \right) = s + O(1),$$

and geometrically tends to a cusp bounded by a horocycle of length 2. Since the complement of the collar neighbourhoods around the $\{\gamma_j\}$ converges to a compact horocycle bordered surface, the complement of the collar neighbourhood must have diameter bounded above by some constant for all t . Therefore, the length of an arbitrary simple closed geodesic on the complement of the $\{\gamma_j\}$ necessarily stabilises and converges as s tends to ∞ . In contrast, for an arbitrary simple closed geodesic γ that intersects the $\gamma_1 \sqcup \dots \sqcup \gamma_k$ by $M > 1$ times, its length grows due to the following two dominant contributing factors:

- the growing width of the collar neighbourhoods around the γ_j , and
- the accumulating Fenchel-Nielsen twists around the $\{\gamma_j\}$ induced by the antistretch.

The former contributing factor tells us that the length of γ as $s \rightarrow \infty$ is bounded below by $2Ms + o(1)$. Coupled with the second contributing factor, we easily obtain the coarse upper bound of $4Ms + o(1)$, which we now improve upon. Observe that the direction of the $\mp 2s + O(1)$ Fenchel-Nielsen twist is orthogonal to the length $2s + O(1)$ geodesic cutting across the collar neighbourhood surrounding each γ_j . The hyperbolic Pythagoras' theorem (see, e.g.: [1, Theorem 2.2.2.i]) therefore tells us that the length of each subsegment of γ that cuts across a the collar neighbourhood of γ_j grows as

$$2 \operatorname{arccosh}(\cosh(s + O(1)) \cosh(t + O(1))) = 2s + O(1),$$

which shows that length of γ_j is bounded above by $2Ms + O(1)$ — which equals the previously established lower bound. The embedding of $x_{-s} = \text{stretch}(x, \Lambda, -s)$ in $\mathbb{P}(\mathbb{R}_{\geq 0}^s)$ can be normalised via division by $2s$ and represented by

$$\phi : \mathcal{S} \rightarrow \mathbb{R}_{\geq 0}, \quad \gamma \mapsto i(\gamma, \gamma_1 + \gamma_2 + \cdots + \gamma_k),$$

where $i(\alpha, \beta)$ denotes the geometric intersection number between α and β . The limit of x_{-s} in $\mathcal{PML}(S)$ is therefore given by the projective class of $\phi \in \mathbb{R}_{\geq 0}^s$, i.e., the projective multicurve $[\gamma_1 + \gamma_2 + \cdots + \gamma_k] \in \mathcal{PML}(S)$. \square

4. STRETCH VECTORS

We now return to the context where $S = S_{g,n}$ is a general orientable surface of genus g with $n \geq 0$ cusps. The aim of this subsection is to derive results pertaining to tangent vectors induced by stretch paths.

Definition 4.1 (stretch vector). Consider a point $x \in \mathcal{T}(S)$ and a complete lamination λ on S . Then Thurston's stretch map construction defines a geodesic ray $\{\text{stretch}(x, \lambda, t)\}_{t \geq 0}$, where $\text{stretch}(x, \lambda, t)$ is obtained from x by e^t -stretching with respect to λ . We refer to the tangent vector

$$v_\lambda := \left. \frac{d}{dt} \right|_{t=0} \text{stretch}(x, \lambda, t) \in T_x \mathcal{T}(S)$$

as the *stretch vector* at x with respect to λ .

We begin by giving a sketch-of-proof for an infinitesimal version of the fact referenced in the second last sentence of Remark 2.5. This claim is implicitly asserted in the proof of [30, Theorem 5.1].

Lemma 4.2 (maximally stretched lamination [30, Theorem 5.1]). *Consider a maximal (geodesic) lamination λ with corresponding stretch vector $v_\lambda \in T_x \mathcal{T}(S)$. For any measured lamination $\mu \in \mathcal{ML}(S)$ such that the geodesic lamination $|\mu|$ supporting μ has a nonempty transverse intersection with λ , we have*

$$v_\lambda(\log \ell(\mu)) < 1.$$

Proof. We may assume without loss of generality that $\ell_x(\mu) = 1$, and hence we compute

$$(21) \quad v_\lambda(\log \ell(\mu)) = \frac{v_\lambda(\ell(\mu))}{\ell_x(\mu)} = v_\lambda(\ell(\mu)) = \lim_{t \rightarrow 0} \frac{\ell_{x_t}(\mu) - \ell_x(\mu)}{t},$$

where $x_t = \text{stretch}(x, \lambda, t)$. Although the precise value of $\ell_{x_t}(\mu)$ is finicking to compute, it suffices for our purposes to find a small enough upper bound for this value so that Eq. (21) approaches something less than 1. Specifically, we compute the length of the quasi-geodesic to which μ deforms with respect to the change of metric parameterised by x_t . To this end, we first show that there exist some $\theta \in (0, \frac{\pi}{2})$ and a subset ν of μ of length $L \in (0, 1]$ such that the geodesics constituting ν intersect the (vertical) geodesic partial foliation for λ at an angle between θ and $\frac{\pi}{2}$. Computing the change in the norm of each unit tangent vector along the geodesic segments of ν as one deforms the

metric $x = x_0$ to x_t shows that the new norm is smaller than $\cos(\theta)e^t$. This suffices to show that Eq. (21) is strictly less than $L \cos(\theta) + (1 - L) < 1$. \square

Lemma 4.3. *Consider a pair of complete geodesic laminations Λ_\pm on $(S, x) \in \mathcal{T}(S)$ which agree everywhere except on a $(1, 1)$ -cusped annulus contained in S with core geodesic γ_0 of length ℓ_0 , whereby Λ_\pm respectively contain λ_\pm (as by the notation of Theorem 3.1). Then, the difference between the stretch vectors v_{Λ_\pm} is expressed as:*

$$(22) \quad v_{\Lambda_+} - v_{\Lambda_-} = \left(\frac{4\ell_0 e^{-\ell_0}}{1 - e^{-\ell_0}} - 4 \log(1 - e^{-\ell_0}) \right) E_{\gamma_0},$$

where E_{γ_0} is the unit Fenchel-Nielsen twist vector with respect to γ_0 . In particular,

$$(23) \quad \|v_{\Lambda_+} - v_{\Lambda_-}\|_{\text{Th}} = (4\ell_0 e^{-\ell_0} + o(\ell_0 e^{-\ell_0})) \|E_{\gamma_0}\|_{\text{Th}}, \text{ as } \ell_0 \rightarrow \infty.$$

Proof. Let $S_0 \subset S$ denote the specified $(1, 1)$ -cusped annulus in S containing γ_0 as its core geodesic. We first observe that since the metric deformations corresponding to λ_- and λ_+ on the boundary of S_0 are identical, the only difference between the metrics $\text{stretch}(x, \Lambda_+, t)$ and $\text{stretch}(x, \Lambda_-, t)$ may be homotoped to occur only within S_0 . In other words, the restriction of the metrics $\text{stretch}(x, \Lambda_\pm, t)$ to $S \setminus S_0$ are isometric. Moreover, since there is only one way to attach a pair of (crowned) half-pants of cuff length $e^t \ell_0$ to each of the two crowned boundaries of $S \setminus S_0$, this further informs us that the geodesically bordered hyperbolic metrics induced by $\text{stretch}(x, \Lambda_\pm, t)$ on $S \setminus \gamma_0$ are also necessarily isometric. Therefore, the vector $v_{\Lambda_+} - v_{\Lambda_-}$ is a multiple of E_{γ_0} , the unit Fenchel-Nielsen twist vector with respect to γ_0 . Theorem 3.1 tells us that this factor is

$$\left. \frac{d}{dt} \right|_{t=0} 4 \left(\log(1 - e^{-e^t \ell_0}) - e^t \log(1 - e^{-\ell_0}) \right) = 4 \left(\frac{\ell_0 e^{-\ell_0}}{1 - e^{-\ell_0}} - \log(1 - e^{-\ell_0}) \right).$$

This gives Eq. (22), and Eq. (23) follows by L'Hôpital's rule. \square

Corollary 4.4. *The tangent vectors v_{Λ_+} and v_{Λ_-} at $T_x \mathcal{T}(S)$ are distinct.*

Proof. Note that

$$\frac{4\ell_0 e^{-\ell_0}}{1 - e^{-\ell_0}} - 4 \log(1 - e^{-\ell_0}) = \text{positive term} - (\text{negative term}) > 0.$$

Lemma 4.3 then tells us that $v_{\Lambda_+} \neq v_{\Lambda_-}$. \square

We can give a much more general version of Corollary 4.4. Let Λ_1 and Λ_2 be two complete geodesic laminations of (S, x) such that $\Lambda_1 \cap \Lambda_2$ contains the boundary of some crowned annulus $T \subset S$, as well as the unique simple closed geodesic γ_T contained in the interior of T . Further require that $\Lambda_1 \cap T$ and $\Lambda_2 \cap T$, which are both complete geodesic laminations on T , are respectively equal to λ_+ and λ_- in the sense of Lemma 3.5. Given this setup, we have the following result:

Lemma 4.5. *Suppose all of the shearing parameters assigned to the bi-infinite geodesic leaves of $\Lambda_1 \cap T$ and $\Lambda_2 \cap T$ which spiral to γ_T are positive. Then, the tangent vectors v_{Λ_1} and v_{Λ_2} at $T_x \mathcal{T}(S)$ are distinct.*

Proof. We can see from Lemma 3.5 and Remark 3.6 that the directional derivatives $v_{\Lambda_1}(\tau_{\gamma_T})$ and $v_{\Lambda_2}(\tau_{\gamma_T})$ are expressible purely in terms of the geometry of $T \subset S$. In particular, the expression of

$$(v_{\Lambda_1} - v_{\Lambda_2})(\tau_{\gamma_T}) = v_{\Lambda_1}(\tau_{\gamma_T}) - v_{\Lambda_2}(\tau_{\gamma_T})$$

is a sum of positive multiples of terms of the form

$$\frac{\ell_{\gamma_T} e^{-\ell_{\gamma_T}}}{1 - e^{-\ell_{\gamma_T}}} - \log(1 - e^{-\ell_{\gamma_T}}),$$

as well as terms of the form

$$\begin{aligned} & \log \left(1 + e^{-s_1} + e^{-(s_1+s_2)} + \dots + e^{-(s_1+\dots+s_{n-1})} \right) \\ & + \frac{s_1 e^{-s_1} + (s_1 + s_2) e^{-(s_1+s_2)} + \dots + (s_1 + \dots + s_{n-1}) e^{-(s_1+\dots+s_{n-1})}}{1 + e^{-s_1} + e^{-(s_1+s_2)} + \dots + e^{-(s_1+\dots+s_{n-1})}}. \end{aligned}$$

Our previous computation for the proof of Corollary 4.4 as well as the assumption of positive shearing parameters then ensures that all of these terms are positive, and hence $(v_{\Lambda_1} - v_{\Lambda_2})(\tau_{\gamma_T}) \neq 0$. Therefore, $v_{\Lambda_1} - v_{\Lambda_2} \neq 0$, and these two vectors are distinct. \square

4.1. Non-chain-recurrent laminations. In order to construct Thurston metric geodesics between two arbitrary points $x, y \in \mathcal{T}(S)$, Thurston shows that any two points in Teichmüller space is joined by a Thurston metric geodesic obtained from concatenating finitely many stretch paths. The initial stretch path segment in the geodesics that Thurston constructs are stretch paths with respect to any complete lamination Λ which extends the maximal ratio-maximising lamination $\mu(x, y)$ (Definition 2.12) between x and y . In particular, one can always take Λ to be the unique complete lamination extending a maximal chain-recurrent lamination containing $\mu(x, y)$. Thurston's geodesic construction [30, Theorem 8.5] tells us to start stretching, with respect to Λ , from x to x_t until such a time that the maximal ratio-maximising lamination $\mu(x_t, y)$ changes in topology. One may then exchange Λ for a complete geodesic lamination containing the updated maximal ratio-maximising lamination. Again, we can choose this new complete geodesic lamination to be a complete lamination extending a maximal chain-recurrent lamination containing the updated maximal ratio-maximising lamination. Repeating this selection procedure suffices to show that one can always join two arbitrary points in Teichmüller space via a concatenation of stretch paths with respect to (complete laminations extending) maximal chain-recurrent laminations. Conversely, when the maximal ratio-maximising lamination $\mu(x, y)$ is a maximal chain-recurrent lamination, there is a unique Thurston geodesic joining x and y given by the stretch path with respect

to (the complete lamination extending) the maximal chain-recurrent lamination containing $\mu(x, y)$.

Based on the above observations, we heuristically expect that stretch paths with respect to maximal chain-recurrent laminations are “extremal” among Thurston metric geodesics, which in turn suggests that stretch vectors for stretch paths with respect to maximal chain-recurrent laminations should constitute extreme points in the unit Thurston norm tangent sphere $\mathbf{S}_x \subsetneq T_x \mathcal{T}(S)$. This is conjectural except in the case when S is a 1-cusped torus or a 4-cusped sphere, but is otherwise unproven. In particular, :

Lemma 4.6. *Let Λ_0 be a complete geodesic lamination that contains λ_0 of Fig. 4 as a sublamination. Then, v_{Λ_0} is not an extreme point in \mathbf{S}_x .*

Proof. Let

- γ_0 denote the simple closed geodesic contained in λ_0 ,
- A denote the convex hull of λ_0 in (S, x) , and
- Λ_{\pm} be the complete geodesic laminations obtained by replacing $\lambda_0 \subset \Lambda_0$ respectively with λ_{\pm} of Fig. 4.

Since Λ_0 , Λ_+ and Λ_- all agree on $S \setminus A$, e^t -stretch maps with respect to these three laminations deform $S \setminus A$ in the same way. This, combined with Theorem 3.1, implies that for any Fenchel-Nielsen coordinate with respect to a pants decomposition containing γ , the coordinates for the e^t -stretch maps with respect to Λ_0 and Λ_{\pm} completely agree except in the twist parameter for γ_0 . Indeed, Theorem 3.1 tells us that the twist parameters for the stretch maps with respect to Λ_{\pm} average to the twist parameter for the stretch map with respect to Λ_0 . Therefore,

$$v_{\Lambda_0} = \frac{1}{2} (v_{\Lambda_+} + v_{\Lambda_-}),$$

where v_{Λ_+} and v_{Λ_-} are the respective stretch vectors with respect to Λ_+ and Λ_- . Therefore, v_{Λ_0} is not an extreme point in \mathbf{S}_x . \square

4.2. Maximal twist lemma. In this subsection, we consider geodesic laminations λ on $S = S_{g,n}$ which contain, as a sublamination, (disjoint) simple closed geodesics $\alpha, \beta_1, \beta_2, \beta_3, \beta_4$ such that the collections $\{\alpha, \beta_1, \beta_2\}$ and $\{\alpha, \beta_3, \beta_4\}$ both bound pairs of pants on S (see Fig. 13). Our goal is to show that stretch vectors for certain λ of the above form maximise and minimise the Fenchel–Nielsen twist parameter for the curve α over the collection of all unit (with respect to the Thurston norm) tangent vectors which “maximally stretch” $\alpha, \beta_1, \beta_2, \beta_3$, and β_4 . This is made precise in Lemma 4.9.

Remark 4.7. We presently only consider the setting where the simple closed geodesics $\alpha, \beta_1, \beta_2, \beta_3, \beta_4$ are all distinct. However, elementary topological arguments suffice to allow us to extend the main results in this subsection also to cases where some of the aforementioned geodesics are the same, namely (up to symmetry):

- $\beta_1 = \beta_2$;

- $\beta_1 = \beta_2$ and $\beta_3 = \beta_4$;
- $\beta_1 = \beta_3$ and $\alpha = \beta_2 = \beta_4$.

A key part of our strategy is to study the lengths of simple closed geodesics which lie in the interior of the 4-holed sphere bounded by $\beta_1, \beta_2, \beta_3, \beta_4$ and which geometrically intersect α precisely twice. Let γ_0 be one shortest such simple closed geodesic, and let γ_m , $m \in \mathbb{Z}$, be the α^m -Dehn twist of γ_0 . The Hausdorff limits λ_{\pm} of γ_m as $m \rightarrow \pm\infty$ are geodesic laminations in the interior of the 4-holed sphere bounded by $\beta_1, \beta_2, \beta_3, \beta_4$ containing α and two other bi-infinite geodesic leaves α^L, α^R spiralling to α in opposite directions from opposite sides (Fig. 13).

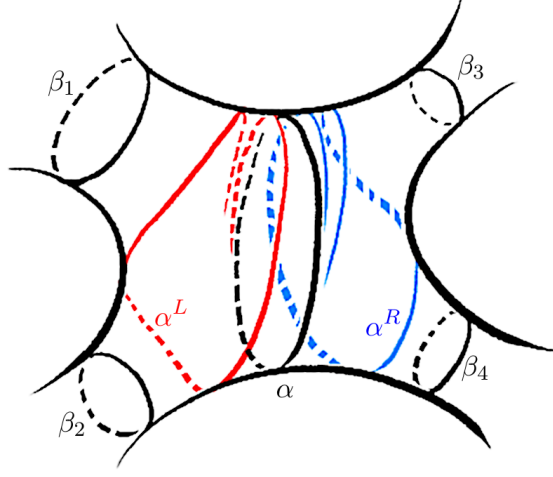


FIGURE 13. Λ_+ is the union of $\alpha, \alpha^L, \alpha^R, \beta_1, \beta_2, \beta_3$ and β_4 .

We complete the lamination $\lambda_+ \cup \{\beta_1, \beta_2, \beta_3, \beta_4\}$ arbitrarily to a complete geodesic lamination Λ_+ on S and likewise complete $\lambda_- \cup \{\beta_1, \beta_2, \beta_3, \beta_4\}$ to Λ_- . Note that the length of γ_m under stretching by Λ_{\pm} depends only on the restriction of Λ_{\pm} to the 4-holed sphere bounded by $\beta_1, \beta_2, \beta_3, \beta_4$. There are (at most) 2^4 different choices (two choices for each β_i) for how one completes λ_{\pm} to a complete lamination on the 4-holed sphere bounded by $\beta_1, \beta_2, \beta_3, \beta_4$, and each of these 2^4 choices come from adding geodesic leaves $\sigma_1, \sigma_2, \sigma_3, \sigma_4$ which spiral from $\beta_1, \beta_2, \beta_3, \beta_4$ to α .

We first study the behaviour of the length of $\gamma = \gamma_m$, for $m \gg 0$, under the K -stretch map with respect to Λ_+ . We approximate γ_m by a homotopic curve ς obtained by concatenating segments along leaves of the stretch-invariant horocyclic foliation and the stretch-invariant geodesic lamination for Λ_+ . Since γ intersects α twice, the latter separates γ into two

geodesic segments γ^L and γ^R . When $m \gg 0$ is large, we know that γ_m well-approximates λ_+ and hence the geodesic segment γ^L is homotopic (whilst fixing endpoints) to a path which traverses along:

- a segment along a leaf of the stretch-invariant horocyclic foliation, so as to go from α to α^L , followed by
- a segment along α^L , followed by
- a segment along a leaf of the stretch-invariant horocyclic foliation, so as to go from α^L back to α .

We homotope γ^R analogously to a horocyclic segment, followed by a α^R segment and then another horocyclic segment. See Fig. 14 where this is depicted on a lift to the universal cover, with labels for each of these six segments along leaves of stretch-invariant foliations.

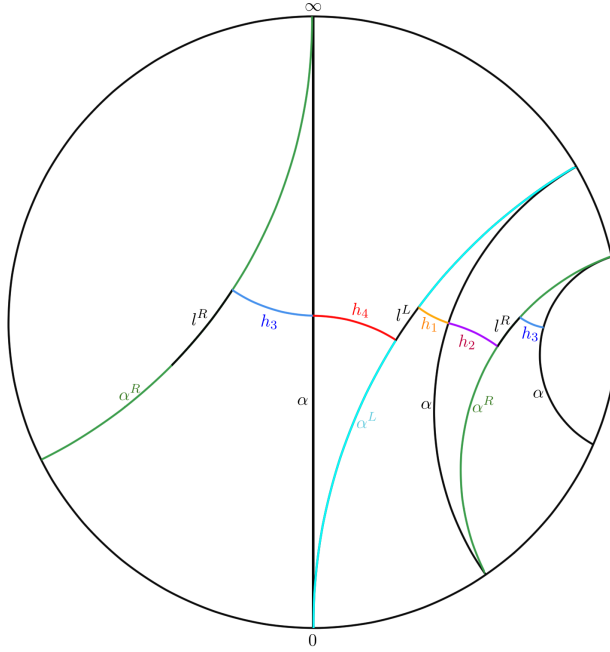


FIGURE 14. The curve γ_m is homotopic to a closed loop ς obtained by concatenating l^L , l^R , h_1 , h_2 , h_3 , and h_4 .

We next express the length of γ in terms of the lengths of these six segments by computing a $\mathrm{PSL}_2(\mathbb{R})$ -conjugacy representative of the holonomy matrix for γ .

Lemma 4.8. *The length of γ is related to lengths of $l^L, l^R, h_1, h_2, h_3, h_4$ (we use the same symbols to represent their lengths) by the following formula:*

$$(24) \quad 2 \cosh\left(\frac{\ell_\gamma}{2}\right) = e^{\frac{1}{2}(l^L + l^R)}(1 + h_1 h_2)(1 + h_3 h_4) + e^{\frac{1}{2}(l^L - l^R)} h_1 h_4 \\ + e^{\frac{1}{2}(l^R - l^L)} h_2 h_3 + e^{-\frac{1}{2}(l^L + l^R)}.$$

Proof. Without loss of generality, we position the lift of α^L depicted in Fig. 13 with endpoints placed at 0 and ∞ and with a lift of the h_4 segment joined to i on the right of the lift of α (see Fig. 14). We obtain the holonomy matrix for γ by composing the matrices corresponding to the following sequence of transformations:

- (1) hyperbolically translate along the lift of α^L by l^L , so that the endpoint of the lift of h_1 is now placed at i ;
- (2) rotate by π around i , so that the lift of h_1 is now placed to the right of i on the lift of α^L ;
- (3) parabolically translate along the lift of h_1 , so that the shared endpoint of h_1 and h_2 is now placed at i ;
- (4) rotate by π so that the lift of h_2 is now to the left of i ;
- (5) parabolically translate along the lift of h_2 , so that the shared endpoint of the respective lifts of h_2 and α^R is positioned at i ;
- (6) hyperbolically translate along the lift of α^R by l^R , so that the endpoint of the lift of h_3 is now placed at i ;
- (7) rotate by π around i , so that the lift of h_3 is now placed to the right of i on the lift of α^R ;
- (8) parabolically translate along the lift of h_3 , so that the shared endpoint of h_3 and h_4 is now placed at i ;
- (9) rotate by π so that the lift of h_4 is now to the left of i ;
- (10) parabolically translate along the lift of h_4 , so that the shared endpoint of the respective lifts of h_4 and α^R is positioned at i . This results in a γ -translate of our initial configuration.

Putting all of these matrices together (ordered from right to left), we obtain:

$$\begin{aligned} & \pm \begin{bmatrix} 1 & h_4 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix} \begin{bmatrix} 1 & -h_3 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix} \begin{bmatrix} e^{\frac{l^R}{2}} & 0 \\ 0 & e^{-\frac{l^R}{2}} \end{bmatrix} \begin{bmatrix} 1 & h_2 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix} \begin{bmatrix} 1 & -h_1 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix} \begin{bmatrix} e^{\frac{l^L}{2}} & 0 \\ 0 & e^{-\frac{l^L}{2}} \end{bmatrix} \\ & = \pm \begin{bmatrix} e^{\frac{1}{2}(l^L+l^R)}(1+h_1h_2)(1+h_3h_4) + e^{\frac{1}{2}(l^L-l^R)}h_1h_4 & * \\ * & e^{\frac{1}{2}(l^R-l^L)}h_2h_3 + e^{-\frac{1}{2}(l^L+l^R)} \end{bmatrix}. \end{aligned}$$

We omit the non-diagonal entries because we only need to know the trace of this matrix, which is equal to $2 \cosh(\frac{\ell_\gamma}{2})$. This yields Eq. (24). \square

The homotopy class of γ is unaffected by the $K = e^t$ -stretching with respect to Λ_+ , and the structure of its homotopy representative in terms of stretch-invariant segments is preserved over the e^t -stretch path x_t with respect to Λ_+ . Therefore, the above formula is both well-defined and true with

$$l_\gamma(t) := \ell_\gamma(x_t), \quad l^L(t) := l^L(x_t), \quad l^R(t) := l^R(x_t)$$

$h_j(t)$ defined as the arc-length of the h_j horocyclic segment for $x_t \in \mathcal{T}(S)$. Since $l^L(t)$ and $l^R(t)$ measure the lengths of segments on a leaf in Λ_+ , they

must be $K = e^t$ stretched:

$$(25) \quad l^L(t) = l^L(0)e^t \text{ and } l^R(t) = l^R(0)e^t.$$

Note in particular that

$$(26) \quad (l^L)'(0) = l^L(0) \text{ and } (l^R)'(0) = l^R(0).$$

The behaviour of the horocyclic segments h_i with respect to t is slightly more nuanced. First observe that the complement of Λ_+ on the pair of pants bounded by $\{\beta_1, \beta_2, \alpha\}$ consists of two ideal triangles Δ_1 and Δ_2 . The set $h_1 \cap \Delta_1$ is comprised of infinitely many connected components all of which are horocyclic segments. These horocyclic segments exponentially decrease in length, the closer a component of $h_1 \cap \Delta_1$ is to α . Let $d_1(t)$ denote the distance between the longest horocyclic subsegment of $h_1 \cap \Delta_1$ and the nearest horocyclic boundary edges of the central stable triangle on Δ_1 (with respect to the metric $x_t \in \mathcal{T}(S)$) when Δ_1 is laid flat on the hyperbolic plane. Then, the lengths of the segments comprising $h_1 \cap \Delta_1$ are given by

$$e^{-d_1(t)}, e^{-d_1(t)-l_\gamma(t)}, e^{-d_1(t)-2l_\gamma(t)}, e^{-d_1(t)-3l_\gamma(t)}, \dots$$

Likewise consider the horocyclic segments which constitute the connected components of $h_1 \cap \Delta_2$, and let d_2 denote the distance between the longest of these segments and nearest horocyclic boundary of the central stable triangle on Δ_2 when Δ_2 is laid flat on the hyperbolic plane. The lengths of the segments comprising $h_1 \cap \Delta_2$ are then given by

$$e^{-d_2(t)}, e^{-d_2(t)-l_\gamma(t)}, e^{-d_2(t)-2l_\gamma(t)}, e^{-d_2(t)-3l_\gamma(t)}, \dots$$

When the surface is e^t -stretched along Λ_+ , the horocyclic heights $d_i(0) + jl_\gamma(0)$ increase to $e^t(d_i(0) + jl_\gamma(0))$ and hence the length of $h_1(t)$, i.e. the sum of the lengths of all these horocyclic segments, is given by

$$(27) \quad h_1(t) = \frac{e^{-e^t d_1(0)} + e^{-e^t d_2(0)}}{1 - e^{-e^t l_\gamma(0)}}.$$

This is a function which monotonically decreases to 0 as $t \rightarrow \infty$. Crucially, we can determine the derivative of h_1 at $t = 0$:

$$(28) \quad h_1'(0) = -\frac{d_1(0)e^{-d_1(0)} + d_2(0)e^{-d_2(0)}}{1 - e^{-l_\gamma(0)}} - \frac{l_\gamma(0)e^{-l_\gamma(0)}(e^{-d_1(0)} + e^{-d_2(0)})}{(1 - e^{-l_\gamma(0)})^2}.$$

Analogous expressions hold for $h_2(t)$, $h_3(t)$ and $h_4(t)$.

Lemma 4.9 (maximal twist lemma). *Given simple closed geodesics $\beta_1, \beta_2, \beta_3$ and β_4 which bound a 4-holed sphere on $x \in \mathcal{T}(S)$ and a simple closed geodesic α in the interior of the aforementioned 4-holed sphere, define*

λ_+ as above and let Λ_+ be any complete geodesic lamination on S which contains $\beta_1, \beta_2, \beta_3, \beta_4$ and λ_+ . Further let $v_{\Lambda_+} \in T_x \mathcal{T}(S)$ denote the stretch vector for Λ_+ . Then there are no Thurston norm unit vectors $w \in \mathbf{S}_x \subsetneq T_x \mathcal{T}(S)$ such that all the following conditions are all satisfied

- (1) $w(\ell_\alpha) = v_{\Lambda_+}(\ell_\alpha) = \ell_\alpha(x)$, where $\ell_\alpha : \mathcal{T}(S) \rightarrow \mathbb{R}$ is the hyperbolic length function for α ,
- (2) for each $j = 1, 2, 3, 4$, $w(\ell_{\beta_j}) = v_{\Lambda_+}(\ell_{\beta_j}) = \ell_{\beta_j}(x)$,
- (3) $w(\tau_\alpha) > v_{\Lambda_+}(\tau_\alpha)$, where $\tau_\alpha : \mathcal{T}(S) \rightarrow \mathbb{R}$ is a Fenchel–Nielsen twist function for α .

Proof. We show that any $w \in T_x \mathcal{T}(S)$ satisfying the three listed properties necessarily has Thurston norm strictly larger than 1. In particular, we show that there is some γ_m for which $\|w\|_{\text{Th}} \geq w(\log \ell_{\gamma_m}) > 1$. Specifically, we take γ_m to be the α^m -Dehn twist of γ_0 : to be the shortest geodesic on the interior of the convex hull of the β_i which intersects α precisely twice.

Since γ_m is in the interior of the 4-holed sphere bounded by $\beta_1, \beta_2, \beta_3, \beta_4$, its geometry is completely determined by $\ell_{\beta_1}, \ell_{\beta_2}, \ell_{\beta_3}, \ell_{\beta_4}, \ell_\alpha$ and τ_α . This implies that $d\ell_{\gamma_m}$ is a linear combination of $d\ell_\alpha, d\ell_{\beta_1}, d\ell_{\beta_2}, d\ell_{\beta_3}, d\ell_{\beta_4}$ and $d\tau_\alpha$. Since the $d\tau_\alpha$ -component of $d\ell_{\gamma_m}$ is equal to $\partial_{\tau_\alpha}(\ell_{\gamma_m})$, we have

$$w(\log \ell_{\gamma_m}) = \frac{w(\ell_{\gamma_m})}{\ell_{\gamma_m}} = \frac{v_{\Lambda_+}(\ell_{\gamma_m}) + \partial_{\tau_\alpha}(\ell_{\gamma_m})\epsilon}{\ell_{\gamma_m}},$$

for $\epsilon = w(\tau_\alpha) - v_{\Lambda_+}(\tau_\alpha) > 0$ (by condition (3)). This is possible because w and v_{Λ_+} agree with regard to their $d\ell_\alpha, d\ell_{\beta_1}, d\ell_{\beta_2}, d\ell_{\beta_3}$, and $d\ell_{\beta_4}$ -components (see conditions (1) and (2)), and only differ in their $d\ell_{\tau_\alpha}$ -components (see condition (3)).

We approximate the right-most term using Wolpert's cosine formula [35, §3]: for $m \gg 0$, the geodesic γ_m intersects α at an angle close to 0 (at both points), and hence $\partial_{\tau_\alpha}(\ell_{\gamma_m}) \approx 2$. We may hence assume, by setting m sufficiently high, that $\frac{\partial_{\tau_\alpha}(\ell_{\gamma_m})\epsilon}{\ell_{\gamma_m}} \approx \frac{2\epsilon}{\ell_{\gamma_m}} > \frac{\epsilon}{\ell_{\gamma_m}}$. Our goal therefore becomes to show that

$$\|w\|_{\text{Th}} \geq w(\log \ell_{\gamma_m}) > \frac{v_{\Lambda_+}(\ell_{\gamma_m}) + \epsilon}{\ell_{\gamma_m}(x)} > 1.$$

To this end, we now turn to computing $v_{\Lambda_+}(\ell_{\gamma_m})$.

Observe that $v_{\Lambda_+}(\ell_{\gamma_m})$ can be determined by taking ℓ_{γ_m} of the e^t -stretch path x_t , based at $x_0 = x$, for Λ_+ and computing the $t = 0$ derivative of $\ell_{\gamma_m}(t) := \ell_{\gamma_m}(x_t)$. Therefore, we differentiate Eq. (24) and substitute in Eq. (26) to replace $(l^L)'$ and $(l^R)'$. Due to the comparative complexity of Eq. (28), we do not insert it directly into our present computations, but reserve this for a later step in our analysis. In any case, differentiating Eq. (24) and rearranging slightly we obtain that

$$\begin{aligned}
v_{\Lambda_+}(\ell_{\gamma_m}) &= \frac{\partial_t 2 \cosh(\frac{1}{2} l_{\gamma_m}(t))}{\sinh(\frac{1}{2} l_{\gamma_m}(t))} \Big|_{t=0} \\
&= \frac{(l^L(0) + l^R(0))}{2 \sinh(\frac{1}{2} l_{\gamma_m}(0))} \left(e^{\frac{1}{2}(l^L(0) + l^R(0))} (1 + h_1(0)h_2(0))(1 + h_3(0)h_4(0)) \right. \\
&\quad + e^{\frac{1}{2}(l^L(0) - l^R(0))} h_1(0)h_4(0) + e^{\frac{1}{2}(l^R(0) - l^L(0))} h_2(0)h_3(0) + e^{-\frac{1}{2}(l^L(0) + l^R(0))} \Big) (*) \\
&\quad + \frac{e^{\frac{1}{2}(l^L(0) + l^R(0))}}{\sinh(\frac{1}{2} l_{\gamma_m}(0))} \left((h'_1(0)h_2(0) + h_1(0)h'_2(0))(1 + h_3(0)h_4(0)) \right. \\
&\quad + (1 + h_1(0)h_2(0))(h'_3(0)h_4(0) + h_3(0)h'_4(0)) \\
&\quad + e^{-l^R(0)}(h'_1(0)h_4(0) + h_1(0)h'_4(0)) + e^{-l^L(0)}(h'_2(0)h_3(0) + h_2(0)h'_3(0)) \\
&\quad - l^L(0)e^{-l^L(0)}h_2(0)h_3(0) - l^R(0)e^{-l^R(0)}h_1(0)h_4(0) \\
&\quad \left. - (l^L(0) + l^R(0))e^{-(l^L(0) + l^R(0))} \right) (**).
\end{aligned}$$

The terms in the first two lines (*) in the above expression are positive, whereas the remaining lines (**) are all negative (since $h'_j(0) < 0$). In fact, the total sum of the terms (*) is precisely equal to $(l^L(0) + l^R(0))\coth(\frac{1}{2}l_{\gamma_m}(0)) > l^L(0) + l^R(0)$. On the other hand, by the triangle inequality, the difference between $l^L(0) + l^R(0)$ and $l_{\gamma_m}(0)$ is less than $h_1(0) + h_2(0) + h_3(0) + h_4(0)$. Moreover, recall from Eq. (27) that

$$h_1(0) = \frac{e^{-d_1(0)} + e^{-d_2(0)}}{1 - e^{-l_{\gamma_m}(0)}}.$$

As $m \rightarrow \infty$, both $d_1(0)$ and $d_2(0)$ grow linearly with order $O(m\ell_\alpha(x))$ from the m Dehn-twists around α , and we see therefore that $h_1(0)$ behaves as $O(e^{-m\ell_\alpha(x)})$ as $m \rightarrow \infty$. This order of growth with respect to increasing m holds for all four $h_j(0)$ terms. One immediate consequence is that for m sufficiently large,

$$(29) \quad (l^L(0) + l^R(0))\coth(\frac{1}{2}l_{\gamma_m}(0)) + \frac{\epsilon}{2} > l^L(0) + l^R(0) + \frac{\epsilon}{2} > l_{\gamma_m}(0).$$

Let us now turn to the terms in (**): the linear growth of d_j with respect to m ensures that $h'_1(0)$, as given by Eq. (28), behaves as $O(me^{-m\ell_\alpha(x)})$. By symmetry, so too do all four $h'_j(0)$ terms. This, coupled with the fact that

$$\frac{e^{\frac{1}{2}(l^L(0) + l^R(0))}}{\sinh(\frac{1}{2}l_{\gamma_m}(0))} \rightarrow 1 \text{ as } m \rightarrow \infty,$$

means that the dominant term in the (**) summands is of order $O(me^{-m\ell_\alpha(x)})$, and this too shrinks to 0 as $m \rightarrow \infty$. Therefore, for sufficiently large m , $\frac{\epsilon}{2}$ suffices to cover the total negativity of the red summands. We have therefore shown that for sufficiently large m ,

$$v_{\Lambda_+}(\ell_{\gamma_m}) + \epsilon = \text{the sum of } (*) + \frac{\epsilon}{2} + \text{the sum of } (**) + \frac{\epsilon}{2} > l_{\gamma_m}(0) := \ell_{\gamma_m}(x),$$

as desired. \square

Remark 4.10. By symmetry, Lemma 4.9 tells us that there are no Thurston norm unit vectors $w \in \mathbf{S}_x \subsetneq T_x \mathcal{T}(S)$ such that of the following conditions are all satisfied

- (1) $w(\ell_\alpha) = v_{\Lambda_-}(\ell_\alpha) = \ell_\alpha(x)$, where $\ell_\alpha : \mathcal{T}(S) \rightarrow \mathbb{R}$ is the hyperbolic length function for α ;
- (2) for each $j = 1, 2, 3, 4$, $w(\ell_{\beta_j}) = v_{\Lambda_-}(\ell_{\beta_j}) = \ell_{\beta_j}(x)$;
- (3) $w(\tau_\alpha) < v_{\Lambda_-}(\tau_\alpha)$, where $\tau_\alpha : \mathcal{T}(S) \rightarrow \mathbb{R}$ is a Fenchel–Nielsen twist function for α .

In particular, this follows directly from Lemma 4.9 by taking α^{-1} as input instead of α .

Definition 4.11 (twist width). We refer to $(v_{\Lambda_+} - v_{\Lambda_-})(\tau_\alpha)$ as the *twist width* of α with respect to $\beta_1, \beta_2, \beta_3, \beta_4$.

Lemma 4.12. *The twist width of α with respect to $\beta_1, \beta_2, \beta_3, \beta_4$ is well-defined, i.e., it does not depend on the choice of twist coordinate τ_α or how one chooses to extend λ_\pm to Λ_\pm .*

Proof. Any two choices of Fenchel–Nielsen twist coordinates for α differ up to addition by some analytic function on Teichmüller space. The derivative of this function is cancelled out when taking the difference between v_{Λ_+} and v_{Λ_-} and hence renders twist width unaffected by the choice of the twist coordinate.

Next, we observe that Lemma 4.9 tells us that for any complete extension Λ_+ , setting

$$\begin{aligned} N &:= N_x([\alpha + \beta_1 + \beta_2 + \beta_3 + \beta_4]) \\ &= \{v \in \mathbf{S}_x \mid \iota_x([\alpha + \beta_1 + \beta_2 + \beta_3 + \beta_4])(v) = \|v\|_{\text{Th}}\}, \end{aligned}$$

we have

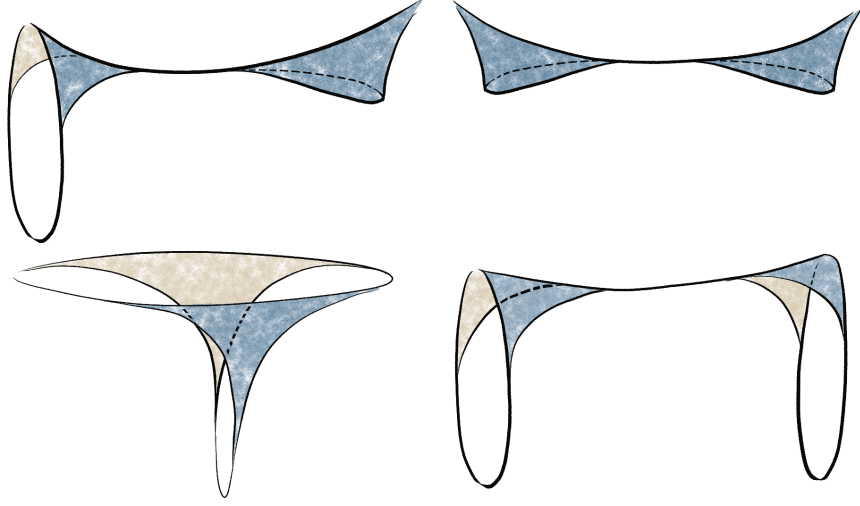
$$v_{\Lambda_+}(\tau_\alpha) = \max_{w \in N} w(\tau_\alpha),$$

and hence the value of $v_{\Lambda_+}(\tau_\alpha)$ is independent of the choice of extension from λ_+ to Λ_+ . Likewise, Remark 4.10 tells us that

$$v_{\Lambda_-}(\tau_\alpha) = \min_{w \in N} w(\tau_\alpha),$$

and hence is also independent. Therefore, the twist width $(v_{\Lambda_+} - v_{\Lambda_-})(\tau_\alpha)$ is independent of the extension of λ_\pm to Λ_\pm , and is well-defined. \square

Definition 4.13 (ϵ -slender pairs of pants). For a pair of pants P , denote the ϵ -neighbourhood of ∂P by $B(\partial P, \epsilon)$. For $\epsilon > 0$, we say that P is ϵ -slender if and only if $P \setminus B(\partial P, \epsilon)$ consists of two connected components, each of which is either a hypercycle-bounded punctured monogon or triangle (see Fig. 15). We further say that a sequence $\{P_i\}_{i \in \mathbb{N}}$ of pairs of pants is *asymptotically slender* if and only if for every $\epsilon > 0$, there is some index $I_\epsilon \in \mathbb{N}$ such that P_i is ϵ -slender for all $i > I_\epsilon$.

FIGURE 15. The four types of ϵ -slender pairs of pants.

Lemma 4.14. *Consider an asymptotically slender sequence of pairs of pants $\{P_i\}_{i \in \mathbb{N}}$ and denote their boundaries as a union of (possibly length 0) curves $\partial P_i = \alpha_i \cup \beta_i \cup \gamma_i$. Then,*

$$\lim_{i \rightarrow +\infty} \min \{ |\ell_{\alpha_i} - \ell_{\beta_i} - \ell_{\gamma_i}|, |\ell_{\alpha_i} - \ell_{\beta_i} + \ell_{\gamma_i}|, |\ell_{\alpha_i} + \ell_{\beta_i} - \ell_{\gamma_i}| \} = +\infty.$$

Proof. For each pair of pants P_i , let σ_{α_i} , σ_{β_i} and σ_{γ_i} respectively denote the (simple) orthogeodesics on P_i which respectively go between α_i and α_i , β_i and γ_i . By taking the endpoints of σ_{β_i} and σ_{γ_i} and consistently dragging them around α_i , β_i and γ_i in one of the two directions and straightening out the resulting paths to geodesics, we obtain (up to) eight complete laminations of P_i . Denote the shearing parameters for the respective bi-infinite geodesics obtained from σ_{α_i} , σ_{β_i} and σ_{γ_i} by s_{α_i} , s_{β_i} and s_{γ_i} . For either of the two choices of spiralling direction for α_i , there are four possibilities for spiralling directions for the β_i, γ_i pair. Invoking well known-results relating shearing parameters and lengths of closed geodesics (see e.g. [2, Section 2.3.1]), we see that there are four possible values for the shearing coordinate corresponding to $|s_{\alpha_i}|$ depending on spiralling, i.e.,

$$\frac{1}{2}|\ell_{\alpha_i} + \ell_{\beta_i} + \ell_{\gamma_i}|, \frac{1}{2}|\ell_{\alpha_i} - \ell_{\beta_i} - \ell_{\gamma_i}|, \frac{1}{2}|\ell_{\alpha_i} - \ell_{\beta_i} + \ell_{\gamma_i}|, \text{ and } \frac{1}{2}|\ell_{\alpha_i} + \ell_{\beta_i} - \ell_{\gamma_i}|.$$

We wish to show that the latter three absolute shearing coordinate values tend to infinity as $i \rightarrow +\infty$ (the first value blows up as a consequence).

Let λ_i be an arbitrary complete geodesic lamination obtained by taking σ_{α_i} , σ_{β_i} and σ_{γ_i} and dragging their ends around each of the boundaries of P_i in a consistent manner. For small $\epsilon > 0$, the two ideal triangles in $P_i \setminus \lambda_i$ approximately agree with the two components of $P_i \setminus B(\partial P_i, \epsilon)$, and hence the centre of mass for the two central stable triangles lie in distinct components

of $P_i \setminus B(\partial P_i, \epsilon)$. See Fig. 15 for a list of all possible configurations (up to permuting α_i, β_i and γ_i) for how these central stable triangles can be positioned relative to each other. We show that these central stable triangles may be set to be arbitrarily far apart as $i \rightarrow +\infty$ and hence obtain that $|s_{\alpha_i}| \rightarrow +\infty$. Roughly speaking, for any $L > 0$, take $\epsilon' > 0$ sufficiently small so that the distance between the central stable triangle and the ϵ' -thin part of each ideal triangle in $P_i \setminus \lambda_i$ is substantially larger than L . For any $\epsilon > 0$, the condition of ϵ -slenderness ensures that for $i > I_\epsilon$, the two components of $P \setminus B(\partial P, \epsilon)$ are separated by one or two orthogeodesics shorter than 2ϵ . Taking ϵ small enough separates the bulk of the ϵ' -thick portion of each ideal triangle in $P_i \setminus \lambda_i$ from each other, thereby ensuring that the distance between their central stable triangles is at least $2L$, and hence that $|s_{\alpha_i}| > L$. \square

Theorem 4.15 (shrinking twist widths). *Consider a sequence of collections of curves $\{\alpha_i, \beta_i^L, \gamma_i^L, \beta_i^R, \gamma_i^R\}_{i \in \mathbb{N}}$ such that each of the two sequences $\{\alpha_i, \beta_i^L, \gamma_i^L\}_{i \in \mathbb{N}}$ and $\{\alpha_i, \beta_i^R, \gamma_i^R\}_{i \in \mathbb{N}}$ constitutes the boundary geodesics for an asymptotically slender sequence of pairs of pants on (S, x) . Then the twist width of α_i with respect to $\beta_i^L, \gamma_i^L, \beta_i^R, \gamma_i^R$ tends to 0.*

Proof. Take two sequences of complete geodesic laminations $\{\Lambda_{+,i}\}_{i \in \mathbb{N}}$ and $\{\Lambda_{-,i}\}_{i \in \mathbb{N}}$ satisfying the conditions of the maximal twist lemma (Lemma 4.9). Our goal is to show that

$$\lim_{i \rightarrow \infty} (v_{\Lambda_{+,i}} - v_{\Lambda_{-,i}})(\tau_{\alpha_i}) = 0.$$

To begin with, we know from Lemma 3.5 that $(v_{\Lambda_{+,i}} - v_{\Lambda_{-,i}})(\tau_{\alpha_i})$ at $x \in \mathcal{T}(S)$ is a linear combination of derivatives (with respect to t) of

- **type one terms**

$$\log(1 - e^{-e^t \ell_{\gamma_S}(x)}) - e^t \log(1 - e^{-\ell_{\gamma_S}(x)}), \text{ and}$$

- **type two terms**

$$\begin{aligned} & e^t \log \left(1 + e^{-s_1(x)} + e^{-(s_1(x)+s_2(x))} + e^{-(s_1(x)+s_2(x)+s_3(x))} \right) \\ & - \log \left(1 + e^{-e^t s_1(x)} + e^{-e^t (s_1(x)+s_2(x))} + e^{-e^t (s_1(x)+s_2(x)+s_3(x))} \right). \end{aligned}$$

The condition that $\{\alpha_i, \beta_i^L, \gamma_i^L\}$ and $\{\alpha_i, \beta_i^R, \gamma_i^R\}$ are ϵ -slender for sufficiently high i tells us that $\ell_{\alpha_i} \rightarrow +\infty$. Otherwise, there is a subsequence of $\{\ell_{\alpha_i}\}$ which converges to something finite, and the only way that this is compatible with the definition of asymptotically slenderness is if $\ell_{\alpha_i} \rightarrow 0$. Due to the discreteness of the (simple) length spectrum of (S, x) , this can only happen if ℓ_{α_i} eventually equals 0. However, this then contradicts α_i being an interior geodesic rather than a boundary cusp.

Now employing the newly established fact that $\ell_{\alpha_i} \rightarrow +\infty$, we see that the derivatives of the first type, which take the form

$$(30) \quad \frac{\ell_{\alpha_i} e^{-\ell_{\alpha_i}}}{1 - e^{-\ell_{\alpha_i}}} - \log(1 - e^{-\ell_{\alpha_i}}),$$

necessarily tend to 0 as $i \rightarrow +\infty$.

We next consider terms of the second type. To begin with, we showed in the proof of Lemma 4.12 that we may use

- (1) any choice of Fenchel–Nielsen twist parameter τ_{α_i} for α_i , provided that we are consistent in setting this choice for computing both $v_{\Lambda_{+,i}}(\tau_{\alpha_i})$ and $v_{\Lambda_{-,i}}(\tau_{\alpha_i})$, and
- (2) any choice of extension of $\lambda_{\pm,i}$ to $\Lambda_{\pm,i}$, provided that $\Lambda_{\pm,i}$ is complete and contains $\lambda_{\pm,i}$ and $\beta_i^L, \gamma_i^L, \beta_i^R, \gamma_i^R$.

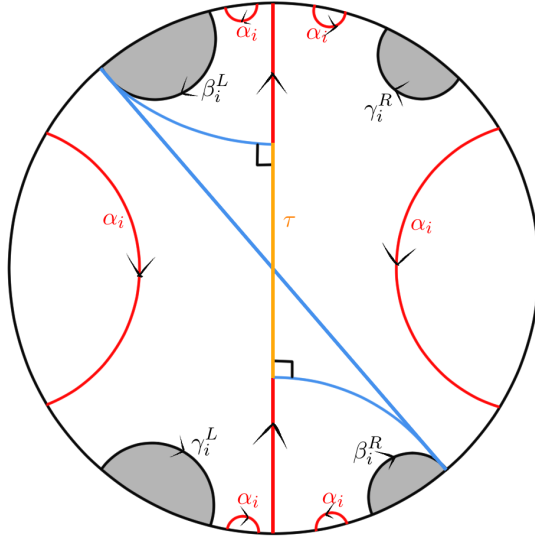


FIGURE 16. The orange segment depicted measures the Fenchel Nielsen twist parameter $\tau = \tau_{\alpha_i}$ for α_i .

For our purposes, we take a Fenchel–Nielsen twist parameter to go between the attractive fixed points of β_i^L and of β_i^R (see Fig. 16) when dealing with $\Lambda_{+,i}$ and to be a Fenchel–Nielsen twist parameter between the attractive fixed points of γ_i^L and of γ_i^R when dealing with $\Lambda_{-,i}$ (this is to ensure some symmetry which simplifies the analysis). We work with the “left” pair of pants bordered by $\alpha_i, \beta_i^L, \gamma_i^L$, the right pair of pants is similarly dealt with. Choose the extension of $\lambda_{\pm,i}$ to $\Lambda_{\pm,i}$ as in Fig. 17. Then,

$$\begin{aligned} s_{+,2}^L &= \ell_{\gamma_i^L}, & s_{+,4}^L &= \ell_{\beta_i^L}, & s_{+,1}^L &= s_{+,3}^L = \frac{1}{2}(\ell_{\alpha_i} - \ell_{\beta_i^L} - \ell_{\gamma_i^L}), \text{ and} \\ s_{-,2}^L &= -\ell_{\gamma_i^L}, & s_{-,4}^L &= -\ell_{\beta_i^L}, & s_{-,1}^L &= s_{-,3}^L = \frac{1}{2}(\ell_{\alpha_i} + \ell_{\beta_i^L} + \ell_{\gamma_i^L}), \end{aligned}$$

where $s_{\pm,k}$ denotes the shearing coordinate for the ideal geodesic shared by the k -th and the $(k+1)$ -th ideal triangle in the complement of $\Lambda_{\pm,i}$.

We first analyse the $\Lambda_{-,i}$ case. For $\Lambda_{-,i}$ (see Fig. 17), the shearing coordinates in question correspond to

$$\begin{aligned} s_{-,1}^L &= \frac{1}{2}(\ell_{\alpha_i} + \ell_{\beta_i^L} + \ell_{\gamma_i^L}), \\ s_{-,1}^L + s_{-,2}^L &= \frac{1}{2}(\ell_{\alpha_i} - \ell_{\beta_i^L} + \ell_{\gamma_i^L}), \\ s_{-,1}^L + s_{-,2}^L + s_{-,3}^L &= \ell_{\alpha_i} + \ell_{\gamma_i^L}. \end{aligned}$$

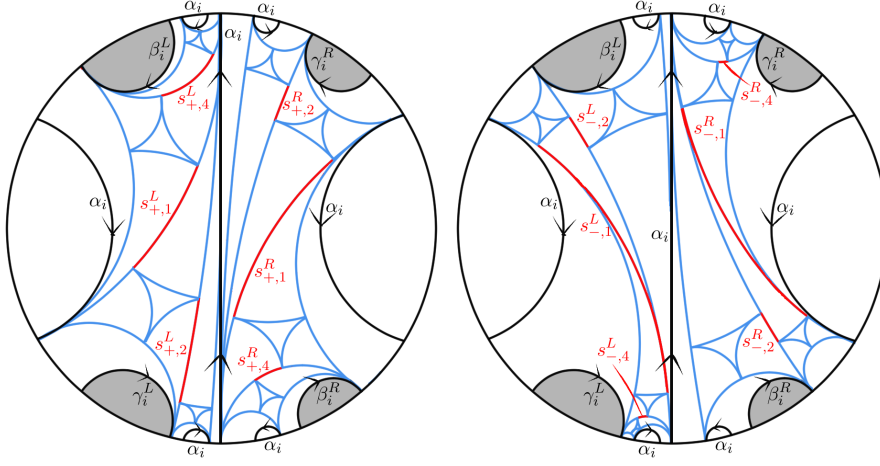


FIGURE 17. A depiction of the shearing parameters $s_{\pm,j}^L$ and $s_{\pm,j}^R$ for $\Lambda_{+,i}$ (left) and $\Lambda_{-,i}$ (right).

Both $s_{-,1}^L$ and $s_{-,1}^L + s_{-,2}^L + s_{-,3}^L$ are positive and tend to $+\infty$ (because $\ell_{\alpha_i} \rightarrow +\infty$). By Lemma 4.14, the absolute value of $s_{-,1}^L + s_{-,2}^L$ tends to ∞ as $i \rightarrow +\infty$. If $s_{-,1}^L + s_{-,2}^L$ is positive, then the derivative of the type two terms for $\Lambda_{-,i}$ takes the form

$$\begin{aligned} (31) \quad & \log \left(1 + e^{-s_{-,1}^L} + e^{-(s_{-,1}^L + s_{-,2}^L)} + e^{-(s_{-,1}^L + s_{-,2}^L + s_{-,3}^L)} \right) \\ & + \frac{s_{-,1}^L e^{-s_{-,1}^L} + (s_{-,1}^L + s_{-,2}^L) e^{-(s_{-,1}^L + s_{-,2}^L)} + (s_{-,1}^L + s_{-,2}^L + s_{-,3}^L) e^{-(s_{-,1}^L + s_{-,2}^L + s_{-,3}^L)}}{1 + e^{-s_{-,1}^L} + e^{-(s_{-,1}^L + s_{-,2}^L)} + e^{-(s_{-,1}^L + s_{-,2}^L + s_{-,3}^L)}}, \end{aligned}$$

and $s_{-,1}^L$, $s_{-,1}^L + s_{-,2}^L$ and $s_{-,1}^L + s_{-,2}^L + s_{-,3}^L$ blow up.

If $s_{-,1}^L + s_{-,2}^L$ is negative, then by dividing the inside of the logarithm of both terms by $e^{-(s_{-,1}^L + s_{-,2}^L)}$, the type two terms take the form

$$\begin{aligned} (32) \quad & e^t \log \left(1 + e^{-(s_{-,1}^L - (s_{-,1}^L + s_{-,2}^L))} + e^{(s_{-,1}^L + s_{-,2}^L)} + e^{-((s_{-,1}^L + s_{-,2}^L + s_{-,3}^L) - (s_{-,1}^L + s_{-,2}^L))} \right) \\ & - \log \left(1 + e^{-e^t (s_{-,1}^L - (s_{-,1}^L + s_{-,2}^L))} + e^{e^t (s_{-,1}^L + s_{-,2}^L)} + e^{-e^t ((s_{-,1}^L + s_{-,2}^L + s_{-,3}^L) - (s_{-,1}^L + s_{-,2}^L))} \right), \end{aligned}$$

and note that $(s_{-,1}^L - (s_{-,1}^L + s_{-,2}^L)), -(s_{-,1}^L + s_{-,2}^L), ((s_{-,1}^L + s_{-,2}^L + s_{-,3}^L) - (s_{-,1}^L + s_{-,2}^L))$ all blow up as the index increases. Thus, the derivative of the type two terms may always be written in the form

$$(33) \quad \log(1 + e^{-u} + e^{-v} + e^{-w}) + \frac{ue^{-u} + ve^{-v} + we^{-w}}{1 + e^{-u} + e^{-v} + e^{-w}},$$

for u, v, w which blow up as $i \rightarrow +\infty$, and hence this derivative term also tends to 0.

We finally consider the $\Lambda_{+,i}$ case, which is somewhat more delicate than the $\Lambda_{-,i}$ scenario. Our present goal is to show also in this setting that the derivative of the type two terms tends to 0. Fig. 17 tells us that

$$\begin{aligned} s_{+,1}^L &= \frac{1}{2}(\ell_{\alpha_i} - \ell_{\beta_i^L} - \ell_{\gamma_i^L}), \\ s_{+,1}^L + s_{+,2}^L &= \frac{1}{2}(\ell_{\alpha_i} - \ell_{\beta_i^L} + \ell_{\gamma_i^L}), \\ s_{+,1}^L + s_{+,2}^L + s_{+,3}^L &= \ell_{\alpha_i} - \ell_{\beta_i^L}. \end{aligned}$$

It is possible for any of these terms to be positive or negative, and so we invoke Lemma 4.14 to assert that precisely one of the following options holds for $i \gg 1$:

- I. $\ell_{\alpha_i} \gg \ell_{\beta_i^L} + \ell_{\gamma_i^L}$;
- II. $\ell_{\beta_i^L} \gg \ell_{\alpha_i} + \ell_{\gamma_i^L}$;
- III. $\ell_{\gamma_i^L} \gg \ell_{\alpha_i} + \ell_{\beta_i^L}$;
- IV. $\ell_{\alpha_i} \ll \ell_{\beta_i^L} + \ell_{\gamma_i^L}$, $\ell_{\beta_i^L} \ll \ell_{\alpha_i} + \ell_{\gamma_i^L}$, and $\ell_{\gamma_i^L} \ll \ell_{\alpha_i} + \ell_{\beta_i^L}$,

and to hence deal with this problem on a case-by-case basis.

For Case I, we leave the expression for the type two terms unchanged. Since all three sums $s_{+,1}^L$, $s_{+,1}^L + s_{+,2}^L$ and $s_{+,1}^L + s_{+,2}^L + s_{+,3}^L$ tend to $+\infty$ as $i \rightarrow +\infty$, which leads to Eq. (31), with the $--$ signs changed to $+$, tending to 0.

For Case II, by letting the inside of the logarithm of both terms be multiplied by $e^{\ell_{\alpha_i} - \ell_{\beta_i^L}}$, the type two terms are re-expressed as:

$$\begin{aligned} &e^t \log \left(1 + e^{-\frac{1}{2}(\ell_{\beta_i^L} - \ell_{\alpha_i} - \ell_{\gamma_i^L})} + e^{-\frac{1}{2}(\ell_{\beta_i^L} - \ell_{\alpha_i} - \ell_{\gamma_i^L}) - \ell_{\gamma_i^L}} + e^{-(\ell_{\beta_i^L} - \ell_{\alpha_i} - \ell_{\gamma_i^L}) - \ell_{\gamma_i^L}} \right) \\ &- \log \left(1 + e^{-\frac{e^t}{2}(\ell_{\beta_i^L} - \ell_{\alpha_i} - \ell_{\gamma_i^L})} + e^{-\frac{e^t}{2}(\ell_{\beta_i^L} - \ell_{\alpha_i} - \ell_{\gamma_i^L}) - e^t \ell_{\gamma_i^L}} + e^{-e^t(\ell_{\beta_i^L} - \ell_{\alpha_i} - \ell_{\gamma_i^L}) - e^t \ell_{\gamma_i^L}} \right). \end{aligned}$$

As before, this is expressed in the form of Eq. (33) with $u, v, w \rightarrow +\infty$ as $i \rightarrow +\infty$, and hence the $t = 0$ derivative of this term tends to 0 as $i \rightarrow +\infty$.

For Case III, we likewise re-express the type two terms as:

$$e^t \log \left(1 + e^{-\frac{1}{2}(\ell_{\gamma_i^L} - \ell_{\alpha_i} - \ell_{\beta_i^L}) - \ell_{\beta_i^L}} + e^{-\frac{1}{2}(\ell_{\gamma_i^L} - \ell_{\alpha_i} - \ell_{\beta_i^L}) - (\ell_{\alpha_i} + \ell_{\beta_i^L})} + e^{-(\ell_{\gamma_i^L} - \ell_{\alpha_i} - \ell_{\beta_i^L}) - (\ell_{\alpha_i} + \ell_{\beta_i^L})} \right) \\ - \log \left(1 + e^{-\frac{e^t}{2}(\ell_{\gamma_i^L} - \ell_{\alpha_i} - \ell_{\beta_i^L}) - e^t \ell_{\beta_i^L}} + e^{-\frac{e^t}{2}(\ell_{\gamma_i^L} - \ell_{\alpha_i} - \ell_{\beta_i^L}) - e^t(\ell_{\alpha_i} + \ell_{\beta_i^L})} + e^{-e^t(\ell_{\gamma_i^L} - \ell_{\alpha_i} - \ell_{\beta_i^L}) - e^t(\ell_{\alpha_i} + \ell_{\beta_i^L})} \right).$$

This again suffices to ensure that the $t = 0$ derivative of this term tends to 0 as $i \rightarrow +\infty$.

Finally, for Case IV, we re-express the type two terms as:

$$e^t \log \left(1 + e^{-\frac{1}{2}(\ell_{\beta_i^L} + \ell_{\gamma_i^L} - \ell_{\alpha_i})} + e^{-\frac{1}{2}((\ell_{\alpha_i} + \ell_{\gamma_i^L} - \ell_{\beta_i^L}) - (\ell_{\beta_i^L} + \ell_{\gamma_i^L} - \ell_{\alpha_i}))} + e^{-\frac{1}{2}(\ell_{\alpha_i} + \ell_{\gamma_i^L} - \ell_{\beta_i^L})} \right) \\ - \log \left(1 + e^{-\frac{e^t}{2}(\ell_{\beta_i^L} + \ell_{\gamma_i^L} - \ell_{\alpha_i})} + e^{-\frac{e^t}{2}((\ell_{\alpha_i} + \ell_{\gamma_i^L} - \ell_{\beta_i^L}) - (\ell_{\beta_i^L} + \ell_{\gamma_i^L} - \ell_{\alpha_i}))} + e^{-\frac{e^t}{2}(\ell_{\alpha_i} + \ell_{\gamma_i^L} - \ell_{\beta_i^L})} \right).$$

And this again suffices to ensure that the $t = 0$ derivative of this term tends to 0 as $i \rightarrow +\infty$.

We therefore see that all the type one and type two terms associated with the left pair of pants tend to 0. By symmetry, this is true also for the terms associated with the right pair of pants. Therefore, the twist width of α_i with respect to $\beta_i^L, \gamma_i^L, \beta_i^R, \gamma_i^R$ tends to 0. \square

5. THE GEOMETRY OF CONVEX BODIES

Our basic setup for this section is the following. Consider a compact convex subset $D \subsetneq \mathbb{V}$ of an m -dimensional vector space \mathbb{V} . Since D is compact and convex, it is homeomorphic to a closed ball of dimension less than or equal to m . In particular, D is a closed m -ball if and only if it has non-empty interior in \mathbb{V} .

In this section, we shall introduce several geometric/combinatorial structures associated with points on the boundary of the convex set D . These include notions such as *dimension*, *face-dimension*, *adherence*, *adherence-dimension*, *codimension* and a few others. These concepts apply to all compact convex bodies in finite dimensional vector spaces, and may be of general interest. We later apply them to points in $\mathbf{S}_x^* := \iota_x(\mathcal{PML}(S)) \subsetneq T_x^*\mathcal{T}(S)$ and its dual unit sphere in the tangent space, thereby laying the groundwork for establishing topological rigidity, Theorem 1.8. At the same time, we shall recall some classical notions from convexity theory. We start by recalling the important notion of support hyperplane for the convex set D .

Definition 5.1 (support hyperplane). For a point P in the boundary ∂D of a convex subset $D \subsetneq \mathbb{V}$, an affine hyperplane $\pi \subset \mathbb{V}$ is a *support hyperplane* to ∂D at P if the following hold:

- π contains P and
- one of the two closed half-spaces bounded by π contains all of D .

5.1. Dimension and faces.

Definition 5.2 (dimension of a point). For any point P on ∂D , let \mathbb{A}_P denote a maximal-dimensional affine subspace of \mathbb{V} such that P is contained in the interior of the set $\mathbb{A}_P \cap \partial D$, where the latter is regarded as a subset of \mathbb{A}_P (endowed with the standard topology on affine subspaces). We define the *dimension of P (with respect to ∂D)* as the dimension of \mathbb{A}_P :

$$\dim_{\partial D}(P) := \dim \mathbb{A}_P.$$

Remark 5.3 (Example). Consider the convex body $D = [0, 2] \times [0, 2] \subset \mathbb{R}^2$. We first note (somewhat trivially) that the dimension of any point on ∂D must be strictly less than 2 as they are not in the interior of D . For the point $P = (0, 1)$, the y -axis $\{x = 0\}$ is a 1-dimensional (affine) subspace such that the intersection $\{x = 0\} \cap \partial D$ is a closed interval containing P as an interior point and hence $\dim_{\partial D}(P) = 1$. For $Q = (0, 0)$, the intersection $\mathbb{L} \cap \partial D$ of every line \mathbb{L} through Q , besides the x and y axes, with ∂D is simply equal to $\{Q\}$, which is a nowhere-dense subset of \mathbb{L} . The two exceptional lines which pass through Q are the x and y axes, and their intersections with ∂D is a closed interval having Q as a boundary (and non-interior) point. Therefore, Q has dimension strictly less than 1, i.e., $\dim_{\partial D}(Q) = 0$.

Lemma 5.4 (alternative formulation of dimension). *Definition 5.2 may be alternatively, and equivalently, formulated by replacing $\mathbb{A}_P \cap \partial D$ with $\mathbb{A}_P \cap D$ in the statement of this definition.*

Proof. Let \mathbb{A} denote a maximal-dimensional affine subspace of \mathbb{V} such that P is contained in the interior of the set $\mathbb{A} \cap \partial D$, where the latter is regarded as a subset of \mathbb{A} . We first note that, by definition, there is an open ball B in \mathbb{A} so that

$$P \in B \subset \mathbb{A} \cap \partial D \subseteq \mathbb{A} \cap D,$$

and P is an interior point of $\mathbb{A} \cap D$ (regarded as a subset \mathbb{A}). Hence the notion of dimension defined as in Definition 5.2 is less than or equal to the alternative proposed here.

Conversely, consider any maximal-dimensional affine subspace \mathbb{A}' containing P as an interior point of $\mathbb{A}' \cap D$, with $\mathbb{A}' \cap D$ topologised as a subset of \mathbb{A}' . This means that there is some open ball $B' \subset \mathbb{A}' \cap D$ in \mathbb{A}' containing P . The existence of a supporting hyperplane to D at $P \in \partial D$ asserts that there is some linear function $w^* : \mathbb{V} \rightarrow \mathbb{R}$ such that $P \in \ker w^*$ and the interior of D is a subset of the preimage $(w^*)^{-1}((0, +\infty))$. This in turn implies that either $\mathbb{A}' \cap D = \mathbb{A}' \cap \partial D$, or $\mathbb{A}' \cap D$ contains an interior point $Q \in D \setminus \partial D$ of D , and hence Q is not an element of the closed affine half-space $(w^*)^{-1}((-\infty, 0])$. This means that P lies in the affine half-space $(w^*)^{-1}((-\infty, 0]) \cap \mathbb{A}'$ but Q does not, thereby indicating that P is a boundary point of $\mathbb{A}' \cap D$. This contradicts the original assumption that P is an interior point of $\mathbb{A}' \cap D$, and hence excludes the latter possibility. Therefore,

we have $P \in B' \subset \mathbb{A}' \cap D = \mathbb{A}' \cap \partial D$, and we see that the two notions of dimension are equal. \square

Definition 5.5 (face for a point). Given a point $P \in \partial D$, we refer to the intersection $\mathbb{A}_P \cap \partial D$ of any affine subspace \mathbb{A}_P (constructed as in Definition 5.2) with ∂D as a *face* (of ∂D) for $P \in \partial D$.

- The dimension of the face $\mathbb{A}_P \cap \partial D$ is defined to be the dimension of \mathbb{A}_P .
- We refer to the collection of faces $\mathbb{A}_P \cap \partial D$ for $P \in \partial D$ as the collection of faces on ∂D .
- We refer to any face that is a subset of another face as a *subface*.
- For each face F of ∂D , we call the subset of points in F which have an open-ball neighbourhood (with respect to the topology induced on the face) contained in F , the *interior* of the face.

We next show that the above notion of a face is equivalent to the one found in convex geometry literature (see, e.g., [27, p. 162]):

Definition 5.6 (face). A convex subset $F \subset D$ is called a *face* of D if and only if for every $x \in F$ and every $y, z \in D$ such that x lies on the open interval between y and z , we have $y, z \in F$.

We shall show that the two notions are equivalent if we ignore the special case where $F = D$, which is permitted under Definition 5.6 but not Definition 5.5.

Proof that the two notions of faces agree. Consider a face $F \subset \partial D$ in the sense of Definition 5.5, and let $P \in F$ be an interior point of F . We assume (for proof by contradiction) that F is not a face in the classical sense (Definition 5.6) and so there is some $x \in F$ and $y, z \in D$ such that x lies on the open interval between y and z , but $y, z \notin F$. Note that it is impossible to have $y \in F$ and $z \notin F$ (or $z \in F$ and $y \notin F$) because the fact that x lies on the line joining y and z would then mean that z lies on a line generated by two elements of F and hence would lie on the affine space \mathbb{A} generated by F and hence would lie in $F = \mathbb{A} \cap D$.

Suppose that neither y nor z is contained in F . Then the convex hull of $F \cup \{y, z\}$ is strictly larger than F . Let \mathbb{A}' be the affine space generated by (the convex hull of) $F \cup \{y, z\}$. Then, $\mathbb{A}' \supsetneq \mathbb{A}$ and hence \mathbb{A}' has dimension strictly greater than \mathbb{A} by 1 (as the line joining y and z transverse to F is only 1-dimensional). This in turn means that the convex hull of $F \cup \{y, z\}$ is a closed ball of one dimension higher than F , and hence $\mathbb{A}' \cap D$, which necessarily contains the convex hull of $F \cup \{y, z\}$ must have the same dimension as \mathbb{A}' . In particular, P must be an interior point of $\mathbb{A}' \cap D$ (or it would be a boundary point of F , thereby contradicting Definition 5.5). This contradicts the maximality of \mathbb{A} , thereby showing that our assumption for contradiction is false, and that F must be a face in the classical sense.

Conversely, given a face $F \subset D$ in the classical sense (Definition 5.6) which is not equal to D itself, let \mathbb{A} denote the affine subspace generated by F . We first note that F is necessarily equal to $\mathbb{A} \cap D$; for otherwise, we would be able to produce an interval with one end point based at an interior point of F and the other point in $\mathbb{A} \cap D \setminus F$, thereby contradicting Definition 5.6. Since F is (by definition) convex and closed (a consequence of the definition), it is a closed ball of the same dimension as \mathbb{A} . Let $P \in F$ be an interior point of F , regarded as a subset of \mathbb{A} . We need to show that there is no affine space \mathbb{A}' strictly larger than \mathbb{A} such that $\mathbb{A}' \cap D$ contains P as an interior point of $\mathbb{A}' \cap D$, the latter regarded as a subset of \mathbb{A}' . If there is such an affine space \mathbb{A}' , then (using the fact that P is such an interior point) we can find a small interval in $\mathbb{A}' \cap D$ which has non-empty transverse intersection with $F = \mathbb{A} \cap D$. This then contradicts the classical definition Definition 5.6 of a face, and hence \mathbb{A} is indeed maximal. Therefore, F is a face for P , in the sense of Definition 5.5. \square

Remark 5.7. It is clear from either Definition 5.5 or Definition 5.6 that any non-trivial intersection of a collection of faces is itself a face.

Definition 5.8 (exposed face (see, e.g., [27, p. 162])). A set which is the intersection of D with one of its support hyperplanes (recalled in Definition 5.1) is called an *exposed face*.

Remark 5.9. When D is a polytope, the notion of a face is equal to that of an exposed face. For general convex bodies, this is false: exposed faces are always faces (see, e.g., [27, p. 162]), but the converse statement is false. Consider the following example: take D to be the ϵ -neighbourhood of any polygon in \mathbb{R}^2 . Then, ∂D is made up of straight edges joined to circular arcs, and a point at which one joins a straight edge and a circular arc is a face, but not an exposed face.

We now show that every $P \in \partial D$ has a unique associated face for P . This will form the basis for the next notion of dimension which we wish to introduce.

Lemma 5.10 (face uniqueness). *For every $P \in \partial D$, there is a unique face for P in the sense of Definition 5.5.*

Proof. Let \mathbb{A}_P and \mathbb{A}'_P denote affine subspaces of \mathbb{V} which satisfy the conditions stated in the alternative formulation of Definition 5.2 as given by Lemma 5.4. Let \mathbb{A}''_P be the convex hull of $\mathbb{A}_P \cup \mathbb{A}'_P$. This is an affine subspace, and we now show that \mathbb{A}''_P also satisfies the conditions for the alternative formulation of dimension. We know by assumption that there exist sets $B \subset \mathbb{A}_P \cap D$ and $B' \subset \mathbb{A}'_P \cap D$ that contain P and are open subsets respectively in \mathbb{A}_P and \mathbb{A}'_P . In particular, we may choose B and B' to be convex (balls). This means that the convex hull of $B \cup B'$ is a convex open ball in \mathbb{A}''_P containing P . Furthermore, since D is convex, it contains necessarily the convex hull of $B \cup B'$, and hence this convex hull is a subset

of $\mathbb{A}_P'' \cap D$, which is open in \mathbb{A}_P'' . Since $P \in B \cup B'$ is contained in this open convex hull, it is an interior point. The maximality of the dimensions of \mathbb{A}_P and \mathbb{A}_P' then implies that $\mathbb{A}_P = \mathbb{A}_P'' = \mathbb{A}_P'$, therefore yielding the uniqueness of \mathbb{A}_P , and hence of the face $\mathbb{A}_P \cap \partial D = \mathbb{A}_P \cap D$. \square

Definition 5.11 (convex stratification). Lemma 5.10 canonically partitions ∂D into the interiors of faces. We refer to this decomposition into strata (i.e., the interiors of faces) as the *convex stratification*. Note that the strata are convex and hence are open cells.

Remark 5.12. The proof of Lemma 5.10 shows that not only is the face for P unique, but the affine subspace \mathbb{A}_P containing the face for P is also unique. In particular, the fact that P is an interior point of $\mathbb{A}_P \cap \partial D$ with respect to the topology on \mathbb{A}_P implies that \mathbb{A}_P is the affine subspace generated by the face for P .

Proposition 5.13. *For any convex set $C \subset \partial D$, there exists a unique smallest face of ∂D containing C .*

Proof. Since C is a convex set, its interior, as a subset of ∂D , is a convex k -dimensional open ball $\text{int}(C)$. By Lemma 5.10, for each point $P \in \text{int}(C)$, there is a unique associated face obtained as an intersection of the form $\mathbb{A}_P \cap \partial D$. For a point Q in $\text{int}(C)$, by the convexity of $\text{int}(C)$, there is a segment on ∂D with endpoint at Q containing P in its interior. Therefore by Definition 5.6, Q must be contained in the face of P . Therefore, there is a unique face F_C which is the face of every interior point $P \in \text{int}(C)$. Since faces are closed, the face F_C must contain all of C . In fact, F_C must be the smallest face of ∂D that contains C as a subset because C contains interior points of F_C , whereas subfaces of F_C are necessarily on the boundary (see Definition 5.6.) \square

Remark 5.14. It is possible to prove Proposition 5.13 without invoking Lemma 5.10. First use the hyperplane separation theorem to show that there must be some exposed face containing C , and hence show that the set of faces containing C is a non-empty set. Then take the intersection of all faces containing C , and show (definitionally) that this intersection is a face. The advantage of the method provided is that we see that C must contain interior points of F_C .

5.2. Face-dimension. The goal of this subsection is to use the unique association of faces to points established in Lemma 5.10 to introduce a new notion of dimension for points in ∂D which we call face-dimension (Definition 5.18).

Definition 5.15 (adherence). We say that a face F of ∂D is *adherent* to a face $F' \supseteq F$ if for any face F'' containing F there is a face \hat{F} which contains both F' and F'' as subfaces. We also introduce the following related notions:

- Each face F is adherent to a unique maximal face (possibly F itself, this is justified below). We refer to this maximal face as the

adherence closure of F . We say that F is *adherence-closed* if its adherence closure is the face F itself. Note that adherence closures are necessarily adherence-closed.

- For an adherence-closed face F , the union of its interior and the interiors of all subfaces whose adherence closures coincide with F is called the *adherence core* of F .
- An adherence-closed face F is said to be *adherence complete* if it is the adherence closure of each of its subfaces. See Fig. 18 for an example.

Proposition 5.16 (adherence closure uniqueness). *Every face F has a unique adherence closure.*

Proof. Let F_1 and F_2 both be adherence closures to F . The fact that F is adherent to F_1 and F_2 contains F means that there's some face in ∂D containing both F_1 and F_2 . Let \hat{F} denote the intersection of every face that contains $F_1 \cup F_2$, and note that \hat{F} is a face. We show that F is adherent closed to \hat{F} . Let F' be an arbitrary face containing F , then there's a face F'' that contains both F' and F_1 . Since F is adherent to F_2 , there is a face F''' that contains both F_2 and F'' and hence $F_1 \cup F_2 \cup F' \subset F'''$. By construction, this means that $\hat{F} \subset F'''$ and $F' \subset F'''$. Having shown that F is adherent to \hat{F} , the fact that F_1 and F_2 are maximal implies that $F_1 = \hat{F} = F_2$. \square

Proposition 5.17. *The adherence closure \hat{F} of a face F is adherence-closed.*

Proof. Let \bar{F} denote the adherence closure of \hat{F} . Consider an arbitrary face F' containing F , and let F'' denote a face containing both F' and \hat{F} . Since \hat{F} is adherent to \bar{F} , this means that there is a face F''' containing both F'' and \bar{F} and hence F''' contains both F' and \bar{F} . This shows that F is adherent to \bar{F} , and by the maximality of \hat{F} , we see that $\hat{F} = \bar{F}$, as desired. \square

Definition 5.18 (face-dimension for points). For a point P on ∂D , let F denote the face for P . We refer to the dimension of the adherence closure of F as the *face-dimension* of P , and denote it by $\text{Fdim}(P)$.

The following is immediately obtained from the definition.

Proposition 5.19 (dimension vs. face-dimension). *For any point $P \in \partial D$,*

$$\text{Fdim}(P) \geq \dim(P),$$

with equality if and only if the face for P is its own adherence closure.

Remark 5.20. It is possible for face-dimension to be strictly greater than dimension. In the left figure in Fig. 18, each of the two vertices is contained in a unique edge, and hence each of them is adherent to the edge, and has face-dimension 1. In the right figure, each vertex has more than one edge containing it, and is not adherent to either edge. Each of these four vertices therefore has face-dimension 0.

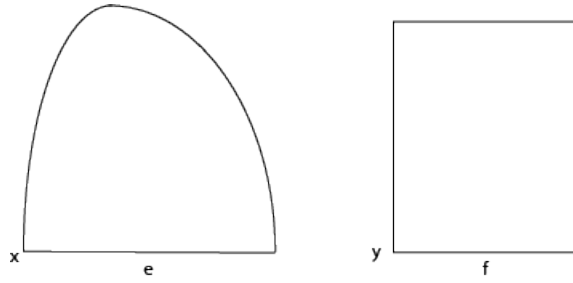


FIGURE 18. Two different examples of 2-dimensional convex sets, with their 1-dimensional boundaries. The adherence closure of x is e (left figure), whereas the vertex y is its own adherence closure (right figure).

5.3. Adherence-dimension.

Definition 5.21 (adherence-dimension for faces). We call a chain of faces

$$F_1 \subsetneq F_2 \subsetneq \dots \subsetneq F_{k-1} \subsetneq F_k$$

an F -dim ascending chain if the face-dimensions of the $\{F_i\}_{i=1,\dots,k}$ are all distinct, i.e. they are strictly increasing. We also introduce the following related notions:

- We call an F -dim ascending chain $\{F_i\}_{i=1,\dots,k}$ *maximal* if there is no other F -dim ascending chain containing $\{F_i\}_{i=1,\dots,k}$ which ends with F_k .
- We define the *adherence height* of F to be the minimum of the lengths of maximal F -ascending chains *ending* with F .
- We define the *adherence depth* of F to be the minimum of the lengths of maximal F -ascending chains *starting* with F .
- We define the *adherence-dimension* $\text{Adim}(F)$ of F to be the sum of its adherence height and its adherence depth subtracted by 2.

Remark 5.22. The underlying concept for adherence height, depth and dimension is motivated by the notion of Krull dimension, and there is a certain level of flexibility in how they might be defined — for example: one can choose to take the *maximum* rather than the minimum length of maximal F -ascending chains starting/ending with F . Our choice is because

- we need to take the minimum rather than maximum length for adherence height as this is needed for the proof of Theorem 7.4,
- the previous choice leads us to also take the minimum rather than the maximum length for defining adherence depth, to ensure that the adherence-dimension of a face F is necessarily greater than or equal to the adherence-dimensions of the subfaces of F .

Definition 5.23 (Adherence-dimension for points). We define the *adherence-dimension* for a point $P \in \partial D$ as the adherence-dimension of the face F for P . In particular, we denote this by $\text{Adim}_{\partial D}(P) := \text{Adim}(F)$.

5.4. Duality and codimension. We have hitherto worked with an arbitrary convex ball $D \subset \mathbb{V}$, but now consider the setting where $\vec{0} \in \mathbb{V}$ is an interior point of D .

Definition 5.24 (dual ball). Define the convex *dual ball* $D^* \subset \mathbb{V}^*$ in the dual space \mathbb{V}^* by

$$\begin{aligned} D^* &:= \{w^* \in \mathbb{V}^* \mid \forall v \in D, w^*(v) \leq 1\} \\ &= \left\{ w^* \in \mathbb{V}^* \mid \sup_{v \in D} w^*(v) \leq 1 \right\} \\ &= \left\{ w^* \in \mathbb{V}^* \mid \sup_{v \in \partial D} w^*(v) \leq 1 \right\}. \end{aligned}$$

We refer to ∂D^* as the dual sphere to the sphere ∂D , and note that positive homothety ensures that

$$\partial D^* = \left\{ w^* \in \mathbb{V}^* \mid \sup_{v \in \partial D} w^*(v) = 1 \right\}.$$

The classical notion of support hyperplane (Definition 5.1) establishes a geometric relation between the two dual pictures:

Lemma 5.25. *Every support hyperplane at P contains the face for P .*

Proof. Let \mathbb{A}_P be the subspace passing through P used to define the face for P , and \mathbb{A}_P^0 its translate under the translation sending P to $\vec{0}$, i.e., $\mathbb{A}_P^0 := \mathbb{A}_P - P$. We claim that any arbitrary support hyperplane π at P necessarily contains \mathbb{A}_P . We also consider the translate $\pi^0 := \pi - P$ of π . We have only to show that π^0 contains \mathbb{A}_P^0 . If not, then there is some vector $v \in \mathbb{A}_P^0 \setminus \pi^0$, and likewise all of its non-zero multiples are in \mathbb{A}_P^0 but not in π^0 . In particular, we may replace v with a small enough scalar multiple so that both $\pm v$ are contained in an open ball around $\vec{0}$ in $F - P = (\mathbb{A}_P \cap D) - P$. However, this contradicts the assumption of π being a support hyperplane, for $P + v$ and $P - v$ lie on different sides of π . Therefore, π contains \mathbb{A}_P , hence the face $\mathbb{A}_P \cap D$ of P . \square

We define the codimension for a point P as the dimension of the space of support hyperplanes at P . To make this precise, we parametrise support hyperplanes via normal vectors:

Definition 5.26 (normal and positive normal vectors). We say that a non-zero vector $w^* \in \mathbb{V}^*$ is *normal* to an affine hyperplane $\pi \subset \mathbb{V}$ if for any two points $P, Q \in \pi$,

$$w^*(P - Q) = 0.$$

When π is a support hyperplane of a convex subset $D \subset \mathbb{V}$, we say that a dual vector $w^* \in \mathbb{V}^*$ is a *positive normal vector* to π if w^* is normal to π and points to the side of π disjoint from D .

Remark 5.27. Elementary linear algebra tells us that every support hyperplane π through $P \in \partial D$ takes the form

$$\pi = \{v \in \mathbb{V} \mid w^*(v) = c\}, \quad c = w^*(P),$$

for some positive normal vector $w^* \in \mathbb{V}^*$. In particular, any two positive normal vectors to the same hyperplane are positive scalar multiples of each other and may be uniquely normalised so that $c = w^*(P) = 1$.

Definition 5.28 (codimension). For any point $P \in \partial D$, we define:

$$N_P := \left\{ w^* \in \mathbb{V}^* \left| \begin{array}{l} w^* \text{ is positive normal to some} \\ \text{support hyperplane to } \partial D \text{ at } P \\ \text{and } w^*(P) = 1 \end{array} \right. \right\}.$$

We refer to the dimension of N_P , as a subset of the $(m-1)$ -sphere ∂D , as the *codimension* of P , and denote it by $\text{codim}_{\partial D}(P)$.

Remark 5.29. The above definition essentially defines the codimension of P as the dimension of the space of support hyperplanes for P , only phrased in terms of unit normal vectors to these hyperplanes.

Corollary 5.30. *For any point P on the boundary ∂D of a convex full-dimensional subset $D \subset \mathbb{V}$, we have*

$$\dim_{\partial D}(P) + \text{codim}_{\partial D}(P) \leq \dim \mathbb{V} - 1.$$

Proof. We again consider to translate the picture by $-P$. Lemma 5.25 may be interpreted as saying that any positive normal vector w^* contains $\mathbb{A}_P^0 := \mathbb{A}_P - P$ in its kernel. The rank-nullity theorem tells us that the space of arbitrary (i.e., not just positive normal) dual vectors which take the value 0 on \mathbb{A}_P^0 is itself of dimension $\dim \mathbb{V} - \dim \mathbb{A}_P^0 = \dim \mathbb{V} - \dim \mathbb{A}_P$, and imposing the normalisation condition for positive normal vectors tells us that N_P has dimension at most $\dim \mathbb{V} - \dim \mathbb{A}_P - 1$, and hence

$$\text{codim}_{\partial D}(P) = \dim N_P \leq \dim \mathbb{V} - \dim \mathbb{A}_P - 1 = \dim \mathbb{V} - \dim_{\partial D}(P) - 1.$$

□

We conclude this subsection by establishing several properties of the sets N_P in Definition 5.28. We first show that inclusion for the N_P satisfies the following contravariance property:

Lemma 5.31. *If F is the face for $P \in \partial D$, i.e., if P is an interior point of F , then for any $Q \in F$, we have*

$$N_P \subseteq N_Q \subset \partial D^*.$$

In particular, if P and Q are interior points to the same face $F \subset \partial D$, then

$$N_P = N_Q.$$

Proof. If a point $P \in \partial D$ is contained in the interior of a face F , then any support hyperplane to ∂D at P is also a support hyperplane to any other point of the face (e.g. $Q \in F$). The claims of the lemma immediately follow from this fact. \square

We further see that each N_P is actually an exposed face in the dual sphere for ∂D .

Theorem 5.32. *For every $P \in \partial D$, the set $N_P \subset \mathbb{V}^*$ is an exposed face (see Definition 5.8) of ∂D^* . In particular, this means that N_P is a (non-empty) convex and compact subset of D^* .*

Proof. We first claim that

$$N'_P := \left\{ w^* \in \mathbb{V}^* \mid \sup_{v \in \partial D} w^*(v) = w^*(P) = 1 \right\}$$

is equal to N_P . To see this, we first note that any $w'^* \in N'_P$ defines an affine hyperplane of the form

$$\{v \in \mathbb{V} \mid w'^*(v) = w'^*(P) = 1\}.$$

The equality $\sup_{v \in \partial D} w'^*(v) = w'^*(P)$ means that ∂D (and hence D) lies on one side of $\{v \in \mathbb{V} \mid w'^*(v) = w'^*(P) = 1\}$. Therefore, if $w'^* \in N'_P$, then it is a positive normal vector to a support hyperplane of D at P , and hence $N'_P \subseteq N_P$.

Conversely, given a positive normal vector $w^* \in N_P$, the fact that ∂D lies on one side of the support hyperplane normal to w^* tells us that

$$w^*(P) \geq \sup_{v \in \partial D} w^*(v) \geq w^*(P), \text{ and hence } \sup_{v \in \partial D} w^*(v) = w^*(P).$$

We also know, from the definition of N_P , that $w^*(P) = 1$, and hence $N_P \subseteq N'_P$, thereby telling us that $N_P = N'_P$.

We next observe that N_P is the intersection of the following hyperplane (in \mathbb{V}^*)

$$\{w^* \in \mathbb{V}^* \mid w^*(P) = 1\}$$

and the sphere

$$\partial D^* = \left\{ w^* \in \mathbb{V}^* \mid \sup_{v \in \partial D} w^*(v) = 1 \right\},$$

and hence N_P is a exposed face of ∂D^* in \mathbb{V}^* . In particular, this is an intersection of two closed and convex sets, where one is bounded, thus their intersection is convex and compact. \square

Remark 5.33. A consequence of the proof of Theorem 5.32 is that each N_P is expressed as

$$N_P = \left\{ w^* \in \mathbb{V}^* \mid \sup_{v \in \partial D} w^*(v) = w^*(P) = 1 \right\}.$$

This has the interpretation that N_P consists of all “unit” dual vectors w^* that attain their maximal value (on ∂D) at P .

5.5. Linear invariants. The goal of this final subsection is to show that the concepts we have defined in this section are invariant under invertible linear maps. To begin with, we point out that the image of a convex set $D \subset \mathbb{V}$ under a linear map $f: \mathbb{V} \rightarrow \mathbb{W}$ is necessarily convex. Moreover, if f is invertible, then the boundary of D maps to the boundary of $f(D)$:

$$f(\partial D) = \partial f(D) \subset \mathbb{W}.$$

Theorem 5.34 (linear invariants). *Consider an invertible linear map $f: \mathbb{V} \rightarrow \mathbb{W}$. For every point $P \in \partial D$,*

- (1) *f maps the face for P on ∂D to the face for $f(P)$ on $\partial f(D)$, and hence takes the convex stratification structure on ∂D to the convex stratification on $\partial f(D)$;*
- (2) *the dimension of P is a linear invariant:*

$$\dim_{\partial D}(P) = \dim_{\partial f(D)}(f(P));$$

- (3) *the relation of adherence is preserved under f ;*
- (4) *the face-dimension of P is a linear invariant:*

$$\text{Fdim}_{\partial D}(P) = \text{Fdim}_{\partial f(D)}(f(P));$$

- (5) *adherence height and adherence depth are linear invariants;*
- (6) *adherence-dimension of P is a linear invariant:*

$$\text{Adim}_{\partial D}(P) = \text{Adim}_{\partial f(D)}(f(P));$$

- (7) *codimension of P is a linear invariant:*

$$\text{codim}_{\partial D}(P) = \text{codim}_{\partial f(D)}(f(P)).$$

Proof. The classical definition of faces (Definition 5.6) combined with the fact that f and its inverse takes intervals to intervals imply (1). More precisely, the linear image of a face F (which is convex) is necessarily convex, and so we need only verify that for every $\xi \in f(F)$ and $\eta, \zeta \in f(D)$ such that ξ lies on the open interval between η and ζ , both η and ζ lie in $f(F)$. In this situation, $f^{-1}(\xi) \in F$ lies on the open interval between $f^{-1}(\eta)$ and $f^{-1}(\zeta)$. Since F is a face, we see that $f^{-1}(\eta), f^{-1}(\zeta) \in D$, and hence $\eta = f(f^{-1}(\eta)), \zeta = f(f^{-1}(\zeta)) \in f(D)$. We see therefore that faces are preserved under invertible linear transformations, and since f is a homeomorphism, any interior point $P \in F$ is sent to an interior point $f(P) \in f(F)$. This suffices to establish (1). The claim (2) then follows immediately from (1).

We next observe that since faces are preserved, if a face F is a subface of F' , then $f(F)$ is a subface of $f(F')$. This in turn means that f preserves the relation of adherence, i.e., the claim (3) is true, and hence the adherence closure of F is mapped to the adherence closure of $f(F)$. This yields claim (4). Furthermore, since adherence is preserved, adherence height, adherence

depth and adherence-dimension must all be preserved: the claims (5) and (6) both hold.

Finally, we shall prove that f preserves codimension. By linearity and invertibility, the map f takes each support hyperplane π at $P \in \mathbb{V}$ to a support hyperplane $f(\pi)$ at $f(P) \in \mathbb{W}$, and hence f induces a map from the face $N_P \subset \mathbb{V}^*$ to the face $N_{f(P)} \subset \mathbb{W}^*$. This is in fact a homeomorphism, with its inverse induced by f^{-1} taking support hyperplanes at $f(P)$ to support hyperplanes at P . Therefore, the dimensions of N_P and $N_{f(P)}$ agree, and hence the codimensions of P and $f(P)$ must be the same, i.e. the claim (7) is true. \square

6. THE CONVEX GEOMETRY OF THURSTON METRIC SPHERES

We now take the various linear invariants developed in Section 5.5 and consider them specifically in the context of convex balls associated with the Thurston metric. In particular, the convex ball D in question will, unless otherwise stated, be the closed ball bounded by

$$\mathbf{S}_x^* := \iota_x(\mathcal{PML}(S)) \subset T_x^*\mathcal{T}(S).$$

The theory of convex bodies (Section 5) we have hitherto seen tells us general properties such as the fact that the dimension of an arbitrary point $P = \iota_x([\lambda])$ is less than its face-dimension (Proposition 5.19):

$$\dim_{\mathbf{S}_x^*}(\iota_x([\lambda])) \leq \text{Fdim}_{\mathbf{S}_x^*}(\iota_x([\lambda])).$$

We now consider the geometry of \mathbf{S}_x^* in this specific context, and in particular establish relationships between the faces on \mathbf{S}_x^* and the topological structure of the support of the measured laminations encoded by each face. We aim to access topological structure from the convex geometry of \mathbf{S}_x^* , and one application is in our proofs of the *topological rigidity* (Section 7.1) of Thurston norm isometries.

As another example, consider the following reformulation of Lemma 5.31 in this specific context. First recall that

$$N_x([\lambda]) := N_{\iota_x([\lambda])} = \{v \in \mathbf{S}_x \mid \iota_x([\lambda])(v) = \|v\|_{\text{Th}}\}.$$

Lemma 6.1. *If $\iota_x([\lambda])$ is contained in the interior of a face F , i.e., if F is the face of $\iota_x([\lambda])$, then for any projective lamination μ with $\iota_x([\mu]) \in F$, we have*

$$N_x([\lambda]) \subseteq N_x([\mu]) \subset T_x\mathcal{T}(S).$$

In particular, if $\iota_x([\lambda])$ and $\iota_x([\mu])$ are in the interior of the same face, then

$$N_x([\lambda]) = N_x([\mu]) \subset T_x\mathcal{T}(S).$$

6.1. Minimal supporting surface and support closure. A measured lamination or a projective lamination is decomposed into components. For each component, we can think of its minimal supporting surface which we define as follows:

We shall use the notion of an *incompressible* subsurface Σ of S . This is a surface with boundary embedded in S such that all its boundary components are essential simple closed curves. We also regard S itself as being incompressible.

Definition 6.2 (minimal supporting surface). Let λ a minimal measured geodesic lamination on S or the projective class of such a lamination. A *minimal supporting surface* for λ is an incompressible subsurface Σ of S containing λ such that for any other incompressible subsurface Σ' containing λ , we have $\Sigma \subset \Sigma'$ after moving Σ by an isotopy. See Fig. 19 for a depiction.

From the above definition, it follows that the minimal supporting surface of λ is unique up to isotopy and depends only on $|\lambda|$. We now define the support closure of a measured/projective measured lamination. This will play an important role later.

Definition 6.3 (support closure). For a measured lamination or a projective measured lamination λ , we consider the minimal supporting surface of every component of λ that is not a simple closed curve. We add to $|\lambda|$ all boundary components of the minimal supporting surfaces and obtain a geodesic lamination. We call this geodesic lamination the *support closure* of λ and denote it by $|\widehat{\lambda}|$ (see Fig. 19).

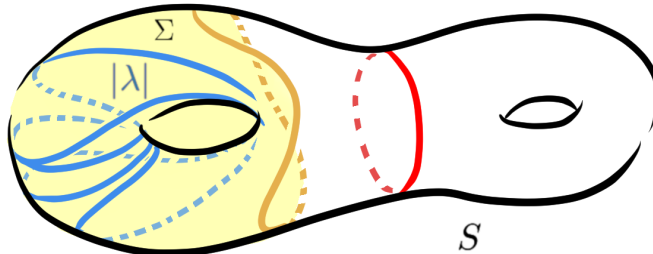


FIGURE 19. The subsurface Σ is the minimal supporting surface for $|\lambda|$, and $|\widehat{\lambda}|$ is obtained by adding the geodesic representative for the boundary curve of Σ to $|\lambda|$.

6.2. Faces for measures of equal support. We first establish some new notation through the following lemma.

Lemma 6.4. *Given a measured lamination λ on S , let $|\lambda|$ denote its support geodesic lamination. The image, under ι_x , of the set of projective measured laminations with support contained in $|\lambda|$ is a convex set, and there is a unique minimal face on \mathbf{S}_x^* containing this convex set. We denote this face by $F_{|\lambda|}$.*

Proof. Instead of working with projective measured laminations $[\mu] \in \mathcal{PML}(S)$ with support in $|\lambda|$, we use a measured lamination representative $\mu \in \mathcal{ML}(S)$. Since $d \log \ell(\mu) = \frac{d\ell(\mu)}{\ell_x(\mu)}$ is homogeneous with respect to $\mu \in \mathcal{ML}(S)$, we assume without loss of generality that μ lies on the unit sphere of $\mathcal{ML}(S)$, i.e., in the subset of $\mathcal{ML}(S)$ defined by $\ell_x(\mu) = 1$. We refer to such measured laminations as *unit measured laminations*.

For two unit measured laminations λ_1 and λ_2 whose supports are contained in $|\lambda|$, the linear combination

$$\lambda_3 = t\lambda_1 + (1-t)\lambda_2, \quad \text{for } t \in [0, 1],$$

defined by taking the transverse measures, is another unit measured lamination, which satisfies that

$$\begin{aligned} \iota_x([\lambda_3]) &= d\ell(t\lambda_1 + (1-t)\lambda_2) \\ &= td\ell(\lambda_1) + (1-t)d\ell(\lambda_2) = t\iota_x([\lambda_1]) + (1-t)\iota_x([\lambda_2]). \end{aligned}$$

Thus the image, under ι_x , of the set of unit measured laminations with support contained in $|\lambda|$, is a convex set. By Proposition 5.13, there is a unique minimal face of \mathbf{S}_x^* containing this convex set. \square

The following lemmas make use of the notation which we introduced in Lemma 6.4, and illustrate some first applications of the notion of support closure (see Definition 6.3). They are key lemmas which will feature in the proofs of results to come in subsequent subsections.

Lemma 6.5. *A recurrent geodesic lamination ξ has non-empty transverse intersection with $|\lambda|$ if and only if it has non-empty transverse intersection with $|\widehat{\lambda}|$.*

Proof. One direction is obvious, and we need only show that if ξ has nonempty transverse intersection with $|\widehat{\lambda}|$, then it also has nonempty transverse intersection with $|\lambda|$. Assume the contrary. Then, first recall that $|\widehat{\lambda}|$ differs from $|\lambda|$ by a collection of simple closed geodesics which themselves serve as the boundary of some minimal supporting surface Σ . The fact that ξ has nonempty intersection with $|\widehat{\lambda}|$ and not $|\lambda|$, combined with the fact that ξ is supported by some recurrent train-track means that there must be some geodesic arc α in $\xi \cap \Sigma$ that starts and ends on the boundary of Σ . Moreover, such a geodesic arc α cannot be homotopic to a boundary arc of Σ as that would form a geodesically bordered hyperbolic bigon. This in turn contradicts the fact that Σ is a minimal supporting surface of ξ , as a component of $\Sigma \setminus \alpha$ retracts onto a surface which supports ξ , but Σ cannot be homotoped into $\Sigma \setminus \alpha$. \square

Lemma 6.6 (non-transversality). *Suppose that $\iota_x([\lambda])$ is contained in the face for $\iota_x([\mu])$. Then $|\mu|$ and $|\widehat{\lambda}|$ cannot have non-empty transverse intersection.*

Proof. Assume that $|\mu|$ and $|\widehat{\lambda}|$ have non-empty transverse intersection. Then by Lemma 6.5, so do $|\mu|$ and $|\lambda|$. We consider a complete geodesic lamination ν containing $|\mu|$. The stretch map along ν on $(S, x) \in \mathcal{T}(S)$ defines an stretch vector $v_\nu \in N_x([\mu]) \subset T_x\mathcal{T}(S)$. By Lemma 6.1, the vector v_ν is contained in $N_x([\mu'])$ for any point $\iota_x([\mu']) \in F_{|\mu|}$ (see Lemma 6.4), and hence in particular we have $v_\nu \in N_x([\lambda])$. However, this is impossible because the infinitesimal Thurston stretch map along ν maximally stretches precisely the laminations which lie in ν (this fact is asserted in [30, Theorem 5.1], see Lemma 4.2 for a sketch-of-proof and Remark 2.5 for some discussion), whereas ν has non-empty transverse intersection with $|\lambda|$ (see Lemma 6.5) and hence v_ν cannot maximally stretch $[\lambda]$. \square

Theorem 6.7 (faces for measures of equal support). *Let $[\lambda] \in \mathcal{PML}(S)$ be a projective measured lamination with support $|\lambda|$. Then, the face $F_{|\lambda|} \subset T_x^*\mathcal{T}(S)$ is an exposed face consisting of precisely the set of ι_x images of projective measured laminations with support contained in $|\lambda|$. In particular, the interior of $F_{|\lambda|}$ consists only of ι_x -images of projective laminations having support equal to $|\lambda|$.*

Proof. Since $(S, x) \setminus |\lambda|$ is a bordered hyperbolic surface, we may endow $(S, x) \setminus |\lambda|$ with a finite geodesic lamination with no closed geodesic leaves. In particular, adding this lamination to $|\lambda|$ yields a complete geodesic lamination ξ on (S, x) which can only support measured laminations on sublaminations of $|\lambda| \subset \xi$. The stretch vector $v_\xi \in \mathbf{S}_x \subset T_x\mathcal{T}(S)$ defines a support hyperplane of the form:

$$\pi_\xi = \left\{ w^* \in T_x^*\mathcal{T}(S) \mid \sup_{v \in \mathbf{S}_x} w^*(v) = w^*(v_\xi) = 1 \right\}.$$

We first observe that the ι_x image of every projective measured lamination supported on $|\lambda|$ lies in this support hyperplane because v_ξ stretches along $|\lambda|$. Also, the minimality of the face $F_{|\lambda|}$ ensures that it is a subset of the exposed face $\pi_\xi \cap \mathbf{S}_x^*$.

Conversely, Lemma 4.2 and the fact that $|\lambda|$ is the largest possible (transverse) measure supporting sublamination of ξ implies that $\pi_\xi \cap \mathbf{S}_x^*$ can only contain ι_x -images of projective measured laminations with support contained in $|\lambda|$. This also shows that $F_{|\lambda|} = \pi_\xi \cap \mathbf{S}_x^*$, and that $F_{|\lambda|}$ is an exposed face as claimed since π_ξ is a support hyperplane.

To finish the proof, we show that the interior points of $F_{|\lambda|}$ consist only of ι_x -images of projective measured laminations whose supports are precisely equal to $|\lambda|$. Consider a projective lamination $[\mu]$ such that $\iota_x([\mu])$ is an interior point for $F_{|\lambda|} = \pi_\xi \cap \mathbf{S}_x^*$. We have shown that $|\mu| \subseteq |\lambda|$, and we now suppose that $|\mu|$ is contained in $|\lambda|$ as a proper subset. Since $\iota_x([\mu])$ is an interior point and $\iota_x([\lambda])$ lies in $F_{|\lambda|}$, Lemma 6.1 tells us that

$$(34) \quad N_x([\mu]) = N_x([\lambda]).$$

Now take ξ to be a complete lamination containing $|\mu|$ that intersects $|\lambda| \setminus |\mu|$ transversely. The stretch vector v_ξ , with respect to ξ , is contained in $N_x([\mu])$, but not in $N_x([\lambda])$ as ξ intersects $|\lambda|$ transversely (Lemma 4.2). This contradicts Eq. (34), and hence $|\mu|$ cannot be a proper subset of $|\lambda|$ and must instead be equal to $|\lambda|$. \square

Remark 6.8. Theorem 6.7 claims that every interior point corresponds to a projective lamination whose support is $|\lambda|$, but it does not say that every lamination whose support is precisely $|\lambda|$ is mapped into the interior of $F_{|\lambda|}$.

Corollary 6.9. *For two arbitrary projective laminations $[\lambda], [\mu] \in \mathcal{PML}(S)$,*

$$|\mu| \subsetneq |\lambda| \text{ if and only if } F_{|\mu|} \subsetneq F_{|\lambda|}.$$

Corollary 6.10. *The dimension of $\iota_x([\lambda])$ is at most equal to the dimension of the cone of transverse measures on $|\lambda|$ subtracted by 1.*

Corollary 6.11. *Given an arbitrary point $\iota_x([\lambda]) \in \mathbf{S}_x^*$, then the dimension $\dim_{\mathbf{S}_x^*}(\iota_x([\lambda]))$ satisfies*

$$\dim_{\mathbf{S}_x^*}(\iota_x([\lambda])) \leq 3g - 4 + n.$$

Proof. By Theorem 6.7, the face of $\iota_x([\lambda])$ is an affine image of the space of projective transverse measures on $|\lambda|$. The dimension of the latter space is bounded by $3g - 4 + n$ as is shown in [23, Corollaire, p. 133]. \square

Remark 6.12. The upperbound on dimension definitionally bounds the face-dimension of every $\iota_x([\lambda])$ above by $3g - 4 + n$.

Corollary 6.13 (embedded curve complex). *The convex stratification of \mathbf{S}_x^* contains the curve complex for S as a subcomplex. We denote this subcomplex as \mathbf{C}_x^* .*

Proof. Following Theorem 6.7, we have a complete topological description of the face of $\iota_x([\Gamma])$ for an arbitrary projective multicurve $[\Gamma]$: when $|\Gamma|$ is comprised of k simple closed geodesics, the face of $\iota_x([\Gamma])$ is a $k - 1$ simplex whose boundary consists of lower-dimensional simplices corresponding to the faces of projective weighted multicurves supported on a proper subset of closed geodesics in $|\Gamma|$. This is precisely a geometric realisation of the curve complex, and we denote it by \mathbf{C}_x^* . \square

Corollary 6.14. *For any projectively weighted multicurve with k components $[\Gamma] \in \mathcal{PML}(S)$, the dimension of $\iota_x([\Gamma])$ is $k - 1$.*

Corollary 6.15 (contravariant labelling). *For two arbitrary projective multicurves $[\Gamma_1], [\Gamma_2] \in \mathcal{PML}(S)$, denote their respective supports by $|\Gamma_1|$ and $|\Gamma_2|$. Then,*

$$(35) \quad |\Gamma_1| \subseteq |\Gamma_2| \text{ if and only if } N_x([\Gamma_2]) \subseteq N_x([\Gamma_1]),$$

with equality on the left if and only if we have equality on the right.

Proof. Since $[\Gamma_1]$ and $[\Gamma_2]$ are projective multicurves, the points $\iota_x([\Gamma_1])$ and $\iota_x([\Gamma_2])$ respectively lie in the interior of $F_{[\Gamma_1]}$ and $F_{[\Gamma_2]}$. Therefore, $[\Gamma_1] \subseteq [\Gamma_2]$ if and only if $\iota_x([\Gamma_1])$ is contained in the face F of $\iota_x([\Gamma_2])$. Lemma 6.1 then yields the desired result. \square

6.3. Closure correspondence.

Theorem 6.16 (closure correspondence). *Let x be a fixed hyperbolic structure on S , and consider a projective lamination $[\lambda]$ with support $|\lambda|$. Let \widehat{F} be the adherence closure of the face F for $\iota_x([\lambda])$. Then the following hold:*

- (a) *The face \widehat{F} is equal to the (exposed) face $F_{\widehat{|\lambda|}}$ consisting of the images of all projective laminations whose supports are contained in $\widehat{|\lambda|}$.*
- (b) *The interior points of \widehat{F} are ι_x -images of projective laminations with support equal to $\widehat{|\lambda|}$.*
- (c) *The ι_x -image of the set of all projective laminations whose support closures are $\widehat{|\lambda|}$ coincides with the adherence core of \widehat{F} .*

Proof. We shall show that the adherence closure \widehat{F} of F coincides with $F_{\widehat{|\lambda|}}$. The fact that $F_{\widehat{|\lambda|}}$ is an exposed face, i.e., that it is the intersection of a hyperplane and \mathbf{S}_x^* (see Theorem 6.7), which contains $\iota_x([\lambda])$, combined with Lemma 5.25 ensures that F is a subface of $F_{\widehat{|\lambda|}}$. We first show that F is adherent to $F' := F_{\widehat{|\lambda|}}$. Suppose that another face F'' contains F as a subface, and take $\iota_x([\mu])$ in the interior of F'' . There are three mutually exclusive possibilities covering all options for the relationship between $\widehat{|\lambda|}$ and $[\mu]$:

- (1) $\widehat{|\lambda|}$ has nonempty transverse intersection with $|\mu|$;
- (2) $\widehat{|\lambda|}$ has no nonempty transverse intersection with $|\mu|$, but $\widehat{|\lambda|} \setminus |\mu| \neq \emptyset$;
- (3) $\widehat{|\lambda|}$ is properly contained in $|\mu|$.

Case (1): If $|\mu|$ intersects $\widehat{|\lambda|}$ transversely, then we take a complete geodesic lamination ξ containing $|\mu|$ and consider the stretch vector v_ξ along ξ . Since $\iota_x([\lambda]) \in F \subset F''$ and F'' contains $\iota_x([\mu])$ in its interior, Lemma 6.1 implies that

$$v_\xi \in N_x([\mu]) \subset N_x([\lambda]).$$

Lemma 6.5 further tells us that $|\mu|$ has non-empty transverse intersection with $|\lambda|$, and hence ξ has non-empty transverse intersection with $|\lambda|$. This then contradicts the non-transversality lemma (Lemma 6.6) and renders Case (1) impossible.

Case (2): Next suppose that $\widehat{|\lambda|} \setminus |\mu| \neq \emptyset$. We choose a complete lamination ξ containing $|\mu|$ which intersects $\widehat{|\lambda|}$ transversely and repeating the argument for Case (1), albeit without needing to invoke Lemma 6.5 as ξ already has

non-empty transverse intersection with $|\lambda|$, we also get a contradiction for Case (2).

Case (3): The only remaining case is when $\widehat{|\lambda|}$ is a proper subset of $|\mu|$. Take a unit measured lamination λ' such that $\iota_x([\lambda'])$ is contained in the interior of $F_{\widehat{|\lambda|}}$, and note that Theorem 6.7 tells us that $|\lambda'| = \widehat{|\lambda|}$. Since both λ' and μ are measured laminations, their supports are both unions of their minimal components. Therefore, we have a decomposition of μ (assumed to be a unit measured lamination) into disjoint sublaminations $\mu_1 \sqcup \mu_2$ such that $|\mu_1| = |\lambda'| = \widehat{|\lambda|}$. Since $|\mu|$ contains $|\lambda'| = \widehat{|\lambda|}$, we can define a path of unit measured lamination $t\lambda' + (1-t)\mu$, for $t \in [0, 1]$, which is mapped to an affine segment on \mathbf{S}_x^* connecting $\iota_x([\mu])$ with $\iota_x([\lambda'])$. Since $|t\lambda' + (1-t)\mu| \subset |\mu_1| \cup |\mu| = |\mu|$, the path is contained in $F_{|\mu|}$. Taking \hat{F} (in Definition 5.15) to be $F_{|\mu|}$, the face \hat{F} contains both $F' = F_{\widehat{|\lambda|}}$ and F'' , hence F satisfies the conditions for adherence to $F' = F_{\widehat{|\lambda|}}$.

We next show that $F_{\widehat{|\lambda|}}$ is maximal among the faces to which F is adherent. If F is adherent to $F' \supsetneq F_{\widehat{|\lambda|}}$, then for any interior point $\iota_x([\mu]) \in F'$, our arguments for Cases (1) and (2) show that $|\mu|$ necessarily contains $\widehat{|\lambda|}$ as a subset. We first consider the possibility that $\widehat{|\lambda|} \subsetneq |\mu|$. Then there is a measured lamination ν whose support contains $\widehat{|\lambda|}$ as a sublamination and which intersects μ transversely. Let $F'' = F_{|\nu|}$, then Theorem 6.7 tells us that $F'' \supset F_{\widehat{|\lambda|}}$. Since $F_{\widehat{|\lambda|}}$ contains F as was explained at the beginning of the proof, F is a subface of F'' . On the other hand, since $|\mu|$ and $|\nu|$ have non-empty transverse intersection, by Theorem 6.7, they cannot lie in the same face in \mathbf{S}_x^* — any point in the interior of such a face would be the ι_x image of a projective measured lamination with both $|\mu|$ and $|\nu|$ in its support, which is impossible. This implies that there is no face (in \mathbf{S}_x^*) containing both F' and $F'' = F_{|\nu|}$, and hence contradicts the assumption that F is adherent to F' . Therefore the possibility that $\widehat{|\lambda|} \subsetneq |\mu|$ is excluded, which implies that $\widehat{|\lambda|} = |\mu|$ and thus $F' = F_{\widehat{|\lambda|}}$, as desired. This means that $F_{\widehat{|\lambda|}}$ is the adherence closure of the face for $\iota_x([\lambda])$ and completes the proof of (a). Theorem 6.7 then gives (b).

We finally turn to proving (c). Consider a projective lamination $[\lambda'] \in \mathcal{PM}\mathcal{L}(S)$ whose support closure is $\widehat{|\lambda|}$. Then part (a) asserts that the adherence closure of the face for $\iota_x([\lambda'])$ is precisely $F_{\widehat{|\lambda|}}$. Therefore, $\iota_x([\lambda'])$ is contained in the adherence core of $\hat{F} = F_{\widehat{|\lambda|}}$. Conversely, suppose that $\iota_x([\mu])$ is in the adherence core of \hat{F} , i.e., it is contained in the interior of a face \tilde{F} whose adherence closure is $F_{\widehat{|\lambda|}}$. Then part (a) shows that the adherence closure of \tilde{F} is $F_{|\mu|}$, hence that $F_{\widehat{|\lambda|}} = F_{|\mu|}$. Part (b) then asserts

that $|\widehat{\mu}| = |\widehat{\lambda}|$, therefore the support closure of $[\mu]$ is indeed equal to $|\widehat{\lambda}|$, as desired. \square

Theorem 6.16 implies the following corollaries.

Corollary 6.17. *For any two projective laminations $[\lambda], [\mu] \in \mathcal{PML}(S)$,*

$$\text{if } |\mu| \subsetneq |\lambda|, \text{ then } F_{|\widehat{\mu}|} \subsetneq F_{|\widehat{\lambda}|}.$$

Corollary 6.18. *The face-dimension of $\iota_x([\lambda])$ is equal to the dimension of the cone of transverse measures on $|\widehat{\lambda}|$ subtracted by 1.*

Corollary 6.19. *If $[\Gamma] \in \mathcal{PML}(S)$ is a projectively weighted multicurve with k components, then the face-dimension of $\iota_x([\Gamma])$ is $k - 1$.*

6.4. Distinguishing laminations. One strength of Theorem 6.16 is that the property of being adherence-closed is linearly invariant. Results in this subsection show us how we might distinguish certain types of projective measured laminations based on linearly invariant properties (such as adherence-closedness) of the face for the ι_x -images of the aforementioned laminations.

Theorem 6.20. *Let $[\lambda]$ be a projective class of a measured lamination. Then the following two properties are equivalent:*

- (1) *$[\lambda]$ is either the projective class of a weighted multicurve or it is maximal and uniquely ergodic;*
- (2) *every subface of the face for $\iota_x([\lambda])$ is adherence-closed.*

Proof. Let F be the face for $\iota_x([\lambda])$. If $[\lambda]$ is a projectively weighted multicurve, then by Theorem 6.16, every subface \tilde{F} of F is the face of a projective multicurve supported on a subset of $|\lambda|$. In particular, since multicurves are their own support closures, \tilde{F} is adherence-closed. On the other hand, if $[\lambda]$ is maximal and uniquely ergodic, then F consists of just one point $\iota_x([\lambda])$ and $F = \{\iota_x([\lambda])\}$ is the only subface of itself. Moreover, since $[\lambda]$ is maximal, then, by Theorem 6.16, there is no face properly containing F , and hence F is again adherence-closed.

Suppose next that $[\lambda]$ is neither a projectively weighted multicurve nor a maximal and uniquely ergodic lamination. We show that there is a subface of F whose adherence closure is not itself. First note that if the adherence closure of F is not F , then we are done, and thus we assume that F is adherence-closed. In particular, since $\iota_x([\lambda])$ is an interior point of F , we know by Theorem 6.16 that $F = F_{|\widehat{\lambda}|}$. We cover the remainder of the proof with two cases:

- (1) $[\lambda]$ is either non-maximal or is disconnected,
- (2) $[\lambda]$ is both maximal and connected.

Case (1): If $[\lambda]$ is either non-maximal or is disconnected, then λ has a component λ_0 which is not a weighted simple closed curve and whose supporting surface is not all of S . This means that the support closure $|\widehat{\lambda_0}|$ properly

contains $|\lambda_0|$, hence Theorem 6.16 tells us that the face \tilde{F} for $\iota_x([\lambda_0])$ is not adherence-closed. Indeed, Theorem 6.16 tells us that the adherence closure of \tilde{F} is $F = F_{|\lambda|}$, hence it gives a subface of the latter which is not adherence-closed. This covers the case when F is either non-maximal or is disconnected.

Case (2): If $[\lambda]$ is both maximal and connected. Then, by the assumption that $[\lambda]$ is not uniquely ergodic, the space of transverse measures on $|\lambda|$ constitutes a convex set whose extreme points correspond to the ergodic measures. (The existence of such extremal points follows from the classical theorem on convex sets which can be found in [16]. The proof of the fact that the ergodic measures constitute the extremal points can be found in [13, 17].) This implies that the image of the entire set of projective transverse measures on $|\lambda|$ cannot be the interior of $F_{|\lambda|}$, therefore there is a projective lamination $[\mu]$ supported on $|\lambda|$ such that the face \tilde{F} for $\iota_x([\mu])$ is not adherence-closed (by Theorem 6.16). Since $F = F_{|\lambda|}$ is the adherence closure of \tilde{F} , the latter is a subface of F which fails to be adherence-closed, as desired. This completes the proof. \square

Proposition 6.21. *If a projective measured lamination $[\lambda]$ is maximal, then the adherence closure of the face for $\iota_x([\lambda])$ is maximal, i.e., it is not a subface of any other face.*

Proof. Since $[\lambda]$ is maximal, its support closure is $|\widehat{\lambda}| = |\lambda|$. Then, Corollary 6.9 ensures that $F_{|\lambda|}$ is maximal among all faces that can be expressed in the form $F_{|\mu|}$ for some projective measured lamination $[\mu]$. Consider a face F containing $F_{|\lambda|}$, and let $\iota_x([\lambda'])$ be an interior point of F . Then Theorem 6.16 tells us that

$$F_{|\lambda|} \subseteq F \subseteq F_{|\lambda'|}.$$

By the maximality of $F_{|\lambda|}$ among faces of the form $F_{|\mu|}$, we see therefore that $F_{|\lambda|} = F_{|\lambda'|}$, and hence $F_{|\lambda|} = F$. Since every face containing $F_{|\lambda|}$ is equal to $F_{|\lambda|}$ itself, the face $F_{|\widehat{\lambda}|} = F_{|\lambda|}$, which is the adherence closure of the face for $\iota_x([\lambda])$, must be maximal. \square

Remark 6.22. The converse to Proposition 6.21 is false: simply take $[\lambda]$ to be a non-maximal lamination such that $|\widehat{\lambda}|$ is maximal.

Theorem 6.23. *A projective lamination $[\lambda]$ is maximal and connected if and only if the adherence closure of the face for $\iota_x([\lambda])$ is maximal (that is, it is not a subface of any other face), and adherence complete (Definition 5.15).*

Proof. We first show the “only if” part. We have already seen from Proposition 6.21 that if $[\lambda]$ is maximal, then the adherence closure $F' = F_{|\widehat{\lambda}|}$ of the face for $\iota_x([\lambda])$ is maximal. We need to prove adherence completeness from connectedness. Take a point $\iota_x([\lambda'])$ lying in the interior of F' . By Theorem 6.16, given the interior point $\iota_x([\mu])$ of an arbitrary face in F' , the

lamination $[\mu]$ has support $|\mu|$ contained in $|\lambda'| = \widehat{|\lambda|} = |\lambda|$ (the last equality follows from the maximality of $|\lambda|$). Since λ is connected, this implies that $|\mu| = |\lambda|$ (see e.g. [8, Cor 1.7.3]), and by Theorem 6.16, the adherence closure of the face for $\iota_x([\mu])$ is also F' . Therefore, F' is adherence complete.

Now, to prove the “if” part, suppose that $[\lambda]$ is either non-maximal or disconnected. Then there is a component of λ whose minimal supporting surface is not the entire S . Let c denote a boundary curve of such a minimal supporting surface. Let F' be the adherence closure of the face for $\iota_x([\lambda])$. Then, by Theorem 6.16, $\iota_x([c])$ is contained in F' . We immediately exclude the case when $|c| = |\lambda|$; for, then F' consists of one point, and cannot be a maximal face (recall that $S \neq S_{1,1}, S_{0,4}$). Thus, $|c| \neq |\lambda|$, and the adherence closure of the face for $\iota_x([c])$ is itself (Theorem 6.16), hence F' is not adherence complete. \square

6.5. Tangential adherence. In part (c) of Theorem 6.16, we described the image of the set of projective measured laminations whose support closures are $\widehat{|\lambda|}$ as a subset on $F_{\widehat{|\lambda|}} \subset \mathbf{S}_x^*$. We now present an analogous result for the set of projective measures supported precisely on $|\lambda|$.

Definition 6.24 (tangential adherence). A subface F of a face $F' \subset \mathbf{S}_x^*$ is said to be *tangentially adherent* to F' if the following condition is satisfied:

Given $[\lambda], [\mu] \in \mathcal{PML}(S)$ such that F is the face for $\iota_x([\lambda])$ and F' is the face for $\iota_x([\mu])$, every stretch vector v , along a complete geodesic lamination, that is contained in $N_x([\lambda])$ is also contained in $N_x([\mu])$.

Proposition 6.25. *Let $[\lambda]$ be a projective measured lamination on S , and consider the face $F_{|\lambda|}$ comprised of all ι_x images of projective measured laminations supported in $|\lambda|$ (see Lemma 6.5 and Theorem 6.7). The union of the interior of $F_{|\lambda|}$ and the interiors of all subfaces that are tangentially adherent to F coincides with the ι_x -image of the set of projective transverse measures supported precisely on $|\lambda|$.*

Proof. Assume without loss of generality that $[\lambda]$ is contained in the interior of $F_{|\lambda|}$. Let $[\lambda']$ be an arbitrary projective lamination with support $|\lambda'| = |\lambda|$. By Theorem 6.7, $\iota_x([\lambda'])$ is contained in $F_{|\lambda|}$, hence the face F for $\iota_x([\lambda'])$ is a subface of the exposed face $F_{|\lambda|}$ (Lemma 5.25). If $F = F_{|\lambda|}$, then $\iota_x([\lambda'])$ (trivially) lies in the interior of a face tangentially adherent to $F_{|\lambda|}$. Now suppose otherwise that the face for $\iota_x([\lambda'])$ is a proper subface F of $F_{|\lambda|}$. Given an arbitrary stretch vector $v_\nu \in N_x([\lambda'])$, the vector v_ν maximally stretches a complete geodesic lamination ν which contains $|\lambda'| = |\lambda|$. Thus, v_ν also maximally stretches λ and so v_ν is also contained in $N_x([\lambda])$. This shows that the face F is tangentially adherent to $F_{|\lambda|}$, hence an arbitrary point $\iota_x([\lambda'])$ with support equal to $|\lambda|$ always lies in the interior of some tangentially adherent subface of $F_{|\lambda|}$.

Conversely, suppose that the face F for some point $\iota_x([\mu])$ is tangentially adherent to $F_{|\lambda|}$. By Theorem 6.7, $|\mu|$ is a subset of $|\lambda|$. If $|\mu|$ is properly

contained in $|\lambda|$, then we can find a complete geodesic lamination ν containing $|\mu|$ but intersecting $|\lambda|$ transversely. Then the stretch vector v_ν along ν is contained in $N_x([\mu])$ but not in $N_x([\lambda])$, which contradicts the assumption that F is tangentially adherent to $F_{|\lambda|}$. We conclude that every projective lamination $[\mu]$ whose image under ι_x is contained in the interior of F has support precisely equal to $|\lambda|$. \square

Corollary 6.26. *For any projective measured lamination $[\lambda]$,*

$$\widehat{F_{|\lambda|}} = F_{\widehat{|\lambda|}}.$$

Proof. By Proposition 6.25, there exists $[\lambda'] \in \mathcal{PML}$ on the interior of $F_{|\lambda|}$ such that $|\lambda'| = |\lambda|$. Then, $F_{|\lambda|}$ is the face for $\iota_x([\lambda'])$ and by the closure correspondence theorem (Theorem 6.16), we obtain:

$$\widehat{F_{|\lambda|}} = \widehat{F_{|\lambda'|}} = F_{\widehat{|\lambda'|}} = F_{\widehat{|\lambda|}}.$$

\square

As previously highlighted, one major advantage of Theorem 6.16 is the linear invariance of the property of adherence (and adherence closures, by extension). In contrast, we are unable to show directly that tangential adherence is linearly invariant. We note that it does follow as a consequence of the infinitesimal rigidity theorem Theorem 1.1. In any case, we establish the following relationship between tangential adherence and adherence:

Lemma 6.27. *A subface tangentially adherent to F' is adherent to F' .*

Proof. Suppose that F is tangentially adherent to F' , and that F is the face for $\iota_x([\lambda])$ whereas F' is the face for $\iota_x([\mu])$ as in the statement of Definition 6.24. Let F'' be a face containing F as a proper subface, and let $\iota_x([\nu])$ be an interior point for F'' . The exposed face $F_{|\nu|}$ contains F'' , and hence also $\iota_x([\mu])$. It follows that $|\nu|$ contains $|\mu|$ (Theorem 6.16). Suppose that $|\nu|$ is not contained in $|\lambda|$. Then we can choose a complete geodesic lamination ξ containing $|\nu|$ but intersecting $|\lambda|$ transversely. The stretch vector v_ξ lies in $N_x([\mu])$ because ξ contains $|\mu|$. However, v_ξ is not an element of $N_x([\lambda])$ since ξ intersects $|\lambda|$ transversely, and this contradicts the condition of Definition 6.24.

We may therefore conclude that $|\nu|$ is contained in $|\lambda|$, and hence the exposed face $F_{|\lambda|}$ (Theorem 6.7) contains both F' and F'' (Lemma 5.25). This affirms the property of adherence. \square

6.6. Codimension of multicurves. We have already determined the dimension and face-dimension of ι_x -images of multicurves (Corollary 6.14 and Corollary 6.19), and now determine their codimension.

Theorem 6.28 (codimension of multicurves). *For any projectively weighted multicurve $[\Gamma] \in \mathcal{PML}(S)$ having k components and for any $x \in \mathcal{T}(S)$, the codimension of $\iota_x([\Gamma])$ is equal to $6g - 6 + 2n - k$. Note in particular that this codimension is independent of x .*

We first establish upper bound for $\text{codim}(\iota_x([\Gamma]))$:

Lemma 6.29 (codimension upper bound). *For any projectively weighted multicurve $[\Gamma]$ having k components,*

$$\text{codim}_{\mathbf{S}_x^*}(\iota_x([\Gamma])) \leq 6g - 6 + 2n - k.$$

Proof. We know from Corollary 6.14 that $\dim_{\mathbf{S}_x^*}(\iota_x([\Gamma]))$ is equal to $k - 1$. The lemma then follows immediately from Corollary 5.30. \square

Remark 6.30. In the specific case when $[\Gamma]$ is a closed geodesic, Theorem 6.28 is a corollary of Thurston's work [30, Theorem 10.1].

We next show that the upper bound on codimension is tight by constructing sufficiently many linearly independent vectors in the length increasing cone. We begin with the special case when the support of λ is a pants decomposition of S .

Theorem 6.31 (codimension for pants decompositions). *Every projectively weighted pants decomposition $[\Gamma] \in \mathcal{PML}(S)$ of $S = S_{g,n}$ has codimension $3g - 3 + n$.*

Proof. After Lemma 6.29, we need only show that there are $3g - 3 + n$ linearly independent vectors contained in $N_x([\Gamma])$. In particular, since $N_x([\Gamma])$ is a convex set, this is equivalent to showing that the affine subspace generated by $N_x([\Gamma])$ is $3g - 3 + n$ -dimensional.

Let $\gamma_1, \dots, \gamma_{3g-3+n}$ denote the simple closed geodesics constituting the support of $[\Gamma]$. For each of the γ_i , let $\Lambda_{+,i}$ and $\Lambda_{-,i}$ denote complete geodesic laminations on S such that

- $\Lambda_{\pm,i}$ contain $\gamma_1, \dots, \gamma_{3g-3+n}$,
- $\Lambda_{\pm,i}$ contain geodesic leaves which bound the same $(1, 1)$ -cusped annulus (see Fig. 3) $A_i \subset S$, which in turn contains γ_i ,
- $\Lambda_{+,i}$ and $\Lambda_{-,i}$ have the same geodesic leaves, except in the interior of A_i , where the two bi-infinite geodesic leaves on A_i in $\Lambda_{+,i}$ and $\Lambda_{-,i}$ which spiral toward γ_i do so in opposite directions (see left-most and right-most pictures in Fig. 4).

Note that it is a simple exercise to see that it is always possible to construct such $\Lambda_{\pm,i}$.

We denote the stretch vectors for these laminations by $v_{\pm,i} \in N_x([\Gamma])$. By Corollary 4.4, we see that $v_{+,i} - v_{-,i}$ is precisely the unit Fenchel–Nielsen twist vector for γ_i , and hence the vectors

$$\{v_{+,i} - v_{-,i}\}_{i=1, \dots, 3g-3+n}$$

correspond to (non-zero multiples of) Fenchel–Nielsen twists for distinct simple closed geodesics, hence they are all linearly independent. Moreover, $v_{+,1}$ is linearly independent from the $v_{+,i} - v_{-,i}$ since stretching along $\Lambda_{+,1}$

increases the lengths of $\gamma_1, \dots, \gamma_{3g-3+n}$. Hence, the following collection of $3g-2+n$ vectors

$$v_{+,1}, (v_{+,1} - v_{-,1}), (v_{+,2} - v_{-,2}), \dots, (v_{+,3g-3+n} - v_{-,3g-3+n})$$

are linearly independent and generate a $(3g-2+n)$ -dimensional vector space $\mathbb{A} \leq T_x \mathcal{T}(S)$. This is necessarily a subspace of the vector space generated by $N_x([\Gamma])$, and hence $\mathbb{A} \cap N_x([\Gamma])$ has the same dimension as $\mathbb{A} \cap \{v \in T_x \mathcal{T}(S) \mid \iota_x([\Gamma])(v) = 1\}$, which in turn has dimension $3g-2+n-1 = 3g-3+n$. Therefore, $N_x([\Gamma])$ is at least $3g-3+n$ dimensional, and this gives us the desired lower bound. \square

We now adapt this argument to deal with general multicurves:

Proof of Theorem 6.28. Given an arbitrary projective class of a weighted multicurve $[\Gamma] \in \mathcal{PML}(S)$, let $\gamma_1, \dots, \gamma_k$ denote the simple closed geodesics which make the support $|\Gamma|$. We first complete $\{\gamma_i\}_{i=1,\dots,k}$ to a pants decomposition $\{\gamma_i\}_{i=1,\dots,3g-3+n}$, whose union is represented by a weighted multicurve Γ' . In the proof of Theorem 6.31, we showed that there is a vector space of dimension $3g-2+n$ generated by stretch vectors contained in $N_x([\Gamma'])$. Therefore, we can choose linearly independent stretch vectors $\{v_j\}_{j=1,\dots,3g-2+n} \in N_x([\Gamma'])$. Since $|\Gamma|$ is contained in $|\Gamma'|$ these vectors also maximally stretch Γ , and hence they are contained in $N_x([\Gamma])$.

We next replace γ_{3g-3+n} by a different simple closed geodesic δ_1 such that $\{\gamma_1, \dots, \gamma_{3g-4+n}, \delta_1\}$ is a new pants decomposition, and we let w_1 denote the stretch vector for some complete geodesic lamination containing $\{\gamma_1, \dots, \gamma_{3g-4+n}, \delta_1\}$.

We show that v_1, \dots, v_{3g-2+n} , and w_1 are linearly independent as follows.

If not, then there exist coefficients $a_1, \dots, a_{3g-2+n}, b_1 \in \mathbb{R}$ such that

$$a_1 v_1 + a_2 v_2 + \dots + a_{3g-2+n} v_{3g-2+n} + b_1 w_1 = \vec{0} \in T_x \mathcal{T}(S).$$

However, the vector w_1 is the only one among this collection which maximally stretches δ_1 , and this would then cause $b_1 = d\ell_{\delta_1}(a_1 v_1 + a_2 v_2 + \dots + a_{3g-2+n} v_{3g-2+n} + b_1 w_1) = 0$. The linear independence of v_1, \dots, v_{3g-2+n} then implies that $a_1 = \dots = a_{3g-2+n} = 0$. Hence $v_1, \dots, v_{3g-3+n}, w_1$ are linearly independent as claimed.

We continue this procedure by replacing γ_{3g-4+n} with a simple closed curve δ_2 such that $\{\gamma_1, \dots, \gamma_{3g-5+n}, \delta_1, \delta_2\}$ is a new pants decomposition and similarly produce a new tangent vector w_2 . The same coefficient argument tells us that $\{v_1, \dots, v_{3g-2+n}, w_1, w_2\}$ is a collection of linearly independent vectors. To clarify, if

$$a_1 v_1 + a_2 v_2 + \dots + a_{3g-2+n} v_{3g-2+n} + b_1 w_1 + b_2 w_2 = \vec{0},$$

then we use the fact that *only* w_2 maximally stretches δ_2 to show that $b_2 = 0$ by considering the value under $d\ell_{\delta_2}$. The linear independence of $v_1, \dots, v_{3g-2+n}, w_1$ then implies that $a_1 = \dots = a_{3g-2+n} = b_1 = 0$.

We iterate the above argument until we construct linearly independent tangent vectors $v_1, \dots, v_{3g-2+n}, w_1, \dots, w_{3g-3+n-k}$. Since these vectors

maximally stretch $\gamma_1, \dots, \gamma_k$, they must lie in $N_x([\Gamma])$, hence the codimension of $\iota_x([\Gamma])$ is at least $3g - 2 + n + 3g - 3 + n - k - 1 = 6g - 6 + 2n - k$. This gives the desired lower bound. \square

Corollary 6.32 (embedded dual curve complex). *For every $x \in \mathcal{T}(S)$, the convex stratification on \mathbf{S}_x contains an ideal cell-complex \mathbf{C}_x which is dual to the curve complex in the sense that*

- *the support $|\Gamma|$ for any projective multicurve $[\Gamma] \in \mathcal{PML}$ is assigned to the face $N_x([\Gamma]) \subsetneq \mathbf{S}_x$;*
- *the subset-relation for cells in the curve complex inverts to the superset-relation for faces in \mathbf{C}_x ;*
- *the dimension of each cell in the curve complex is equal to the codimension of its corresponding face in \mathbf{C}_x as a subset of \mathbf{S}_x .*

Proof. Corollary 6.15 asserts that the assignment of the face $N_x([\Gamma])$ to the multicurve $[\Gamma]$ is purely dependent on the support $|\Gamma|$ and is hence well-defined. Moreover, Corollary 6.15 ensures that this is an injective assignment and Lemma 6.1 ensures that the superset-relation in \mathbf{C}_x is inverted. Let us denote the collection of faces of the form $N_x([\Gamma])$, for some multicurve $[\Gamma]$, by \mathbf{C}_x .

All that remains is to demonstrate the codimension condition on \mathbf{C}_x . Theorem 6.28 (definitionally) asserts that for any projective multicurve $[\Gamma]$ comprised of k components, the dimension of $N_x([\Gamma])$ is $6g - 6 + 2n - k$. Therefore, its codimension as a subset of \mathbf{S}_x , is $6g - 7 + 2n - (6g - 6 + 2n - k) = k - 1$ — this is precisely the dimension of the cell in the curve complex corresponding to $|\Gamma|$. \square

7. INFINITESIMAL RIGIDITY OF THE THURSTON METRIC

The goal of this section is to prove the infinitesimal rigidity theorem, i.e., Theorem 1.1. Before proceeding, we clarify an important technical point. We actually have two ways in which mapping classes induce maps between tangent (and cotangent) spaces on the Teichmüller space as we now explain. A diffeomorphism $h: S \rightarrow S$ induces a mapping class $[h]$ which acts diffeomorphically on $\mathcal{T}(S)$ by the pushing-forward of metrics, as well as the natural homeomorphism $h: \mathcal{PML}(S) \rightarrow \mathcal{PML}(S)$ given by taking $[\lambda]$ to $[h(\lambda)]$. Then, for any $x \in \mathcal{T}(S)$,

- (1) we obtain a map (using the fact that $[h]$ is an analytic isometry and hence differentiable) dual to the derivative map $d[h]$

$$d^*[h]: T_{[h](x)}^* \mathcal{T}(S) \rightarrow T_x^* \mathcal{T}(S),$$

- (2) as well as a map

$$\iota_x \circ h^{-1} \circ \iota_{[h](x)}^{-1}: \mathbf{S}_{[h](x)}^* \rightarrow \mathbf{S}_x^*,$$

which we extend by non-negative homothety to a map

$$\iota_x \circ h^{-1} \circ \iota_{[h](x)}^{-1}: T_{[h](x)}^* \mathcal{T}(S) \rightarrow T_x^* \mathcal{T}(S).$$

Lemma 7.1. *The two maps $\iota_x \circ h^{-1} \circ \iota_{[h](x)}^{-1}$ and $d^*[h]$ agree.*

Proof. Any tangent vector $v \in T_x \mathcal{T}(S)$ at x can be represented as the derivative $\frac{d}{dt}c(0)$ at 0 of a differentiable curve $c(t) : (-\epsilon, \epsilon) \rightarrow \mathcal{T}(S)$. Then,

$$d[h](v) = \frac{d}{dt}([h] \circ c)(0).$$

For any measured lamination μ and any point $y \in \mathcal{T}(S)$, we have $\ell_{[h](y)}(h(\lambda)) = \ell_y(\mu)$, and hence

$$\ell_{[h] \circ c(t)}(h(\mu)) = \ell_{c(t)}(\mu).$$

This in turn implies that $d \log \ell(h(\mu))(\frac{d}{dt}([h] \circ c)(0)) = d \log \ell(\mu)(\frac{d}{dt}(c(0)))$, hence

$$(\iota_{[h](x)} h([\mu]))(d[h](v)) = (\iota_x([\mu]))(v).$$

Note that for an arbitrary covector $w^* \in \mathbf{S}_{[h](x)}^* \subset T_{[h](x)}^* \mathcal{T}(S)$, the property $d^*[h](w^*)(v) = w^*(d[h])(v)$ for every vector $v \in T_x \mathcal{T}(S)$ characterises $d^*[h](w^*)$. Setting w^* to be $\iota_x([\mu])$ for some $[\mu]$, we see that $d^*[h] \equiv \iota_x h^{-1} \iota_{[h](x)}^{-1}$ on $\mathbf{S}_{[h](x)}^*$ and hence on $T_{[h](x)}^* \mathcal{T}(S)$. \square

Corollary 7.2. *A diffeomorphism $h : S \rightarrow S$ representing a mapping class $[h] : \mathcal{T}(S) \rightarrow \mathcal{T}(S)$ preserves linear invariants such as the dimension, the face-dimension, adherence-dimension and codimension of points. In other words, for any $[\lambda] \in \mathcal{PML}(S)$, we have*

$$\begin{aligned} \dim_{\mathbf{S}_x^*}(\iota_x([\lambda])) &= \dim_{\mathbf{S}_{[h](x)}^*}(\iota_{[h](x)}(h([\lambda]))), \\ \text{Fdim}_{\mathbf{S}_x^*}(\iota_x([\lambda])) &= \text{Fdim}_{\mathbf{S}_{[h](x)}^*}(\iota_{[h](x)}(h([\lambda]))), \\ \text{Adim}_{\mathbf{S}_x^*}(\iota_x([\lambda])) &= \text{Adim}_{\mathbf{S}_{[h](x)}^*}(\iota_{[h](x)}(h([\lambda]))), \\ \text{and } \text{codim}_{\mathbf{S}_x^*}(\iota_x([\lambda])) &= \text{codim}_{\mathbf{S}_{[h](x)}^*}(\iota_{[h](x)}(h([\lambda]))). \end{aligned}$$

Proof. This follows from Theorem 5.34 applied to the observation that $d^*[h]$ is an invertible linear map. \square

7.1. Topological rigidity. Recall from §1.10 that every Thurston-norm-preserving isometry $f : T_x \mathcal{T}(S) \rightarrow T_y \mathcal{T}(S)$ defines a self-homeomorphism of projective lamination space

$$f_{y,x} := \iota_x^{-1} \circ f^* \circ \iota_y : \mathcal{PML}(S) \rightarrow \mathcal{PML}(S),$$

where $f^* : T_y^* \mathcal{T}(S) \rightarrow T_x^* \mathcal{T}(S)$ is the dual isometry. We now show that any such $f_{y,x}$ comes from a mapping class:

Theorem 7.3 (Topological rigidity, Theorem 1.8). *The map $f_{y,x} : \mathcal{PML}(S) \rightarrow \mathcal{PML}(S)$ associated with a Thurston-norm isometry $f : T_x \mathcal{T}(S) \rightarrow T_y \mathcal{T}(S)$ is necessarily induced by a diffeomorphism $h : S \rightarrow S$.*

As previously mentioned in §1.10, we shall give two different proofs for Theorem 1.8. The first proof utilises the notion of adherence-closedness previously introduced in Definition 5.15 and developed in Theorem 6.16;

and the second proof uses adherence-dimension to quantify the properties of *chains of faces*'.

7.1.1. Method 1: adherence-closedness of subfaces. The first proof of topological rigidity relies on Theorem 6.20. Suppose that $f : T_x\mathcal{T}(S) \rightarrow T_y\mathcal{T}(S)$ is a Thurston norm isometry, then its dual Thurston co-norm isometry $f^* : T_y^*\mathcal{T}(S) \rightarrow T_x^*\mathcal{T}(S)$ takes each face of $\mathbf{S}_y^* := \iota_y(\mathcal{PML}(S))$ to that of $\mathbf{S}_x^* := \iota_x(\mathcal{PML}(S))$. By Theorem 5.34, f^* preserves the relation of dim, adherence, and adherence-closedness.

Let c be a weighted multicurve, and F the face for $\iota_y([c]) \in \mathbf{S}_y^*$. Due to Theorem 6.16 and Theorem 6.20, we see that $\iota_y([c])$ necessarily lies on the interior of $F_{|c|}^\wedge = F_{|c|}$. Therefore, $F_{|c|}$ is the face for $\iota_y([c])$ hence $F = F_{|c|}$, which is to say that F consists of the image of the set of projective measures supported by $|c|$. By Theorem 6.20, its image $f(F_{|c|})$ is a face for either a projective weighted multicurve or a maximal and uniquely ergodic projective lamination. In the latter case, the face $f(F_{|c|})$ is 0-dimensional and maximal. The maximality condition is preserved under (invertible) linear transformations, hence $F_{|c|}$ is also maximal. However, $[c]$ would then need to be a projectively weighted pants decomposition, and would have strictly positive dimension, as follows from Corollary 6.14, thereby contradicting the 0-dimensionality. Therefore $f(F_{|c|})$ must be a face for some projectively weighted multicurve $\iota_x([d]) \in \mathbf{S}_x^*$, and hence coincides with $F_{|d|}$.

Corollary 6.9 says that if $|c_1| \subsetneq |c_2|$ for some two projectively weighted multicurves $[c_1]$ and $[c_2]$, then we have $F_{|c_1|} \subsetneq F_{|c_2|}$. Coupling this with the fact that inclusion between faces is preserved under f^* , we see that f^* induces a simplicial isomorphism on the curve complex for S . By the work of Ivanov, Luo and Korkmaz [12, 18, 15], the map on the curve complex given by f is induced by an extended mapping class $[h]$. Since the set of multicurves is dense in $\mathcal{PML}(S)$, we see that $f_{y,x}$ coincides with the homeomorphism induced by this mapping class $[h]$.

7.1.2. Method 2: Characterising multicurves via adherence-dimension. We now give the second proof of Theorem 1.8, which uses adherence-dimension (Definition 5.21).

Theorem 7.4. *Every face F on \mathbf{S}_x^* has adherence-dimension $\leq 3g - 3 + n$, with equality if and only if F is the face for the ι_x -image of a projective weighted multicurve $[c] \in \mathcal{PML}(S)$.*

Proof. By Remark 6.12, the longest possible sequence of faces with distinct face-dimensions has corresponding face-dimensions valued from 0 to $3g - 3 + n$, and hence adherence-dimension is bounded above by $3g - 3 + n$. In particular, it is easy to verify that equality holds for any face corresponding to a projectively weighted multicurve $[c]$ by considering a projective measured lamination supported on a pants decomposition containing $|c|$. This shows the “if” part of the theorem.

Now, we turn to the “only if” part. Let F be a face for $\iota_x([\lambda])$ such that $\text{Adim}_{\mathbf{S}_x^*}(\iota_x([\lambda])) = 3g - 3 + n$. Then there must be two ascending sequence of faces

$$F_1 \subsetneq \dots \subsetneq F_k = F \text{ and } F = F_k \subsetneq \dots \subsetneq F_{3g-3+n} \subsetneq F_{3g-2+n},$$

which respectively represent the ascending sequences used to define the adherence height and the adherence depth of F . In particular, these $\{F_i\}_{i=0, \dots, 3g-2+n}$ all have distinct face-dimensions, and since dimension itself can only take value between 0 and $3g - 3 + n$ (Corollary 6.11), we see that $\dim(F_i) = \text{Fdim}(F_i) = i - 1$ for every $i = 1, \dots, 3g - 2 + n$ by Proposition 5.19. This in turn means that each F_i is adherence-closed, and by Theorem 6.16, each face takes the form $F_i = F_{|\lambda_i| = |\widehat{\lambda_i}|}$ and

$$|\widehat{\lambda_1}| \subsetneq |\widehat{\lambda_2}| \subsetneq \dots \subsetneq |\widehat{\lambda_{3g-4+n}}| \subsetneq |\widehat{\lambda_{3g-3+n}}|.$$

This is possible only if every component of λ_{3g-3+n} is either a simple closed curve or a measured lamination whose minimal supporting surface is either a torus with one hole or a sphere with four holes. Moreover, to get the length of $3g - 3 + n$, whenever a non-multicurve component ν appears in $|\widehat{\lambda_j}|$ for the first time, all the boundary components of the minimal supporting surface for ν must be already contained in $|\widehat{\lambda_{j-1}}|$.

Now suppose that $|\lambda_k| = |\lambda|$ contains a non-multicurve component ν . Then the maximal F-dim ascending chain whose first term is the face $F_{|\nu|}$ is necessarily shorter than $F_1 \subsetneq \dots \subsetneq F_k$ since its face-dimension is equal to $\dim F_{|\nu|} \geq 2$. This contradicts our assumption that $F_1 \subsetneq \dots \subsetneq F_k$ attains the adherence height of F_k . Therefore $[\lambda]$ must be a projectively weighted multicurve. \square

Theorem 5.34 says that $f^*: T_y^* \mathcal{T}(S) \rightarrow T_x^* \mathcal{T}(S)$ preserves the relation of adherence, and hence the adherence-dimension. Theorem 7.4 shows that if $F \subsetneq \mathbf{S}_y^*$ is a face for a projectively weighted multicurve, then so is $f^*(F) \subsetneq \mathbf{S}_x^*$. The rest of the proof is the same as in Method 1.

7.2. Geometric rigidity. Theorem 1.8 illustrates a form of “mapping class realisability” for Thurston-norm isometries on $\mathcal{PML}(S)$. We first consider the consequences of such a property on length-increasing cones.

Proposition 7.5. *Consider the following setup:*

- $x, y \in \mathcal{T}(S)$ and $[\lambda] \in \mathcal{PML}(S)$,
- a Thurston-norm isometries

$$f: T_x \mathcal{T}(S) \rightarrow T_y \mathcal{T}(S),$$

- and its corresponding extended mapping class $[h] \in \text{Mod}^*(S)$ (see Theorem 1.8).

Then, we have:

$$f(N_x([\lambda])) = N_y(h([\lambda])).$$

Proof. Theorem 1.8 precisely tells us that $f^*(\iota_y([\lambda])) = \iota_x(h([\lambda]))$, and hence

$$\begin{aligned}
& f(N_x([\lambda])) \\
&= f\left(\left\{v \in T_x\mathcal{T}(S) \mid \begin{array}{l} \|v\|_{\text{Th}} = 1 \text{ and } v \text{ is positive normal} \\ \text{to a support hyperplane at } \iota_x([\lambda]) \end{array} \right\}\right) \\
&= \left\{f(v) \in T_y\mathcal{T}(S) \mid \begin{array}{l} \|f(v)\|_{\text{Th}} = 1 \text{ and } f(v) \text{ is positive normal} \\ \text{to a support hyperplane at } \iota_y(h([\lambda])) \end{array} \right\} \\
&= N_y(h([\lambda])),
\end{aligned}$$

where the first equality holds by definition (Definition 5.28), and the second equality comes from the fact that f is a linear isometry and hence preserves norms and takes support hyperplanes to support hyperplanes. \square

Theorem 7.6. *Consider a sequence of projectively weighted pants decompositions $\{[\Gamma_i]\}_{i \in \mathbb{N}}$ whose sequence of supports $\{|\Gamma_i|\}_{i \in \mathbb{N}}$ converges to a maximal chain-recurrent geodesic lamination Λ in the Hausdorff topology. Then, the sequence of faces $\{N_x([\Gamma_i])\}_{i \in \mathbb{N}}$ converges to the stretch vector v_Λ .*

Proof. We show that the diameter of the faces $N_x([\Gamma_i])$, with respect to the Thurston norm, tend to 0 as $i \rightarrow \infty$. Theorem 6.31 asserts that each of these faces is $(3g - 3 + n)$ -dimensional. In particular, with respect to the Fenchel–Nielsen coordinates for $|\Gamma_i|$, the length component of the vectors $v \in N_x([\Gamma_i])$ are all identical by the definition of $N_x([\Gamma_i])$, and the vectors there differ from each other only in the $(3g - 3 + n)$ -dimensional subspace of twists. It therefore suffices to show that the diameter of the set of possible twist values $\{v(\tau_\gamma) \mid v \in N_x([\Gamma_i])\}$ tends to 0 as $i \rightarrow \infty$ for each γ in $|\Gamma_i|$; for this would ensure that $N_x([\Gamma_i])$ is trapped in a $(3g - 3 + n)$ -dimensional product of intervals with the width of each product interval shrinking to 0 as $i \rightarrow \infty$, and thereby giving us the claim that

$$\lim_{i \rightarrow \infty} \text{diam}(N_x([\Gamma_i])) = 0.$$

We shall show this by employing Theorem 4.15. To this end, we first show that any sequence of pairs of pants $\{P_i\}_{i \in \mathbb{N}}$, where P_i is a component of $S \setminus |\Gamma_i|$, is asymptotically slender (Definition 4.13). Suppose not. Then by passing to a subsequence, we may assume that none of the P_i are ϵ -slender, and hence the length of the shorter orthogeodesic on $\{P_i\}_{i \in \mathbb{N}}$ remains bounded from below by some $(2 + \epsilon)\epsilon > 0$.

Now cover each P_i with two disjoint (except along their boundaries) ideal triangles Δ_i^1, Δ_i^2 . We invoke the fact that the collection of embedded ideal triangles on S is compact [10, Prop. 3.8], and by changing indices we produce $\{\Delta_i^1\}_{i \in \mathbb{N}}$ which converges to an embedded ideal triangle Δ^1 . We can assume without loss of generality that a lift $\tilde{\Delta}^1$ of Δ^1 has ideal vertices placed at $\{0, 1, \infty\}$. Then there is a sequence of lifts $\{\tilde{\Delta}_i^1\}_{i \in \mathbb{N}}$ for $\{\Delta_i^1\}_{i \in \mathbb{N}}$ with vertices close to and limiting to $\{0, 1, \infty\}$. Taking a further subsequence and possibly applying a permutation of $\{0, 1, \infty\}$ by a Möbius transformation, we can find a sequence of lifts $\{\tilde{\Delta}_i^2\}_{i \in \mathbb{N}}$ for $\{\Delta_i^2\}_{i \in \mathbb{N}}$ so that the edge shared

between $\tilde{\Delta}_i^1$ and $\tilde{\Delta}_i^2$ tends to the geodesic joining 0 and ∞ . We denote by $\xi_i \in (0, \infty)$ the ideal vertex of $\tilde{\Delta}_i^2$ which is not shared with $\tilde{\Delta}_i^1$. The fact that P_i is not ϵ -slender means that ξ_i cannot converge to either 0 or ∞ . This in turn means that the geodesic joining 0 and ∞ is a shared edge between $\tilde{\Delta}^1$ and $\tilde{\Delta}^2$ which is disjoint from Λ . This contradicts the fact that Λ is a maximal chain-recurrent geodesic lamination because then the only possibility is that the geodesic joining 0 and ∞ is the unique isolated simple bi-infinite geodesic lying on some punctured monogon (Remark 2.9); but this is impossible since such a geodesic bounds the same ideal triangle on both sides. The same observation that governs Remark 4.7 tells us that Theorem 4.15 applies also to cases where one (or more) geodesic γ in $|\Gamma_i|$ lies in a component of $S \setminus (|\Gamma_i| \setminus \gamma)$ which is a one-holed torus. We have therefore shown that the diameter of the faces $N_x(|\Gamma_i|)$ tend to 0.

The compactness of the (Thurston-norm) unit sphere \mathbf{S}_x ensures that the sequence of faces $\{N_x(|\Gamma_i|)\}_{i \in \mathbb{N}}$ has subsequences which converge to a vector $v \in \mathbf{S}_x$. We aim to show that this limit v must be equal to v_Λ . To this end, let $\{\Lambda_i\}$ be a sequence of maximal chain-recurrent geodesic laminations respectively containing $|\Gamma_i|$, and consider the metrics $\text{stretch}(x, \Lambda_i, t) \in \mathcal{T}(S)$ (see Remark 2.9). By construction, the maximal ratio-maximising lamination $\mu(x, \text{stretch}(x, \Lambda_i, t))$ is equal to Λ_i . Lemma 2.10 shows that the sequence $\{\Lambda_i\}$ converges, and its Hausdorff limit is Λ . Therefore, [30, Theorem 8.4] tells us that every convergent subsequence (which necessarily exist due to the fact that the Thurston metric sphere of radius t around x is compact) of $\text{stretch}(x, \Lambda_i, t)$ necessarily converges to a metric $y \in \mathcal{T}(S)$ such that $\mu(x, y) = \Lambda$. Since Λ is a maximal chain-recurrent lamination, the only metric which satisfies this criterion is $y = \text{stretch}(x, \Lambda, t)$, and hence we have

$$(36) \quad \forall t \geq 0, \quad \lim_{i \rightarrow \infty} \text{stretch}(x, \Lambda_i, t) = \text{stretch}(x, \Lambda, t).$$

To finish the proof, we observe that since the diameter of $N_x(|\Gamma_i|)$ tends to 0, we need only show that an arbitrary convergent sequence of stretch vectors of the form $\{v_i \in N_x(|\Gamma_i|)\}_{i \in \mathbb{N}}$ limits to v_Λ . Eq. (36) implies that for each $k \in \mathbb{N}$, we may choose $i_k \in \mathbb{N}$ large enough so that the distance between $\text{stretch}(x, \Lambda_{i_k}, 2^{-k})$ and $\text{stretch}(x, \Lambda, 2^{-k})$, with respect to the standard Euclidean metric for any a priori fixed analytically compatible global coordinate chart on $\mathcal{T}(S)$, is less than 2^{-2k} . Then, thanks to the analyticity of the stretch map with respect to the stretching time [30, Corollary 4.2], we have:

$$\begin{aligned} x + 2^{-k} v_{i_k} + O(2^{-2k}) &= \text{stretch}(x, \Lambda_{i_k}, 2^{-k}) \\ &= \text{stretch}(x, \Lambda, 2^{-k}) + O(2^{-2k}) \\ &= x + 2^{-k} v_\Lambda + O(2^{-2k}), \end{aligned}$$

and hence

$$\lim_{i \rightarrow \infty} v_i = \lim_{k \rightarrow \infty} v_{i_k} = v_\Lambda.$$

□

Corollary 7.7 (equivariant stretch vectors). *Given an arbitrary Thurston-norm isometry $f: T_x\mathcal{T}(S) \rightarrow T_y\mathcal{T}(S)$, let $[h] \in \text{Mod}^*(S)$ (see Theorem 1.8) denote its inducing extended mapping class. Let Λ denote an arbitrary maximal chain-recurrent lamination of S , then the stretch vector $v_\Lambda \in T_x\mathcal{T}(S)$ (see Remark 2.9) satisfies the following equivariance property:*

$$f(v_\Lambda) = v_{h(\Lambda)} \in T_y\mathcal{T}(S).$$

Proof. The chain-recurrence of Λ means that it is arbitrarily closely approximated, in the Hausdorff topology, by a sequence of simple closed multicurves $\{\gamma_i\}_{i \in \mathbb{N}}$ (see Lemma 2.7). Extend each multicurve γ_i to a pants decomposition Γ_i . Lemma 2.10 shows that the $\{\Gamma_i\}_{i \in \mathbb{N}}$ has a Hausdorff limit and converges to Λ . Having satisfied the conditions for Theorem 7.6, we know that any sequence of unit vectors $\{v_i \in N_x([\Gamma_i])\}_{i \in \mathbb{N}}$ necessarily converges to v_Λ . Since f is continuous, the sequence $\{f(v_i)\}$ converges to $f(v_\Lambda)$. However, Proposition 7.5 tells us that there is a mapping class h such that

$$f(v_i) \in f(N_x([\Gamma_i])) = N_y(h([\Gamma_i])).$$

Theorem 7.6 ensures that the sequence $\{f(v_i)\}$ necessarily converges to $v_{h(\Lambda)}$. Therefore, $f(v_\Lambda) = v_{h(\Lambda)}$, as desired. □

Given an arbitrary simple closed geodesic γ on S , take α_0 to be any simple closed geodesic which has non-empty transverse intersection with γ , and define a sequence of simple closed curves $\{\alpha_m\}_{m \in \mathbb{Z}}$, where α_m is obtained from α_0 by m Dehn twists around γ . For each α_m , let $\Lambda_{+,m}$ and $\Lambda_{-,m}$ denote chain-recurrent complete geodesic laminations such that

- $\Lambda_{\pm,m}$ contain geodesic leaves which bound the same $(1,1)$ -cusped annulus (see Fig. 3) $A_m \subset S$, which in turn contains α_m ;
- $\Lambda_{+,m}$ and $\Lambda_{-,m}$ have the same geodesic leaves, except on the interior of A_m , where the two bi-infinite geodesic leaves on A_i in $\Lambda_{+,m}$ and $\Lambda_{-,m}$ which spiral toward α_m do so in opposite directions (see left-most and right-most pictures in Fig. 4).

Proposition 7.8 (length extraction). *Let $v_{\pm,m}$ denote the respective stretch vectors for the complete laminations $\Lambda_{+,m}$ and $\Lambda_{-,m}$. Then,*

$$\ell_x(\gamma) = \lim_{m \rightarrow \pm\infty} \frac{-\log \|v_{+,m} - v_{-,m}\|_{\text{Th}}}{i(\gamma, \alpha_0) \cdot |m|},$$

where $i(\gamma, \alpha_0)$ denotes the geometric intersection number between γ and α_0 .

Proof. From [30, Theorem 5.2], we know that the unit earthquake vector map $E_{(\cdot)}: \text{PML}(S) \rightarrow T_x\mathcal{T}(S)$ given by

$$[\lambda] \mapsto \frac{1}{\ell_x(\lambda)} E_\lambda|_x,$$

where λ is an arbitrary measured lamination representing $[\lambda]$ and $E_\lambda|_x$ is the earthquake vector at x corresponding to λ , defines an embedding. Since

$\mathcal{PML}(S)$ is compact and $\|E_{(\cdot)}\|_{\text{Th}}$ is a continuous map, this means that there are constants $c_1, c_2 > 0$ such that

$$(37) \quad c_1 \ell_x(\lambda) \leq \|E_\lambda|_x\|_{\text{Th}} \leq c_2 \ell_x(\lambda).$$

Then, Eq. (23), combined with Eq. (37) tells us that

$$c'_1 \ell_x(\alpha_m)^2 e^{-\ell_x(\alpha_m)} \leq \|v_{+,m} - v_{-,m}\|_{\text{Th}} \leq c'_2 \ell_x(\alpha_m)^2 e^{-\ell_x(\alpha_m)},$$

for some positive constants c'_1, c'_2 . Applying the logarithm to the above formula and using the fact that $\ell_x(\alpha_m)$ asymptotically grows as $|m|i(\gamma, \alpha_0)\ell_x(\gamma)$ as $m \rightarrow \pm\infty$ yields the desired result. \square

Theorem 7.9 (Geometric rigidity, Theorem 1.15). *Let $f: T_x\mathcal{T}(S) \rightarrow T_y\mathcal{T}(S)$ be a linear isometry as given in Theorem 1.1, and let h be a diffeomorphism representing the mapping class as in Theorem 1.8. Then for every simple closed curve $\gamma \in \mathcal{S}$ on S , we have*

$$\ell_x(\gamma) = \ell_y(h(\gamma)).$$

Proof. Define α_m , $\Lambda_{\pm,m}$ and $v_{\pm,m}$ as in the proof of Proposition 7.8. Since f is an isometry, we see that

$$\|v_{+,m} - v_{-,m}\|_{\text{Th}} = \|f(v_{+,m}) - f(v_{-,m})\|_{\text{Th}} = \|w_{+,m} - w_{-,m}\|_{\text{Th}},$$

where $w_{\pm,m} \in T_y\mathcal{T}(S)$ are stretch vectors with respect to the laminations $h(\Lambda_{\pm,m})$ on y (Corollary 7.7). Then,

$$\begin{aligned} \ell_x(\gamma) &= \lim_{m \rightarrow \infty} \frac{-\log \|v_{+,m} - v_{-,m}\|_{\text{Th}}}{i(\gamma, \alpha_0) \cdot |m|} \\ &= \lim_{m \rightarrow \infty} \frac{-\log \|w_{+,m} - w_{-,m}\|_{\text{Th}}}{i(h(\gamma), h(\alpha_0)) \cdot |m|} = \ell_y(h(\gamma)), \end{aligned}$$

as desired. \square

7.3. Infinitesimal rigidity. We now prove Theorem 1.1: the infinitesimal rigidity of the Thurston metric.

Proof of Theorem 1.1. By Theorem 1.8, for an isometry f in Theorem 1.1, there is a diffeomorphism h representing some mapping class group which induces the same homeomorphism on $\mathcal{PML}(S)$ as $f_{y,x} := \iota_x^{-1} \circ f^* \circ \iota_y$. By Theorem 1.15, for any simple closed curve γ on S , the length of $h(\gamma)$ with respect to the hyperbolic metric at y coincides with that of γ at x . Since the length functions of the simple closed curves define an embedding of $\mathcal{T}(S)$ to $\mathbb{R}^{\mathcal{S}}$, we see that h takes (S, x) to (S, y) . Since dh takes the cotangent vector corresponding to $d_y \log \ell(\gamma)$ to $d_x \log \ell(h(\gamma))$, we see that

$$dh \equiv \iota_y \circ h \circ \iota_x^{-1} \equiv f.$$

\square

There are various corollaries to Theorem 1.1, including the local rigidity (Corollary 1.17 and global rigidity (Corollary 1.18) results stated in the introduction. For example, we can strengthen Corollary 7.7 as follows:

Corollary 7.10. *Given a Thurston-norm isometry $f: T_x\mathcal{T}(S) \rightarrow T_y\mathcal{T}(S)$, consider its corresponding extended mapping class $[h] \in \text{Mod}^*(S)$ (see Theorem 1.8) and let Λ denote an arbitrary complete lamination of S . The stretch vector $v_\Lambda \in T_x\mathcal{T}(S)$ satisfies the following equivariance property*

$$f(v_\Lambda) = v_{h(\Lambda)} \in T_y\mathcal{T}(S).$$

To clarify, we have removed every assumption on Λ apart from completeness.

Corollary 7.10 implies that the notion of tangential adherence introduced in Definition 6.24 is a linear invariant. This then strengthens Proposition 6.25 as a linearly invariant characterisation of faces of the form $F_{|\lambda|} \subset \mathbf{S}^*$.

7.4. Further questions. Despite our increased understanding of the infinitesimal and global structure of Thurston's Finsler metric on Teichmüller space, there is still much that remains mysterious. We list some of the questions we have under consideration:

Question 7.11. The notion of *maximal ratio-maximising chain-recurrent lamination* [30, Theorem 8.2] and its role in Thurston's theory of stretch maps as being (intuitively speaking) “extremal” stretch maps [30, Theorem 8.4] suggests that stretch vectors with respect to maximal ratio-maximising chain-recurrent laminations should constitute extreme points of the faces of \mathbf{S}_x . This is unproven except when S is the 1-cusped torus or the 4-cusped sphere, and/or possibly as a statement which is “generically” true (and even this is not without subtleties). Is this, in fact, true?

Question 7.12. What are the dimensions, codimensions, face-dimensions and adherence-dimensions of arbitrary faces in \mathbf{S}_x and \mathbf{S}_x^* ?

Question 7.13. Corollary 6.9 can be paraphrased in a category-theoretic language as saying that $F_{(\cdot)}$ induces a covariant functor between two full subcategories $|\mathcal{ML}|$ and \mathcal{F} of the category of sets where

- the objects of $|\mathcal{ML}|$ are given by the supporting geodesic laminations $|\lambda|$ of measured laminations $\lambda \in \mathcal{ML}$, and
- the objects of \mathcal{F} are given by the faces of \mathbf{S}_x^* .

Moreover, Corollary 6.17 tells us that there is a natural transformation η from the functor $F_{(\cdot)}$ to itself that takes each object $|\lambda| \in \text{ob}(|\mathcal{ML}|)$ to the morphism $\eta_{|\lambda|}$ of inclusion: i.e., $F_{|\lambda|} \subset F_{\widehat{|\lambda|}}$. Is this structure part of some general phenomenon?

Question 7.14. On a related note to Question 7.13, it seems highly plausible that the map $N_x(\cdot)$ taking a measured lamination $\lambda \in \mathcal{ML}$ to the corresponding face should not require the information of the measure on λ and only uses x as an auxiliary metric. Philosophically speaking, then, this map should induce a contravariant functor $N_{(\cdot)}$ from $|\mathcal{ML}|$ to the full subcategory \mathcal{N} of the category of sets whose objects are given by the faces

of \mathbf{S}_x . The role of such a purported contravariant functor $N_{(\cdot)}$ is reminiscent of the theory of algebraic varieties in algebraic geometry. Assuming that such a conjectural theory for $N_{(\cdot)}$ holds, is there perhaps a comparable theory of convex “varieties” for convex geometry that applies beyond the confines of our work? And is it dual in a structured sense to what we see in Question 7.13?

REFERENCES

- [1] P. Buser, *Geometry and Spectra of Compact Riemann Surfaces*, Progress in Mathematics, vol. 106. Birkhäuser Boston Inc., Boston, 1992.
- [2] D. Calegari, *3-Manifolds*, (in preparation, to be published by) Cambridge University Press.
- [3] R. Canary, D. Epstein and P. Green, Notes on notes of Thurston, London Math. Soc. Lecture Note Ser. **328**, Fundamentals of hyperbolic geometry: selected expositions, pp. 1-115, Cambridge Univ. Press, Cambridge, 2006.
- [4] A. J. Casson and S. A. Bleiler, *Automorphisms of surfaces after Nielsen and Thurston*, No. 9., Cambridge University Press, 1988.
- [5] A. Calderon and J. Farre, Shear-shape cocycles for measured laminations and ergodic theory of the earthquake flow, arxiv:2102.13124
- [6] D. Dumas, A. Lenzhen, K. Rafi and J. Tao, Coarse and fine geometry of the Thurston metric, arxiv:1610.07409
- [7] A. Fathi, F. Laudenbach and V. Poénaru, *Travaux de Thurston sur les surfaces*, Astérisque 66-67, Société Mathématique de France, 1979. English transl. by D. Kim and D. Margalit, Mathematical Notes 48. Princeton, NJ, Princeton University Press, 2012.
- [8] R. C. Penner and J. L. Harer, *Combinatorics of Train Tracks. (AM-125)*, Volume 125. Princeton University Press, 2016.
- [9] Y. Huang, Moduli Spaces of Surfaces. PhD thesis, The University of Melbourne (2014).
- [10] Y. Huang and S. Zhe, McShane identities for higher Teichmüller theory and the Goncharov-Shen potential, arXiv:1901.02032v4 (to appear in the Memoirs of the Amer. Math. Soc.)
- [11] Y. Huang and A. Papadopoulos, Optimal Lipschitz maps on one-holed tori and the Thurston metric theory of Teichmüller space, *Geometriae Dedicata* 214 (2021), 465–488.
- [12] N. Ivanov, Automorphism of complexes of curves and of Teichmüller spaces, *Internat. Math. Res. Notices* 14 (1997), 651–666.
- [13] A. B. Katok, Invariant measures of flows on oriented surfaces, *Soviet Math. Dokl.*, 14, (1973), 1104–1108
- [14] S. Kerckhoff, Earthquakes are analytic, *Comment. Math. Helv.* 60 (1985), 17–30.
- [15] M. Korkmaz, Automorphisms of complexes of curves on punctured spheres and on punctured tori, *Topology Appl.* 95 (1999), 85–111.
- [16] M. Krein and D. Milman, On extreme points of regular convex sets, *Studia Math* 9 (1940), 133–138.
- [17] G. Levitt, Feuilletages des surfaces, *Ann. Inst. Fourier (Grenoble)* 32 (1982), 179–217.
- [18] F. Luo, Automorphisms of the complex of curves. *Topology* 39 (2000), 283–298.
- [19] H. Masur, Interval exchange transformations and measured foliations. *Annals of Mathematics* 115(1), (1982), 169–200.
- [20] S. Mazur, and S. Ulam, Sur les transformations isométriques d’espaces vectoriels normés, *CR Acad. Sci. Paris*, 1932, no. 194, 116.

- [21] K. Ohshika, Teichmüller spaces and the rigidity of mapping class group actions, to appear in *Surveys in Geometry I*, ed. by A. Papadopoulos, Springer, 2021.
- [22] H. Pan, Local Rigidity of Teichmüller space with Thurston Metric, arXiv:2005.11762 [math.GT], 2020.
- [23] A. Papadopoulos, Deux remarques sur la géométrie symplectique de l'espace des feuilletages mesurés sur une surface, *Ann. Inst. Fourier (Grenoble)* 36 (1986), 127–141.
- [24] A. Papadopoulos, Sur le bord de Thurston de l'espace de Teichmüller d'une surface non compacte, *Math. Ann.* 282 (1988), 353–359.
- [25] A. Papadopoulos, On Thurston's boundary of Teichmüller space and the extension of earthquakes. *Topology Appl.* 41 (1991), 147–177.
- [26] R. C. Penner, The decorated Teichmüller space of punctured surfaces, *Comm. Math. Phys.* 113(2) (1987), 299–339.
- [27] R. T. Rockafellar, *Convex analysis*. Princeton University Press, 2015.
- [28] H. L. Royden, Automorphisms and isometries of Teichmüller space, in: *Advances in the theory of Riemann surfaces*, Ann. of Math. Studies, vol. 66 (1971), 317–328.
- [29] G. Théret, On elementary antistretch lines, *Geom. Dedicata* 136 (2008), 79–93.
- [30] W. P. Thurston, *Minimal stretch maps between hyperbolic surfaces*, arXiv:9801039 [math.GT], 1986.
- [31] W. P. Thurston, On the geometry and dynamics of diffeomorphisms of surfaces, *Bull. Amer. Math. Soc* 19(2) (1988), 417–431.
- [32] W. P. Thurston, *Three-dimensional geometry and topology*, ed. by Sylvio Levy, vol. I, Princeton University Press, Princeton, N.J., 1997.
- [33] W. P. Thurston, *The geometry and topology of 3-manifolds*, Mimeographed Course Notes, Princeton University, 1977–78, to appear as Volume IV of Thurston's Collected Works, AMS, Vol. IV, 2022.
- [34] C. Walsh, *The horoboundary and isometry group of Thurston's Lipschitz metric*, Handbook of Teichmüller theory, ed. A. Papadopoulos, Vol. IV, Eur. Math. Soc., Zurich, 2014.
- [35] S. Wolpert, An elementary formula for the Fenchel–Nielsen twist, *Comment. Math. Helv.* 56 (1981), 132–135.

YI HUANG, JINCHUNYUAN WEST BUILDING, YAU MATHEMATICAL SCIENCES CENTER, TSINGHUA UNIVERSITY, HAIDIAN DISTRICT BEIJING 100084, CHINA. EMAIL: yihuang-math@tsinghua.edu.cn

KEN'ICHI OHSHIKA, DEPARTMENT OF MATHEMATICS, GAKUSHUIN UNIVERSITY, MEJIRO, TOSHIMA-KU, TOKYO, JAPAN. EMAIL: ohshika@math.gakushuin.ac.jp

ATHANASE PAPADOPOULOS, INSTITUT DE RECHERCHE MATHÉMATIQUE AVANCÉE (UNIVERSITÉ DE STRASBOURG ET CNRS), 7 RUE RENÉ DESCARTES, 67084 STRASBOURG CEDEX FRANCE. EMAIL: papadop@math.unistra.fr