

The Morse–Smale Property of the Thurston Spine

Ingrid IRMER

International Center for Mathematics, Department of Mathematics,
Southern University of Science and Technology, Shenzhen, P.R. China

E-mail: ingridmary@sustech.edu.cn

URL: <https://www.sustech.edu.cn/en/faculties/ingrid-irmer.html>

Received February 07, 2024, in final form April 10, 2025; Published online April 23, 2025

<https://doi.org/10.3842/SIGMA.2025.027>

Abstract. The Thurston spine consists of the subset of Teichmüller space at which the set of shortest curves, the systoles, cuts the surface into polygons. The systole function is a topological Morse function on Teichmüller space. This paper studies the local properties of the Thurston spine, and the smooth pieces out of which it is constructed. Some of these local properties are shown to have global consequences, for example that the Thurston spine satisfies properties defined in terms of the systole function analogous to that of Morse–Smale complexes of (smooth) Morse functions on compact manifolds with boundary.

Key words: Thurston spine; moduli space; handle decomposition; systole; topological Morse function; unstable manifold; mapping class group

2020 Mathematics Subject Classification: 57R99; 57M99

1 Introduction

The systoles on a closed, connected, orientable hyperbolic surface S_g of genus g are the set of shortest geodesics. This is a finite nonempty set that varies as the hyperbolic structure varies in Teichmüller space \mathcal{T}_g . The Thurston spine \mathcal{P}_g is a mapping class group-equivariant CW complex embedded in \mathcal{T}_g . Setwise it consists of the points of \mathcal{T}_g where the systoles cut the corresponding surface into polygons. In [21], a construction of a mapping class group-equivariant deformation retraction of \mathcal{T}_g onto \mathcal{P}_g is given. Although Morse theory was not explicitly used, the construction is analogous to the deformation retraction of a compact manifold with boundary onto the Morse–Smale complex defined by a Morse function.

A key role is played by the systole function $f_{\text{sys}}: \mathcal{T}_g \rightarrow \mathbb{R}_+$ taking a point in \mathcal{T}_g to the length of the systoles. The systole function is a piecewise-smooth topological Morse function, [1, 19].

The purpose of this paper is to relate the structure of \mathcal{P}_g to that of the equivariant handle decomposition of \mathcal{T}_g coming from f_{sys} . This is done by proving three results, outlined below.

To begin with, all critical points of f_{sys} are contained in \mathcal{P}_g [21]. As f_{sys} is not smooth, the notion of gradient of f_{sys} at x is replaced by a “cone of increase” consisting of the elements of $T_x \mathcal{T}_g$ specifying directions in which f_{sys} is increasing. Similarly, as \mathcal{P}_g is not a manifold, the notion of tangent space to \mathcal{P}_g at a point x is replaced by the “tangent cone” to \mathcal{P}_g at x . The following theorem states informally that restricting f_{sys} to \mathcal{P}_g does not create critical points.

Theorem 1.1. *Suppose a point $x \in \mathcal{P}_g$ has the property that the cone of increase of f_{sys} at x is full. Then the intersection of this cone with the tangent cone of \mathcal{P}_g is nonempty.*

In other words, at a point at which f_{sys} is increasing in some direction to first order, there exists a direction in the tangent cone to \mathcal{P}_g in which the lengths of all the systoles are increasing.

The next two results ensure that the Thurston spine contains a choice of unstable manifolds¹ of the critical points of f_{sys} .

Theorem 1.2. *Where a stratum of \mathcal{P}_g is locally top-dimensional it is balanced. If $x \in \mathcal{P}_g$ is not balanced, there are points directly above x in \mathcal{P}_g that are balanced.*

Technical terms, such as stratum, tangent cone and cone of increase, are defined in Section 2. A balanced stratum has properties in common with an unstable manifold of the systole function, in the way the systole function increases towards such a stratum.

Proposition 1.3. *Fix a critical point $p \in \mathcal{P}_g$ of index j with set of systoles C . A neighbourhood of p in \mathcal{P}_g contains a piecewise smooth cell of codimension j on which f_{sys} is increasing to second order away from the critical point. The tangent space of this cell at p has codimension j and is given by $\{\nabla L(c)(p) \mid c \in C\}^\perp$. The subspace spanned by $\{\nabla L(c)(p) \mid c \in C\}$ consists of directions in which the systole function is decreasing away from p .*

The proofs given in this paper do not assume Thurston’s equivariant deformation retraction of \mathcal{T}_g onto \mathcal{P}_g , and are elementary in the sense that they only require Thurston’s Lipschitz map construction from [21], Wolpert’s formula for Dehn twists, convexity of length functions, calculus and linear algebra. They will be used to prove geometric properties of gradients of length functions, to justify part of the construction of [21], and in forthcoming papers to relate geometric, topological, analytic, algebraic and number theoretic properties of Teichmüller space and the mapping class group.

As f_{sys} is not smooth, there can potentially exist regular points at which there is no open cone of directions in which the systole function is increasing. These points are called *boundary points* of the systole function. A geometric model for \mathcal{P}_g around critical points and boundary points of f_{sys} is given in Section 5 and an explicit example of a critical point is worked through in Section 7. Whether or not boundary points actually exist in Teichmüller space of surfaces without boundary is not known. An example of a boundary point of f_{sys} for a surface with boundary is given in [19, p. 434, point (ii)].

The Morse theory used in this paper is not the usual (smooth) Morse theory, but the Morse theory of so-called *topological Morse functions* (defined in Section 2) developed in [15]. It was shown in [1] that f_{sys} is a topological Morse function. The critical points necessarily occur where f_{sys} is not smooth. Studying the pieces of the piecewise smooth structure is a key ingredient of most of the proofs in this paper. These pieces, the strata, are contained in so-called “loci”. These are solutions to the set of constraints that the lengths of curves in a given set are equal, or differ by a fixed amount encoded in a tuple d . Of particular interest are the loci described by sets of curves that do not contain any proper filling subsets; “minimal filling sets”.

Corollary 1.4. *Let C be a minimal filling set of curves. When the locus $E(C, d)$ is nonempty, it is a connected, embedded submanifold of \mathcal{T}_g . The restriction to $E(C, d)$ of the lengths of the curves in C is strictly convex and has a unique minimum.*

Typically, locally top-dimensional strata of \mathcal{P}_g are open sets of loci described by Corollary 1.4, and important properties of such loci are shown to generalise to locally top-dimensional strata of \mathcal{P}_g .

The reference to Morse–Smale in the title comes from a CW-complex embedded in a manifold with boundary that is often referred to as the “Morse–Smale complex” in computer science literature. This complex is obtained as the image of a deformation retraction constructed

¹As the systole function is only a topological Morse function, the stable/unstable manifolds are not uniquely defined. This is one reason for wanting to work with the Thurston spine instead of the unstable manifolds. The notion of stable manifolds will be replaced by Schmütz’s notion of sets of minima, as made rigorous in an upcoming paper [9].

using the flow of the gradient vector field of a Morse function. The cells of the complex consist of the unstable manifolds of the critical points. This CW-complex should not be confused with the chain complex of the same name used for computing homology of the manifold. An analogue of the Morse–Smale complex for topological Morse functions is described in [15].

Organisation of the paper. In an effort to be self-contained, a fairly detailed Section 2 is given, surveying the basic properties of the objects studied in this paper and introducing the notation and assumptions that will be used. This includes a subsection on Schmutz’s sets of minima; a construction that is perhaps not as well-known in the literature as it deserves to be. In order to understand the properties of the stratification of \mathcal{T}_g , the topological properties of the loci are studied in Section 3. An important concept appears to be that of a balanced locus, which is introduced in Section 3.2. Section 4 proves the two theorems stated at the beginning of the introduction, and similar properties are established for critical and boundary points of f_{sys} in Section 5. The implications of the theorems in Sections 4 and 5 for the structure of a vector field studied by Thurston are discussed in Section 6. In Section 7, Proposition 1.3 is applied to an important example of a critical point due to Schmutz.

2 Definitions and conventions

The unique closed topological surface of genus g for $g \geq 2$ will be denoted by S_g . When S_g has the additional structure of a marked hyperbolic structure, i.e., it represents a point of Teichmüller space \mathcal{T}_g of S_g , it will be denoted by S_g .

Curves on S_g are nontrivial isotopy classes of embeddings of S^1 into S_g , where S^1 is the unpointed, unoriented 1-sphere. A curve on S_g is assumed to inherit a marking, and by abuse of notation, the word curve on S_g will often be confused with the image of its unique geodesic representative. The term geodesic will always refer to a curve except in the case of a geodesic arc. The set of curves on S_g will be denoted by \mathcal{C}_g . A curve will usually be denoted by lowercase c with a subscript, while uppercase C is reserved for finite sets of curves. A geodesic on S_g of smallest length will be called a *systole*. The *systole function* f_{sys} is the piecewise smooth function $\mathcal{T}_g \rightarrow \mathbb{R}_+$ whose value at a point x is given by the length of the systoles. Throughout the paper, it will always be assumed that the elements of a set C of curves have pairwise geometric intersection number at most one, mirroring a property of sets of systoles.

A set of curves on S_g *fills* the surface if the complement of a set of representatives in minimal position consists of a set of topological disks. The Thurston spine \mathcal{P}_g is therefore the set of all points in \mathcal{T}_g at which the systoles fill.

In this paper, a cell is homeomorphic to an open ball. The standard definition of a cell in a CW complex is closed; these will be understood to be the closures of the cells discussed in this paper. By abuse of notation, vertices, edges and faces of a cell will be used to refer to the vertices, edges and faces of the closure of the cell.

The symbol \subset will be used to denote subset. For example, $\{1\} \subset \{1, 2\}$ and $\{1, 2\} \subset \{1, 2\}$, whereas \subsetneq is used to rule out the possibility that the sets are equal.

Length functions. Length functions generalise the notion of the length of a curve. A curve c on S_g determines an analytic function $L(c): \mathcal{T}_g \rightarrow \mathbb{R}_+$ whose value at a point x is equal to the length of the geodesic representative of c at x . For an ordered set of curves $C = (c_1, \dots, c_k)$ and a corresponding set of real, positive weights $A = (a_1, \dots, a_k)$, the length function $L(A, C): \mathcal{T}_g \rightarrow \mathbb{R}_+$ is the analytic function given by

$$L(A, C)(x) = \sum_{j=1}^k a_j L(c_j)(x).$$

When all the weights in A are equal to 1, the map $L(A, C)$ from \mathcal{T}_g to \mathbb{R}^C will be written $L(C)$ or $L(c)$ when $C = \{c\}$. Length functions are so useful because they satisfy many convexity properties, as shown in [23] for the Weil–Petersson metric, in [10] for earthquake paths and [5] for Fenchel–Nielsen coordinates.

The *sublevel set* of $L(c)$ is the set of all points of \mathcal{T}_g for which $L(c)$ is less than or equal to the constant $l \in \mathbb{R}_+$, and will be denoted by $L(c)_{\leq l}$. The boundary of $L(c)_{\leq l}$ is then the level set $L(c)^{-1}(l)$. Similarly, the *superlevel set* $\{x \in \mathcal{T}_g \mid L(c)(x) \geq l\}$ will be denoted by $L(c)_{\geq l}$.

Local finiteness. The lengths of elements of \mathcal{C} is known to satisfy a property that will be referred to as local finiteness. It can be proven from the collar lemma that for any $x \in \mathcal{T}_g$, the set of curves of length within ϵ of $f_{\text{sys}}(x)$ is finite. Consequently, there is a neighbourhood of $x \in \mathcal{T}_g$ in which the systoles are contained in the set of curves that are realised as systoles at x .

Definition 2.1 (systole stratum). A *systole stratum* $\text{Sys}(C)$ is the subset of \mathcal{T}_g on which C is the set of systoles.

Here the term stratification is used in a minimal sense, following Thurston, to mean a disjoint union of locally closed subsets, for which each point is on the boundary of only finitely many of these locally closed subsets (strata).

The existence of a triangulation of \mathcal{T}_g compatible with the stratification follows from [12]. A triangulation of this type is assumed when referring to the boundary of a stratum, or the tangent vector or 1-sided limit of tangent vectors to strata of \mathcal{P}_g .

At first sight it might seem surprising that \mathcal{P}_g is nonempty. This is a consequence of the following proposition, and of the theorem due to Bers’ [4] that f_{sys} is bounded from above.

Proposition 2.2 ([21, Proposition 0.1], [3, Lemma 4]). *Let C be a set of curves on \mathcal{S}_g that do not fill. Then at any point x of \mathcal{T}_g , there is an open cone of derivations in $T_x \mathcal{T}_g$ whose evaluation on every element of $\{L(c) \mid c \in C\}$ is strictly positive at x .*

The following is a short proof of Proposition 2.2 suggested by an anonymous referee.

Proof. Since C does not fill, there exists a curve c^* disjoint from all the curves in C . Dehn twisting around c^* takes x to a different point x' without changing the lengths of curves in C . Since length functions are convex along Weil–Petersson geodesics [23], the tangent vector v at x to the Weil–Petersson geodesic from x to x' is a direction in which the lengths of all the curves in C are decreasing to first order. It follows that v is contained in an open cone of vectors satisfying the proposition. ■

Twist flow. In addition to a length function $L(c)$, a curve c on \mathcal{S}_g determines a flow $\mathbb{R} \times \mathcal{T}_g \rightarrow \mathcal{T}_g$. Choose any set of Fenchel–Nielsen coordinates on \mathcal{T}_g defined using a set of curves containing c . Then a point (r, x) is mapped to the point in \mathcal{T}_g with Fenchel–Nielsen coordinates equal to the coordinates at x , except for the twist parameter around c , which is incremented by r . A flowline will be called a *twist path* and denoted by $\gamma_{c,x}$, where $x = \gamma_{c,x}(0)$. A twist path is a special case of an earthquake path, for which it is known that length functions are convex [10].

This flow will be implicit in many of the arguments in this paper, for the simple reason that for any c_1 and c_2 disjoint, c_1 determines a flow that leaves the level sets of $L(c_2)$ invariant.

Topological Morse functions. Suppose M is an n -dimensional topological manifold. A continuous function $f: M \rightarrow \mathbb{R}_+$ will be called a *topological Morse function* if the points of M are either regular or critical points with respect to f . A point $x \in M$ is regular if there is an open neighbourhood U containing x , for which U admits a homeomorphic parametrisation by n parameters, one of them being f . If the point p is a critical point, there exists

a $j \in \mathbb{Z}$, $0 \leq j \leq n$, called the *index* of p , and a homeomorphic parametrisation of U by parameters $\{x_1, \dots, x_n\}$ such that for every point in U

$$f(x) - f(p) = \sum_{i=1}^{i=n-j} x_i^2 - \sum_{i=n-j+1}^{i=n} x_i^2.$$

A reference for topological Morse functions is [15].

Whenever a metric is needed in this paper, the Weil–Petersson metric will be assumed, and denoted by \mathbf{g} . None of the important concepts are dependent on a specific choice of mapping class group-equivariant metric; in most instances, any other choice would work just as well.

The orthogonal complement of a set of gradients of length functions, $\{\nabla L(c) \mid c \in C\}$ will be denoted by $\{\nabla L(c) \mid c \in C\}^\perp$. Although a metric is needed to define gradients and orthogonal complements, $\{\nabla L(c) \mid c \in C\}^\perp$ is independent of a choice of metric. This is because it can be defined as a tangent space to the intersection of level sets of length functions.

Tangent cones and cones of increase. There are many places throughout this paper at which objects similar to polyhedra are used. For a polyhedron P in \mathbb{R}^d , the *tangent cone* at a point $x \in \partial P$ is the set of $v \in T_x \mathbb{R}^d$ such that $v = \dot{\gamma}(0)$ for a smooth oriented path $\gamma(t)$ with $\gamma(0) = x$ and $\gamma(\epsilon)$ in P for sufficiently small $\epsilon > 0$. For background and definitions relating to polyhedra, see [6] or [2].

As \mathcal{P}_g can be triangulated, the tangent cone of \mathcal{P}_g at a point x can be defined analogously. When x is on the boundary of more than one simplex of \mathcal{P}_g , the tangent cone is the union of the tangent cones of the simplices with x on the boundary. Tangent cones of strata are defined analogously. For a triangulation compatible with the stratification, the tangent cone to $\text{Sys}(C)$ at x is the union of the tangent cones of the simplices with interior contained in $\text{Sys}(C)$ with x on the boundary.

Since f_{sys} is only piecewise smooth, it does not have a gradient at every point. The level sets of f_{sys} are a locally finite intersection of level sets of curves in \mathcal{T}_g . On a neighbourhood of $x \in \mathcal{T}_g$ at which C is the set of systoles, by local finiteness the level set of f_{sys} passing through x lies on the boundary of the intersection $I(C, x)$ defined as follows $I(C, x) := \bigcap_{c \in C} L(c)_{\geq f_{\text{sys}}(x)}$. Each of these sublevel sets is linear with respect to some set of Fenchel–Nielsen coordinates, but as the systoles can intersect, $I(C, x)$ is not always an intersection of half-spaces linear with respect to the same set of coordinates. Nevertheless, it is possible to define a tangent cone in the same way as for a polyhedron, and the *cone of increase* of f_{sys} at a point x is defined to be the tangent cone of $I(C, x)$.

The cone of increase of f_{sys} will be called *full* if it has dimension equal to that of \mathcal{T}_g , namely $6g - 6$. Since f_{sys} is only a topological, but not smooth, Morse function, it is possible that there exist regular points at which the cone of increase of f_{sys} is not full. If x is one such point, then f_{sys} can increase away from x to at most second order. Following Schmutz, these points will be called *boundary points of f_{sys}* . It was shown in [19] that the set of critical and boundary points of f_{sys} are a discrete subset of \mathcal{P}_g . An alternative definition will be given in Section 2.1, and these points will be studied in Section 5.

Wherever it is nonempty, $\partial I(C, x)$ is a piecewise smooth, codimension 1 submanifold of \mathcal{T}_g , and as such, it makes sense to speak of $y \neq x$ on $\partial I(C, x)$. A point y on $\partial I(C, x)$ will be called a *vertex* of $I(C, x)$ if the tangent cone to $I(C, x)$ at y contains no vector subspaces of $T_y \mathcal{T}_g$. Otherwise, y will be said to be on a *face*, where the face corresponds to a maximal vector subspace of the tangent cone of $I(C, x)$ at y , and two such subspaces at points y_1 and y_2 determine the same face if y_1 and y_2 are connected by a path in $\partial I(C, x)$ along which the vector subspace has fixed dimension and varies smoothly. When the face has codimension 1 in \mathcal{T}_g , it will be called a *facet*.

When $I(C, x)$ is replaced by the intersection $D(C, t) := \bigcap_{c \in C} L(c)_{\leq t}$ tangent cones, facets, faces and vertices are defined analogously.




Figure 1. Thurston's Lipschitz map construction, where m is shown in black and A in blue in the upper leftmost figure.

The fact that f_{sys} is a topological Morse function [1] implies that the set of unit tangent vectors in the cone of increase at a regular point of f_{sys} is a topological ball [15]. Removing the zero vector from the cone of increase does not give a disconnected set.

Eutactic. In the literature, a point $x \in \mathcal{T}_g$ in $\text{Sys}(C)$ is defined to be *eutactic* if for every derivation v in $T_x \mathcal{T}_g$, either the evaluations $\{vL(c)(x) \mid c \in C\}$ are all zero, or there is at least one $c \in C$ for which $vL(c)(x)$ is strictly positive and another for which it is strictly negative. In this paper the term eutactic will be generalised. A set of curves C (not necessarily systoles) will be called eutactic at a point x in \mathcal{T}_g if for every derivation v in $T_x \mathcal{T}_g$, either the evaluations $\{vL(c)(x) \mid c \in C\}$ are all zero, or there is at least one $c \in C$ for which $vL(c)(x)$ is strictly positive and another for which it is strictly negative. Similarly, for a fixed subspace $V \in T_x \mathcal{T}_g$, the set of curves C will be called V -eutactic at x if for every derivation v in V , either the evaluations $\{vL(c)(x) \mid c \in C\}$ are all zero, or there is at least one $c \in C$ for which $vL(c)(x)$ is strictly positive and another for which it is strictly negative.

Thurston's Lipschitz map construction. A construction of a map $\mathcal{T}_g \rightarrow \mathcal{T}_g$ will now be surveyed. This is explained in detail in [16]. The input of this construction is (m, A) , where m is a geodesic multicurve on S_g and A is a set of pairwise disjoint geodesic arcs with endpoints on m , meeting m at right angles, and with the property that A has nonzero geometric intersection number with every boundary component of a tubular neighbourhood of m . The output of the construction is a family of maps $\mathcal{T}_g \rightarrow \mathcal{T}_g$ which are used to control the change in the lengths of some curves relative to others. This construction was used in [21] to prove Proposition 2.2, with m isotopic to the boundary of the subsurface filled by C , and A a set of arcs that intersect every curve in C .

Start with a surface S_g representing a point of \mathcal{T}_g . Cut S_g along the geodesic representative of m to get a union of hyperbolic surfaces with boundary. On every boundary component of each of these surfaces, attach a hyperbolic annulus, to obtain a surface without boundary, as shown in Figure 1. Each of the geodesic representatives of arcs in A can be extended to a bi-infinite

geodesic in these extended surfaces. Cut the extended surfaces along each bi-infinite geodesic, and glue in a strip of the hyperbolic plane bounded by two infinite geodesics. The widths of the geodesic strips is chosen in such a way that for every curve c in m , the curves c_L and c_R in the extended surfaces corresponding to c have the same length after the strips have been glued in. For each curve c in m , cut the extended surfaces along the geodesic representatives of c_L and c_R , and glue the surface together along the resulting boundary components. This gives a new hyperbolic structure. Clearly there are many choices involved in this construction.

In this paper, as in [21], limits are taken as the widths of the strips are allowed to approach zero. The resulting 1-parameter families of maps do not change the length of any curve disjoint from the geodesic representatives of the elements of m and A . If the widths are chosen to increase to first order with the parameter, the tangent vectors to these 1-parameter families lengthen to first order every curve in m and every curve with geodesic representative disjoint from m that intersects one or more of the arcs in A .

2.1 Sets of minima

As an attempt at parameterising \mathcal{T}_g without reference to punctures or marked points, Schmutz introduced the concept of cell decompositions of Teichmüller space parametrised by length functions. Schmutz’s construction will be surveyed here as it provides insight into the differential topology of \mathcal{T}_g .

A length function $L(A, C)$ with fixed $A \in \mathbb{R}_+^{|C|}$ is convex with respect to the Weil–Petersson metric [23]. A necessary and sufficient condition for x to represent the minimum of $L(A, C)$ in \mathcal{T}_g is therefore that the gradient of $L(A, C)$ is zero at x . When C is eutactic at x , there exists an A for which the unique minimum of $L(A, C)$ occurs at x . If C is not a filling set, Proposition 2.2 ensures that $L(A, C)$ has no minimum in the interior of \mathcal{T}_g .

Definition 2.3 ($\text{Min}(C)$). The subset $\text{Min}(C)$ of \mathcal{T}_g is the set of all points at which $L(A, C)$ has a minimum for some A with strictly positive entries. Alternatively, $\text{Min}(C)$ is the set of all $x \in \mathcal{T}_g$ at which C is eutactic.

In Section 3, minima of sets of curves will be defined. This should not be confused with minima of length functions in the definition of $\text{Min}(C)$.

There is a notion of closure of $\text{Min}(C)$ that will be denoted by $\overline{\text{Min}(C)}$.

Definition 2.4 (modified version of [20, Proposition 2]). Let C be a finite set of closed geodesics that fills S_g . A point x of \mathcal{T}_g is in $\overline{\text{Min}(C)}$ if and only if there does not exist a derivation in $T_x\mathcal{T}_g$ whose evaluation on each length function of a curve in C is strictly positive. Then $\partial\text{Min}(C)$ is defined to be the set of points in $\overline{\text{Min}(C)} \setminus \text{Min}(C)$.

Lemma 2.5 (consequence of [19, Lemma 14 and Theorem 15]). *Any point $x \in \mathcal{T}_g$ in $\partial\text{Min}(C)$ is in $\text{Min}(C_i)$, for some filling $C_i \subsetneq C$. If C_i is the largest subset of C that is eutactic at x , there exists a derivation v in $T_x\mathcal{T}_g$ for which $vL(c)(x) = 0$ for all $c \in C_i$ and $vL(c)(x) < 0$ for all $c \in C \setminus C_i$.*

Lemma 1 of [19] showed that when C fills, any choice of A with strictly positive entries determines a length function $L(A, C)$ with a minimum realised somewhere in \mathcal{T}_g . It therefore follows from Proposition 2.2 that $\text{Min}(C)$ is nonempty if and only if C is a filling set of curves. There exists therefore a surjective map $\phi_C: \mathbb{R}_+^{|C|} \rightarrow \text{Min}(C)$, where $\phi_C(A)$ is the point in $\text{Min}(C)$ at which $L(A, C)$ has its minimum. Since ϕ_C is not injective, the entries of the k -tuple A are parameters, not coordinates. For $C = \{c_1, \dots, c_k\}$, let $F(C): \mathcal{T}_g \rightarrow \mathbb{R}_+^k$ be the differentiable function $x \mapsto (L(c_1), \dots, L(c_k))$. When the rank of $F(C)$ is constant on $\text{Min}(C)$, $\text{Min}(C)$ is a cell, otherwise it is “pinched” at places where the rank drops.

For a filling set C , some points on the boundary of \mathbb{R}_+^k also determine length functions with minima in \mathcal{T}_g ; Schmutz called these admissible boundary points. The union of \mathbb{R}_+^k with the admissible boundary points will be denoted by A_C . In [19, Lemma 2], it was shown that there exists an open set B_C in \mathbb{R}^k containing A_C , with the property that every tuple (b_1, \dots, b_k) in B_C gives a function $b_1 L(c_1) + \dots + b_k L(c_k)$ with a unique minimum in \mathcal{T}_g . The image under ϕ_C of $B_C \setminus A_C$ will be denoted by $\text{Conv}(C)$, where “Conv” stands for convex.

Boundary points of the systole function. It was shown in [1] that a critical point of f_{sys} occurs exactly where, for some set C , $\text{Sys}(C)$ intersects $\text{Min}(C)$. Schmutz defined a boundary point of f_{sys} to be a point at which, for some C , $\text{Sys}(C)$ intersects $\partial \text{Min}(C)$. By Lemma 2.5, this is a point at which $\text{Sys}(C)$ intersects $\text{Min}(C')$, for $C' \subsetneq C$. There are reasons for calling such points boundary points of f_{sys} , but boundary points of f_{sys} should not be confused with the usual topological notion of boundary. The structure of critical and boundary points of f_{sys} will be discussed in Section 5. By [19, Corollary 22, Theorem 23], critical points and boundary points of f_{sys} are isolated, so there are only finitely many of them modulo the action of the mapping class group.

3 Minimal filling sets of curves

Suppose C is a filling set of curves on S_g , any two of which intersect in at most one point. The set C will be called a *minimal filling set of curves* if no proper subset of C fills. This section begins by proving some key lemmas about minimal filling sets of curves, and ends by discussing a property that generalises many of the important properties of minimal filling sets of curves.

A basic property of minimal filling sets of curves is the following, which will be assumed from now on without comment.

Lemma 3.1. *Suppose C is a minimal filling set of curves. For any curve $c \in C$, there is a curve c^* that intersects c but does not intersect any of the other curves in C .*

Proof. Suppose every curve on S_g that intersects $c \in C$ also intersects a curve in $C \setminus \{c\}$. Then $C \setminus \{c\}$ fills, contradicting the assumption that C is a minimal filling set. ■

All known examples of locally top-dimensional strata of \mathcal{P}_g have systoles that are minimal filling sets. The difficulty in proving this is always the case is due to a very specific configuration of linearly dependent length functions discussed in Remark 4.7. It will be shown that loci of minimal filling sets of curves have beautiful properties that generalise to locally top-dimensional strata of \mathcal{P}_g .

3.1 Key lemmas

This subsection proves some key lemmas about minimal filling sets of curves and their loci, which will now be defined.

Definition 3.2 (locus $E(C)$ labelled by C). The locus of points in \mathcal{T}_g at which the lengths of all curves in C are equal will be denoted by $E(C)$, and called the locus labelled by C . Alternatively, $E(C) = \{x \in \mathcal{T}_g \mid L(c)(x) = L(c')(x) \forall c, c' \in C\}$.

Note that $\text{Sys}(C) \subset E(C)$.

The set $E(C)$ is called a locus to emphasise the similarity with the locus of points equidistant from a finite set of points $\{p_1, \dots, p_k\}$ in \mathbb{R}^n . The analogy is drawn by replacing the set of points in \mathbb{R}^n by points in the completion of \mathcal{T}_g with respect to the Weil–Petersson metric, where p_i for $i = 1, \dots, k$ is a point at which the curve $c_i \in C$ has been pinched. Note that by convexity, $L(c_i)$ increases monotonically along a Weil–Petersson geodesic emerging from p_i .

Recall that \mathcal{C} is the set of curves on \mathcal{S}_g . Denote by $F(\mathcal{C})$ the map from \mathcal{T}_g to $\mathbb{R}^{\mathcal{C}}$ taking x to the collection of lengths $\{L(c)(x) \mid c \in \mathcal{C}\}$, and by $\overline{F}(\mathcal{C}): \mathcal{T}_g \rightarrow \mathbb{R}^{\mathcal{C}}/\mathbb{R}$ the composition of this map with the projection from $\mathbb{R}^{\mathcal{C}}$ (arbitrary maps from \mathcal{C} to \mathbb{R}) to its quotient by \mathbb{R} (the constant maps). Fix C and $d \in \mathbb{R}^{\mathcal{C}}/\mathbb{R}$. Then the locus $E(C, d)$ is defined to be the set of points in \mathcal{T}_g for which the lengths of the curves in C realise the specified differences in $d|_C$, in other words

$$E(C, d) = \{x \in \mathcal{T}_g \mid L(c)(x) + d|_c = L(c')(x) + d|_{c'} \forall c, c' \in C\}. \quad (3.1)$$

When every element of $d|_C$ is zero, d will be omitted.

The next lemma will be used in proving results about loci of minimal filling sets of curves.

Lemma 3.3. *For a filling set C of curves, suppose $C \setminus \{c\}$ is not filling for some $c \in C$. Then for any $x \in \mathcal{T}_g$, there exists a derivation $v \in T_x \mathcal{T}_g$ such that $vL(c)(x) < 0$ and $vL(c')(x) > 0$ for all $c' \in C \setminus \{c\}$.*

Definition 3.4 (C -generic). Suppose C is a (not necessarily proper) subset of a minimal filling set of curves. A point $x \in \mathcal{T}_g$ is C -generic if for every $c \in C$, the sum of the cosines of the angles of intersection of c with c^* is nonzero, for some choice of c^* . Alternatively, the length function $L(c)$ is strictly convex along the twist path $\gamma_{c^*, x}$ and hence has a unique minimum. The point x is C -generic if for every $c \in C$, x is not the unique minimum of $L(c)$ on $\gamma_{c^*, x}$.

Proof. Wolpert’s twist formula [22] states that the rate of change of the length of $L(c)$ with respect to the twist parameter around c^* is equal to the sum of the cosines of the angles of intersection of c with c^* . For points of \mathcal{T}_g that are $\{c\}$ -generic, the lemma follows from Wolpert’s twist formula and Proposition 2.2. For points that are not $\{c\}$ -generic, the proof uses Thurston’s Lipschitz map construction from Section 2. The geodesic multicurve m is given by c^* , where, as usual, c^* is a curve intersecting c and disjoint from any curve in $C \setminus \{c\}$. It will now be explained how to obtain the set of arcs A .

If c intersects c^* just once, and the cosine of the Weil–Petersson angle of intersection is zero, then $c \cap (S_g \setminus c^*)$ is an arc with endpoints on c^* making an angle $\pi/2$ with c^* . Choose this arc to be an element of A and add more arcs to the set, subject to the constraints in the definition of A , until each curve in the set $C \setminus \{c\}$ intersects at least one arc in A . As explained below, this choice of m and A then determines a 1-parameter family of maps, with the property that the tangent to this 1-parameter family gives a derivation that is strictly negative when evaluated on $L(c)$ and strictly positive when evaluated on the length of any curve in $C \setminus \{c\}$.

When c intersects c^* more than once, it is necessary to keep in mind that the intersections between the arcs in A and the curve c that govern whether or not a given Lipschitz map increases or decreases the length of c are not the usual topological intersections, but intersections between the geodesic representative of the curve c , and the geodesic representatives of the homotopy classes of arcs relative to the geodesic representative of c^* , i.e., geodesic arcs orthogonal at their endpoints to the geodesic c^* .

Choose an orientation for c^* . This makes it possible to define a “right side” \mathcal{N}_R of c^* in a tubular neighbourhood \mathcal{N} of c^* and a “left side” \mathcal{N}_L of c^* in \mathcal{N} . Moving around c^* and measuring angles to the right of c^* , there must be an arc e_1 of $c^* \cap (S_g \setminus \{c\})$ at whose endpoints the signs of the cosines of the angles of intersection either go from $+$ to $-$, 0 to $-$, $+$ to 0 or 0 to 0 . Here the angles to the right of c^* at the endpoints of the arc e_1 are taken to be the angles at the vertices of the connected component of $S_g \setminus \{c, c^*\}$ to the right of e_1 . In Figure 2(a), these are the angles $\pi - \theta_1$ and θ_2 . Similarly, moving around c^* and measuring angles to the left of c^* , there must be an arc e_2 at which the signs of the cosines of the angles of intersection either go from $+$ to $-$, 0 to $-$, $+$ to 0 or 0 to 0 . The arc e_1 joins two adjacent vertices of a connected component of $S_g \setminus \{c, c^*\}$ to the right of e_1 , for which the inner angles at

the two vertices are both at most $\frac{\pi}{2}$; see Figure 2 (a). The arc e_2 does the same for a connected component of $S_g \setminus \{c, c^*\}$ to the left of e_2 . These arcs will be used to construct arcs on both sides of the geodesic c^* , called central arcs, whose geodesic representatives are disjoint from the geodesic c .




Figure 2. In this figure, the blue arcs lie along the geodesic c^* and the red arcs lie along the geodesic c in the hyperbolic surface S_g . In (a), the green arc is a central arc disjoint from c that only intersects \mathcal{N} in \mathcal{N}_R , i.e., it is to the right of c^* . In (b), when e_1 and e_2 are contained in the same cycle, the central arcs are constructed as shown in green. These central arcs intersect both \mathcal{N}_R and \mathcal{N}_L .

A set of cycles $A_i := \{a'_1, \dots, a'_m\}$ is obtained, consisting of the connected components of the boundary of a normal neighbourhood of the union of the geodesics c and c^* . By choosing c^* to have the smallest possible geometric intersection number with c , it is possible to assume without loss of generality that each element of A_i is homotopically nontrivial. This is because the curves in the set C intersect pairwise at most once, and c is the only curve in C to intersect c^* . A curve in $C \setminus \{c\}$ that intersects a homotopically trivial element a'_k of A_i must intersect c at least twice, which is not possible. When a'_k is homotopically trivial, the curve c^* can then be replaced by a homotopically nontrivial connected component of the boundary of a normal neighbourhood of $c^* \cup a'_k$ that intersects c .

Each cycle in A_i has edges alternating between edges lying along c and c^* . Suppose e_1 and e_2 are not contained in the same cycle. From each cycle a'_l , remove a single edge lying along c^* to obtain an arc a_l with endpoints on c^* . Whenever possible, the arc removed should be one of the arcs e_1 or e_2 . This is shown in Figure 2 (a). An arc a_l obtained by deleting e_1 or e_2 will be called *central*. When e_1 and e_2 are contained in the same cycle, a pair of central arcs are constructed as shown in Figure 2 (b). A central arc has the property that the shortest representative of its homotopy class with endpoints on c^* is disjoint from the geodesic representative of c .

Choose a set A of pairwise disjoint arcs as above, containing $\{a_1, \dots, a_m\}$.

If A only contained central arcs, it follows that the resulting Lipschitz map construction gives a direction in which $L(c)$ is strictly decreasing and the lengths of curves in $C \setminus \{c\}$ are increasing. That the lengths of curves in $C \setminus \{c\}$ are increasing in this direction is a property of Thurston's Lipschitz map construction. Trigonometry implies that $L(c)$ decreases in this direction, as follows: Gluing in strips corresponding to the central arcs, as shown in Figure 3, takes the geodesic c^* to a piecewise geodesic, with angles less than π on the side of c^* that is not the flaring end. The geodesic representative of the piecewise geodesic is obtained by shifting the piecewise geodesic to the side with angles less than π , thereby decreasing $L(c)$.

A Lipschitz map constructed from m and A , for which the widths of the strips corresponding to the central arcs are sufficiently large relative to the widths of any strips corresponding to arcs that are not central, gives a 1-parameter family of maps whose tangent vector at x corresponds to a derivation that is strictly negative when evaluated on $L(c)$ and strictly positive when evaluated on the length of any curve in $C \setminus \{c\}$. It is possible to simultaneously satisfy the




Figure 3. The geodesic c^* , the central arc a_1 extended out into the flaring end (blue) and an arc of $c \cap (S_g \setminus c^*)$ (green) are shown on the left. Gluing in a strip as shown turns c^* into a piecewise smooth curve, which is shortened by shifting in the direction of the red arrows.

constraints in Thurston’s Lipschitz map construction and choose the strips corresponding to the central arcs to be wider than the other strips, because it was shown that there are at least 2 central arcs; one on either side of c^* , or both intersecting both connected components of $\mathcal{N} \setminus c^*$. ■

Given a set C of curves, and a point $x \in \mathcal{T}_g$, denote by $d(x)$ the value of d for which $E(C, d(x))$ passes through x . Choose $c \in C$, then a point in $E(C, d(x))$ at which $L(c)|_{E(C, d(x))}$ has a minimum is a point at which the length of every curve in C , when restricted to $E(C, d(x))$, has a minimum. A point $x \in \mathcal{T}_g$ will be called a *minimum* of C if $L(c)|_{E(C, d(x))}$ has a minimum at x for every $c \in C$.

Lemma 3.5. *Let C be a minimal filling set of curves. Then the gradients of the lengths of the curves in C are linearly independent everywhere in \mathcal{T}_g away from minima of C .*

Remark 3.6. Note that this lemma also holds for subsets of minimal filling sets of curves.

Proof. As in the proof of Lemma 3.3, there are two cases. When x is C -generic, the lemma follows from Wolpert’s twist formula. To prove the lemma for C -nongeneric points that are not minima of C , it will be shown that C -generic points satisfy a geometric property that implies linear independence of the gradients. It will then be shown that away from minima, Lemma 3.3 implies this property can only break down on open sets. As the set of C -nongeneric points are not open, the lemma then follows.

Let $L(C, t, d) := \{x \in \mathcal{T}_g \mid L(c)(x) = t + d_c \forall c \in C\}$ be a level set of $E(C, d)$ from (3.1) on which the length of each $c \in C$ is given by $t + d|_c$. Then $D(C, t, d)$ is defined to be the intersection of the sublevel sets $\{L(c) \leq t + d|_c \mid c \in C\}$. Recall the definition of faces, facets and tangent cones from Section 2. Since each sublevel set is strictly convex with respect to, for example, the Weil–Petersson metric [23], the intersection $D(C, t, d)$ is convex, in particular connected. Each facet of $D(C, t, d)$ is contained in a level set $L(c)^{-1}(t + d|_c)$ for some $c \in C$. Lower-dimensional faces of $D(C, t, d)$ are intersections of level sets.

The assumption that $E(C)$ is nonempty guarantees that for sufficiently large t the intersection $L(C, t) = \{x \in \mathcal{T}_g \mid L(c)(x) = t \forall c \in C\}$ will be nonempty, and is then the lowest-dimensional face of $D(C, t)$.

Let $\{x_1, \dots, x_n\}$ be a standard set of coordinates on \mathbb{R}^n , and denote by

$$\mathfrak{H}^n = \{(x_1, \dots, x_n) \in \mathbb{R}^n \mid x_1 x_2 \cdots x_n = 0\} \subset \mathbb{R}^n$$

a union of hyperplanes passing through the origin of \mathbb{R}^n . A property of level sets of curves called the \mathfrak{H} -property will be defined, so-called because the level sets satisfying this property intersect like a subset of the hyperplanes in \mathfrak{H}^n . The level sets of $\{L(c) \mid c \in C\}$ will be said to have the \mathfrak{H} -property at $x \in \mathcal{T}_g$ if $T_x \mathcal{T}_g \simeq \mathbb{R}^{6g-6}$ has a basis $\{\frac{\partial}{\partial x_1}, \dots, \frac{\partial}{\partial x_n}\}$ that includes $\{\nabla L(c)(x) \mid c \in C\}$. Note that, when specifying the isomorphism between $T_x \mathcal{T}_g$ and \mathbb{R}^{6g-6} , a metric \mathbf{g} on \mathcal{T}_g is assumed, and the vector space isomorphism between $T_x \mathcal{T}_g$ and \mathbb{R}^{6g-6} is an isometry with respect to the standard metric of \mathbb{R}^{6g-6} . An equivalent definition of the \mathfrak{H} -property is that when the level sets of $\{L(c) \mid c \in C\}$ have the \mathfrak{H} -property at x , the following conditions are satisfied:

- (1) Each $C' \subset C$ determines exactly one face of $D(C, t, d)$ adjacent to x contained in $L(C', t, d|_{C'})$. This will be called the face labelled by C' .
- (2) The face from part 1 labelled by C' has connected unit tangent cone.
- (3) The tangent cone at x of the face labelled by $C' \subset C$ has codimension equal to $|C'| - 1$.

Understanding the difference between linear independence of $\{\nabla L(c) \mid c \in C\}$ and level sets of $\{L(c) \mid c \in C\}$ satisfying the \mathfrak{H} -property is slightly subtle, because the two properties are only different in dimensions at least 4. An example will be given later to illustrate this.

A tangent cone of a face of $D(C, t, d)$ labelled by C' will be called *full* if it has codimension $|C'| - 1$, which is the codimension of $E(C', d|_{C'})$ when the lengths of curves in C' have linearly independent gradients. A tangent cone to $D(C, t, d)$ is full if it has dimension equal to the dimension of \mathcal{T}_g .

Claim. *At a C -generic point, the level sets of $\{L(c) \mid c \in C\}$ have the \mathfrak{H} -property.*

The claim will be proven using twist paths. For each $c \in C$, at x , both $D(C \setminus \{c\}, t, d)$ and $L(C \setminus \{c\}, t, d) \cap \partial D(C \setminus \{c\}, t, d)$ have a product structure. One term of the product (called an \mathbb{R} -factor below) is diffeomorphic to \mathbb{R} , and corresponds to the parameter on the c^* -twist paths. The intersection of the codimension 1 submanifold $L(c)^{-1}(t + d|_c)$ with $L(C \setminus \{c\}, t, d) \cap \partial D(C \setminus \{c\}, t, d)$ cuts through this \mathbb{R} factor; on one side, c is longer than the other curves in C and on the other it is shorter. The existence of an \mathbb{R} -factor corresponding to $c \in C$ guarantees the existence of a face labelled by $C \setminus \{c\}$; the face is the union of subintervals of $\gamma_{c^*, y}$ with $y \in L(C, t, d) \cap \partial D(C, t, d)$ on which $L(c) + d|_c < L(c_i) + d|_{c_i}$ for all $c_i \in C \setminus \{c\}$. The C -generic assumption guarantees this is not empty, because there is a direction along $\gamma_{c^*, y}$ for $y \in L(C, t, d) \cap \partial D(C, t, d)$ in which $L(c)$ is decreasing and the lengths of the other curves in C are stationary. For c_i and $c_j \in C$, the face labelled by $C \setminus \{c_i, c_j\}$ lies between the faces labelled by $C \setminus \{c_i\}$ and $C \setminus \{c_j\}$. Faces labelled by smaller subsets are constructed similarly. Connectivity of unit tangent cones to faces is a consequence of linear independence, as is fullness of the tangent cones. This concludes the proof of the claim.

At all C -generic points that are not minima of C , the dimension of the tangent cone to the face labelled by $C' \subset C$ remains constant. This is a consequence of linear independence of $\{\nabla L(c) \mid c \in C\}$ at C -generic points.

The assumption that x is C -generic will now be dropped. Suppose x is a point of $L(C, t, d)$ on $\partial D(C, t, d)$ that is not a minimum. This assumption ensures that $D(C, t, d)$ has nonempty interior; t can be decreased slightly and $L(C, t, d) \cap \partial D(C, t, d)$ does not become empty.

When C does not fill, Proposition 2.2 implies that at x , the tangent cone to $D(C, t, d)$ is full; these are directions in which the lengths of all the curves in C are decreasing. When C fills, x is not in $\text{Min}(C)$ as this would imply x is a minimum of C . Hence by Definition 2.4 the tangent cone to $D(C, t, d)$ is also full. For any $x \in L(C, t, d) \cap \partial D(C, t, d)$ and any $c \in C$, Lemma 3.3 gives a full tangent cone of the intersection of $D(C \setminus \{c\}, t, d)$ and $L(c)_{\geq t+d|_c}$. This is the cone of directions in which $L(c)$ is increasing and the curves in $C \setminus \{c\}$ are decreasing. Every $c \in C$ labels a facet of $D(C, t, d)$ with x on the boundary. This is because, at x , the tangent cones of $D(C, t, d)$ and of the intersection of $D(C \setminus \{c\}, t, d)$ and $L(c)_{\geq t+d|_c}$ for every $c \in C$ are full, and by convexity the unit tangent cone of $D(C \setminus \{c\}, t, d)$ is connected for every $c \in C$.

If the level sets $\{L(c) \mid c \in C\}$ have the \mathfrak{H} -property at $x \in \mathcal{T}_g$, this implies that the set of vectors $\{\nabla L(c)(x) \mid c \in C\}$ is linearly independent. Note that the converse is not true. One example is given by the gradients in \mathbb{R}^4 of the functions

$$f_1 = e^{x_1} + e^{-x_2}, \quad f_2 = e^{x_2} + e^{-x_3}, \quad f_3 = e^{-x_3} + e^{x_4}, \quad f_4 = e^{-x_4} + e^{x_1}.$$

In this example, the gradients are linearly independent, except where $x_2 = x_4$. Moreover, this set of functions has the property that at every point there is an open cone of directions in which

all four functions are increasing, and for each f_i , $i = 1, \dots, 4$, there is an open cone of directions in which f_i is decreasing and the other functions increasing. Define $D(x)$ to be the intersection of the sublevel sets $\{f_i(x) \leq f_i(p) \mid i = 1, 2, 3, 4\}$. As for $D(C, t, d)$, there is a facet of $D(x)$ along a level set of each of the functions. This means that when $x_2 = x_4$ at x , $\partial D(x)$ has no face labelled by $\{f_1, f_2\}$; the tangent space at x to the submanifold on which $f_1 = f_2$ consists of directions in which one of $\{f_3, f_4\}$ is increasing, and the other decreasing. Similarly for f_3 and f_4 . On a neighbourhood of the submanifold on which $x_2 = x_4$, the level sets also cannot have the \mathfrak{H} -property.

As in the example, away from minima, Lemma 3.3 and Proposition 2.2 imply that if the level sets of $\{L(c) \mid c \in C\}$ do not have the \mathfrak{H} -property at x , this is also the case on a neighbourhood of x . However, the set of C -nongeneric points has measure zero, and the level sets of $\{L(c) \mid c \in C\}$ have the \mathfrak{H} -property at C -generic points. This concludes the proof of the lemma. ■

The key property of a minimal filling set of curves C used to prove the previous lemma was that for every $c \in C$, $C \setminus \{c\}$ fills a strictly smaller subsurface of \mathcal{S}_g than C . The next lemma is a generalisation of Lemma 3.5.

Lemma 3.7. *Let C be a filling set of curves (not necessarily minimal) and suppose $c \in C$ has the property that $C \setminus \{c\}$ is not filling. Then $\nabla L(c)$ is linearly independent from $\{\nabla L(c') \mid c' \in C \setminus \{c\}\}$ everywhere in \mathcal{T}_g away from minima of C .*

Proof. To prove this lemma, a “relative version” of the \mathfrak{H} -property will be defined. After it is shown that C satisfies the $C \setminus \{c\}$ -relative version of the \mathfrak{H} -property at points that are $\{c^*\}$ -generic, the lemma follows from the same argument as in the proof of Lemma 3.5.

Suppose $x \in \mathcal{T}_g$ is not a minimum of C . Recall the three conditions given under the definition of the \mathfrak{H} -property. The curve c will be said to satisfy the $C \setminus \{c\}$ -relative \mathfrak{H} -property if

- (1) For every $C' \subset (C \setminus \{c\})$ that is the label of a face of $D(C \setminus \{c\}, t, d(x))$ adjacent to x contained in a level set of $E(C', t, d(x))$, there is exactly one face of $D(C, t, d(x))$ adjacent to x contained in a level set of $E(C' \cup \{c\}, t, d(x))$.
- (2) The face from (1) labelled by $C' \cup \{c\}$ has unit tangent cone with the same number of connected components as the corresponding face of $D(C \setminus \{c\}, t, d(x))$.
- (3) For nonempty C' , the tangent cone at x of the face of $D(C, t, d(x))$ labelled by $C' \cup \{c\}$ has dimension exactly one less than the corresponding face of $D(C \setminus \{c\}, t, d(x))$ labelled by C' . When C' is the empty set, $D(C, t, d(x))$ has a facet contained in a level set of c and whose tangent cone has codimension 1 in $T_x \mathcal{T}_g$.

At points that are $\{c^*\}$ -generic, to prove that c satisfies the $C \setminus \{c\}$ -relative \mathfrak{H} -property, conditions (1), (2) and the first part of (3) are established using the same \mathbb{R} -factors as in the proof of the claim in Lemma 3.5. Lemma 3.3 gives the existence of the facet in the second part of (3). As in the proof of Lemma 3.5, the assumption that x is not a minimum of C implies x is not in $\overline{\text{Min}(C)}$ hence the tangent cone to $D(C, t, d(x))$ is full. The tangent cone to $D(C \setminus \{c\}, t, d(x)) \setminus D(C, t, d(x))$ is full by Lemma 3.3. That the tangent cone to the facet labelled by $\{c\}$ has codimension 1 in $T_x \mathcal{T}_g$ is then a consequence of convexity of $D(C \setminus \{c\}, t, d(x))$. This concludes the proof that c satisfies the $C \setminus \{c\}$ -relative \mathfrak{H} -property at x when x is $\{c^*\}$ -generic.

As in the proof of Lemma 3.5, away from minima, c satisfies the $C \setminus \{c\}$ -relative \mathfrak{H} -property except on the complement of open sets. The lemma then follows from the same proof by contradiction as Lemma 3.5. ■

The next lemma will be assumed from now on without comment.

Lemma 3.8. *If C is a filling set of curves, whenever $E(C, d)$ is nonempty, $E(C, d)$ has a minimum of C .*

Proof. Whenever C fills, it contains a minimal filling set C_m . It was shown in Corollary 3.10 that $E(C_m, d|_{C_m})$ has a minimum, and since $E(C, d) \subset E(C_m, d|_{C_m})$, the lengths of curves on $E(C, d)$ cannot be less than the value at the minimum on $E(C_m, d|_{C_m})$. Since the intersection of $E(C, d)$ with the thick part of \mathcal{T}_g is closed, the infimum is realised. As each connected component of $E(C, d)$ is closed, this also shows that each connected component of $E(C, d)$ has a minimum. ■

Lemma 3.9. *Let C be a minimal filling set of curves and $C' \subsetneq C$. The only minima of C occur in the intersection of $E(C, d)$ with $\text{Min}(C)$, and C' has no minima.*

Proof. It was shown in the proof of Lemma 3.5 that for $c \in C$ there are no saddle points of $L(c)|_{L(C \setminus \{c\}, t, d)}$ and for $c \in C'$ there are no saddle points of $L(c)|_{L(C' \setminus \{c\}, t, d)}$. It follows that each connected component of $E(C, d)$ and each connected component of $E(C', d)$ has at most one minimum of the restriction of the lengths of curves in C respectively C' . By Lemma 3.5, the gradients of the lengths of curves in C , respectively C' , are linearly independent everywhere away from this minimum.

Every point x in \mathcal{T}_g is in $E(C, d)$ for some $d(x)$. It will now be shown that every $E(C, d)$ intersects $\text{Min}(C)$ in a unique point.

That $E(C, d)$ intersects $\text{Min}(C)$ in at most one point is an immediate consequence of the definition of $\text{Min}(C)$. If $E(C, d)$ intersected $\text{Min}(C)$ in two points, p_1 and p_2 , then either the lengths of the curves in C at both points are equal, or the lengths of the curves at one of the points, call it p_1 , are smaller. In the first case, there are length functions with minima realised at both points, contradicting convexity. In the second case, a contradiction is obtained, because the length function with global minimum at p_2 is smaller at p_1 .

Since C is a minimal filling set of curves, by [19, Lemma 9, Theorem 10], $\text{Min}(C)$ is a continuously differentiable cell with empty boundary. As each $E(C, d)$ intersects $\text{Min}(C)$ in at most one point, there is a differentiable bijection from $\text{Min}(C)$ to a subset of elements of $\mathbb{R}^{|C|}$ whose entries sum to zero; a point $x \in \text{Min}(C)$ is mapped to the tuple d for which $E(C, d)$ passes through x .

Claim. *The set of loci $E(C, d)$ that intersect $\text{Min}(C)$ with varying d foliate \mathcal{T}_g .*

To prove the claim, suppose $\{y_i\}$ is a sequence of points, each of which is in a stratum $E(C, d(y_i))$ intersecting $\text{Min}(C)$, but with limit y in a stratum $E(C, d(y))$ that does not intersect $\text{Min}(C)$. Let $\{x_i\}$ be a sequence of points in \mathcal{T}_g converging to a point x , where x_i is in the same locus as y_i for every $i \in \mathbb{N}$. Since d varies smoothly over \mathcal{T}_g , $\{d(y_i)\}$ approaches $d(y)$ and $\{d(x_i)\}$ approaches $d(x) = d(y)$.

Note that there is a neighbourhood of $\text{Min}(C)$ on which the claim holds. If the points of the sequence $\{x_i\}$ are chosen sufficiently close to $\text{Min}(C)$, x and every x_i is in an $E(C, d(x_i))$ that intersects $\text{Min}(C)$. Moreover, it is possible to choose the sequences in such a way that each x_i and y_i are arbitrarily close together, giving that x and y are in the same component of $E(C, d(x))$. The claim then follows by contradiction. This concludes the proof of the lemma in the case of filling sets of curves.

For a nonfilling subset C' of C , any minima of the lengths of the curves in C' restricted to $E(C', d)$ must occur within $\text{Min}(C)$; otherwise this would give a minimum on some $E(C, d)$ away from $\text{Min}(C)$. However, by [19, Lemma 5], the assumption that C is a minimal filling set of curves implies that the gradients of the lengths of curves in C' are everywhere linearly independent on $\text{Min}(C)$. This rules out the existence of minima on $E(C', d)$, as these can only occur at a point at which $L(c)^{-1}(t + d|_c)$ intersects $L(C' \setminus \{c\}, t, d)$ nontransversely for any $c \in C'$. ■

Corollary 3.10. *Let C be a minimal filling set of curves. When $E(C, d)$ is nonempty, it is a connected, embedded submanifold of \mathcal{T}_g . The restriction to $E(C, d)$ of the lengths of the curves in C is strictly convex and has a unique minimum.*

Proof. On a connected component of $E(C, d)$ that intersects $\text{Min}(C)$, the gradients of the lengths of curves are linearly independent away from $\text{Min}(C)$. Since $C = \{c_1, \dots, c_k\}$ is a minimal filling set of curves, it follows from [19, Lemma 5] that $\{\nabla L(c_2) - \nabla L(c_1), \dots, \nabla L(c_k) - \nabla L(c_1)\}$ is linearly independent everywhere on $\text{Min}(C)$. Repeated applications of the pre-image lemma (see, for example, [11, Corollary 5.14]) then implies that each $E(C, d)$ passing through $\text{Min}(C)$ is an embedded submanifold.

It was also shown in the proof of Lemma 3.9 that on each connected component of $E(C, d)$, the only nonregular point of the restriction of the lengths of curves in C is a unique minimum where the connected component passes through $\text{Min}(C)$ and that $E(C, d)$ can intersect $\text{Min}(C)$ at most once. This proves convexity and connectivity. ■

3.2 Balanced loci

This subsection discusses a property of loci. This concept is related to the way loci of filling sets of curves intersect corresponding sets of minima, leading to a trichotomy of three fundamentally different types of loci.

Suppose now that C is any finite set of curves on \mathcal{S}_g that intersect pairwise at most once. It follows from Lojasiewicz’s theorem, [12] that $E(C, d)$ admits a triangulation, i.e., it is a simplicial complex. Consequently, $E(C, d)$ from equation (3.1) has a dimension, and in the interior of each simplex, it has a tangent space. In what follows, v_C will always denote a vector or vector field that specifies a direction in which the lengths of the curves in C are all increasing at the same rate. Whenever possible, v_C is chosen to be contained in or on the boundary of the convex hull of $\{\nabla L(c) \mid c \in C\}$.

Definition 3.11 ($E(C, d)$ is balanced, semi-balanced or unbalanced at x). Suppose x is a point of $E(C, d)$. Then

- $E(C, d)$ is *balanced* at x if x is in $\text{Min}(C)$ or there is a nonzero vector v_C in $T_x E(C, d)$ in the interior of the convex hull of $\{\nabla L(c)(x) \mid c \in C\}$. Equivalently, if V is the orthogonal complement of $v_C(x)$, C is V -eutactic.
- $E(C, d)$ is *semi-balanced* at x if it is not balanced but x is in $\text{Min}(C')$ for $C' \subsetneq C$ or there is a nonzero vector v_C on the boundary of the convex hull of $\{\nabla L(c)(x) \mid c \in C\}$.
- $E(C, d)$ is *unbalanced* at x if there exists a nonzero vector v_C at x and $E(C, d)$ is neither balanced nor semi-balanced.

As will be explained later, the property of being balanced is independent of the choice of metric.

When $E(C, d)$ is balanced at x and x is not in $\text{Min}(C)$, this implies that there is a direction v_C in $T_x E(C, d)$ in which the lengths of the curves in $C = \{c_1, \dots, c_k\}$ are increasing. While $a_1 \nabla L(c_1) + \dots + a_k \nabla L(c_k)$, $a_1, \dots, a_k \in \mathbb{R}_+$ could be a direction in which some of the lengths of curves in C are decreasing, when such a vector is tangent to $E(C, d)$, the constraints defining $E(C, d)$ force this to be a direction in which the lengths of all the curves in C are increasing.

Remark 3.12. When C is a minimal filling set of curves, it was shown in Corollary 3.10 that $E(C, d) \subset \mathcal{T}_g$ is an embedded submanifold, in which case it has a tangent space at every point. The statement that $E(C, d)$ is balanced at every point therefore makes sense.

Proposition 3.13. *Let C be a minimal filling set of curves. Then $E(C, d)$ is balanced at every point.*

Proof. The proof is essentially an application of Lagrange multipliers and the characterisation of $\text{Min}(C)$ from [19]. Fix $C = \{c_1, \dots, c_k\}$. For x not in $\text{Min}(C)$, the locus $E(C, d)$ is balanced at x if and only if it is possible to find a k -tuple (a_1, \dots, a_k) in \mathbb{R}_+^k , with entries normalised to sum to one, and such that v_C can be normalised to give $v_C = a_1 \nabla L(c_1) + \dots + a_k \nabla L(c_k)$ in $T_x \mathcal{T}_g$.

Claim. *For $x \in E(C, d) \setminus \text{Min}(C)$, $v_C = a_1 \nabla L(c_1) + \dots + a_k \nabla L(c_k)$ where (a_1, \dots, a_k) are the parameters of the point in $\text{Min}(C)$ through which the locus $E(C, d)$ passes.*

To prove the claim, recall that on every $E(C, d)$, by Corollary 3.10, there is a unique point $p(C, d)$ in the intersection $E(C, d) \cap \text{Min}(C)$; the point $p(C, d)$ is the point at which the restrictions to $E(C, d)$ of the lengths of curves in C have their minimum. Define

$$L(t) := \bigcup_d \{y \in E(C, d) \mid L(c)(y) = L(c)(p(C, d)) + t \forall c \in C\}.$$

Choose t so that the connected component of $L(t)$ passes through the point x . Theorem 10 of [19] showed that $\text{Min}(C)$ is a C^1 -submanifold. Since $L(t)$ is the boundary of a tubular neighbourhood of a C^1 -submanifold, it follows from a classical result, see [7], that $L(t)$ is also C^1 . Lemma 2.5 implies that $\text{Min}(C)$ has empty boundary.

Recall that the entries of the tuple (a_1, \dots, a_k) are normalised to sum to one. By construction, x is a local minimum of the length function $a_1 L(c_1) + \dots + a_k L(c_k)$ along any path in $L(t)$ crossing $E(C, d(x))$ transversely. It follows from the method of Lagrange multipliers that $a_1 \nabla L(c_1) + \dots + a_k \nabla L(c_k)$ is normal to $L(t)$ everywhere in the intersection $E(C, d(x)) \cap L(t)$, and hence parallel to the normal vector field $\frac{\partial}{\partial t}$ on $L(t)$. In $E(C, d(x))$, the common gradient of the length functions, a vector field parallel to v_C , is also always orthogonal to the level sets of the lengths of curves in C . This proves the claim.

As this construction works for any value of $t > 0$, it follows that the k -tuple $d(x)$ defining the locus $E(C, d(x))$ passing through x is the value of $d|_C$ for which $E(C, d(x))$ passes through the point in $\text{Min}(C)$ representing the minimum of the length function $a_1 L(c_1) + \dots + a_k L(c_k)$, and v_C is parallel to $a_1 \nabla L(c_1) + \dots + a_k \nabla L(c_k)$ everywhere on $E(C, d(x))$. This concludes the proof. \blacksquare

Proposition 3.14. *Let C' be a subset of a minimal filling set, or a multicurve possibly containing separating curves. Then $E(C')$ is balanced at every point.*

Remark 3.15. Unlike the case for minimal filling sets, $E(C', d)$ is not balanced for every possible d . However, since $\{\nabla L(c)(x) \mid c \in C'\}$ is linearly independent at $x \in E(C')$, this linear independence holds on a neighbourhood of x in \mathcal{T}_g , as does the existence of a $v_{C'}$ in the convex hull of $\{\nabla L(c) \mid c \in C'\}$. It follows that points in loci $E(C', d)$ sufficiently close to $E(C')$ are also balanced.

Proof. When C' is a subset of a minimal filling set, it follows from Lemmas 3.5 and 3.9 that the loci $E(C', d)$ are embedded submanifolds that foliate \mathcal{T}_g with varying d . When C' is a multicurve, this follows from the fact that the lengths of curves in C' determine a subset of Fenchel–Nielsen coordinates defined everywhere on \mathcal{T}_g . The proof of Proposition 3.13 therefore generalises as soon as a replacement for the set of minima $\text{Min}(C)$ can be found.

By Proposition 2.2, $\text{Min}(C')$ is empty in \mathcal{T}_g . However, nonfilling subsets can be understood to have sets of minima in the metric completion $\overline{\mathcal{T}}_g$ of \mathcal{T}_g with respect to the Weil–Petersson metric. Points on the boundary of \mathcal{T}_g with respect to this completion can be understood to be noded surfaces that have been pinched along multicurves, see, for example, [24]. In [19], sets of

minima were also defined for noded surfaces. It follows from Riera’s formula [18] that a necessary condition for a length function $L(A, C')$ to have a minimum is that the multicurve consisting of the boundary of the subsurface filled by C' has been pinched. Within such noded surfaces, C' is a minimal filling set, and it is possible to define a set of minima that will be denoted by $\text{Min}(C')^\infty$. Not every $E(C', d)$ will necessarily have a point in $\text{Min}(C')^\infty$ in its completion; the proof consequently reduces to the problem of showing that $E(C')$ does.

Corollary 3.16 ([24, Corollary 21]). *Denote by $\text{Sys}(m)^\infty$ a stratum in $\overline{\mathcal{T}}_g \setminus \mathcal{T}_g$ corresponding to noded surfaces on which the multicurve m has been pinched. Denote by $L(m)$ the length function $\sum_{c \in m} L(c)$. The distance $d(x, \text{Sys}(m)^\infty)$ of a point $x \in \mathcal{T}_g$ to $\text{Sys}(m)^\infty$ has the expansion*

$$d(x, \text{Sys}(m)^\infty) = \sqrt{2\pi L(m)} + \mathcal{O}(L(m)^2).$$

When C' is a multicurve m , $\text{Min}(C')^\infty$ is contained in $\text{Sys}(m)^\infty$, which is in the metric completion of $E(C')$. This proves the proposition for multicurves.

When C' is not a multicurve, denote by m a multicurve contained in C' with largest cardinality, and c_1 a curve in $C' \setminus m$. Then by the collar lemma, $L(c_1)$ is zero on $\text{Sys}(\{c_1\})^\infty$ and infinite on $\text{Sys}(m)^\infty$, similarly $L(m)$ is zero on $\text{Sys}(m)^\infty$ and infinite on $\text{Sys}(\{c_1\})^\infty$. Again by Riera’s formula, $\text{Min}(m \cup \{c_1\})^\infty$ must be in a noded surface pinched along the multicurve on the boundary of the subsurface filled by $m \cup \{c_1\}$. This multicurve is disjoint from both m and c_1 , so this analogue of a set of minima has the analogues of sets of minima of subsets of $m \cup \{c_1\}$ on its boundary, generalising Lemma 2.5.

Let γ be a path in the thin part of \mathcal{T}_g near $\text{Min}(m \cup \{c_1\})^\infty$ with one endpoint close to $\text{Sys}(\{c_1\})^\infty$ and the other close to $\text{Sys}(m)^\infty$. Using the estimate in Corollary 3.16, it follows from the intermediate value theorem applied to γ that $E(m \cup \{c_1\})$ has a point in $\text{Min}(m \cup \{c_1\})^\infty$ in its metric completion. This proves the proposition when $C' = m \cup \{c_1\}$. Now let c_2 be a curve in $C' \setminus (m \cup \{c_1\})$ chosen such that c_1 intersects a curve in $m \cup \{c_1\}$; this is always possible because m was assumed to be a multicurve of maximal cardinality. The same argument proves the proposition for $C' = m \cup \{c_1, c_2\}$. This construction is then iterated further if necessary. ■

Remark 3.17. When using Proposition 3.14, it can be helpful to know that any pair $\{c_1, c_2\}$ of curves intersecting exactly once is contained in a minimal filling set. This follows from the observations that the action of Γ_g on curves takes minimal filling sets to minimal filling sets, Γ_g acts transitively on pairs of nonseparating curves intersecting exactly once, and any minimal filling set contains a pair of curves that intersect once.

Corollary 3.18. *A stratum $\text{Sys}(C)$ is on the boundary of $\text{Sys}(\{c\})$ for every $c \in C$.*

Proof. Let x be a point in $\text{Sys}(C)$. It follows from Proposition 3.14 and Remark 3.17 that for every $c_1, c_2 \in C$, $\mathbf{g}(\nabla L(c_2)(x), \nabla L(c_1)(x)) < \mathbf{g}(\nabla L(c_1)(x), \nabla L(c_1)(x))$. For every $c \in C$, it follows from local finiteness that the vector $-\nabla L(c)(x)$ is in the tangent cone of $\text{Sys}(\{c\})$. ■

Proposition 3.19. *The property of being balanced at x is independent of the choice of metric.*

Proof. The proof is basically a repeat of the arguments given in Propositions 3.13 and 3.14. Suppose $E(C, d)$ is balanced at x with respect to a metric \mathbf{g} . Setting $C = \{c_1, \dots, c_k\}$, one obtains that v_C is proportional to $\sum_{i=1}^k a_i \nabla L(c_i)$ where $a_i \in \mathbb{R}_+$, $i = 1, \dots, k$ and $\sum_{i=1}^k a_i = 1$. Define $L := \sum_{i=1}^k a_i L(c_i)$. The proposition will be proven by showing that $\nabla L(x)$ is proportional to $v_C(x)$ for any metric.

The functions $\mathcal{T}_g \rightarrow \mathbb{R}_+$ obtained by taking components of the k -tuple d will be denoted by d_1, \dots, d_k . On a neighbourhood of x , it is possible to obtain a subset of independent coordinates $\{d'_1, \dots, d'_n\}$ with the property that each d'_i is a linear combination of d_1, \dots, d_k , and $\text{Span}\{\nabla d'_1, \dots, \nabla d'_n\} = \text{Span}\{\nabla d_1, \dots, \nabla d_k\}$.

The vectors $v_C(x)$ and a spanning set for $\{\nabla L(c)(x) \mid c \in C\}^\perp$ in $T_x \mathcal{T}_g$ span the subspace $\{\nabla d_1, \dots, \nabla d_k\}^\perp$ of $T_x \mathcal{T}_g$ on which the tuple d is stationary. Note that this property is independent of \mathbf{g} , as the subspace of $T_x \mathcal{T}_g$ on which d is stationary can be defined in terms of tangents to intersections of level sets, as can $\{\nabla L(c)(x) \mid c \in C\}^\perp$. Taking d'_1, \dots, d'_n, L to be a subset of independent coordinates on a neighbourhood of x in \mathcal{T}_g , calculating differentials in this coordinate chart gives that $dL(x)$ is contained in $(\{\nabla d'_1, \dots, \nabla d'_n\}^\perp)^\flat$. Also by construction, $\nabla L(x) = \sum_{i=1}^k a_i \nabla L(c_i)(x)$. Consequently, $dL(x) = v_C^b(x)$, as required. ■




Figure 4. Inner versus outer loci. In the figure on the right, the locus should be regarded as coming out of the page. A filling set of curves that intersect pairwise at most once must have cardinality at least four. This figure represents a lower-dimensional analogue, where $\text{Sys}(\{c_i\})^\infty$ denotes a point in the completion of \mathcal{T}_g with respect to the Weil–Petersson metric at which the curve c_i has been pinched.

Once again allowing C to be any filling set of curves that intersects pairwise at most once, now that $\text{Min}(C)$ can have a boundary, a new phenomenon occurs.

Inner, outer and borderline loci. A connected component of $E(C, d)$ is

- *inner* if it passes through a point of $\text{Min}(C)$,
- *borderline* if it passes through a point of $\partial \text{Min}(C)$,
- *outer* if it is disjoint from $\overline{\text{Min}(C)}$.

Lower-dimensional analogues of inner and outer loci are shown in Figure 4. As $E(C, d)$ can intersect $\overline{\text{Min}(C)}$ in at most one point, each locus is exactly one of these three types.

An inner locus $E(C, d)$ has the property that at the global minimum point of the lengths of curves in C restricted to $E(C, d)$, there is no direction in which the lengths of all the curves in C are decreasing, whereas there is an open cone of such directions at the minimum point when $E(C, d)$ is outer.

The loci $E(C, d)$ studied here so far were all embedded submanifolds of dimension greater than one that foliate \mathcal{T}_g for varying d . The arguments of Proposition 3.13 suggest that for filling C , such loci are

- balanced if they are inner
- semi-balanced if they are borderline
- unbalanced if they are outer.

It would be interesting to know to what extent this generalises.

For a filling set C , if $\text{Sys}(C) \subset E(C)$ is not balanced at some point x , then by Proposition 3.13, C strictly contains a minimal filling set of curves. The next proposition can be used to characterise a filling subset of C .

Proposition 3.20. *Suppose C is a filling set of curves, and at a point x that is not a minimum of C , a nonzero $v_C(x)$ exists which cannot be chosen to be in the interior of the convex hull of $\{\nabla L(c)(x) \mid c \in C\}$. Then the vertices of the face f of the convex hull of $\{\nabla L(c)(x) \mid c \in C\}$ closest to or containing $v_C(x)$ are the gradients of a filling set of curves $c(f)$.*

Remark 3.21. The set of curves $c(f)$ from Proposition 3.20 has the following property in common with a minimal filling set of curves: A nonzero vector $v_{c(f)}(x)$ exists in the interior of the convex hull of $\{\nabla L(c)(x) \mid c \in c(f)\}$. Identifying $T_x \mathcal{T}_g$ with \mathbb{R}^{6g-6} by a vector space isomorphism that preserves the inner product given by \mathbf{g} , the vector $v_{c(f)}(x)$ is obtained by adding to $v_C(x)$ a multiple of the normal vector to f representing the shortest path between the vertex of $v_C(x)$ and the convex hull.

Proof of Proposition 3.20. The motivation for the proof is to compare the assumption on the gradients with the statement of Lemma 2.5. In the projection to the orthogonal complement of $v_{c(f)}(x)$ in $T_x \mathcal{T}_g$, the gradients $\{\nabla L(c)(x) \mid c \in C\}$ look like the gradients at a point in $\text{Min}(c(f)) \subset \overline{\text{Min}(C)}$, where $\text{Min}(c(f))$ is only nonempty if $c(f)$ fills. The intuition is that $c(f)$ corresponds to the set of minima on the boundary of $\text{Min}(C)$ closest to x ; such a boundary point is reached by decreasing the lengths of curves in C .

For every $c \in (C \setminus c(f))$, the inequality

$$\mathbf{g}(\nabla L(c), v_{c(f)}) \geq \mathbf{g}(v_{c(f)}, v_{c(f)}) \quad (3.2)$$

is assumed in the statement of the proposition to hold at x , where the assumption is made that $v_{c(f)}$ is normalised in such a way that it represents a point in \mathbb{R}^{6g-6} contained in the polytope with vertices given by $\{\nabla L(c) \mid c \in c(f)\}$.

Denote by C_{ineq} the connected component containing x of the subset of $E(c(f), d(x))$ on which inequality (3.2) holds. A point at which $v_{c(f)} = 0$ is assumed to be in C_{ineq} if it is a limit point of a sequence on which inequality (3.2) holds.

The argument is greatly simplified at a number of points by arguing that pathological cases (usually to do with linear dependence of gradients) can only occur when one or more curves can be discarded from the set without changing the subsurface of \mathcal{S}_g filled by the set. To begin with, it will be shown that the following assumptions can be made without loss of generality.

- (1) The set of vectors $\{\nabla L(c) \mid c \in (C \setminus c(f))\}$ is linearly independent in \mathcal{T}_g away from minima of $C \setminus c(f)$, and each one is linearly independent from $\{\nabla L(c) \mid c \in c(f)\}$ away from minima of C .
- (2) The set of vectors $\{\nabla L(c)(x) \mid c \in c(f)\}$ is linearly independent everywhere in C_{ineq} where $v_{c(f)}$ is nonzero. A nonzero vector $v_{c(f)}$ whose existence is guaranteed by this linear independence can be chosen to be contained in the convex hull of $\{\nabla L(c) \mid c \in c(f)\}$ and is in the tangent cone to $E(c(f), d(x))$.
- (3) Strict inequality in inequality (3.2) holds at x .

To see that assumption (1) can be made without loss of generality, suppose some of the curves $C \setminus c(f)$ have lengths with linearly dependent gradients at a point that is not a minimum of $C \setminus c(f)$. It follows from Lemma 3.7 that the proof of the proposition can be simplified by removing these curves from C . Doing so does not destroy the property that inequality (3.2) holds for the remaining curves in $C \setminus c(f)$. Similarly, if there is a curve in $C \setminus c(f)$ for which the gradient of the length is linearly dependent on $\{\nabla L(c) \mid c \in c(f)\}$ away from minima of C ,

discarding this curve from C still results in a filling set. If all the curves $C \setminus c(f)$ have lengths with gradients linearly dependent on $\{\nabla L(c) \mid c \in c(f)\}$ at a point that is not a minimum of C , it follows that $c(f)$ fills as required. This concludes the proof that the first assumption can be made without loss of generality.




Figure 5. If no nonzero $v_{c(f)}(y)$ exists because the dimension of $\text{Span}(\{\nabla L(c) \mid c \in c(f)\})$ drops at y (part (a)), then at points near y where a nonzero $v_{c(f)}$ exists, this is not in the convex hull of $\{\nabla L(c) \mid c \in c(f)\}$ (part (b)), viewed from the side.

At every point y in C_{ineq} at which $v_{c(f)}$ is nonzero, it is possible to assume without loss of generality that the set of vectors $\{\nabla L(c')(y) \mid c' \in c(f)\}$ is linearly independent. Otherwise, by Lemma 3.7 there is a subset $c'(f)$ of $c(f)$ that fills the same subset of \mathcal{S}_g as $c(f)$ and for which $v_{c'(f)}(y) = v_{c(f)}(y)$. It is then possible to prove the proposition with $c'(f)$ in place of $c(f)$. The same argument shows that it is possible to assume without loss of generality that if a nonzero $v_{c(f)}(y)$ exists for $y \in C_{\text{ineq}}$, it is in the tangent cone to $E(c(f), d(x))$ at y . It is also possible to assume without loss of generality that $v_{c(f)}$ is contained in the convex hull of $\{\nabla L(c') \mid c' \in c(f)\}$ at every point in C_{ineq} where $v_{c(f)}$ is nonzero because otherwise there is a point y' satisfying the conditions of the proposition with $c'(f) \subset c(f)$ in place of $c(f)$. This concludes the proof that the second assumption can be made without loss of generality at points of C_{ineq} where a nonzero $v_{c(f)}$ exists.

Strict convexity of length functions with respect to the Weil–Petersson metric, proven in [23], implies that if equality holds in inequality (3.2) for a curve $c \in (C \setminus c(f))$, shifting x an arbitrarily small amount in the direction of $\nabla L(c)(x)$ gives a point x' at which strict inequality holds. By assumption (1), it is then possible to work in the locus containing x' instead of $E(C, d(x))$, showing that assumption (3) can also be made without loss of generality.

To understand what happens when $v_{c(f)}(y) = 0$ for $y \in C_{\text{ineq}}$, suppose y is on the boundary of the subset of C_{ineq} on which $v_{c(f)}$ is zero. It follows from assumption (2) applied to points of C_{ineq} at which $v_{c(f)}$ is nonzero that y is either in $\overline{\text{Min}(c(f))}$ or is a point at which linear independence of $\{\nabla L(c') \mid c' \in c(f)\}$ breaks down. The former implies $c(f)$ fills, so suppose from now on that it is the latter. Definition 2.4 implies that there is an open cone of directions at y in which the lengths of all the curves in $c(f)$ are decreasing. By assumption, y is a limit point of a sequence on which linear independence of $\{\nabla L(c') \mid c' \in c(f)\}$ holds and on which there is a nonzero $v_{c(f)}$ in the convex hull of $\{\nabla L(c') \mid c' \in c(f)\}$. As illustrated in Figure 5, this prevents linear independence of $\{\nabla L(c') \mid c' \in c(f)\}$ from breaking down at y , because for all points sufficiently close to y at which linear independence holds, $v_{c(f)}$ could not be in the convex hull of $\{\nabla L(c') \mid c' \in c(f)\}$. This shows that $v_{c(f)} = 0$ on C_{ineq} implies that $c(f)$ fills. It will therefore be assumed that $v_{c(f)}$ is nonzero on C_{ineq} .

Next, assume to begin with that $|C \setminus c(f)| = 1$. The locus $E(C, d(x))$ is a separating subset of $E(c(f), d(x))$. One side of $E(c(f), d(x)) \setminus E(C, d(x))$ contains the loci in which the curve c in $C \setminus c(f)$ is longer than the curves in $c(f)$, and the other contains the loci on which it is shorter.

The loci $E(C, d)$ for varying d contained in $E(c(f), d(x))$ are shown schematically in Figure 6, as well as other objects that will be defined in this proof.




Figure 6. A schematic drawing of $E(c(f), d(x))$. The subset C_{ineq} of $E(c(f), d(x))$ is shaded. The concentric circles are level sets of $L(c')$ for $c' \in c(f)$, and the curved vertical lines are loci $E(C, d)$ for different values of d .

In what follows, within $E(c(f), d(x))$ it will be convenient to confuse the tuple d with a function $d: E(c(f), d(x)) \rightarrow \mathbb{R}$, $x \mapsto L(c)(x) - L(c')(x)$ for a choice (it does not matter which) of curve $c' \in c(f)$.

Claim. *Under the assumptions listed above, if strict inequality in inequality (3.2) holds at x , it holds everywhere on the connected component $E'(C, d(x))$ of $E(C, d(x))$ containing x .*

The claim will be proven by contradiction. If the claim were false, it follows from the linear independence assumptions that there is a point y_0 on $E(C, d(x))$ at which equality in inequality (3.2) holds. Choose a point y on $E'(C, d(x))$ close to y_0 at which inequality in inequality (3.2) holds, and for which y is not a minimum of C . Setting $C = \{c_1, \dots, c_k\}$, $c = c_1$, there is a $b_1 < 0$, $b_i > 0$ for $i = 2, \dots, k$ with $\sum_{i=1}^k b_i = 1$ such that $v_C(y) = \sum_{i=1}^k b_i \nabla L(c_i)(y) \neq 0$. The point y is chosen such that $|b_1|$ is sufficiently small that the function $\sum_{i=1}^k b_i L(c_i)$ is convex on a neighbourhood of y containing y_0 .

An argument similar to the proof of Proposition 3.13 will now be given, but using an analogue of $\text{Conv}(C)$ (defined in Section 2.1) in place of $\text{Min}(C)$. The general idea is to use the linear independence assumptions and convexity to derive a contradiction to the existence of y_0 .

Within some neighbourhood of y , the flow-box theorem guarantees the existence of a continuously differentiable codimension 1 submanifold $L(y)$ embedded in $E(c(f), d(x))$ orthogonal to the vector field $v_C = v_{\{c\} \cup c(f)}$ and passing through y . On a neighbourhood of y , a locus $E(C, d)$ near $E'(C, d(x))$ intersects $L(y)$ along a level set of $L(C)|_{E(C, d)}$. Define

$$L(y, t) = \bigcup_{d \in \mathcal{N}(d(x))} \{z \in E(C, d) \mid L(c_i)(z) = t + L(c_i)(z_0) \forall c_i \in C, \\ \text{where } z_0 \in E(C, d) \cap L(y)\},$$

where t takes values in a neighbourhood of 0 in \mathbb{R} , and $\mathcal{N}(d(y))$ is a neighbourhood of $d(x) = d(y)$ in \mathbb{R} .

It follows from convexity of the length functions $L(c)$ and $L(c_2)$ and local convexity of $\sum_{i=1}^k b_i L(c_i)$ that $E'(C, d(x)) \cap L(y)$ is the minimum of the restriction of the function $\sum_{i=1}^k b_i L(c_i)$ to $L(y)$. Similarly, $E'(C, d(x)) \cap L(y, t)$ is the minimum of the restriction of the function $\sum_{i=1}^k b_i L(c_i)$ to $L(y, t)$. The method of Lagrange multipliers then shows that on a neighbourhood of y containing y_0 ,

$$v_C|_{E'(C, d(x))} = \sum_{i=1}^k b_i \nabla L(c_i) \neq 0,$$

contradicting the assumption that equality in inequality (3.2) occurs at y_0 . This concludes the proof of the claim.

The minima of C on $E'(C, d(x))$ are a connected set, because the linear independence assumptions rule out any nonregular points of $L(C)|_{E'(C, d(x))}$ that are not minima. Within C_{ineq} the loci $E(C, d)$ for varying d make up the leaves of a codimension 1 foliation. This is because every connected component of $E(C, d)$ in C_{ineq} is noncompact, as otherwise there would be loci in C_{ineq} with maxima of C , contradicting the linear independence assumptions. Since C fills, level sets of $L(C)|_{E(C, d)}$ are compact. Consequently, there is no connected component of a locus $E(C, d)$ in C_{ineq} contained in a level set of $L(C)$. The function d can only take regular values in the interior of C_{ineq} because otherwise there would be a nonregular value d^* of d such that assumption (1) breaks down at every point of a connected component $E'(C, d^*)$ of $E(C, d^*)$ in C_{ineq} that is not a minimum of C .

Remark 3.22. Since assumptions (1) and (2) hold on a neighbourhood of C_{ineq} , and v_{c_f} is assumed to be nonzero in C_{ineq} , it also follows that there is a neighbourhood of C_{ineq} in $E(c(f), d(x))$ on which the function d can only take regular values.

Let γ be a continuous path traced out by minima of C on $E(c(f), d(x))$, starting with the minimum of C in $E(C, d(x)) \subset E(c(f), d(x))$ and passing through the loci in the direction of $-v_{c(f)}$. Since d only takes regular values on C_{ineq} , strict inequality in inequality (3.2) holds along the intersection of γ with the interior of C_{ineq} . Moreover, by assumption (2) and Remark 3.22, a point at which γ ends or leaves C_{ineq} is a point at which $v_{c(f)} = 0$. The proposition will be proven by showing that γ contained in C_{ineq} leads to a contradiction.

If γ is contained in C_{ineq} , γ approaches a noded surface in \overline{T}_g on which a multicurve has been pinched. For example, when $d(x) = 0$, as discussed in the proof of Lemma 3.14, this is the noded surface on which the multicurve on the boundary of the subsurface filled by $c(f)$ has been pinched. Since C fills, this multicurve intersects a curve in C . It follows from the collar lemma that γ cannot be a path along which the lengths of curves in C decrease monotonically, giving the required contradiction.

When $|C| > |c(f)| + 1$, assumption (1) makes it possible to replace $L(c)$ in the previous argument by the length function

$$L(C \setminus c(f)) := \frac{1}{|C \setminus c(f)|} \sum_{c' \in (C \setminus c(f))} L(c').$$

Along γ , the fact that the length function $L(C \setminus c(f))$ is decreasing faster in the direction of $v_{c(f)}$ than the lengths of curves in $c(f)$ does not mean that the length of every curve in $C \setminus c(f)$ is decreasing faster in the direction of $-v_{c(f)}$ than the lengths of curves in $c(f)$. If, for example, $C \setminus c(f) = \{c_1, c_2\}$, and at the point p at which γ leaves C_{ineq} the inequality (3.2) holds for c_1 but not c_2 , then γ passes through a locus on which the assumptions of the proposition hold with $c(f)$ replaced by $c(f) \cup \{c_2\} \subset C$. The argument then shows $c(f) \cup \{c_2\} \subset C$ fills, and can subsequently be applied with $c(f) \cup \{c_2\}$ in place of C to show that $c(f)$ fills. This argument can be generalised to the case with $|C \setminus c(f)| > 2$. \blacksquare

4 Proof of the Morse–Smale properties

The purpose of this section and the next is to show that \mathcal{P}_g contains a choice of unstable manifolds of the critical points of f_{sys} . This section studies the behaviour of \mathcal{P}_g away from critical points and boundary points of f_{sys} , i.e., at points at which the cone of increase is full. It will be shown that the locally top-dimensional strata making up the cells of \mathcal{P}_g are balanced, and that restricting f_{sys} to \mathcal{P}_g does not create critical points. In this way, the intuition of Thurston that f_{sys} is increasing towards \mathcal{P}_g is encoded.

The main results of this section will be proven by generalising the properties of strata of minimal filling sets to the locally top-dimensional strata of \mathcal{P}_g . The simple case is described in the next lemma.

Lemma 4.1. *Suppose $x \in \text{Sys}(C)$ and the set of vectors $\{\nabla L(c)(x) \mid c \in C\}$ is linearly independent. Then $\text{Sys}(C)$ is on the boundary of $\text{Sys}(C')$ for every nonempty $C' \subsetneq C$.*

Proof. For any fixed $C' \subset C$, $E(C) \subset E(C')$, so $E(C')$ is not empty. The linear independence assumption then ensures that the codimension of $E(C')$ in \mathcal{T}_g is equal to $|C'| - 1$, and that the tangent space at x to the level set of $E(C')$ passing through x is given by the orthogonal complement of $\text{Span}\{\nabla L(c) \mid c \in C'\}$. Choose $c_1 \in C$. It follows from local finiteness that the nonzero vector obtained by projecting $\nabla L(c_1)(x)$ onto the orthogonal complement of $\text{Span}\{\nabla L(c) \mid c \in C' \setminus \{c_1\}\}$ is in the tangent cone to $\text{Sys}(C \setminus \{c_1\})$. At a point $y \in \text{Sys}(C \setminus \{c_1\})$ that can be made arbitrarily close to x , for any choice of $c_2 \in (C \setminus \{c_1\})$, a stratum $\text{Sys}(C \setminus \{c_1, c_2\})$ adjacent to y , and hence also to x , can be constructed similarly. This construction can be iterated to prove the lemma. ■

It follows from Lemma 4.1 that when $x \in \text{Sys}(C)$ and $\text{Sys}(C)$ is a locally top-dimensional stratum of \mathcal{P}_g with $\{\nabla L(c)(x) \mid c \in C\}$ linearly independent, C is a minimal filling set of curves. Also under these assumptions, by Corollary 3.10 and local finiteness the cone of increase at x has nonempty intersection with the tangent cone to \mathcal{P}_g at x .

A stratum will be said to be *above* a point x if the stratum contains a sequence $\{x_i\}$ of points converging to x such that $f_{\text{sys}}(x_i) > f_{\text{sys}}(x)$. Similarly, a level set of f_{sys} is above x if the value of f_{sys} on this level set is larger than $f_{\text{sys}}(x)$.

Recall that $\text{Sys}(C) \subset E(C)$ for any set of systoles C .

Definition 4.2 (accidental systole). Suppose there exists an open set U in \mathcal{T}_g such that $\text{Sys}(C) \cap U$ is given by $E(C') \cap U$ for $C' \subsetneq C$ and f_{sys} is nonconstant on $\text{Sys}(C) \cap U$. The curves in $C \setminus C'$ will be called accidental systoles at points of $\text{Sys}(C) \cap U$.

The existence of accidental systoles means the linear independence condition in Lemma 4.1 breaks down.

Remark 4.3. In the presence of accidental systoles, the vector $v_{C'}$ can be set equal to v_C , and since the convex hull of $\{\nabla L(c)(x) \mid c \in C'\}$ is contained in the convex hull of $\{\nabla L(c)(x) \mid c \in C\}$, if $E(C')$ is balanced at x , this implies that $\text{Sys}(C)$ is balanced at x .

Remark 4.4. The author is not aware of any examples of accidental systoles. Possibly they do not exist.

Theorem 4.5. *Suppose a point $x \in \mathcal{P}_g$ has the property that the cone of increase of f_{sys} at x is full. Then the intersection of this cone of increase with the tangent cone of \mathcal{P}_g is nonempty.*

Proof. The theorem is proven by constructing a stratum above x in \mathcal{P}_g . The first step is to show the existence of a nonzero vector $v_{C'}(x)$ in the cone of increase, where C' is a filling subset of C , and for every $c \in C \setminus C'$, $L(c)$ is increasing faster in the direction of $v_{C'}(x)$ than the length of any curve $c' \in C'$. Suppose to begin with that x is not a minimum of C in $E(C)$.

Let n be the dimension of $\text{Span}\{\nabla L(c)(x) \mid c \in C\}$. Identify the subspace of $T_x\mathcal{T}_g$ spanned by $\{\nabla L(c)(x) \mid c \in C\}$ with \mathbb{R}^n via a vector space isomorphism that preserves the inner product on $T_x\mathcal{T}_g$ induced by \mathbf{g} . Find a hyperplane in \mathbb{R}^n containing the points represented by at least n of the vectors $\{\nabla L(c)(x) \mid c \in C\}$ and with the property that all other vectors in the set $\{\nabla L(c)(x) \mid c \in C\}$ represent points on the opposite side of the hyperplane from the zero vector. The vector in $T_x\mathcal{T}_g$ corresponding to the closest point on the hyperplane to the zero vector is $v_{C'}(x)$. It follows from Lemma 3.7 that C' fills. The set of curves C' cannot be eutactic at x , as this would contradict the assumption that the cone of increase at x is full. Consequently, $v_{C'}(x)$ is not zero.

When $\{\nabla L(c)(x) \mid c \in C'\}$ are not linearly independent, as a result of second and higher order rates of change of lengths, it does not necessarily follow that $v_{C'}(x)$ is in the tangent cone to a stratum $\text{Sys}(C')$ at x . In order to show that there does, nevertheless, exist a stratum in \mathcal{P}_g above x , a ‘‘Voronoi decomposition’’ of the unit tangent cone $UT_x\mathcal{T}_g$ will be defined, and properties of this decomposition invariant under arbitrarily small deformations in length will be used.

The decomposition² $\mathcal{V}(C, x)$ of $UT_x\mathcal{T}_g$ has a top-dimensional cell labelled by a curve $c \in C$ consisting of all $v \in UT_x\mathcal{T}_g$ for which

$$\mathbf{g}(v, -\nabla L(c)(x)) > \mathbf{g}(v, -\nabla L(c')(x)) \forall c' \in C \setminus \{c\}.$$

Smaller-dimensional cells of the decomposition are defined analogously, and labelled by subsets of C . Note that it follows from Proposition 3.14 and Remark 3.17 that every curve in C determines a top-dimensional cell of $\mathcal{V}(C, x)$.

In what follows, it will only be necessary to study the intersection of the cell decomposition with the unit vectors in the span of $\{\nabla L(c)(x) \mid c \in C\}$. This will be denoted by $\mathcal{V}(C, x)^{\text{Span}}$. It follows from Lemma 3.7 that zero-dimensional cells of $\mathcal{V}(C, x)^{\text{Span}}$ are labelled by filling subsets of C . It was shown that the intersection of $\mathcal{V}(C, x)^{\text{Span}}$ with the cone of increase contains a 0-dimensional cell labelled by C' .

When the dimension of $\text{Span}\{\nabla L(c)(x) \mid c \in C\}$ stays constant on a neighbourhood of x , define $\mathcal{V}(C, x, \epsilon)^{\text{Span}}$ as follows: scale the unit vectors in $UT_x\mathcal{T}_g \cap \text{Span}\{\nabla L(c)(x) \mid c \in C\}$ by ϵ , and take the image under the exponential map to obtain an embedded n -dimensional sphere \mathcal{S}_ϵ . Assume ϵ is small enough that local finiteness ensures that the systoles everywhere on \mathcal{S}_ϵ are contained in the set C . $\mathcal{V}(C, x, \epsilon)^{\text{Span}}$ is the cell decomposition of \mathcal{S}_ϵ induced by the stratification; a cell is labelled by a subset C'' of C if the points of the cell are contained in $\text{Sys}(C'')$.




Figure 7. A cell of small diameter in $\mathcal{V}(C, x, \epsilon)^{\text{Span}}$ can shrink to a point as ϵ is decreased to zero.

There is a sense in which $\mathcal{V}(C, x, \epsilon)^{\text{Span}}$ can be made arbitrarily close to $\mathcal{V}(C, x)^{\text{Span}}$ for sufficiently small ϵ . For small ϵ , the cells of $\mathcal{V}(C, x, \epsilon)^{\text{Span}}$ and $\mathcal{V}(C, x)^{\text{Span}}$ correspond except when a cell of small diameter in $\mathcal{V}(C, x, \epsilon)^{\text{Span}}$ shrinks to a point as ϵ is decreased to zero.

²It will not be shown that the decompositions used in this proof truly are cell decompositions, as this is cumbersome, and not essential. If the reader is concerned, it is possible to simply assume that unspecified subdivisions have been performed to obtain cell decompositions.

A lower-dimensional analogue of this is illustrated in Figure 7. In particular, there exists a 0-dimensional cell of $\mathcal{V}(C, x, \epsilon)^{\text{Span}}$ corresponding to a vector in the cone of increase of f_{sys} at x . This shows that at x , there is nonempty intersection of the cone of increase with the tangent cone of \mathcal{P}_g , as required.

When the dimension of $\text{Span}\{\nabla L(c) \mid c \in C\}$ does not stay constant on a neighbourhood of x , i.e., the dimension of the span increases, defining $\mathcal{V}(C, x, \epsilon)^{\text{Span}}$ involves constructing a cross section with tangent space that will be identified with $\text{Span}\{\nabla L(c)(x) \mid c \in C\}$. A cross section will be used that does not “see” the increase in the dimension of the span. The purpose of this is to obtain a stratum labelled by a subset of C the elements of which have lengths whose gradients at x have the same span as $\{\nabla L(c)(x) \mid c \in C\}$. On a neighbourhood of x in \mathcal{T}_g , choose $\text{Croppedspan}\{\nabla L(c) \mid c \in C\}$ to satisfy the following conditions:

- For some fixed $C' \subset C$ independent of y , $\text{Croppedspan}\{\nabla L(c)(y) \mid c \in C'\}$ is an n -dimensional subspace of $T_y \mathcal{T}_g$ spanned by $\{\nabla L(c)(y) \mid c \in C'\}$.
- When $y = x$, $\text{Croppedspan}\{\nabla L(c)(y) \mid c \in C'\} = \text{Span}\{\nabla L(c)(x) \mid c \in C'\}$.
- The projection of $\nabla L(c)(y)$ to the orthogonal complement of $\text{Croppedspan}\{\nabla L(c)(y) \mid c \in C'\}$ in $T_y \mathcal{T}_g$ for every $c \in C$ has length bounded from above by a function of distance from x that approaches zero as y approaches x .

The existence of $\text{Croppedspan}\{\nabla L(c) \mid c \in C\}$ will now be shown. Choose and fix any $C' \subset C$ satisfying the first two conditions. The third condition then follows from the fact that every curve $c \in C \setminus C'$ has length whose gradient becomes linearly dependent on $\{\nabla L(c)(x) \mid c \in C'\}$ at x .

Suppose \mathcal{B}'_α has the property that $\mathcal{B}'_\alpha \setminus x$ is foliated by \mathcal{S}_ϵ for $\epsilon \in (0, \alpha]$, and let $\phi: \mathcal{D}^n \rightarrow \mathcal{B}'_\alpha$ be a diffeomorphism of the n -dimensional unit ball \mathcal{D}^n into \mathcal{B}'_α . There is an isotopy of ϕ to a map with image a ball \mathcal{B}_α centered on x with tangent space at every point y given by $\text{Croppedspan}\{\nabla L(c)(y) \mid c \in C'\}$. Then $\mathcal{V}(C, x, \epsilon)^{\text{Span}}$ is defined to be the cell decomposition induced by the stratification on the sphere of radius ϵ in \mathcal{B}_α .

Using $\mathcal{V}(C, x, \epsilon)^{\text{Span}}$, it is then possible to argue as in the previous case that at x , there is nonempty intersection of the cone of increase with the tangent cone of \mathcal{P}_g . This is because there is a 0-dimensional cell in $\mathcal{V}(C, x, \epsilon)^{\text{Span}}$ above x labelled by a subset C' of C for which $\text{Span}\{\nabla L(c)(x) \mid c \in C'\} = \text{Span}\{\nabla L(c)(x) \mid c \in C\}$.

It remains to prove the theorem in the case that x is a minimum of C on $E(C)$ and the cone of increase at x is full. Let x_1, x_2, \dots be a sequence of points in \mathcal{T}_g with $\lim_{i \rightarrow \infty} x_i = x$ and such that no x_i is a minimum of C . The sequence is also assumed to be contained within a neighbourhood \mathcal{N} of x on which the systoles are all contained in C and every point has an open cone of directions in which the lengths of all curves in C are strictly increasing. Within \mathcal{N} , define a relative stratum $E^{\text{rel}}(C', d)$ to be the set of all points on $E(C', d)$ at which

$$L(c') + d|_{c'} < L(c) + d|_c \quad \forall c' \in C' \quad \text{and} \quad \forall c \in C \setminus C'.$$

The set of curves C' will be called relative systoles. The previous arguments show the existence of a relative stratum $E^{\text{rel}}(C'(x_i), d(x_i))$ above each point x_i , where $C'(x_i) \subset C$ is a filling set, and $d(x_i)$ is the value of d for which $E(C, d(x_i))$ contains x_i . Assume, in addition, that the points of the sequence have been chosen such that they all have relative strata given by the same subsets of C . Every point on such a relative stratum in the interior of \mathcal{N} has another relative stratum with filling relative systoles above it, defined using the tuple $d(x_i)$. In the limit as x_i approaches x , $d(x_i)$ approaches $d(x) = 0$, and these relative strata limit to strata in \mathcal{P}_g above x . This concludes the proof of the theorem. \blacksquare

Theorem 4.6. *Where a stratum of \mathcal{P}_g is locally top-dimensional it is balanced. If $x \in \mathcal{P}_g$ is not balanced, there are points directly above x in \mathcal{P}_g that are balanced.*

Proof. The theorem is proven by assuming x is in a locally top-dimensional stratum $\text{Sys}(C)$ of \mathcal{P}_g and using Proposition 3.20 to derive a contradiction when $\text{Sys}(C)$ is not balanced at x .

Suppose $x \in \text{Sys}(C)$ is in a locally top-dimensional stratum in \mathcal{P}_g that is not balanced at x . It follows that x is not a critical point of f_{sys} . If x were a boundary point of f_{sys} , as discussed in Section 5, these are known to have higher-dimensional balanced strata in \mathcal{P}_g above them. It is therefore possible to assume that the cone of increase at x is full.

On a neighbourhood of x , $\text{Sys}(C)$ is not contained in a level set of f_{sys} . If it were, Theorem 4.5 would give a stratum $\text{Sys}(C')$ above x in \mathcal{P}_g with $C' \subset C$ and for which there exists a $v_{C'}$ in the tangent cone to $\text{Sys}(C')$ at points sufficiently close to x . This shows it is possible to construct a stratum in \mathcal{P}_g directly above x and with larger dimension than $\text{Sys}(C)$, contradicting the assumption that x is in a locally top-dimensional stratum.

Denote by $c(f)$ the set of curves corresponding to the vertices of f , where f is the face of the convex hull of $\{\nabla L(c)(x) \mid c \in C\}$ from Proposition 3.20. The existence of a larger-dimensional stratum above $x \in \text{Sys}(C)$ is a consequence of Proposition 3.20 and the fact that every facet f_i of the convex hull of $\{\nabla L(c)(x) \mid c \in C\}$ determines a stratum $\text{Sys}(c'(f_i))$ adjacent to x , where $c'(f_i) \subset c(f_i)$ fills the same subsurface of \mathcal{S}_g as $c(f_i)$. To prove this fact about facets, note that the span of $\{\nabla L(c)(x) \mid c \in c(f_i)\}$ has a 1-dimensional orthogonal complement in $\text{Span}\{\nabla L(c)(x) \mid c \in C\}$. Denote by o_{f_i} a tangent vector in $T_x \mathcal{T}_g$ to this orthogonal complement. The fact that f_i is a facet is used to ensure that there is a choice of o_{f_i} with strictly positive inner product with every element of $\{\nabla L(c)(x) \mid c \in C \setminus c(f_i)\}$. To first order, $o(f_i)$ appears to be in the tangent cone at x to a stratum $\text{Sys}(c(f_i))$ adjacent to x . The construction in the proof of Theorem 4.5 then shows that o_{f_i} is in the tangent cone at x to a stratum $\text{Sys}(c'(f_i))$, where $c'(f_i) \subset c(f_i)$ fills the same subsurface as $c(f_i)$. The stratum $\text{Sys}(c'(f_i))$ has higher dimension than $\text{Sys}(C)$ at x , giving the required contradiction. ■

Remark 4.7. There are two obstacles in trying to prove that sets of systoles on locally top-dimensional strata of \mathcal{P}_g are all minimal filling sets. The first one is in ruling out accidental systoles. The second is that the converse of Lemma 2.5 does not hold. Not every filling subset C' of C determines a set of minima $\text{Min}(C')$ contained in $\partial \text{Min}(C)$. This could give rise to higher-dimensional analogues of the following: Let $C = \{c_1, c_2, c_3, c_4\}$ where $\{\nabla L(c)(x) \mid c \in C\}$ are linearly dependent and $v_C(x)$ is nonzero and in the convex hull of $\{\nabla L(c)(x) \mid c \in C\}$. Suppose also that $\{c_1, c_2\}$ is a minimal filling set, and c_1 and c_2 are the labels of a pair of vertices of the convex hull of $\{\nabla L(c)(x) \mid c \in C\}$ diagonally opposite each other. There seems to be no reason to expect that any of the faces of the convex hull determine a set of filling systoles.

5 Critical points and boundary points of the systole function

The purpose of this section is to study the structure of \mathcal{P}_g around critical and boundary points of f_{sys} . Section 7 later gives an example that is interesting in its own right, illustrating the theorem proven in this section. The exposition begins with critical points.

Recalling the definition of a topological Morse function, on a neighbourhood of a critical point p of index j there is a homeomorphism ψ that fixes the critical point. This homeomorphism has the property that there exists a smooth chart on a neighbourhood of p such that

$$f_{\text{sys}} \circ \psi(x) - f_{\text{sys}}(p) = -x_1^2 - \cdots - x_j^2 + x_{j+1}^2 + \cdots + x_n^2.$$

In this sense, the level sets of f_{sys} on a neighbourhood of a critical point of index j are the same, up to homeomorphism, as for a critical point of index j of a smooth Morse function. It was shown in [1] that critical points of f_{sys} occur exactly where $\text{Sys}(C)$ intersects $\text{Min}(C)$. The index j is then equal to the dimension of the span of $\{\nabla L(c) \mid c \in C\}$ at the critical point.

Suppose p is a critical point of index j in the intersection of $\text{Sys}(C)$ with $\text{Min}(C)$. It follows from local finiteness that there is a neighbourhood of p in \mathcal{T}_g on which the systoles are contained in the set C . Convexity of length functions implies that the subspace $\{\nabla L(c) \mid c \in C\}^\perp$ is contained in the cone of increase of f_{sys} at p . Within $\text{Min}(C)$, in every direction there is by definition a curve or curves in C whose length is decreasing away from p to first order. For a point q in $\text{Min}(C)$ near p , the systoles are the subset of curves in C whose lengths are decreasing fastest away from p in the direction of q . It follows that $j \leq |C| - 1$ and on a neighbourhood of p , the level set of f_{sys} containing p intersects $\text{Min}(C)$ in the single point p . This is illustrated schematically in Figure 8 for the toy example with two systoles.




Figure 8. The level sets of the systole function passing through the critical point are shown in green. A level set at higher value of f_{sys} is shown in violet. The red arrows show the direction in which f_{sys} is increasing. The simplices σ_1 and σ_2 are shown in blue.

For p in $\overline{\text{Sys}(C)}$, the tangent cone to $\text{Sys}(C)$ at p will be denoted by $TC_p\text{Sys}(C)$.

Proposition 5.1. *Fix a critical point $p \in \mathcal{P}_g$ of index j with set of systoles C . A neighbourhood of p in \mathcal{P}_g contains a piecewise smooth cell of codimension j on which f_{sys} is increasing to second order away from the critical point. The tangent space of this cell at p is $\{\nabla L(c)(p) \mid c \in C\}^\perp$. The subspace spanned by $\{\nabla L(c)(p) \mid c \in C\}$ consists of directions in which the systole function is decreasing away from p .*

Proof. By local finiteness, convexity of length functions and the definition of $\text{Min}(C)$, the subspace spanned by $\{\nabla L(c)(p) \mid c \in C\}$ consists of directions in which f_{sys} is decreasing away from p to first order and all the directions in which f_{sys} is increasing at p to second order are exactly given by $\{\nabla L(c)(p) \mid c \in C\}^\perp$.

The subspace $\{\nabla L(c)(p) \mid c \in C\}^\perp$ of $T_p\mathcal{T}_g$ is tangent to the level set of f_{sys} containing p . Recall that there exists a triangulation of \mathcal{P}_g compatible with the stratification. Denote by $\sigma = \bigcup_{i=1}^n \sigma_i$ the union of simplices for which $TC_p\sigma = \{\nabla L(c)(p) \mid c \in C\}^\perp$. The existence of σ is guaranteed by two observations: firstly, every vector in the cone of increase at p is in $\{\nabla L(c) \mid c \in C\}^\perp$ and hence tangent to the level sets passing through p . Secondly, the restriction of f_{sys} to the interior of a stratum or a simplex is smooth, so the restriction of the level sets to the interiors of the simplices can only contain the smooth pieces of the piecewise smooth level sets. This is illustrated schematically in Figure 8.

It will now be shown that, on a neighbourhood of p , σ is contained in \mathcal{P}_g . This starts by showing that each of the locally top-dimensional simplices $\sigma_1, \dots, \sigma_n$ of σ is necessarily contained in locus that is balanced near p . Let x be a point in the interior of a simplex σ_i in σ where the systoles on σ_i are given by $C_i \subset C$, and let V be the subspace of $T_x\mathcal{T}_g$ given by the orthogonal complement of $v_{C_i}(x)$. From the known geometry of the level sets near p , it follows that any vector in V is either tangent to a level set of f_{sys} restricted to σ_i or to a geodesic γ with $\gamma(0) = x$ passing through level sets below x for $t \in (-\epsilon, \epsilon)$. In Figure 8, γ is approximately parallel to the y axis. Local finiteness then implies that C_i is V -eutactic, and hence σ_i is balanced.




Figure 9. The intersection of level sets near a boundary point p of f_{sys} . The vector $v_{\{c_1, c_2, c_3\}}$ shows a direction in which f_{sys} is increasing away from p to second order.

Consider a sequence of points in σ_i approaching p from above. In the limit as p is approached, $f_{\text{sys}}|_{\sigma_i}$ is increasing only to second order. At points in the sequence, the balanced sets of gradients $\{\nabla L(c) \mid c \in C_i\}$ therefore become arbitrarily close to being eutactic. This implies that C_i is a filling set, since otherwise a contradiction is obtained by invoking Proposition 2.2 and observing that in the limit as p is approached, the open cone of directions in which the lengths of curves in C_i are increasing is empty. Since \mathcal{P}_g is closed, this concludes the proof that $\sigma \subset \mathcal{P}_g$. ■

Remark 5.2. Proposition 5.1 implies that for a critical point $p \in \text{Sys}(C)$, codimension of $\mathcal{P}_g \leq |C| - 1$ or dimension of $\mathcal{P}_g \geq 6g - |C| - 5$. This does not require Thurston's construction from [21] but as explained above, is a consequence of Akrouit's characterisation of critical points of f_{sys} and Proposition 2.2. This has not been well understood in the literature, with conjectures made and lengthy proofs given of special cases of this statement. When p is not a critical point, but is contained in a locally top-dimensional stratum of \mathcal{P}_g , by Theorem 4.6, the cardinality of the systole set similarly gives an upper bound on the codimension of the stratum and hence of \mathcal{P}_g .

Boundary points of the systole function and their structure. Recall that a boundary point of f_{sys} is a point at which $\text{Sys}(C)$ intersects $\text{Min}(C')$, for $C' \subsetneq C$. Since $\text{Min}(C')$ is empty unless C' fills, boundary points of f_{sys} are all necessarily contained in \mathcal{P}_g . Boundary points of f_{sys} were also shown in [19] to be isolated and hence there are at most finitely many modulo the action of Γ_g .

A boundary point of f_{sys} is not a critical point, but it also has no open cone in which f_{sys} is increasing. The level sets of the curves in C intersect as illustrated in Figure 9. The function f_{sys} increases away from a boundary point of f_{sys} only to second order.

Ignoring the curves in $C \setminus C'$, around p , it is possible to construct the same simplices $\sigma_1, \dots, \sigma_n$ as in Proposition 5.1. For any vector in the tangent cone to σ at p giving a direction in which the length of a curve in $C \setminus C'$ is decreasing away from p , the set of systoles is in $C \setminus C'$. In all other directions near p , the behaviour of f_{sys} is identical to that near a critical point. Note that all directions in which f_{sys} is increasing away from p are in the tangent cone to \mathcal{P}_g at p .

6 Thurston's vector field

In the 80s, Thurston gave a talk at Princeton on the construction of a Γ_g -equivariant deformation retraction of \mathcal{T}_g onto \mathcal{P}_g . This was done by showing the existence of a vector field that generates a flow used in the construction of the deformation retraction. Notes from this talk are given in [21]. The purpose of this section is to explain what Theorems 4.6, 4.5 and Proposition 5.1 imply about this vector field.

The vector field from [21] will first be defined. Choose and fix an $\epsilon > 0$. For *finite* C , define

$$U_C = \{x \in \mathcal{T}_g \mid C \text{ is exactly the set of curves of length less than } f_{\text{sys}} + |C|\epsilon\},$$

where $|C|$ is the cardinality of the set C . It was shown that when ϵ is sufficiently small, the set

$$\{U_C \mid C \text{ is a finite set}\} \tag{6.1}$$

covers \mathcal{T}_g . This uses the well-known result (see, for example, [14]) that the number $n(L)$ of curves on a hyperbolic surface corresponding to a point $x \in \mathcal{T}_g$ of length less than or equal to L increases faster than linearly with L . Consequently, ϵ must be made small otherwise a set C of curves of length less than $f_{\text{sys}}(x) + |C|\epsilon$ quickly becomes infinite, and the set in equation (6.1) is not a cover.

On U_C , for nonfilling C , a vector field X_C is defined to be the vector of unit length such that

$$\sum_{c \in C} \log X_C(L(c))$$

is maximised. When C fills, X_C is defined to be zero. Note that this formula assumes real logs, i.e., X_C is a derivation whose evaluation on each of the functions $\{L(c) \mid c \in C\}$ is strictly positive. A smooth vector field X_ϵ was obtained with the help of a partition of unity subordinate to the cover. This gives a vector field, nonvanishing away from a neighbourhood of \mathcal{P}_g , and in a direction of increasing f_{sys} . For any $x \in \mathcal{T}_g \setminus \mathcal{P}_g$, by choosing ϵ sufficiently small, it can be ensured that $X_\epsilon(x) \neq 0$. More details on the construction of X_ϵ can be found in [21].

Definition 6.1 (gap function g_{sys}). The *gap function* $g_{\text{sys}}: \mathcal{T}_g \rightarrow \mathbb{R}_+$ is given by

$$x \mapsto \inf\{r \in \mathbb{R} \mid \text{the curves of length at most } f_{\text{sys}}(x) + r \text{ fill}\}.$$

By local finiteness, g_{sys} is zero only on \mathcal{P}_g , and increases away from \mathcal{P}_g on some neighbourhood of \mathcal{P}_g .

Remark 6.2. Close to \mathcal{P}_g , g_{sys} behaves like a measure of distance from \mathcal{P}_g . Globally it is not an easy function to work with. For example, g_{sys} is not a topological Morse function. This follows from the observation that at some points, the cone of increase of g_{sys} is obtained by taking a *union* of cones of increase corresponding to different choices of longest curve in a filling set. The cone of increase can therefore become disconnected when the zero vector is removed.

Proposition 6.3. *There is a neighbourhood \mathcal{N}_g of \mathcal{T}_g such that for every point $x \in \mathcal{N}_g \setminus \mathcal{P}_g$, choosing ϵ small enough to ensure that $X_\epsilon(x) \neq 0$ ensures that g_{sys} is strictly decreasing in the direction of $X_\epsilon(x) \neq 0$.*

Proof. The neighbourhood \mathcal{N}_g of \mathcal{P}_g is required to satisfy the following conditions

- \mathcal{N}_g is a regular, Γ_g -equivariant neighbourhood of \mathcal{P}_g .
- \mathcal{N}_g is small enough so that g_{sys} is increasing radially away from \mathcal{P}_g on \mathcal{N}_g .
- \mathcal{N}_g is small enough so that local finiteness guarantees that the systoles on \mathcal{P}_g determine the systoles within \mathcal{N}_g .

For sufficiently small \mathcal{N}_g , near critical points and boundary points of f_{sys} , the proposition follows from Proposition 5.1 and the characterisation of boundary points of f_{sys} given at the end of Section 5. Away from critical points and boundary points, for sufficiently small \mathcal{N}_g the proposition follows from Theorems 4.5 and 4.6. In \mathcal{N}_g the gradients of the systoles in the top-dimensional strata of \mathcal{T}_g point radially inwards towards \mathcal{P}_g , and the same is then true of the linear combinations of these gradients where the top-dimensional strata meet along lower-dimensional strata. Where cells of \mathcal{P}_g come together as shown schematically in Figure 10, Theorem 4.6 rules out examples such as in the left hand side of Figure 10, where there is a conflict between increasing f_{sys} and decreasing g_{sys} . Theorems 4.5 and 4.6 also rule out a vector field pointing away from the boundary of \mathcal{P}_g at x . ■




Figure 10. The black lines/shaded region represent points on \mathcal{P}_g , and the arrows represent the direction of increasing f_{sys} . Where cells of \mathcal{P}_g come together, Theorem 4.6 ensures there is no conflict between increasing f_{sys} and decreasing g_{sys} ; the behaviour schematically represented by the figure on the left is not possible.




Figure 11. The set of 6 systoles at the critical point in the example are shown on the right. The left side of the figure shows a fundamental domain with edges lying along the systoles, and with numbers indicating the gluing maps (not the labels of the curves). This figure is taken from [8].

7 Example: Schmutz's critical point

The purpose of this subsection is to discuss an application of Proposition 5.1 in understanding the geometry of \mathcal{P}_2 . This critical point is due to Schmutz [19] and is the first of a family of examples, one in each genus, of critical points of coindex equal to the virtual cohomological dimension of the mapping class group. This family of examples plays a key role in the study of the Steinberg module, [8].

If p is a critical point of index j , and C is the set of systoles at p , a *descendent* C' of C is a subset of C with the following properties:

- (1) C' is eutactic at p .
- (2) The span of $\{\nabla L(c) \mid c \in C'\}$ is equal to the span of $\{\nabla L(c) \mid c \in C\}$.
- (3) C' does not have a proper subset satisfying the first two conditions.

Note that the cardinality of a descendent C' is $j + 1$.

Let C be the set of curves on the right hand side of Figure 11. It was shown in [19, Theorem 36] that there is a critical point p of index 3 at which this set of curves is realised as the set of systoles. At p , the systoles intersect at right angles, cutting the surface into four regular hexagons. In [19], it was shown that $\text{Min}(C)$ is a cell of dimension 3 with empty boundary.

Deleting any pair of intersecting curves in C gives a minimal filling set of curves, all of which are descendents of C . The complete set of descendents is

$$\begin{aligned} C_1 &= \{c_1, c_2, c_3, c_4\}, & C_2 &= \{c_2, c_3, c_4, c_5\}, & C_3 &= \{c_3, c_4, c_5, c_6\}, \\ C_4 &= \{c_4, c_5, c_6, c_1\}, & C_5 &= \{c_5, c_6, c_1, c_2\}, & C_6 &= \{c_6, c_1, c_2, c_3\}. \end{aligned}$$




Figure 12. The arrangement of strata in \mathcal{P}_2 on a neighbourhood of the critical point p in the example. The 3-sphere is piecewise smooth, and has a symmetry group given by the dihedral group D_6 that acts transitively on the top-dimensional strata of \mathcal{P}_2 adjacent to p .

For each descendent C_i , $i = 1, \dots, 6$, the locus $E(C_i)$ is an embedded submanifold of dimension 3, intersecting $\text{Min}(C) = \text{Min}(C_i)$ in the single point p , with $T_p E(C_i)$ given by $\text{Span}\{\nabla L(c) \mid c \in C_i\}^\perp = \text{Span}\{\nabla L(c) \mid c \in C\}^\perp$. This follows from the results in Section 3, or on a neighbourhood of p by the rank calculations in [19].

Near p , \mathcal{P}_2 could potentially lie along any $E(C_i)$; this is determined by second order behaviour. It is a consequence of the main construction in [13] that $E(C) = \bigcap_{i=1}^6 E(C_i)$ has a connected component consisting of a one-dimensional embedded submanifold of \mathcal{T}_2 containing p . On a neighbourhood of p , by local finiteness, the points of $E(C)$ are in $\text{Sys}(C)$. The same construction from [13] also implies that along $E(C) \setminus \{p\}$ the set of gradients $\{\nabla L(c) \mid c \in C\}$ is linearly independent. It follows from the argument in Lemma 4.1 that every proper filling subset of C determines a stratum of \mathcal{P}_2 adjacent to $E(C)$; these are arranged as shown in Figure 12.

The automorphism group of the hyperbolic surface corresponding to p contains a subgroup of the mapping class group isomorphic to the dihedral group D_6 , and the locus $E(C)$ passing through p is the fixed point set of this subgroup. It follows that this connected component of $E(C)$ is a geodesic with respect to any mapping class group-equivariant metric on \mathcal{T}_2 , such as the Teichmüller metric or the Weil–Petersson metric.

The set of minima $\text{Min}(C)$ also has a simple parameterisation. In genus 2, $\text{Min}(C)$ is the set of points at which the curves in C intersect at right angles.

Theorem 7.1 ([17, Theorem 3.5.14]). *For any $a, b, c \in \mathbb{R}_+$, there is a right-angled hyperbolic convex hexagon, unique up to congruence, with alternate sides of length a, b and c .*

Using Theorem 7.1, from the left side of Figure 11, it is not difficult to verify that by Theorem 7.1 the lengths of 3 alternate edges of any one of the hexagons of $S_g \setminus C$ completely parameterise $\text{Min}(C)$.

Acknowledgements

The author would like to thank Stavros Garoufalidis, Don Zagier and a number of very helpful anonymous referees for comments that led to significant improvements.

References

- [1] Akrouit H., Singularités topologiques des systoles généralisées, *Topology* **42** (2003), 291–308.
- [2] Barvinok A., Lattice points, polyhedra, and complexity, in Geometric Combinatorics, *IAS/Park City Math. Ser.*, Vol. 13, [American Mathematical Society](#), Providence, RI, 2007, 19–62.
- [3] Bers L., Nielsen extensions of Riemann surfaces, *Ann. Acad. Sci. Fenn. Ser. A I Math.* **2** (1976), 29–34.
- [4] Bers L., An inequality for Riemann surfaces, in Differential Geometry and Complex Analysis, [Springer](#), Berlin, 1985, 87–93.
- [5] Bestvina M., Bromberg K., Fujiwara K., Souto J., Shearing coordinates and convexity of length functions on Teichmüller space, *Amer. J. Math.* **135** (2013), 1449–1476, [arXiv:2013.0049](#).
- [6] Brion M., Vergne M., Lattice points in simple polytopes, *J. Amer. Math. Soc.* **10** (1997), 371–392.
- [7] Foote R.L., Regularity of the distance function, *Proc. Amer. Math. Soc.* **92** (1984), 153–155.
- [8] Irmer I., Explicit generators of the Steinberg module of the mapping class group, [arXiv:2312.08721](#).
- [9] Irmer I., Schmutz–Thurston duality, in preparation.
- [10] Kerckhoff S.P., The Nielsen realization problem, *Ann. of Math.* **117** (1983), 235–265.
- [11] Lee J.M., Introduction to smooth manifolds, 2nd ed., *Grad. Texts in Math.*, Vol. 218, [Springer](#), New York, 2013.
- [12] Lojasiewicz S., Triangulation of semi-analytic sets, *Ann. Scuola Norm. Sup. Pisa Cl. Sci.* **18** (1964), 449–474.
- [13] Mathieu O., Estimating the dimension of the Thurston spine, [arXiv:2310.15618](#).
- [14] Mirzakhani M., Growth of the number of simple closed geodesics on hyperbolic surfaces, *Ann. of Math.* **168** (2008), 97–125.
- [15] Morse M., Topologically non-degenerate functions on a compact n -manifold M , *J. Anal. Math.* **7** (1959), 189–208.
- [16] Papadopoulos A., Th  ret G., Shortening all the simple closed geodesics on surfaces with boundary, *Proc. Amer. Math. Soc.* **138** (2010), 1775–1784, [arXiv:0905.4661](#).
- [17] Ratcliffe J., Foundations of hyperbolic manifolds, 2nd ed., *Grad. Texts in Math.*, Vol. 149, [Springer](#), New York, 2006.
- [18] Riera G., A formula for the Weil–Petersson product of quadratic differentials, *J. Anal. Math.* **95** (2005), 105–120.
- [19] Schmutz Schaller P., Systoles and topological Morse functions for Riemann surfaces, *J. Differential Geom.* **52** (1999), 407–452.
- [20] Schmutz Schaller P., Riemann surfaces with longest systole and an improved Voronoi algorithm, *Arch. Math. (Basel)* **76** (2001), 231–240.
- [21] Thurston W., A spine for Teichm  ller space, Preprint, 1985.
- [22] Wolpert S.A., An elementary formula for the Fenchel–Nielsen twist, *Comment. Math. Helv.* **56** (1981), 132–135.
- [23] Wolpert S.A., Geodesic length functions and the Nielsen problem, *J. Differential Geom.* **25** (1987), 275–296.
- [24] Wolpert S.A., Geometry of the Weil–Petersson completion of Teichm  ller space, in Surveys in Differential Geometry, *Surv. Differ. Geom.*, Vol. 8, [International Press](#), Somerville, MA, 2003, 357–393, [arXiv:math.DG/0502528](#).