HIERARCHICAL HYPERBOLICITY OF ADMISSIBLE CURVE GRAPHS AND THE BOUNDARIES OF MARKED STRATA

AARON CALDERON AND JACOB RUSSELL

Abstract. We show that for any surface of genus at least 3 equipped with any choice of framing, the graph of non-separating curves with winding number 0 with respect to the framing is hierarchically hyperbolic but not Gromov hyperbolic. We also describe how to build analogues of the curve graph for marked strata of abelian differentials that capture the combinatorics of their boundaries, analogous to how the curve graph captures the combinatorics of the augmented Teichmüller space. These curve graph analogues are also shown to be hierarchically, but not Gromov, hyperbolic.

1. Introduction

The moduli space $\mathbb{P}\Omega\mathcal{M}_g$ of projectivized genus g Abelian differentials is a bundle over the usual moduli space $\mathcal{M}_g$ of genus g Riemann surfaces. It decomposes into *strata*, subvarieties which parametrize differentials with a given number and order of zeros and which are the ambient theatre for Teichmüller dynamics. The overall structure of strata is still poorly understood, and recent work has been largely guided by the following:

Question 1.1. *How similar are strata and $\mathcal{M}_g$?*

There has been a great deal of success constructing compactifications of strata akin to the Deligne–Mumford compactification of $\mathcal{M}_g$ [EMZ03, BCG$^+$18, BCG$^+$19]. The structure of these boundaries can then be used to compute constants of dynamical interest [EMZ03], perform intersection theory on strata [CMSZ20], and compute their Euler characteristics [CMZ22], among many other things.

Another version of Question 1.1 deals with their fundamental groups. Recall that $\mathcal{M}_g$ is an (orbifold) $K(\pi, 1)$ for the usual mapping class group $\mathrm{Mod}(S)$, the group of homeomorphisms of the surface up to homotopy. By analogy, Kontsevich predicted that each connected component of a stratum should be a $K(\pi, 1)$ for "some mapping class group" [KZ].

In [CS22], the first author and Salter showed that the fundamental groups of strata surject onto *framed mapping class groups* $\mathrm{FMod}(S, \phi)$, the stabilizers inside $\mathrm{Mod}(S)$ of trivializations $\phi : TS \cong S \times \mathbb{R}^2$ (see §2 for a formal definition). Apisa, Bainbridge, and Wang subsequently showed that certain strata of twisted 1-forms are $K(\pi, 1)$'s for framed mapping class groups [ABW23], and while this article was under revision, Salter showed that the fundamental groups of many strata are indeed framed mapping class groups [Sal25].

A group-theoretic analogue of Question 1.1 is thus:

Question 1.2. *How similar are* $\mathrm{FMod}(S, \phi)$ *and* $\mathrm{Mod}(S)$*?*

1.1. Curve graphs and strata.

This paper initiates the study of Questions 1.1 and 1.2 from the coarse-geometric perspective by analyzing the geometry of certain curve graphs.

The classical curve graph $\mathscr{C}(S)$ has a vertex for each isotopy class of essential simple closed curve on an orientable surface and an edge when two curves can be realized disjointly [Har81]. In addition to this topological interpretation, this graph also plays the role of (the 1-skeleton of) a Tits building for Teichmüller space $\mathcal{T}_g$, recording the incidences of top-dimensional boundary strata of the *augmented Teichmüller space*, a certain bordification of $\mathcal{T}_g$ that "lifts" the Deligne–Mumford compactification of $\mathcal{M}_g$ (see §5.2).

Masur and Minsky famously proved that $\mathscr{C}(S)$ is Gromov hyperbolic [MM98]. This marquee result has far-reaching implications for the coarse geometry of the mapping class group [Iva97, MM00], the geometry of Teichmüller space [MM98, Raf05], and the structure of hyperbolic 3-manifolds [Min10, BCM12]. More generally, the geometry of curve graphs has proven useful in a variety of settings; examples of this paradigm include relationships between the pants graph/the Weil–Petersson metric on $\mathcal{T}_g$ [Bro03, BF06], the Torelli complex and separating curve graph/the Torelli subgroup and the Johnson kernel [FI05, BM04], and the disk graph/the handlebody group and Heegaard splittings [Hen20, MS13].

As a first step towards Question 1.2, we study a topological analogue of $\mathscr{C}(S)$ that takes the framing into account. Any framing $\phi\colon TS \cong S \times \mathbb{R}^2$ can be used to measure the winding number of a smooth, oriented curve in S by lifting the curve to TS via its tangent vector, projecting to the second coordinate, then measuring the winding number of the image about $0 \in \mathbb{R}^2$. A simple closed curve on S is *admissible* for ϕ if it is nonseparating and has zero winding number, and the *admissible curve graph* $\mathscr{C}_{\mathrm{adm}}(S, \phi)$ is the subgraph of $\mathscr{C}(S)$ spanned by admissible curves.

The framed mapping class group $\mathrm{FMod}(S, \phi)$ preserves the winding number of every curve, hence acts on $\mathscr{C}(S)$ with infinitely many orbits of vertices. In contrast, $\mathrm{FMod}(S, \phi)$ acts on $\mathscr{C}_{\mathrm{adm}}(S, \phi)$ with finitely many orbits of vertices and edges (Proposition 2.10), indicating that the admissible curve graph is better adapted to study $\mathrm{FMod}(S, \phi)$.

Our first main result is that the admissible curve graph is *not* Gromov hyperbolic, but does possess a generalized notion of hyperbolicity.

Theorem A. *For any surface $S = S_{g,n}$ of genus $g \geq 3$ and any framing ϕ of S, the admissible curve graph $\mathscr{C}_{\mathrm{adm}}(S, \phi)$ is hierarchically hyperbolic but not Gromov hyperbolic.*

Hierarchical hyperbolicity was introduced by Behrstock, Hagen, and Sisto to unify similarities between the coarse geometry of mapping class groups, Teichmüller spaces, and right-angled Artin groups [BHS17b]. Briefly, this framework allows one to understand the geometry of a space by projecting it onto a collection of Gromov hyperbolic spaces. The presence of "orthogonal" projections leads to quasi-isometrically embedded flats, hence a failure of Gromov hyperbolicity.

We can also define a geometric analogue of $\mathscr{C}(S)$ that captures the intersection pattern of the boundary of a marked stratum. More precisely, since holomorphic differentials are determined up to scaling by the order and position of their zeros, any stratum component $\mathcal{H} \subset \mathbb{P}\Omega^1\mathcal{M}_g$ may be identified with a subvariety of $\mathcal{M}_{g,n}$, the moduli space of genus g Riemann surfaces with n marked points. Let us conflate $\mathcal{H}$ with this subvariety.

Take any non-hyperelliptic stratum component $\mathcal{H} \subset \mathcal{M}_{g,n}$ and let $\mathcal{H}_\phi$ be any component of the preimage of $\mathcal{H}$ in $\mathcal{T}_{g,n}$. By the main theorem of [CS22], so long as $g \geq 5$, then $\mathcal{H}_\phi$

consists of those marked abelian differentials such that the pullback of any directional vector field along the marking is isotopic to a fixed framing ϕ; see Section 5.1. Consider the closure $\overline{\mathcal{H}_\phi}$ of $\mathcal{H}_\phi$ in the augmented Teichmüller space $\overline{\mathcal{T}_{g,n}}$. We now define a graph $\mathscr{C}(\overline{\mathcal{H}_\phi})$ whose vertices are those multicurves γ such that $\overline{\mathcal{H}_\phi} \cap \mathcal{T}_{g,n}(\gamma) \neq \emptyset$, where $\mathcal{T}_{g,n}(\gamma)$ is the boundary stratum of $\overline{\mathcal{T}_{g,n}}$ in which γ is pinched. Two vertices γ and δ are connected by an edge if the closure of $\overline{\mathcal{H}_\phi} \cap \mathcal{T}_{g,n}(\gamma)$ in $\overline{\mathcal{T}_{g,n}}$ meets $\mathcal{T}_{g,n}(\delta)$ or vice-versa. The intricate structure of the boundary of $\overline{\mathcal{H}_\phi}$ means there are various other natural ways to define this graph (see Section 5.4), but they almost all turn out to be quasi-isometric to $\mathscr{C}(\overline{\mathcal{H}_\phi})$.

The geometry of $\mathscr{C}(\overline{\mathcal{H}_\phi})$ is closely linked to that of $\mathscr{C}_{\text{adm}}(S,\phi)$, and using Theorem A plus structural results about compactifications of strata [BCG+18, BCG+19], we prove:

Theorem B. *For any non-hyperelliptic stratum component $\mathcal{H} \subset \mathbb{P}\Omega^1\mathcal{M}_g$ with $g \geq 5$, the graph $\mathscr{C}(\overline{\mathcal{H}_\phi})$ is hierarchically hyperbolic but not Gromov hyperbolic.*

Remark 1.3. As shown in [CS22, Corollary 1.2], admissible curves are exactly the core curves of cylinders on surfaces in $\mathcal{H}_\phi$. One can also construct a partial bordification of $\mathcal{H}_\phi$ in which only cylinders are allowed to degenerate; the combinatorics of how this space meets $\partial\overline{\mathcal{T}_{g,n}}$ then correspond to $\mathscr{C}_{\text{adm}}(S,\phi)$. Thus Theorem A can also be interpreted as a statement about the coarse geometry of $\mathcal{H}_\phi$.

Remark 1.4. Our restriction to non-hyperelliptic components is because the hyperelliptic ones do not exhibit new phenomena. Indeed, hyperelliptic stratum components are essentially strata of quadratic differentials on $\mathbb{CP}^1$, which are in turn parametrized by their poles and zeros. Thus we can understand compactifications of hyperelliptic stratum components entirely in terms of the Deligne–Mumford compactification of $\mathcal{M}_{0,n}$.

Remark 1.5. The restriction to $g \geq 3$ in Theorem A is because for $g = 1, 2$ the admissible curve graph is not necessarily connected. The restriction to $g \geq 5$ in Theorem B comes from the fact that the main theorem of [CS22] relating $\pi_1(\mathcal{H})$ and $\text{FMod}(S,\phi)$ only applies for $g \geq 5$. In Section 5 we give a (slightly circuitous) definition of $\mathscr{C}(\overline{\mathcal{H}_\phi})$ that agrees with the one given above for $g \geq 5$ and for which Theorem B holds in genus 3 and 4. In particular, all of the proofs in this paper hold for $g \geq 3$.

Remark 1.6. The global residue condition of [BCG+19] does not play a role in our proof. In light of the connection between logarithmic and generalized multi-scale differentials [CGH+25], it would be interesting to know if there is a corresponding version of Theorem B in the logarithmic setting.

Curve graph techniques have been used successfully to study certain $\mathsf{SL}_2\mathbb{R}$-invariant subvarieties of of the moduli space of abelian differentials: [Tan21] proved that Veech groups are undistorted in $\text{Mod}(S)$, [RS09] proved a similar result for covering constructions, and [AHW24] used curve graphs to study the geometry of totally geodesic subvarieties of Teichmüller space. It is our hope that the tools developed in this paper will yield insights into both the intrinsic and extrinsic geometry of framed mapping class groups and strata. For example, we ask:

Question 1.7. *Is $\text{FMod}(S,\phi)$ distorted in $\text{Mod}(S)$? Are strata distorted in $\mathcal{M}_{g,n}$?*

1.2. **Outline of proof and paper.** To prove Theorems A and B, we need to exhibit projections from $\mathscr{C}_{\mathrm{adm}}(S,\phi)$ and $\mathscr{C}(\overline{\mathcal{H}_\phi})$ to Gromov hyperbolic spaces. In both settings, we use Masur and Minsky's subsurface projection maps to the curve graphs of *witnesses* — subsurfaces of S that intersect every admissible curve. This approach was inspired by work of Vokes, who showed that a wide variety of graphs of curves are hierarchically hyperbolic using their subsurface projection maps to witnesses [Vok22]. Vokes first uses the set of witnesses to build a hierarchically hyperbolic "model graph" $\mathcal{K}$, then shows that if the graph of curves admits a cocompact action of $\mathrm{Mod}(S)$ then it is quasi-isometric to $\mathcal{K}$.

To prove Theorem A, we construct a hierarchically hyperbolic model $\mathcal{K}$ for $\mathscr{C}_{\mathrm{adm}}(S,\phi)$ à la Vokes (Section 3). However, we cannot employ her quasi-isometry as $\mathscr{C}_{\mathrm{adm}}(S,\phi)$ does not admit an action by all of $\mathrm{Mod}(S)$ and the action of $\mathrm{FMod}(S,\phi)$ on $\mathcal{K}$ is not sufficiently cofinite to adapt her argument. Instead, we construct a novel quasi-isometry $\mathcal{K} \to \mathscr{C}_{\mathrm{adm}}(S,\phi)$ via the graph $\mathcal{G}$ of genus-separating curves (Section 4). The graph $\mathcal{G}$ can be quasi-isometrically realized as a "blow-up" of $\mathcal{K}$, while $\mathscr{C}_{\mathrm{adm}}(S,\phi)$ is quasi-isometric to a "cone-off" of $\mathcal{G}$. To build the map $\mathcal{K} \to \mathscr{C}_{\mathrm{adm}}(S,\phi)$, we show that the blown-up subsets from $\mathcal{K} \to \mathcal{G}$ coarsely match the coned-off subsets from $\mathcal{G} \to \mathscr{C}_{\mathrm{adm}}(S,\phi)$. This step requires some fairly delicate computations with curves on surfaces.

Theorem B follows by constructing a quasi-isometric model for $\mathscr{C}(\overline{\mathcal{H}_\phi})$ entirely in terms of framing data. This requires unpacking some of the finer structure of the space of multiscale differentials [BCG$^+$19]. These steps are accomplished in Section 5. In this section, we also build a suite of graphs whose definitions capture different facets of the incidences of components of $\partial\overline{\mathcal{H}_\phi}$ and topological properties of the indexing multicurves.

In the final Section 6, we show that $\mathscr{C}(\overline{\mathcal{H}_\phi})$ and the graphs from Section 5 are all quasi-isometric, and that they are quasi-isometric to a Vokes model graph $\overline{\mathcal{K}}$. Again, there is not sufficient transitivity to apply Vokes's methods, and the construction of a quasi-isometry is quite subtle. The graph $\overline{\mathcal{K}}$ is an $\mathrm{FMod}(S,\phi)$–equivariant cone-off of the model $\mathcal{K}$ for $\mathscr{C}_{\mathrm{adm}}(S,\phi)$, and the inclusion $\mathscr{C}_{\mathrm{adm}}(S,\phi) \hookrightarrow \mathscr{C}(\overline{\mathcal{H}_\phi})$ is also an equivariant cone-off. As in the case of Theorem A, the main difficulty is then showing that these two cone-offs coarsely match.

A common theme running throughout this paper is that if one understands the $\mathrm{FMod}(S,\phi)$ action on configurations of curves and subsurfaces well enough, then many surface-topological arguments can be adapted to the framed setting with a little extra care and effort. As such, we prove a number of transitivity results (Propositions 2.10, 6.2, and 6.3) for the $\mathrm{FMod}(S,\phi)$ action that may be of broader interest.

Acknowledgments. We thank Martin Möller and Frederik Benirschke for helping us to understand the structure of the boundary of strata. We also thank Dan Margalit for helpful comments on an early draft of this paper. We heartily thank the referee for a close reading of this paper and a number of suggestions which helped streamline our arguments.

This paper grew out of conversations begun in Spring 2022 when JR was attending the trimester program "Groups acting on fractals, hyperbolicity and self-similarity." We gratefully acknowledge financial support from the Institut Henri Poincaré (UAR 839 CNRS-Sorbonne Université) and LabEx CARMIN (ANR-10-LABX-59-01), and thank them for their hospitality. AC acknowledges support from NSF grants DMS-2005328 and DMS-2202703. JR acknowledges support from NSF grant DMS-2103191.

2. Surfaces, curves, and framings

Let us first recall some basic surface-topological notions and set our notation for the rest of the paper. Let $S = S_{g,n}$ denote an orientable surface with genus g and n punctures. The *complexity* of $S = S_{g,n}$ is $\xi(S) = 3g - 3 + n$. By a *curve* on S we mean an isotopy class of an essential (i.e., non-nulhomotopic), non-peripheral (i.e., not homotopic to a puncture), simple closed curve on S. An *arc* on S is an isotopy class of essential, non-peripheral simple arcs running between the punctures. Curves and arcs are unoriented unless we say otherwise. By a *subsurface* of S, we mean an isotopy class of an essential, non-peripheral, open subsurface of S. If U is a subsurface of S we say that a curve is *U-peripheral* if it is homotopic to a puncture or boundary component of U; note that U-peripheral curves may or may not be S-peripheral. For two subsurfaces U and V, we say $U \subseteq V$ if U and V can be realized such that U is contained in V.

Two curves and/or subsurfaces are *disjoint* if their isotopy classes can be realized disjointly. Otherwise, we say they *intersect*. A *multicurve* on S is a collection of distinct, disjoint curves on S. Throughout the paper, we use lowercase Latin letters to refer to curves, Greek letters to multicurves and arcs, and uppercase letters to subsurfaces.

Given two multicurves α, β on S, we let $i(\alpha, \beta)$ denote their *geometric intersection number*. If α and β are oriented curves, then $\langle \alpha, \beta \rangle$ will denote their *algebraic intersection number*. If a multicurve α intersects a subsurface $W \subseteq S$, then $\alpha \cap W$ is the (proper) isotopy class of curves and arcs obtained by taking the intersection of W with a representative for α that realizes $i(\alpha, \partial W)$. Two arcs α_1, α_2 on the subsurface W are *parallel* if they are properly isotopic.

If α is a multicurve on S, then $S \setminus \alpha$ will denote the open subsurface obtained by removing each curve in α from S. Similarly, if W is a subsurface of S, then $S \setminus W$ is the open subsurface obtained by the closure of W from S. We denote the genus of a subsurface $W \subseteq S$ by $g(W)$.

The (pure) mapping class group, $\mathrm{Mod}(S)$, is the group of homeomorphisms of S that fix each of its punctures, modulo isotopy. The mapping class group is generated by *Dehn twists*: for any simple closed curve c, let T_c denote the homeomorphism obtained by cutting open S along c, twisting one of the boundary components of $S \setminus c$ once to the left, and then regluing.

2.1. Framings and winding numbers.

A *framing* of a surface S is a trivialization of its tangent bundle $\phi : TS \xrightarrow{\sim} S \times \mathbb{R}^2$. For surfaces of genus not equal to 1, the existence of a framing requires S to have punctures and/or boundary. Throughout this paper we will think of S as having punctures.

We are interested in the set of framings up to isotopy; these were called "absolute framings" in [CS22]. Isotopy classes of framings can be described by the discrete invariant of a "winding number function" as follows. Given any C^1 immersed curve $\gamma : [0, 1] \to S$, the tangent framing (γ, γ') gives a curve in $TS \cong S \times \mathbb{R}^2$. Projecting into the second factor gives a loop in $\mathbb{R}^2 \setminus \{0\}$ and so one can measure the *winding number* $\phi(\gamma)$ of γ' about 0. This number is an invariant of the isotopy class of framing as well as the isotopy class of γ (though not its homotopy class), and so to every framing ϕ we have an associated winding number function of the same name

$$\phi : \mathcal{S} \to \mathbb{Z},$$

where $\mathcal{S}$ denotes the set of isotopy classes of oriented simple closed curves. It is not hard to show that the function ϕ is actually a complete invariant of the isotopy class of the framing [RW14, Proposition 2.4], and so for the remainder of the paper we will conflate a(n isotopy class of) framing and its associated winding number function.

Remark 2.1. In a previous version of this paper, we considered surfaces with boundary where the framing was allowed to vary on the boundary. These are equivalent to the absolute framings we now consider; see [CS22, Section 6.2].

Winding number functions have two very important properties, which were first elucidated by Humphries and Johnson [HJ89]. As a consequence, a framing is completely determined (up to isotopy) by its values on a basis for homology.

Lemma 2.2 (Humphries–Johnson)**.** *Any winding number function ϕ associated to a framing satisfies the following properties.*

1. *(Twist-linearity) Let $a, b \subset S$ be oriented simple closed curves. Then*
$$\phi(T_a(b)) = \phi(b) + \langle b, a\rangle \phi(a),$$
where $\langle \cdot, \cdot \rangle : H_1(S;\mathbb{Z}) \times H_1(S;\mathbb{Z}) \to \mathbb{Z}$ denotes the algebraic intersection pairing.
2. *(Homological coherence) Let $U \subset S$ be a subsurface and let $c_1, \ldots, c_k$ denote the set of U-peripheral curves, oriented such that U lies to the left of each c_i. Then*
$$\sum_{i=1}^{k} \phi(c_i) = \chi(U),$$
where $\chi(U)$ denotes the Euler characteristic.

Let $\Delta_1, \ldots, \Delta_n$ denote the set of S-peripheral curves, i.e., small loops about the punctures of S (oriented with the surface on their left). The *signature* of a framing ϕ is the tuple

$$\mathrm{sig}(\phi) := (\phi(\Delta_1), \ldots, \phi(\Delta_n)) \in \mathbb{Z}^n.$$

A framing is said to be of *holomorphic type* if every $\phi(\Delta_i)$ is negative; this terminology comes from the fact that the horizontal vector fields of holomorphic abelian differentials give rise to such framings (compare Section 5.1).

Remark 2.3. We note that *not* every framing of holomorphic type comes from a holomorphic abelian differential. This is the case for framings on surfaces of genus at least 3, but the following families of framings do not come from abelian differentials due to certain low-complexity strata being empty (see just below for the definitions of Arf_1 and Arf).

- $g = 1$, $b = 1$, and $\mathrm{Arf}_1(\phi) \neq 0$.
- $g = 2$, $b = 1$, and $\mathrm{Arf}(\phi) = 0$.

The peripheral curves Δ_i span an $n-1$ dimensional subspace of $H_1(S)$, so we can construct all framings with a given signature by specifying the values on $2g$ other homologically independent curves [CS22, Remark 2.7]. That is, so long as the signature satisfies homological coherence, the set of framings with a given signature is a torsor for the first integral homology of the surface $\overline{S}$ obtained by filling in the punctures of S. See [Kaw18, §2.1], for example.

The following is one particularly nice configuration of $2g$ homologically independent curves.

Definition 2.4. A collection of simple closed curves $\mathcal{B} = \{a_1, b_1, \ldots, a_g, b_g\}$ on S is called a *geometric symplectic basis* (GSB) if $i(a_i, b_i) = 1$ for all i and all other pairs of curves from $\mathcal{B}$ are disjoint.

In the sequel, we will often specify a framing of S by specifying its signature and the winding numbers of a GSB.

2.2. **Framed mapping class groups.** The *framed mapping class group* $\mathrm{FMod}(S, \phi)$ associated to a framing ϕ is the stabilizer of ϕ in $\mathrm{Mod}(S)$ up to isotopy. Equivalently, and more usefully, $f \in \mathrm{FMod}(S, \phi)$ if and only if it preserves all winding numbers, i.e.,

$$(f \cdot \phi)(a) := \phi(f^{-1}(a)) = \phi(a)$$

for every $a \in \mathcal{S}$. In light of Lemma 2.2, in order to check if an element $f \in \mathrm{Mod}(S)$ actually preserves ϕ, it suffices to show that f preserves the ϕ–winding numbers of all curves of a GSB.

Throughout the paper, a particularly important role will be played by the set of non-separating simple closed curves with $\phi(a) = 0$ (note that this does not depend on orientation); these curves are said to be *admissible*. By twist-linearity (Lemma 2.2.1), Dehn twists in admissible curves are always in $\mathrm{FMod}(S, \phi)$, and in [CS22] it is shown (for $g \geq 5$) that $\mathrm{FMod}(S, \phi)$ is generated up to finite index by admissible twists.

Since each orbit of $\mathrm{Mod}(S)$ on the set of framings has infinite size (this is an immediate consequence of Lemma 2.2) and $\mathrm{FMod}(S, \phi)$ is a stabilizer, it is an infinite-index subgroup. Along the same lines, understanding the possible conjugacy classes of $\mathrm{FMod}(S, \phi)$ for different ϕ is equivalent to listing the $\mathrm{Mod}(S)$ orbits. To state this "classification of framed surfaces" [Kaw18] (see also [RW14] for the relatively framed version), we first need to recall the definitions of the Arf invariant and its genus 1 version; see [CS22, §2.2], [Kaw18, §2.4], and [RW14, §2.4] for more detailed discussions.

Suppose first that $g = g(S) \geq 2$ and that every $\phi(\Delta_i)$ is odd. In this case, we say that ϕ is of *spin type.* [1] Fix a geometric symplectic basis $\{a_1, b_1, \ldots, a_g, b_g\}$ on S. Then the *Arf invariant* of ϕ is defined to be

$$\mathrm{Arf}(\phi) := \sum_{i=1}^{g} (\phi(a_i) + 1)(\phi(b_i) + 1) \mod 2. \tag{1}$$

This invariant turns out to only be well-defined when each $\phi(\Delta_i)$ is odd, and in this setting it does not depend on our choice of GSB. If $g = 1$, then there is an $\mathbb{Z}$-valued refinement of the Arf invariant which we denote by

$$\mathrm{Arf}_1(\phi) := \gcd(\phi(c), \phi(\Delta_1) + 1, \ldots, \phi(\Delta_n) + 1 \mid c \text{ is a non-separating simple closed curve}).$$

Theorem 2.5. *Two framings ϕ and ϕ' of S are in the same $\mathrm{Mod}(S)$ orbit if and only if*

($g = 0$) *$sig(\phi) = sig(\phi')$*

[1] In this case, the framing induces a (2-)spin structure on the closed surface obtained by capping off all boundary components, and the Arf invariant of the framing coincides with the parity of the spin structure.

($g = 1$) *$sig(\phi) = sig(\phi')$ and* $\mathrm{Arf}_1(\phi) = \mathrm{Arf}_1(\phi')$
($g \geq 2$) *$sig(\phi) = sig(\phi')$ and if ϕ and ϕ' are of spin type, then* $\mathrm{Arf}(\phi) = \mathrm{Arf}(\phi')$.

In particular, for genus at least 2 there are only ever at most 2 distinct conjugacy classes of framed mapping class groups of a given signature.

The Arf invariant interacts in a complicated way with taking subsurfaces $V \subset S$; sometimes the Arf invariant of $\phi|_V$ is forced by the topology of V, and sometimes it can vary for different V and V' of the same topological type. For later use, we record an example of this phenomenon below. See also the proofs of Propositions 2.10 and 6.2.

Lemma 2.6. *Suppose that $V \subset S$ is a connected subsurface of full genus.*

(1) If $g(S) \geq 2$ and ϕ is of spin type, then $\mathrm{Arf}(\phi) = \mathrm{Arf}(\phi|_V)$.
(2) $g(S) = 1$ and ϕ is of holomorphic type, then $\mathrm{Arf}_1(\phi) = \mathrm{Arf}_1(\phi|_V)$

Proof. When S has genus at least 2, this is an immediate consequence of (1). In the case when S has genus 1, homological coherence together with holomorphic type imply that two curves which differ by sliding over a boundary component must have the same winding number. Thus for any simple closed curve c on S, there is some $c' \subset V$ with $\phi(c) = \phi(c')$, and hence their genus-1 Arf invariants must agree. □

Note that statement (2) is false if one does not assume holomorphic type.

2.3. **Framed change-of-coordinates.** The standard change-of-coordinates principle for the entire mapping class group roughly states that given two multicurves γ and δ, there is some $f \in \mathrm{Mod}(S)$ taking γ to δ if and only if $S \setminus \gamma$ and $S \setminus \delta$ have the same topological type and are glued together in the same way. This technique is often used in surface topology to show the existence of certain configurations of curves with prescribed intersection pattern and to show the transitivity of the $\mathrm{Mod}(S)$ action on such configurations. Its proof is a corollary of the classification of surfaces: one uses the classification to build a homeomorphism between the complements then extends that to a self-homeomorphism of S.

In the framed setting, we can similarly use Theorem 2.5 to show the existence of configurations with certain intersection pattern and winding number (compare [CS22, Proposition 2.5]). For example, we can quickly show that (sub)surfaces with genus always contain admissible curves. Essentially the same statement appears as Corollary 4.3 of [Sal], but we include a proof as we will repeatedly use this statement throughout the paper.

Lemma 2.7. *For any framing ϕ on a surface S of positive genus, there is some non-separating simple closed curve $a \subset S$ with $\phi(a) = 0$.*

Proof. Fix a GSB $\{a_1, \ldots, b_g\}$ on S. Then by stipulating winding numbers on our GSB we can build a framing ψ such that

- $\mathrm{sig}(\phi) = \mathrm{sig}(\psi)$
- $\psi(a_1) = 0$, and
- if $g(S) = 1$ then $\mathrm{Arf}_1(\psi) = \mathrm{Arf}_1(\phi)$, or
- if $g(S) \geq 2$ and ϕ is of spin type then $\mathrm{Arf}(\psi) = \mathrm{Arf}(\phi)$.

Now by Theorem 2.5 there is some homeomorphism $f \in \mathrm{Mod}(S)$ taking ψ to ϕ, and the curve $f(a_1)$ is our desired admissible curve. □

Along the same lines, one can show that S always admits a GSB with given winding numbers so long as those winding numbers yield the correct Arf invariant; the proof is left to the reader. See also the proof of the first part of [CS22, Proposition 2.15].

Lemma 2.8. *Let ϕ be a framing of a surface S of genus $g \geq 1$ and fix any tuple of integers $(x_1, y_1, \ldots, x_g, y_g)$ such that*

- *if $g = 1$, then $\gcd(x_1, y_1, \phi(\Delta_1) + 1, \ldots, \phi(\Delta_n) + 1) = \mathrm{Arf}_1(\phi)$,*
- *if $g \geq 2$ and ϕ is of spin type, then*

$$\sum_{i=1}^{g}(x_i + 1)(y_i + 1) = \mathrm{Arf}(\phi) \mod 2$$

- *if $g \geq 2$ and ϕ is not of spin type, then we impose no conditions on the tuple.*

Then there is a GSB $\mathcal{B} = \{a_1, b_1, \ldots, a_g, b_g\}$ on S such that $\phi(a_i) = x_i$ and $\phi(b_i) = y_i$.

In particular, any surface of genus at least 2 contains both nonseparating curves of arbitrary winding number as well as genus 1 subsurfaces with $\mathrm{Arf}_1 = 1$.

Example 2.9. We caution the reader that if S has genus 1 and $\mathrm{Arf}_1(\phi) = 1$, then genus 1 subsurfaces $T \subset S$ may have $\mathrm{Arf}_1(\phi|_T) \neq 1$. For example, consider a framing ϕ of $S = S_{1,3}$ with signature $(3, -4, -2)$. Because $\gcd(4, -3, -1) = 1$, we see $\mathrm{Arf}_1(\phi) = 1$. Using Lemma 2.8, pick a GSB with $x_1 = 0$ and $y_1 = 4$. Pick a curve disjoint from (x_1, y_1) that together with Δ_2 and Δ_3 bounds a pair of pants, and let T denote the genus 1 component complementary to this pair of pants. By construction, T contains (x_1, y_1) and ϕ_T has signature $(3, -5)$, so we compute $\mathrm{Arf}_1(\phi|_T) = 4$.

The classification of framed surfaces can also be used to easily obstruct transitivity of the $\mathrm{FMod}(S, \phi)$ action. For example, $\mathrm{FMod}(S, \phi)$ does not act transitively on the set of curves that separate off a genus 1 subsurface with one boundary component, even though those curves all have the same topological type and same winding number. The reason is that the induced framing on the subsurface may have different Arf_1 invariant.

We caution the reader that Theorem 2.5 does not imply transitivity on the set of multicurves of the same topological type that induce homeomorphic framings on each subsurface. Indeed, suppose that some $\phi(\Delta_i)$ is even so ϕ does not have an induced Arf invariant. If we consider the set of multicurves $\gamma = c \cup d$ where c cuts off a genus 1 subsurface with one puncture and d is an admissible curve on that subsurface, then the paragraph above implies that $\mathrm{FMod}(S, \phi)$ does not act transitively on this set, even though there is only one $\mathrm{Mod}(S \setminus \gamma)$ orbit of framing on $S \setminus \gamma$. At issue is what happens when we try to glue together framings on subsurfaces to a framing on the entire surface; this can be dealt with by using *relative* framings and being careful about boundary conditions (compare the proof of Lemma 5.3 in [CS22]). Since such arguments require a fair amount of delicacy and are beyond what we need in this paper, we will restrict ourselves to proving those transitivity results we will need in the sequel.

Proposition 2.10. *Let ϕ be a framing of a surface S of genus at least 3. Then $\mathrm{FMod}(S, \phi)$ acts transitively on pairs of disjoint admissible curves of a given topological type. That is, if γ, γ' are pairs of admissible curves and there is some $g \in \mathrm{Mod}(S)$ taking γ to γ', then there is also some $f \in \mathrm{FMod}(S, \phi)$ taking γ to γ'.*

Before proving Proposition 2.10, we first record a useful lemma that allows us to adjust the winding numbers of curves in a configuration without changing their intersection properties. A similar statement appears as Corollary 4.4 of [Sal].

Lemma 2.11. *Let ϕ be a framing of a surface S and let $\{c_1, \ldots, c_k, d\}$ be a collection of simple closed curves. Assume there is some subsurface $T \subset S$ disjoint from all of these curves and such that*

$$g(T) = 1 \text{ and } \mathrm{Arf}_1(\phi|_T) = 1.$$

Suppose also that the union of the c_i does not separate d from T. Then for any $z \in \mathbb{Z}$, there is a simple closed curve d_z such that $\phi(d_z) = z$ and $i(c_i, d_z) = i(c_i, d)$ for all i.

Proof. Our assumption on $\mathrm{Arf}_1(\phi|_T)$ implies (via Lemma 2.8) that there is some GSB (a, b) on T with $\phi(a) = 1$ and $\phi(b) = -z - \phi(d) - 1$. Since d is not separated from T and b is nonseparating, we may choose an arc ε connecting the right-hand sides of b and d. Let d_z be the resulting curve; then by homological coherence (Lemma 2.2.2) we have that

$$\phi(d_z) + \phi(d) + \phi(b) = -1$$

and so d_z is our desired curve. It clearly has the same intersection pattern as d with each c_i since we have only altered d away from c_i, and is essential and non-peripheral as $[d_z] = [d]+[e]$ in homology. □

One particularly important consequence is that we can complete any admissible curve to a partial GSB while specifying the winding number of the transverse curve.

Corollary 2.12. *For any surface of genus at least 2, any admissible a, and any $z \in \mathbb{Z}$, there is a curve b with $i(a, b) = 1$ and $\phi(b) = z$.*

Proof. The subsurface $S \setminus a$ has two boundary components with winding number 0 and so $\mathrm{Arf}_1(S \setminus a) = 1$. Applying Lemma 2.8 we can pick some GSB on $S \setminus a$ with coprime winding numbers; let T denote the subsurface filled by this pair of curves. We can now pick any curve b' disjoint from T with $i(a, b') = 1$. Since b' does not meet T and $\mathrm{Arf}_1(\phi|_T) = 1$, we can apply Lemma 2.11 to adjust $\phi(b')$ at will. □

With these results in hand, we can now prove the desired transitivity statements.

Proof of Proposition 2.10. Obviously transitivity on single curves follows from the result for pairs, but since the proof for pairs requires a bit of casework we will prove the result for single curves first as a demonstration of our techniques.

Single curves. Suppose first that $a, a' \subset S$ are both admissible. Complete a to a GSB $a = a_1, b_1, \ldots, a_g, b_g$ of S. Using Corollary 2.12, there is some b_1' on S with $i(a', b_1') = 1$ and $\phi(b_1') = \phi(b_1)$. Now take the subsurface Y' filled by a' and b_1' and consider its complement. If $\phi|_{S \setminus Y'}$ is of spin type, then the additivity of the Arf invariant [RW14, Lemma 2.11] implies that

$$\mathrm{Arf}(\phi|_{S \setminus Y'}) = \mathrm{Arf}(\phi) - (\phi(a') + 1)(\phi(b_1') + 1) = \sum_{i=2}^{g} (\phi(a_i) + 1)(\phi(b_i) + 1) \mod 2.$$

Otherwise, it is not of spin type; in either case we can now apply Lemma 2.8 to find a GSB $a_2', b_2', \ldots, a_g', b_g'$ on $S \setminus Y'$ with

$$\phi(a_i) = \phi(a_i') \text{ and } \phi(b_i) = \phi(b_i') \text{ for all } i.$$

By the usual change-of-coordinates principle (compare Lemma 2.3 of [Sal]), there is some $f \in \mathrm{Mod}(S)$ taking a to a', each a_i to a_i', and each b_i to b_i'. Since f preserves the winding numbers of the curves of a GSB, it preserves the winding numbers of all simple curves (Lemma 2.2), and thus we see that $f \in \mathrm{FMod}(S, \phi)$.

Nonseparating pairs. If $g \geq 4$ and the admissible curves a_1, a_2 together do not separate S, then we can just repeat our argument for transitivity on single admissible curves: extend a_1, a_2 to an arbitrary GSB, use Corollary 2.12 and 2.8 to extend a_1', a_2' to a GSB with the same winding numbers, and then use the transitivity of the mapping class group action on GSBs to find some f (necessarily in $\mathrm{FMod}(S, \phi)$) taking one GSB to the other.

If $g = 3$ then we must be slightly more clever about how we choose our intial GSB since our choice of transverse curves b_1 and b_2 may constrain the winding numbers of the remaining curves a_3 and b_3 due to the Arf_1 invariant. Suppose first that ϕ is of spin type. Using Corollary 2.12 twice, we can choose disjoint curves b_1 and b_2, each meeting their respective a_i and disjoint from the other, such that

$$\mathrm{Arf}(\phi) + \phi(b_1) + \phi(b_2) = 0 \mod 2.$$

In particular, this implies that if we let Y denote the (disconnected) subsurface obtained by taking a regular neighborhood of $a_1 \cup a_2 \cup b_1 \cup b_2$, then the contribution to $\mathrm{Arf}(\phi)$ of $\phi|_{S \setminus Y}$ must be 0, hence for any GSB (a_3, b_3) on $S \setminus Y$ at least one of $\phi(a_3)$ or $\phi(b_3)$ must be odd. Now we observe that

$$\mathrm{sig}(\phi|_{S \setminus Y}) = (\mathrm{sig}(\phi), +1, +1)$$

and so $\mathrm{Arf}_1(\phi|_{S \setminus Y})$ is the gcd of an odd number and 2, i.e., is 1.

If ϕ is not of spin type then choose any disjoint b_1 and b_2, each meeting their respective a_i and disjoint from the other, and define Y similarly. Then since some $\phi(\Delta_i)$ is even, the signature of $\phi|_{S \setminus Y}$ contains both an even number and $+1$, and so we see that $\mathrm{Arf}_1(\phi|_{S \setminus Y}) = 1$. Therefore, no matter whether ϕ is of spin type or not, we can choose our b_1 and b_2 such that $\phi|_{S \setminus Y}$ has fixed Arf_1, and so by Lemma 2.8 admits a GSB a_3, b_3 with $\phi(a_3) = 0$ and $\phi(b_3) = 1$. We can now finish the proof by inserting a prime in all of the arguments above to get another GSB on S with the same winding number data and then concluding as in the $g \geq 4$ case.

Separating pairs. Finally, suppose that $a_1 \cup a_2$ separates S into two subsurfaces T and U. In this case, neither of the complementary components to $a_1 \cup a_2$ is of spin type, so if ϕ is of spin type then we will need be somewhat clever about our choice of GSB to deal with the emergence of the Arf invariant.

Since at least one of T or U has genus at least 2 or genus 1 with $\mathrm{Arf}_1 = 1$, we can choose a genus 1 subsurface V disjoint from $a_1 \cup a_2$ with a single boundary component and $\mathrm{Arf}_1(\phi|_V) = 1$ — simply take a symplectically dual pair of curves with winding numbers $(0, 1)$ and take a neighborhood. Pick an arbitrary curve b_1 meeting a_1 and a_2 each exactly once which is disjoint from V, and use Lemma 2.11 to turn this curve into an admissible b_1 that also meets each of a_1 and a_2 exactly once.

Choose GSBs

$$\mathcal{B}_T := s_1, t_1, \ldots, s_{g(T)}, t_{g(T)} \text{ for } T \text{ and } \mathcal{B}_U := u_1, v_1, \ldots, u_{g(U)}, v_{g(U)} \text{ for } U$$

that are disjoint from b_1; then $\{a_1, b_1\} \cup \mathcal{B}_T \cup \mathcal{B}_U$ is a GSB for S.

Since (a_1, a_2) and (a_1', a_2') are in the same mapping class group orbit, there is a correspondence between their complementary components; let T' and U' denote the two components of $a_1' \cup a_2'$ corresponding to T and U. Since neither component is of spin type (having a boundary component with even winding number) or, if they have genus 1, have $\mathrm{Arf}_1 = 1$ with an admissible boundary component, Lemma 2.8 implies that both T' and U' admit GSBs with any given tuples of winding numbers. We may therefore choose GSBs $\mathcal{B}_{T'}$ and $\mathcal{B}_{U'}$ with the same winding numbers as those for $\mathcal{B}_T$ and $\mathcal{B}_U$. To extend these to a GSB of S, we just need to find an admissible curve corresponding to b_1.

Suppose ϕ is of spin type, and pick some b_1' that meets each of a_1' and a_2' exactly once and is disjoint from $\mathcal{B}_T \cup \mathcal{B}_U$. Then

$$(\phi(a_1)+1)(\phi(b_1)+1) + \sum_{g(T)} (\phi(s_i)+1)(\phi(t_i)+1) + \sum_{g(U)} (\phi(u_i)+1)(\phi(v_i)+1) = \mathrm{Arf}(\phi)$$

$$= (\phi(a_1')+1)(\phi(b_1')+1) + \sum_{g(T')} (\phi(s_i')+1)(\phi(t_i')+1) + \sum_{g(U')} (\phi(u_i')+1)(\phi(v_i')+1) \mod 2$$

which simplifies to $\phi(b_1) = \phi(b_1') \mod 2$ by our choices of $\mathcal{B}_{T'}$ and $\mathcal{B}_{U'}$. Thus $\phi(b_1')$ must be even. Now choose a curve c on either T' or U' that

- is disjoint from $\mathcal{B}_{T'} \cup \mathcal{B}_{U'}$,
- meets b_1' exactly once, and
- together with a_1' bounds a surface of genus 1 with 2 boundary components.

Such a c can be obtained, for example, by taking the boundary of a regular neighborhood of $u_1' \cup v_1'$ and then connect summing that curve with a_1'. See Figure 1. By homological coherence (Lemma 2.2.2), it must be that $\phi(c) = \pm 2$ (where sign depends on orientation). Twist-linearity (Lemma 2.2.1) then implies that some twist of b_1' about c will be admissible. Thus the configurations of curves

$$a_1, b_1, a_2, \mathcal{B}_T, \mathcal{B}_U \text{ and } a_1', T_c^{-\phi(b_1')/2}(b_1'), a_2', \mathcal{B}_{T'}, \mathcal{B}_{U'}$$

have the same topological type, so there is a mapping class taking one to the other, and since all of the corresponding curves have the same winding number, any such mapping class must preserve ϕ.

If ϕ is not of spin type, then we can conclude by picking an arbitrary b_1' disjoint from $\mathcal{B}_{T'} \cup \mathcal{B}_{U'}$. We then note that since ϕ is not of spin type, then there is some peripheral curve Δ_i with even winding number. Choose c as before and let d be a curve disjoint from all of the listed curves except b_1', obtained by taking the connect sum of a_2 with this Δ_i; by homological coherence again, its winding number must be odd. See Figure 1. Thus, by twisting around c and d we can change the winding number of b_1' by any amount (while keeping all other winding numbers fixed) and so in particular $T_c^m T_d^n(b_1')$ is admissible for some m, n. We can then conclude as in the spin case. □

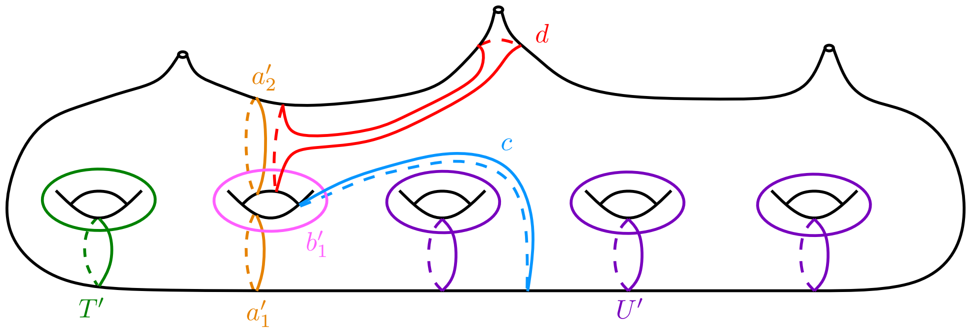

FIGURE 1. GSBs and auxiliary curves as in the proof of Proposition 2.10.

3. The admissible curve graph and its geometric model

A *graph of multicurves* for a surface S is any graph whose vertices are multicurves on S. The simplest and most influential example is the *curve graph* $\mathscr{C}(S)$. The curve graph has all curves on S as vertices and edges between two curves if and only if they intersect the fewest number of times possible for a pair of curves on S. If $\xi(S) > 1$ then edges correspond with disjointness, and when $\xi(S) = 1$ the minimal intersection number is either 1 or 2.

We will focus on the following subset of the curve graph: given a framing ϕ of S, the *admissible curve graph* $\mathscr{C}_{\mathrm{adm}}(S, \phi)$ relative to ϕ is the subgraph of $\mathscr{C}(S)$ spanned by the non-separating curves that are admissible with respect to ϕ.

Proposition 2.10 implies that the framed mapping class group $\mathrm{FMod}(S, \phi)$ acts transitively on the vertices of $\mathscr{C}_{\mathrm{adm}}(S, \phi)$ and with finitely many orbits on its edges. As a consequence of Lemma 2.7, every vertex of $\mathscr{C}(S)$ is distance 1 from a vertex of $\mathscr{C}_{\mathrm{adm}}(S, \phi)$ when $g(S) \geq 2$. When $g(S) \geq 3$, Lemma 2.7 also allows us to copy Salter's "hitchhiking argument" in the case of r-spin structures [Sal, Lemma 3.11] to show $\mathscr{C}_{\mathrm{adm}}(S, \phi)$ is connected.

Lemma 3.1. *If $g(S) \geq 3$, then $\mathscr{C}_{\mathrm{adm}}(S, \phi)$ is connected for any framing ϕ.*

Proof sketch. The graph of genus 1 subsurfaces (with edges for disjointness) is connected [Put08]. Since each genus 1 subsurface contains an admissible curve, the paths in this graph can be upgraded to a path in $\mathscr{C}_{\mathrm{adm}}(S, \phi)$. □

Given a graph of multicurves $\mathcal{X}$, a connected subsurface $W \subseteq S$ is a *witness* for $\mathcal{X}$ if every vertex of $\mathcal{X}$ intersects W and W is not homeomorphic to $S_{0,3}$. We let $\mathrm{Wit}(\mathcal{X})$ denote the set of all witness for $\mathcal{X}$. For the admissible curve graph, the witnesses are all subsurfaces whose complement has no genus and whose peripheral curves' winding numbers do not satisfy a particular set of linear equations.

Lemma 3.2. *Let $S = S_{g,n}$ with $g \geq 3$ and $n \geq 1$. Fix a framing ϕ of S.*

(1) Suppose $Z \subseteq S$ is a genus 0 subsurface and let $z_1, \dots, z_k$ denote the set of Z-peripheral curves, oriented such that Z is to the left of each z_i. Then Z contains a (possibly S-peripheral) curve of winding number 0 if and only if there exists $I \subsetneq \{z_1, \dots, z_k\}$ such that

$$\sum_{z \in I} \phi(z) = 1 - |I|. \tag{2}$$

(2) *A subsurface W of S is a witness for $\mathscr{C}_{\mathrm{adm}}(S,\phi)$ if and only if each curve in ∂W is not admissible and each component Z of $S \setminus W$ is a genus 0 subsurface with the following property: enumerate the Z-peripheral curves as in the previous item. Then there is no I such that equation (2) holds and both I and $\{z_i\}_{i=1}^k \setminus I$ contain curves of ∂W.*

(3) *If $V, W \in \mathrm{Wit}(\mathscr{C}_{\mathrm{adm}}(S,\phi))$ are disjoint, then each is a genus 0 subsurface that does not contain any admissible curves, and there does not exist $Z \in \mathrm{Wit}(\mathscr{C}_{\mathrm{adm}}(S,\phi))$ that is disjoint from both V and W.*

Proof. The first item is an immediate consequence of homological coherence (Lemma 2.2.(2)) and the fact that every curve on a genus 0 surface separates it. The second item follows from the first plus Lemma 2.7's guarantee that every subsurface with genus contains an admissible curve plus the fact that W must be connected (so each z_i is either S-peripheral or W-peripheral). Note that the condition that ∂W meets both I and $\{z_i\}_{i=1}^k \setminus I$ indicates whether or not a curve cutting off the boundaries $\{z_i\}_{i\in I}$ separates S or not. The third item is an immediate consequence of the second. □

Paralleling [Vok22], we now use the witnesses of a graph of multicurves to construct a "model graph," which is in some sense the largest graph of multicurves that has the same witness set as the starting graph.

Definition 3.3. Let $\mathfrak{S}$ be a collection of subsurfaces of S. We say $\mathfrak{S}$ is a set of *valid witnesses* if for all $W \in \mathfrak{S}$,

(1) W is connected;
(2) $\xi(W) \geq 1$;
(3) if Z is a connected subsurface with $W \subseteq Z$, then $Z \in \mathfrak{S}$;

Definition 3.4. Let $\mathfrak{S}$ be a set of valid witnesses for the surface S. If $\mathfrak{S} = \emptyset$, define $\mathcal{K}_{\mathfrak{S}}(S)$ to be a single point. Otherwise, define $\mathcal{K}_{\mathfrak{S}}(S)$ to be the graph such that:

- each vertex is a multicurve γ on S with the property that each component of $S \setminus \gamma$ is *not* an element of $\mathfrak{S}$;
- two multicurves γ and δ are joined by an edge if either
  (1) γ differs from δ by either adding or removing a single curve, or
  (2) γ differs from δ by "flipping" a curve in some subsurface of S, that is, δ is obtained from γ by replacing a curve $c \subset \gamma$ by a curve d, where c and d are contained in the same component Y_c of $S \setminus (\gamma \setminus c)$ and are adjacent in $\mathscr{C}(Y_c)$.

By construction, the set of witness for $\mathcal{K}_{\mathfrak{S}}(S)$ is precisely $\mathfrak{S}$. Moreover, the vertex set of $\mathcal{K}_{\mathfrak{S}}(S)$ is the maximal collection of multicurves whose set of witnesses is $\mathfrak{S}$. Thus, if $\mathcal{X}$ is a graph of multicurves with $\mathrm{Wit}(\mathcal{X}) = \mathfrak{S}$, then the vertices of $\mathcal{X}$ are a subset of $\mathcal{K}_{\mathfrak{S}}(S)$. In the case of the admissible curve graph, this inclusion is Lipschitz.

Lemma 3.5. *If $\mathfrak{S} = \mathrm{Wit}(\mathscr{C}_{\mathrm{adm}}(S,\phi))$, then the inclusion $\mathscr{C}_{\mathrm{adm}}(S,\phi) \to \mathcal{K}_{\mathfrak{S}}(S)$ is 1-Lipschitz.*

Proof. If a, b are a pair of disjoint admissible curves, then there is a "flip" edge of $\mathcal{K}_{\mathfrak{S}}(S)$ between a and b because they are adjacent in $\mathscr{C}(S)$ and $S = S \setminus (a \setminus a)$. □

Vokes studied the family of $\mathcal{K}_{\mathfrak{S}}(S)$ as quasi-isometric models for graphs of multicurves. Specifically, she showed that if $\mathcal{X}$ is a graph of multicurves on S with a cocompact action of $\mathrm{Mod}(S)$ and no annular witnesses, then the inclusion $\mathcal{X} \hookrightarrow \mathcal{K}_{\mathfrak{S}}(S)$ for $\mathfrak{S} = \mathrm{Wit}(\mathcal{X})$ is a quasi-isometry. The advantage of using $\mathcal{K}_{\mathfrak{S}}(S)$ as a model is that it is a *hierarchically hyperbolic space* in a natural way. In particular, the coarse geometry of $\mathcal{K}_{\mathfrak{S}}(S)$ can be well understood using the subsurface projection machinery of Masur and Minsky and the relations between the subsurfaces in $\mathfrak{S}$; see [BHS17b, BHS19, Vok22] for full details.

We note that while Vokes states her results in the case of an action of the full mapping class group, the only actual use of the action is in establishing the quasi-isometry described above. In particular, the proof in Section 3 of [Vok22] as written demonstrates that $\mathcal{K}_{\mathfrak{S}}(S)$ is a hierarchically hyperbolic space, even in the case where $\mathfrak{S}$ is not invariant under the mapping class group.

One consequence of Vokes's hierarchically hyperbolic structure is that Gromov hyperbolicity of the the graph is encoded in the disjointness of the witnesses.

Theorem 3.6 (Corollary 1.5 of [Vok22])**.** *The graph $\mathcal{K}_{\mathfrak{S}}(S)$ is Gromov hyperbolic if and only if $\mathfrak{S}$ does not contain a pair of disjoint subsurfaces.*

4. A quasi-isometry with the model

Vokes's proof of the quasi-isometry between graphs of multicurves and their models relies on the action of the mapping class group in a fundamental way. Specifically, given any connected graph of multicurves $\mathcal{X}$ that has no annular witnesses and has a cocompact action by $\mathrm{Mod}(S)$, she uses the "change-of-coordinates" principle and curve surgery arguments to build a quasi-isometry from $\mathcal{K}_{\mathfrak{S}}(S)$ to $\mathcal{X}$, where $\mathfrak{S}$ is the set of witnesses of $\mathcal{X}$.

In our setting, we only have access to the (weaker) framed versions of these techniques. Moreover, there are infinitely many $\mathrm{FMod}(S, \phi)$ orbits of curves and of witnesses, so we cannot employ standard change-of-coordinates arguments of the form "make a choice for each orbit, then propagate that choice around using the group action to get finiteness" (e.g., [Vok22, Claim 4.3] or Lemma 4.4 below).

Instead of relying on change-of-coordinates, we build our quasi-isometry $\mathcal{K}_{\mathfrak{S}}(S) \to \mathscr{C}_{\mathrm{adm}}(S, \phi)$ by going through an intermediary graph $\mathcal{G}$, which admits a coarsely Lipschitz map Π onto $\mathscr{C}_{\mathrm{adm}}(S, \phi)$ (Lemma 4.5). One can then define a map Ψ from $\mathcal{K}_{\mathfrak{S}}(S)$ to subsets of $\mathcal{G}$; while this map is not coarsely Lipschitz or even coarsely well-defined, the composition $\Pi \circ \Psi$ turns out to be (Proposition 4.13).

The utility of this approach is that $\mathcal{G}$ admits an action of the entire mapping class group, so we can use standard change-of-coordinates arguments. A fruitful comparison is the "hitching a ride" argument we used to show the connectivity of $\mathscr{C}_{\mathrm{adm}}(S, \phi)$ in Lemma 3.1.

For the remainder of the section, $S = S_{g,n}$ will be a surface with genus $g \geq 3$ and $n \geq 1$ punctures, and $\mathfrak{S}$ will be the set of witnesses for $\mathscr{C}_{\mathrm{adm}}(S, \phi)$ with respect to a fixed framing ϕ. Since we will only be considering theses graphs for the surface S, we will use $\mathscr{C}_{\mathrm{adm}}$ and $\mathcal{K}$ to denote $\mathscr{C}_{\mathrm{adm}}(S, \phi)$ and $\mathcal{K}_{\mathfrak{S}}(S)$ respectively.

4.1. **Coarse maps and quasi-isometries.** Let X, Y be metric spaces. A map $f \colon X \to 2^Y$ is *coarsely well-defined* if $f(x)$ has uniformly bounded diameter for every $x \in X$. It is *coarsely*

Lipschitz if there are constants $K \geq 1$ and $C \geq 0$ such that

$$\operatorname{diam}_Y(f(x) \cup f(x')) \leq K d_X(x, x') + C$$

for every $x, x' \in X$. In particular, note that coarsely Lipschitz maps are in particular coarsely well-defined. Prototypical examples are the inclusion of a connected subgraph into a connected graph, the subsurface projection map from the marking graph to $\mathscr{C}(W)$ where $W \subseteq S$ is a subsurface, or the systole map that sends a point in Teichmüller space to its hyperbolic systole(s).

When X is a graph, one can simply define a map $f \colon X \to 2^Y$ on the vertices and assume that the image of any point on an edge is the union of the images of the endpoints of that edge. In this case, to show f is coarsely Lipschitz, it suffices to show that

(1) $f(x)$ is uniformly bounded for all vertices x of X, and
(2) if x and x' are two vertices joined by an edge of X, then $\operatorname{diam}(f(x) \cup f(x'))$ is uniformly bounded.

Two spaces are *quasi-isometric* if there exist two coarsely Lipschitz maps $f \colon X \to 2^Y$ and $\overline{f} \colon Y \to 2^X$ such that $d_X(x, \overline{f} \circ f(x))$ and $d_Y(y, f \circ \overline{f}(y))$ are uniformly bounded for all $x \in X$ and $y \in Y$. In this case, f is a *quasi-isometry* from X to Y and $\overline{f}$ is the *quasi-inverse* of f. We remark that the uniform boundedness of $d_Y(y, f \circ \overline{f}(y))$ is equivalent to every point of Y being contained in a uniform neighborhood of the image of f.

4.2. **The genus-separating curve graph.** We begin building our quasi-isometry from $\mathcal{K}$ to $\mathscr{C}_{\text{adm}}$ by defining the intermediate graph $\mathcal{G}$ that we use throughout this section. We say that a separating curve $c \subseteq S$ is *genus-separating* if each component of $S \setminus c$ has positive genus.

Definition 4.1. The *genus-separating curve graph* $\mathcal{G} = \mathcal{G}(S)$ is the graph whose vertices are genus-separating curves, and where two vertices are connected by an edge if the corresponding curves are disjoint.

Putman's argument that the separating curve graph is connected in the closed case also shows that $\mathcal{G}$ is connected [Put08]. The key commonality is that every vertex of $\mathcal{G}$ is adjacent to a genus-separating curve that cuts off a torus with one boundary component.

Lemma 4.2. *The graph $\mathcal{G}$ is connected so long as $g(S) \geq 3$.*

Since every subsurface with genus contains an admissible curve, we see that for any $c \in \mathcal{G}$ both components of $S \setminus c$ are not witnesses for $\mathscr{C}_{\text{adm}}$. Thus $\mathcal{G}$ is a subgraph of $\mathcal{K}$.

Remark 4.3. While we will not use this in the sequel, we can in fact relate the geometries of $\mathcal{G}$ and $\mathcal{K}$ by considering their sets of witnesses. The witnesses for $\mathcal{G}$ are exactly those subsurfaces that have genus 0 complements, which form a strict superset of the witnesses for $\mathcal{K}$ (characterized in Lemma 3.2). Using the "factored space" construction from [BHS17a], we can thus view $\mathcal{K}$ as being obtained from $\mathcal{K}_{\text{Wit}(\mathcal{G})}(S)$ by coning off regions corresponding to the non-shared witnesses.

As for the usual curve graph, intersection number bounds distance in $\mathcal{G}$.

Lemma 4.4. *For each $n \geq 0$ there exists $N = N(n) \geq 0$ such that for any two genus-separating curves $c, d \in \mathcal{G}$, if $i(c,d) \leq n$, then $d_{\mathcal{G}}(c,d) \leq N$.*

Proof. By the change-of-coordinates principle in $\mathrm{Mod}(S)$, there exist finitely many pairs $\{(c_i, d_i)\}_{i=1}^k$ of genus-separating curves such that every pair of genus-separating curves that intersect at most n times is in the $\mathrm{Mod}(S)$–orbit of some (c_i, d_i). Setting $N = \max\{d_{\mathcal{G}}(c_i, d_i) : 1 \leq i \leq k\}$, the fact that $\mathrm{Mod}(S)$ acts by isometries on $\mathcal{G}$ implies any two genus-separating curves that intersect at most n times are at most N far apart in $\mathcal{G}$. □

4.3. **From genus-separating to admissible curves.** Define a map

$$\Pi \colon \mathcal{G} \to 2^{\mathscr{C}_{\mathrm{adm}}}$$

by sending a genus-separating curve to the collection of admissible curves disjoint from it. This set is always non-empty by Lemma 2.7.

Lemma 4.5. *The map Π is coarsely Lipschitz.*

Proof. As remarked above, it suffices to check that the diameters of the images of vertices and edges are both bounded.

Let $c \in \mathcal{G}$ be any genus-separating curve and let U, V denote the components of $S \setminus c$. Let a be any admissible curve in $\Pi(c)$, and assume without loss of generality that $a \subset U$. Every admissible curve in V is distance 1 from a, and likewise every admissible curve in U is disjoint from any curve in V. Thus $\Pi(c)$ has diameter 2 as a subgraph of $\mathscr{C}_{\mathrm{adm}}$.

Now suppose c and d in $\mathcal{G}$ are disjoint; this implies that one of the (positive genus) components of $S \setminus c$ is nested inside a component of $S \setminus d$. In particular, this implies that $\Pi(c)$ and $\Pi(d)$ overlap, and since each has bounded diameter their union does as well. □

The map Π is defined such that if $a \in \mathscr{C}_{\mathrm{adm}}$ and $c \in \mathcal{G}$ with $i(a,c) = 0$, then

$$d_{\mathscr{C}_{\mathrm{adm}}}(a, \Pi(c)) = 0.$$

Below, we prove a generalization of this fact that allows us to bound the distance between a and $\Pi(c)$ by bounding the geometric intersection number $i(a,c)$. While we end up not using this in the proof of Theorem A, these sorts of bounds are fundamental whenever using curve graph techniques, and indicate that change-of-coordinates style arguments have a chance of working in our setting.

Lemma 4.6. *For any $m \geq 0$, there exists $M = M(m) \geq 0$ such that for any admissible curve a and any genus-separating curve c with $i(a,c) \leq m$, we have $d_{\mathscr{C}_{\mathrm{adm}}}(a, \Pi(c)) \leq M$.*

Proof. If a is disjoint from c, then $a \in \Pi(c)$ and we are done. Otherwise, we will surger c along a to produce a new genus-separating curve c' disjoint from c that intersects a strictly fewer times. By Lemma 4.4, this will allow us to decrease the intersection number of a and c at the cost of moving c a fixed distance in $\mathcal{G}$. Since Π is a coarsely Lipschitz map, this procedure moves the projection a uniformly bounded amount in $\mathscr{C}_{\mathrm{adm}}$, proving the desired statement.

Since S has genus at least 3, there is at least one component $U_c \subset S \setminus c$ of genus at least 2. Consider an arc α of $a \cap U_c$. The regular neighborhood of $c \cup \alpha$ forms a pair of pants P_α, one of whose boundaries is c; label the other two by d and e. Because any strand of $a \cap U_c$

that meets d or e must travel through P_α while avoiding α, any such strand must exit P_α through c. Thus, we have

$$i(a,d) + i(a,e) \leq i(a,c) - 2.$$

If either d or e is separating, then the other one is either separating or S-peripheral (they cannot both be peripheral as c is genus-separating). Since U_c has positive genus, at least one of d and e is genus-separating; we then take c' to be whichever is, completing the proof in this case.

In the other case, d and e are individually non-separating. Let $V_c \subset U_c$ denote the connected subsurface of $U_c \setminus (d \cup e)$ not containing α. Choose an arc β in V_c connecting d and e that is disjoint from $a \cap V_c$. Such an arc always exists because either $a \cap V_c$ contains such an arc, or it does not, in which case one can take an arbitrary arc from d to e and surger it along its intersections with $a \cap V_c$ to make it disjoint; see Figure 2.

The curve c' obtained from a regular neighborhood of $d \cup e \cup \beta$ forms a pair of pants P_β with d and e. Since any arc of a that enters P_β through c' cannot intersect β, that arc must exit through either d or e. Thus

$$i(c', a) \leq i(a,d) + i(a,e) < i(a,c).$$

Since c' is constructed to cut off a genus $g(U_c) - 1 \geq 1$ subsurface, we see that c' is still genus-separating and is clearly disjoint from c. This completes the proof. □

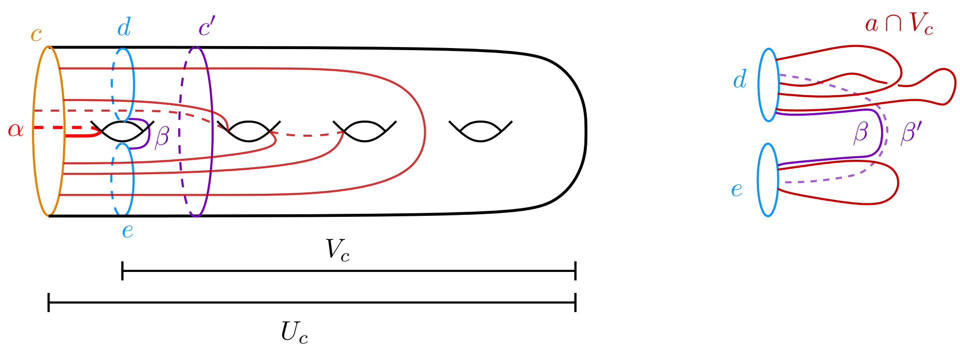

FIGURE 2. On the left, the subsurfaces involved in the proof of Lemma 4.6. On the right, surgering an arbitrary arc β' from d to e along $a \cap V_c$ to obtain a disjoint arc β.

Finally, we use the genus-separating curves to prove that the inclusion $\mathscr{C}_{\mathrm{adm}}(S,\phi) \to \mathcal{K}$ is coarsely onto.

Lemma 4.7. *There exists $D \geq 0$, depending only on S, such that for all vertices $\mu \in \mathcal{K}$, there exists an admissible curve $a \in \mathscr{C}_{\mathrm{adm}}(S,\phi)$ with $d_{\mathcal{K}}(\mu, a) \leq D$*

Proof. Let $\mathcal{P}(S)$ be the pants graph of S. This is the graph that has a vertex for each pants decomposition of S and an edge between two vertices if they differ by flipping a curve in a 4-holed sphere or 1-holed torus subsurface as described in Definition 3.4. Since pants are never witnesses (as their complement always has genus), $\mathcal{P}(S)$ is a subgraph of $\mathcal{K}$. Moreover, any vertex of $\mathcal{K}$ can be joined to a vertex of $\mathcal{P}(S)$ by adding at most $\xi(S) - 1$ curves.

Since the natural $\mathrm{Mod}(S)$-action on $\mathcal{P}(S)$ has compact quotient, there exists a number $D_0 > 0$ depending only on S so that every vertex of $\mathcal{P}(S)$ is within D_0 of a vertex of $\mathcal{P}(S)$ that contains a genus-separating curve. Hence, given any vertex $\mu \in \mathcal{K}$, we can connect it to an admissible curve a in $\mathcal{K}$ with at most $D = 2\xi(S) + D_0$ edges as follows: connect μ to a pants decomposition ρ with at most $\xi(S) - 1$ "add curve" edges; use at most D_0 "flip" edges and at most $\xi(S) - 1$ "delete curve" edges to connect ρ to a genus-separating curve c; connect c to an admissible curve $a \in \Pi(c)$ with 1 "add curve" edge and 1 "delete curve" edge. □

4.4. **A quasi-inverse.** We now construct a map Ψ that sends vertices of $\mathcal{K}$ to sets of genus-separating curves so that the composition $\Pi \circ \Psi$ is a quasi-inverse of the inclusion $\mathscr{C}_{\mathrm{adm}} \to \mathcal{K}$. The idea to is assign a multicurve $\alpha \in \mathcal{K}$ to the set of genus-separating curves that intersect the components of $S \setminus \alpha$ in a particularly nice way. This is always possible by the following lemma.

Lemma 4.8. *For any multicurve α on S, there exists a genus-separating curve c so that for each component Y of $S \setminus \alpha$, we have exactly one of the following:*

(1) c is disjoint from Y,
(2) $c \subseteq Y$,
(3) $c \cap Y$ is a single arc with both endpoints on the same curve of ∂Y, or
(4) $c \cap Y$ is a pair of parallel arcs that both go from one curve $y_1 \in \partial Y$ to a different curve $y_2 \in \partial Y$.

Proof. If a component of $S \setminus \alpha$ has positive genus, then the lemma is true using a separating curve cutting off that genus. Otherwise, the dual graph D of α on S must contain a cycle. We can use the dual graph to build such a separating curve c as follows:

(1) Take any cycle $v_1, \ldots, v_n$ in the dual graph D that meets any vertex of D at most once. Let a_i be the curve of α/edge in the dual graph connecting v_i to v_{i+1} (where indices are taken mod n).
(2) On (the closure of) each subsurface Y_i of $S \setminus \alpha$ corresponding to a vertex v_i of the cycle, choose an arc β_i connecting a_{i-1} to a_i.
(3) The concatenation of the β_i is now a curve b that meets each a_i exactly once.
(4) Set c to be a regular neighborhood of $b \cup a_n$.

By construction $c \cap Y_i$ is a pair of arcs parallel to β_i for each $i \neq 1, n$, and it follows by inspection that $c \cap Y_1$ (and $c \cap Y_n$) is a single arc with both endpoints on a_1 (and a_{n-1}, respectively). See Figure 3. □

In light of Lemma 4.8, we define a map

$$\Psi \colon \mathcal{K} \to 2^{\mathcal{G}}$$

by setting $\Psi(\alpha)$ to be the set of genus-separating curves c that satisfy the conclusion of Lemma 4.8.

Our discussion in Remark 4.3 shows that this map is rather poorly behaved. Viewing $\mathcal{K}$ as (quasi-isometric to) the cone-off of (the model $\mathcal{K}_{\mathrm{Wit}(\mathcal{G})}(S)$ for) $\mathcal{G}$, this map sends cone

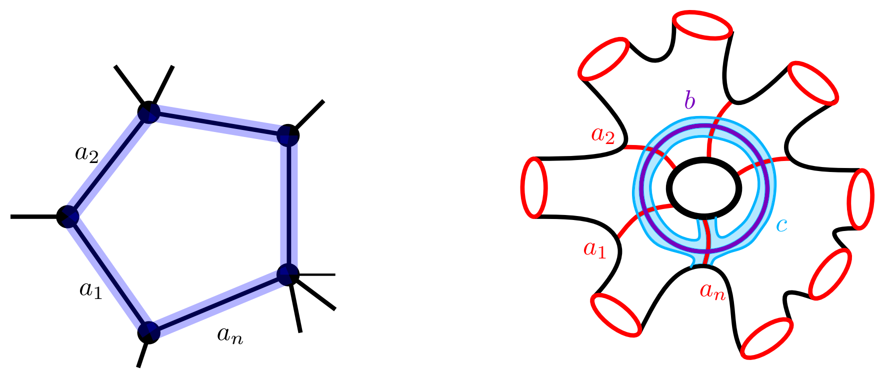

FIGURE 3. Building a genus-separating curve out of a cycle in the dual graph.

points to entire product regions. In particular, the diameter of $\Psi(\alpha)$ need not be bounded. Nevertheless, we will show that the composition $\Pi \circ \Psi$ is coarsely Lipschitz and is hence a quasi-inverse of the inclusion $\mathscr{C}_{\text{adm}} \to \mathcal{K}$.

The key technical step is the next lemma, which provides an "anchor" that allows us to modify c inside a component of $S \setminus \alpha$ without large changes in the eventual composition $\Pi \circ \Psi(\alpha)$. It is in this lemma where we need the finer control over the genus-separating curves in $\Psi(\alpha)$ ensured by Lemma 4.8 as opposed to defining $\Psi(\alpha)$ to be all genus-separating curves that intersect each curve of α some fixed number of times.

Lemma 4.9. *Let α be a multicurve in $\mathcal{K}$ and $c \in \Psi(\alpha)$. For each component Y of $S \setminus \alpha$ that c intersects, consider a component Z of $S \setminus Y$ that contains an admissible curve. Then either:*

- *there is an admissible curve $a_Y \subset Z$ disjoint from c;*
- *$g(Z) \leq 1$ and for* every *admissible curve $a \subset Z$, we have $d_{\mathscr{C}_{\text{adm}}}(a, \Pi(c)) \leq 1$.*

Remark 4.10. The subsurface Z can be an annulus whose core curve is admissible.

Proof. Since Y is not a witness for $\mathscr{C}_{\text{adm}}$ by the definition of $\mathcal{K}$, some component Z of $S \setminus Y$ contains an admissible curve. Pick any such Z. If c is disjoint from Z, then c is disjoint from the admissible curve on Z and we are done.

So suppose that c intersects Z; then $c \cap Z$ separates Z since c is separating. Since c satisfies the requirement of Lemma 4.8, c intersects either 1 or 2 of the boundary curves of Z. Moreover, if $z \in \partial Z$ intersects c, then $i(c, z) = 2$. This implies that if c intersects a single boundary curve of Z, then $c \cap Z$ must be a single separating arc on Z. Hence, $Z \setminus (c \cap Z)$ must contain genus and there is an admissible curve a_Y on Z disjoint from c.

Else, c intersects two boundary curves of $z_1, z_2 \in \partial Z$. In this case $c \cap Z$ must be a pair of essential arcs c_1, c_2 such that each c_i has both of its endpoints on z_i. If either $c_i \cap Z$ separates Z, then both must, and in particular we see that $Z \setminus (c \cap Z)$ contains genus as before, and we are done. Otherwise, each $c_i \cap Z$ is nonseparating by itself, and so $Z \setminus (c \cap Z)$ has exactly 2 components.

If $g(Z) \geq 2$, the fact that $Z \setminus (c \cap Z)$ has exactly two components implies that it has positive genus. This can be seen, for example, by noting that all 4 boundary curves of the tubular neighborhoods of $(c_i \cap Z) \cup z_i$ are homologous in Z, so their complement must have genus. Hence as in the previous case, we let a_Y be the admissible on Z which is disjoint from c.

Finally, we have reduced to the case where $g(Z) \leq 1$ and $Z \setminus (c \cap Z)$ has exactly 2 components. Since every Z-peripheral curve is either S-peripheral or Y-peripheral, we see that $S \setminus Z$ is connected and therefore has genus at least 2. By the definition of Ψ, our curve c must intersect $S \setminus Z$ in a pair of parallel arcs that separate $S \setminus Z$ into two components, one of which is a disk. Thus $(S \setminus Z) \setminus c$ has positive genus and hence contains an admissible curve $a_c \in \Psi(c)$. This curve a_c is now disjoint from *every* admissible curve on Z, and we are done. □

We now prove that $\Pi \circ \Psi(\alpha)$ has uniformly bounded diameter for each $\alpha \in \mathcal{K}$. The proof will use Lemma 4.9 to anchor the image of $\Pi \circ \Psi(\alpha)$ while we modify the genus-separating curves on the components of $S \setminus \alpha$ to reduce intersection numbers.

Proposition 4.11. *There are $L \geq 0$ such that for any $\alpha \in \mathcal{K}$ and $c, d \in \Psi(\alpha)$, there is $c' \in \Psi(\alpha)$ with*

(1) $i(c', d) \leq L \cdot |\chi(S)|$ and

(2) The diameter of $\Pi(c) \cup \Pi(c')$ in $\mathscr{C}_{\mathrm{adm}}$ is at most $4|\chi(S)|$.

In particular, $\Pi \circ \Psi(\alpha)$ has uniformly bounded diameter for all $\alpha \in \mathcal{K}$.

Proof. Throughout the proof, we fix representatives of the isotopy classes of all of the curves involved such that c and d are each in minimal position with respect to α, and such that no points of $c \cap d$ lie on α. This allows us to give meaning to statements like "c and d intersect on a component Y of $S \setminus \alpha$" even though there is no canonical minimal position for triples of isotopy classes of curves.

Having fixed representatives, the proposition will follow by inductively applying the following claim.

Claim 4.12. *There is an $L \geq 0$ depending only on $\xi(S)$ so that for and component Y of $S \setminus \alpha$ on which c and d intersect, then there exists $c_Y \in \Psi(\alpha)$ such that c_Y and d intersect at most L times on Y and c_Y agrees with c on $S \setminus Y$.*

Proof. By construction, each of $c \cap Y$ and $d \cap Y$ is either a single arc connecting a boundary component to itself (which necessarily separates Y) or a pair of parallel arcs connecting different boundary components (and neither of these arcs can individually separate Y).

Since $\mathrm{Mod}(Y)$ acts with finitely orbit of pairs of isotopy classes of arcs, there exist a constant L and arc(s) γ on Y such that

- $i(\gamma, d \cap Y) \leq L$ and L depends only on $\xi(Y) < \xi(S)$,
- there is a homeomorphism $h \colon Y \setminus (c \cap Y) \to Y \setminus \gamma$ that is the identity on the boundary curves of Y.

The desired curve c_Y is then obtained by replacing $c \cap Y$ with γ. □

To prove Proposition 4.11, let $Y_1, \ldots Y_k$ be the components of $S \setminus \alpha$ on which c and d intersect. Applying Claim 4.12 to Y_1, we get a genus-separating curve $c_1 \in \Psi(\alpha)$ that intersects d at most L times in Y_1 and agrees with c outside of Y_1. Using Lemma 4.9, we have two options:

- There exists an admissible curve a that is disjoint from both c and Y_1. Thus a is also disjoint from c_1.
- There is a genus 0 or 1 component Z of $S \setminus Y_1$ that contains an admissible curve and every admissible curve a on Z has $d_{\mathscr{C}_{\text{adm}}}(a, \Pi(c)) \leq 1$ and $d_{\mathscr{C}_{\text{adm}}}(a, \Pi(c_1)) \leq 1$.

In both cases, we have at least one admissible curve a_1 with $d_{\mathscr{C}_{\text{adm}}}(a_1, \Pi(c)) \leq 1$ and $d_{\mathscr{C}_{\text{adm}}}(a, \Pi(c_1)) \leq 1$. Thus $d_{\mathscr{C}_{\text{adm}}}(\Pi(c), \Pi(c_1)) \leq 2$.

Repeating this argument, we produce a sequence of genus-separating curves $c = c_0, c_1, \dots, c_k$ in $\Psi(\alpha)$ such that $\Pi(c_i)$ and $\Pi(c_{i+1})$ are 2-close in $\mathscr{C}_{\text{adm}}$ and $i(c_k, d)$ is at most L times the number of components of $S \setminus \alpha$, which is at most $|\chi(S)|$. The final curve c_k is the desired curve c'. Now $\Pi(c) \cup \Pi(c')$ has diameter at most $4|\chi(S)|$, because k is bounded by $|\chi(S)|$, each $\Pi(c_i)$ has diameter 2 in $\mathscr{C}_{\text{adm}}$, and each $\Pi(c_i)$ and $\Pi(c_{i+1})$ are 2-close.

Finally, c' and d have uniformly bounded intersection number by construction, so by Lemma 4.4 they have uniformly bounded distance in $\mathcal{G}$. Since Π is coarsely Lipschitz (Lemma 4.5), we see that $\Pi(c') \cup \Pi(d)$ also has uniformly bounded diameter. The last statement of Proposition 4.11 now follows by the triangle inequality. □

We now show that the admissible curve graph $\mathscr{C}_{\text{adm}}$ is quasi-isometric to the model $\mathcal{K}$. Since the inclusion $\mathscr{C}_{\text{adm}} \to \mathcal{K}$ is 1-Lipschitz and coarsely onto (Lemmas 3.5 and 4.7), this statement is implied by the following:

Proposition 4.13. *The map $\Pi \circ \Psi \colon \mathcal{K} \to \mathscr{C}_{\text{adm}}$ is a quasi-inverse to the inclusion $\mathscr{C}_{\text{adm}} \to \mathcal{K}$.*

Proof. We first check that for all $a \in \mathscr{C}_{\text{adm}}$, the image $\Pi \circ \Psi(a)$ is uniformly close to a in $\mathscr{C}_{\text{adm}}$. Since $g(S) \geq 3$, there must exists a genus-separating curve c disjoint from a. Hence $c \in \Psi(a)$ and $a \in \Pi(c)$. Thus $a \in \Pi \circ \Psi(a)$ as desired.

We now show that $\Pi \circ \Psi$ is coarsely Lipschitz; this will complete the proof of Proposition 4.13. We have already shown in Proposition 4.11 that the image of every vertex of $\mathcal{K}$ has uniformly bounded diameter, so it suffices to do the same for every edge. That is, if $\alpha, \alpha' \in \mathcal{K}$ are two vertices joined by an edge, then we must show that

$$\operatorname{diam}(\Pi \circ \Psi(\alpha) \cup \Pi \circ \Psi(\alpha'))$$

is uniformly bounded.

If the edge from α to α' corresponds to adding a curve to α to achieve α', then $\Psi(\alpha') \subseteq \Psi(\alpha)$ by definition. This implies $\Pi \circ \Psi(\alpha') \subseteq \Pi \circ \Psi(\alpha)$; the desired diameter bound then follows from Proposition 4.11.

Now assume the edge from α to α' corresponds to a flip move. Let $x \in \alpha$ and $x' \in \alpha'$ such that x is flipped to x'. If x and x' are disjoint, then $\alpha \cup x'$ is a vertex of $\mathcal{K}$ as adding curves to a vertex of $\mathcal{K}$ always produces a new vertex of $\mathcal{K}$. Now $\alpha \cup x'$ is joined by an edge to both α and α' as removing x' produces α and removing x produces α'. The desired bound now follows from the proceeding paragraph about add/remove edges.

If x and x' are not disjoint, then the component Y of $S \setminus (\alpha \setminus x)$ that contains x has $\xi(Y) = 1$. If Y is not a witness, then $\alpha \setminus x = \alpha' \setminus x'$ is a vertex of $\mathcal{K}$ that is joined by an add/remove-edge to both α and α'. As before this establishes the bound.

If Y is a witness, then Lemma 3.2 requires $S \setminus Y$ has no genus. Since $\xi(Y) = 1$ and $g(S) \geq 3$, some topological casework shows this is only possible if $g(S) = 3$, Y is a 4-holed sphere, every Y-peripheral curve is not S-peripheral, and Y does not separate S. In this case, x and x' intersect twice in the 4-holed sphere Y. Thus, flipping α to α' corresponds to moving from the dual graph D for α to the dual graph D' for α' by performing a "Whitehead move" where one collapses the edge of D dual to x and then expands an edge dual to x'; see Figure 4. Since no Y-peripheral curves are separating or S-peripheral, the dual graph D contains a cycle C with an edge dual to x such that after performing the Whitehead move to produce D', the cycle C becomes a cycle C' of D' that does not include the edge dual to x'. There is therefore a genus-separating curve c built from C that will be disjoint from x', which implies $c \in \Psi(\alpha) \cap \Psi(\alpha')$. Since $\Pi(c)$ will then be contained in $\Pi \circ \Psi(\alpha) \cap \Pi \circ \Psi(\alpha')$, we have that $\mathrm{diam}(\Pi \circ \Psi(\alpha) \cup \Pi \circ \Psi(\alpha'))$ is uniformly bounded by Proposition 4.11. □

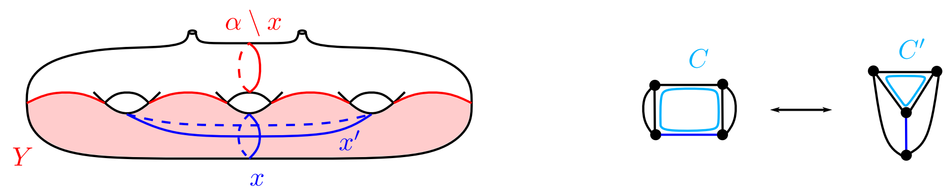

FIGURE 4. On the left, the subsurface Y where x is flipped to x'. On the right, the Whitehead move on the dual graph corresponding to flipping x to x'. The cycle C is sent to the cycle C' under this move.

Proof of Theorem A. Lemma 3.5 and Proposition 4.13 together show that $\mathscr{C}_{\mathrm{adm}}$ is quasi-isometric to the hierarchically hyperbolic space $\mathcal{K}$. Since hierarchical hyperbolicity can be passed along quasi-isometries, $\mathscr{C}_{\mathrm{adm}}$ is also hierarchically hyperbolic.

As Gromov hyperbolicity is also a quasi-isometry invariant, it suffices to to verify that $\mathcal{K}$ is not Gromov hyperbolic. By Corollary 3.6, $\mathcal{K}$ is not Gromov hyperbolic if and only if $\mathscr{C}_{\mathrm{adm}}$ has a pair of disjoint witnesses. Let $\Delta_1, \dots, \Delta_n$ be peripheral curves encircling the punctures of S. Without loss of generality, assume $\phi(\Delta_i) \geq 0$ for $i \in \{1, \dots, k\}$ and $\phi(\Delta_i) < 0$ for $i \in \{k+1, \dots, n\}$. Let α be a multicurve consisting of $g+1$ non-separating curves $a_1, \dots, a_{g+1}$ such that $S \setminus \alpha$ is a pair of genus zero subsurfaces, W^+ and W^-, where W^+ contains $\Delta_1, \dots, \Delta_k$ and W^- contains $\Delta_{k+1}, \dots, \Delta_n$; see Figure 5. Orient each curve of α such that W^+ is to the left.

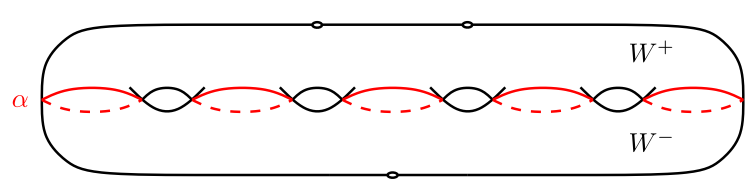

FIGURE 5. The multicurve α whose complement is a pair of witnesses for $\mathscr{C}_{\mathrm{adm}}$.

By homological coherence (Lemma 2.2.2), we have that for any framing ψ of S,

$$\sum_{i=1}^{g+1} x_i + \sum_{j=1}^{k} \psi(\Delta_j) = 1 - g - k \tag{3}$$

where $x_i = \psi(a_i)$. On the other hand, we know from Lemma 3.2 that W^+ contains a ψ-admissible curve (i.e., W is a witness) if and only if there is some subset $\mathcal{C}$ of $\alpha \cup \Delta_1 \cup \ldots \cup \Delta_k$ such that $\mathcal{C} \cap \alpha \neq \emptyset$, $\alpha \not\subseteq \mathcal{C}$, and

$$\sum_{c \in \mathcal{C}} \psi(c) = 1 - |\mathcal{C}|. \tag{4}$$

A similar condition tells us if W^- contains admissible curves/W^+ is a witness.

Now since g of the curves of α are homologically independent, we see that for any $(x_1, \ldots, x_{g+1}) \in \mathbb{Z}^{g+1}$ such that Equation (3) holds, there is a framing ψ of S such that $\psi(a_i) = x_i$ for all i and $\psi(\Delta_j) = \phi(\Delta_j)$ for each $j \in \{1, \ldots, n\}$ (see [CS22, Remark 2.7]). We can choose x_i not to satisfy Equation (4) for any subset $\mathcal{C}$ of $\alpha \cup \Delta_1 \cup \ldots \cup \Delta_k$ that meets α, and similarly for the corresponding equations for W^- since these are all linearly independent by Equation (3). Thus W^+ and W^- are a pair of disjoint witnesses for $\mathscr{C}_{\text{adm}}(S, \psi)$.

The choices in the previous paragraph can all be made explicitly. Set

$$K = \sum |\phi(\Delta_j)| + g + n + 2$$

and choose $x_1, \ldots, x_g$ all to be positive and larger than $2K$ and such that their differences are all larger than $2K$. Set x_{g+1} to satisfy (3), so it will necessarily be very negative. Then for any subset $\mathcal{C}$ of $\alpha \cup \Delta_1, \ldots, \Delta_k$ that intersects but does not contain α, the left-hand side of (4) has magnitude larger than $K \geq |\mathcal{C}| - 1$. Thus W^+ contains no admissible curves, so W^- is a witness. The argument for W^+ is completely analogous.

Finally, we note that in the case that ϕ is of spin type, we can also choose ψ to have the same Arf invariant as ϕ by stipulating the winding numbers on the completion of $a_1, \ldots, a_g$ to a GSB. Theorem 2.5 now provides $f \in \text{Mod}(S)$ such that $\phi = f(\psi)$, and thus $f(W^+)$ and $f(W^-)$ are the desired pair of disjoint witnesses for $\mathscr{C}_{\text{adm}}(S, \phi)$. □

Because Proposition 4.13 proves that the inclusion $\mathscr{C}_{\text{adm}} \to \mathcal{K}$ is a quasi-isometry, the HHS structure on $\mathscr{C}_{\text{adm}}$ is the same as the structure described by Vokes for graphs with a cocompact $\text{Mod}(S)$-action. In particular, the projections maps are the standard subsurface projections maps onto the curve graphs of witnesses and relations are given by nesting and transversality of witnesses. See [Vok22] for full details on the HHS structure.

5. Curve graphs for strata

In this section we define a number of versions of the curve complex that capture the incidence patterns of (bordifications of) strata. After reviewing the relationship between strata and framed mapping class groups (Section 5.1) and how the curve complex can be viewed as a Tits building for the (augmented) Teichmüller space (Section 5.2), we discuss some of the results of [BCG+19], which provides a nice compactification of strata using certain enhanced multicurves. In the final Section 5.4, we use this compactification to build a variety of different complexes which play the analogue of the curve complex. Unlike the classical case, it is much more subtle to determine exactly which nodal surfaces can appear in the boundary, leading us to define a number of different graphs that we will eventually prove are quasi-isometric (Corollary 6.5).

5.1. **Framings and strata.** We start by recalling some of the results of [CS22] on the relationship between strata, markings, and framed mapping class groups. We refer the reader to that paper as well as Sections 2 and 3 of [BSW22] for a more thorough discussion of these relationships, and to [Zor06] for a general reference on abelian differentials.

A *stratum* of abelian differentials is a (quasi-projective) subvariety of the bundle of holomorphic abelian differentials $\Omega\mathcal{M}_g$ on genus g Riemann surfaces defined by conditioning the number and order of zeros. More explicitly, given any partition $\underline{\kappa} = (k_1, \ldots, k_n)$ of $2g-2$ into positive integers, we let $\Omega\mathcal{M}_g(\underline{\kappa}) \subset \Omega\mathcal{M}_g$ denote the stratum parametrizing pairs (X, ω) where X is a Riemann surface and ω is a holomorphic 1-form on X with n distinct zeros of orders $k_1, \ldots, k_n$. Since a holomorphic 1-form is entirely determined up to global scaling by the order and position of its zeros, the projectivization of a stratum $\mathbb{P}\Omega\mathcal{M}_g(\underline{\kappa})$ may be thought of as a subvariety of $\mathcal{M}_{g,n}$. In the sequel, we will freely conflate a stratum and this subvariety in $\mathcal{M}_{g,n}$; we trust this will not cause any confusion.

Strata are not necessarily connected, but each has at most 3 irreducible components [KZ03]. Let $\mathbb{P}\Omega\mathcal{T}_{g,n}(\underline{\kappa})$ denote the full preimage of the stratum $\mathbb{P}\Omega\mathcal{M}_g(\underline{\kappa})$ inside of $\mathcal{T}_{g,n}$. By the orbit-stabilizer theorem, to understand the connected components covering a given component $\mathcal{H}$ of $\mathbb{P}\Omega\mathcal{M}_g(\underline{\kappa})$, it is enough to understand the stabilizer of one of them. Given $g \in \mathrm{Mod}(S_{g,n})$ stabilizing a component of $\mathbb{P}\Omega\mathcal{T}_{g,n}(\underline{\kappa})$ covering $\mathcal{H}$, one can connect a point X and its image gX by a path in that component which projects to a loop in $\mathcal{H}$. Conversely, any loop in $\mathcal{H}$ lifts to a path connecting points in the same component of $\mathbb{P}\Omega\mathcal{T}_{g,n}(\underline{\kappa})$. Thus in order to understand the connected components of $\mathbb{P}\Omega\mathcal{T}_{g,n}(\underline{\kappa})$, one needs to understand the image of the map

$$\rho : \pi_1(\mathcal{H}) \to \pi_1(\mathcal{M}_{g,n}) \cong \mathrm{Mod}(S_{g,n})$$

of orbifold fundamental groups for each component $\mathcal{H}$ of $\mathbb{P}\Omega\mathcal{M}_g(\underline{\kappa})$.

When $\mathcal{H}$ consists entirely of hyperelliptic differentials (which happens only when $\underline{\kappa} = (2g-2)$ or $(g-1, g-1)$ [KZ03]), it is not hard to see that the image of ρ is (conjugate to) a hyperelliptic mapping class group [LM14, Cal20]. The main theorem of [CS22] characterizes the image of ρ for non-hyperelliptic components.

Observe first that a differential ω has an associated horizontal vector field that does not vanish outside the zeros of ω; we denote this by $1/\omega$.

Theorem 5.1 (Theorem A of [CS22])**.** *Let $\mathcal{H}$ be a non-hyperelliptic stratum component and suppose that $g \geq 5$. Then the image of ρ is (conjugate to) the framed mapping class group associated to the framing $1/\omega$.*

We therefore introduce the following notation:

Definition 5.2. Suppose that $\mathcal{H}$ is a non-hyperelliptic stratum component and let $(X, \omega) \in \mathcal{H}$. Choose an arbitrary marking $f : S_{g,n} \to X$ and let ϕ denote the framing corresponding to the vector field $f^*(1/\omega)$. Then we use $\mathcal{H}_\phi$ to denote the subset of $\mathbb{P}\Omega\mathcal{T}_{g,n}(\underline{\kappa})$ parametrizing those marked differentials (X', ω', f') such $(X', \omega') \in \mathcal{H}$ and $(f')^*(1/\omega')$ is isotopic to ϕ.

By Theorem 5.1, if $g \geq 5$ then $\mathcal{H}_\phi$ is just a specified connected component of $\mathbb{P}\Omega\mathcal{T}_{g,n}(\underline{\kappa})$. The reader should think of $\mathcal{H}_\phi$ this way; Definition 5.2 is written as it is only so that we have something that works for all $g \geq 3$.

The Theorem also reveals a relationship between cylinders and admissible curves. Integrating ω induces a singular flat metric on X, and the core curve of any embedded Euclidean cylinder has constant slope with respect to the horizontal vector field $1/\omega$, hence winding number 0. Moreover, since the cylinder has nonzero period with respect to a holomorphic 1-form, the core curve must necessarily be non-separating by Stokes's theorem. Thus the core curve is admissible. Transitivity of the $\mathrm{FMod}(S, \phi)$ action on admissible curves (see Proposition 2.10) now implies that every admissible curve is realized as a cylinder on some differential in $\mathcal{H}_\phi$ [CS22, Corollary 1.2].

In Section 6 below, we will use similar transitivity arguments to understand which multicurves can be pinched in the boundary of $\mathcal{H}_\phi$.

5.2. **The curve complex as a nerve.** Recall that the Deligne–Mumford compactification $\overline{\mathcal{M}_{g,n}}$ of the moduli space of Riemann surfaces is obtained by adjoining moduli spaces corresponding to (stable) nodal surfaces to $\mathcal{M}_{g,n}$. Equivalently, it can also be obtained by taking the completion of $\mathcal{M}_{g,n}$ with respect to the Weil–Petersson metric. A sequence of surfaces X_i degenerates to the boundary if the (extremal or hyperbolic) length of an essential simple closed curve goes to 0; if γ is a topological type of multicurve, then we use $\mathcal{M}_{g,n}(\gamma)$ to denote the boundary stratum where γ is pinched.

One can do a similar thing at the level of Teichmüller space. For any multicurve γ, let $\mathcal{T}_{g,n}(\gamma)$ denote the Teichmüller space of the open subsurface $S \setminus \gamma$. The *augmented Teichmüller space* $\overline{\mathcal{T}_{g,n}}$ is then obtained by adjoining all possible $\mathcal{T}_{g,n}(\gamma)$ to $\mathcal{T}_{g,n}$, marking $S \setminus \gamma$ by the subsurface complementary to γ. Equivalently, $\overline{\mathcal{T}_{g,n}}$ is also the Weil–Petersson metric completion of $\mathcal{T}_{g,n}$. Points in $\mathcal{T}_{g,n}(\gamma)$ can be obtained as geometric limits: for example, if $\mathcal{T}_{g,n} \ni X_i \to X_\infty \in \mathcal{T}_{g,n}(\gamma)$ then the hyperbolic length of γ on X_i goes to 0, so the X_i develop a long collar that limits to a pair of cusps in X_∞.

We direct the reader to [HK14] and its extensive bibliography for a thorough discussion of the history and construction of these spaces.

Remark 5.3. It is useful (though not quite correct) to think of $\overline{\mathcal{T}_{g,n}}$ as covering $\overline{\mathcal{M}_{g,n}}$. There is a surjective map $\overline{\mathcal{T}_{g,n}} \to \overline{\mathcal{M}_{g,n}}$, which when restricted to any stratum $\mathcal{T}_{g,n}(\gamma)$ is a covering onto $\mathcal{M}_{g,n}(\gamma)$, but the overall map is not a covering. This is because $\mathcal{T}_{g,n}$ is infinitely ramified around the boundary stratum $\mathcal{T}_{g,n}(\gamma)$ (and likewise $\mathcal{T}_{g,n}(\gamma)$ is infinitely ramified around its boundary, etc).

The collar lemma implies that the nerve of the (closures of the) top-dimensional boundary strata of $\overline{\mathcal{T}_{g,n}}$ is exactly given by the usual curve complex $\mathscr{C}(S)$ (with vertices given by simple closed curves and simplices given by disjointness). The barycentric subdivision of the curve complex is the *multicurve complex*, which is both the order complex of the poset of multicurves on S as well as the order complex of the poset of the (closures of) all boundary strata of $\overline{\mathcal{T}_{g,n}}$, both ordered by inclusion.

5.3. **Multi-scale differentials and enhanced multicurves.** In this and the next section, we perform a similar construction for (projectivized, marked) strata of abelian differentials.

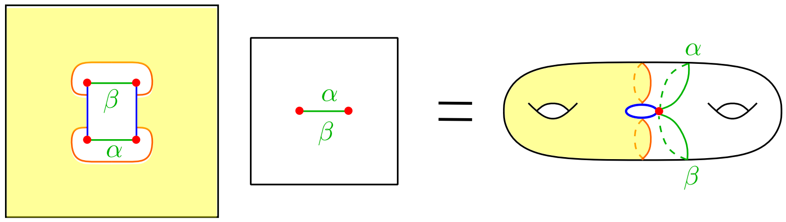

FIGURE 6. A 2-level multicurve consisting of two homologous curves. Neither α nor β is admissible, and neither separates S by itself, so by Theorem 5.8 below, $\overline{\mathcal{H}_\phi}$ does not meet $\mathcal{T}_{g,n}(\alpha)$ or $\mathcal{T}_{g,n}(\beta)$.

Our discussion is complicated by a number of factors, one of which is that some curves on S cannot be pinched by themselves (since any abelian differential is in particular a cohomology class). In fact, if $\overline{\mathcal{H}_\phi} \cap \mathcal{T}_{g,n}(\gamma) \neq \emptyset$ and γ is a single simple closed curve, then it must either be admissible or separating. See Section 5.4 just below.

Example 5.4 (Pinching a multicurve but not its components)**.** Consider the surface shown in Figure 6. In this example, curves α and β are homologous and so their periods must be equal. Crushing the right-hand torus to have 0 area degenerates into the boundary stratum $\overline{\mathcal{H}_\phi} \cap \mathcal{T}_{g,n}(\alpha \cup \beta)$. However, it is impossible to pinch either α or β individually while remaining in $\overline{\mathcal{H}_\phi}$.

Specifying $\mathcal{H}_\phi \subset \mathcal{T}_{g,n}$ as in Definition 5.2, let $\overline{\mathcal{H}_\phi} \subset \overline{\mathcal{T}_{g,n}}$ denote its closure (equivalently, its Weil–Petersson metric completion). Exactly which boundary strata $\overline{\mathcal{H}_\phi}$ meets is a very intricate question, and is related to subtle properties of certain compactifications of $\mathcal{H}$. Let $\overline{\mathcal{H}}$ be the closure of $\mathcal{H}$ inside of $\overline{\mathcal{M}_{g,n}}$ (without markings). The structure of its boundary is determined by the so-called "incidence variety compactification" (IVC) of $\mathcal{H}$ [BCG$^+$18]. A point in the IVC consists of a "level graph" and a "twisted differential" compatible with the level graph; forgetting the differential and remembering only the underlying complex structure yields a surjective map from the IVC onto $\overline{\mathcal{H}}$ [BCG$^+$18, Corollary 1.4].

It turns out that the IVC is highly singular, and in [BCG$^+$19], the IVC is refined into the moduli space of "multi-scale differentials" $\mathbb{P}\Xi\mathcal{H}$ which has nicer geometric properties (e.g., its boundary is a normal crossing divisor). We will not give the precise definition of multi-scale differentials here, and direct the reader to the original papers, especially Section 5 of [BCG$^+$19] and Section 3 of [CMZ22]. Instead, we simply record the relevant topological data of multicurves and winding numbers. We keep the numbering conventions of [BCG$^+$19].

Definition 5.5. Let $S = S_{g,n}$.

(1) An *ordered multicurve* is a multicurve Λ together with a partition of $S \setminus \Lambda$ into (nonempty, but possibly disconnected) subsurfaces $Y_0, \ldots, Y_{-N}$. We say that such a multicurve has $N+1$ levels, or N level passages.

(2) The individual curves of Λ are called *horizontal* if they connect subsurfaces in the same level, and *vertical* otherwise. We say Λ is horizontal/vertical if all of its constituent curves are. Every Λ can be partitioned into its horizontal and vertical pieces, which

we denote by α and β, respectively. We will refer to a vertical multicurve with 2 levels as a *BIC* (for bi-colored graph, see [CMZ23]).

(3) An *enhanced multicurve* (also denoted by Λ) is an ordered multicurve together with a choice of positive number $\kappa(b)$ for each vertical $b \subset \beta$.

(4) Given a framing ϕ of S, an enhanced multicurve Λ is *ϕ-compatible* if:
   - For each horizontal curve $a \subset \alpha$, we have $\phi(a) = 0$.
   - For each vertical curve $b \subset \beta$, we have $\phi(b) = -\kappa(b)$ when b is oriented such that the surfaces on its left and right are in Y_i and Y_j, where $i > j$.

(5) Given a framing ϕ of S, we say that an (unordered, unenhanced) multicurve γ is ϕ-compatible if there exists some enhancement Λ of γ which is ϕ-compatible. We similarly say γ is N-level, horizontal, vertical, or a BIC if there is a ϕ-compatible enhancement with these properties.

Throughout, we use Λ to denote an enhanced multicurve and γ an unenhanced one.

Remark 5.6. It is possible that an (unenhanced) multicurve may be compatible with N-level enhancements for multiple different N (compare the "cherry divisors" from [BCG$^+$19, §14.3]).

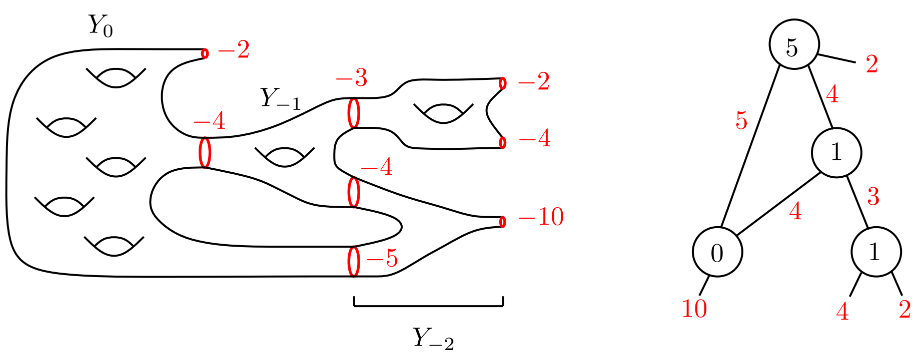

FIGURE 7. An enhanced multicurve with 3 levels and corresponding enhanced level graph.

When ϕ has holomorphic type and β is vertical, every boundary component and peripheral curve of its top (that is, 0^{th}) level has negative winding number. Homological coherence (Lemma 2.2.2) then implies that each of these components must have positive genus. This corresponds to the fact that the top level of a multi-scale differential in the boundary of a holomorphic stratum must itself be holomorphic if its level graph has no horizontal edges. We record this for later use:

Fact 5.7. *Let ϕ be a framing of holomorphic type and let β be a vertical multicurve for ϕ. Then each component of Y_0 has positive genus.*

The moduli space $\mathbb{P}\Xi\mathcal{H}$ of (projectivized) multi-scale differentials is obtained by constructing a bordification of the entire (projectivized) marked stratum $\mathbb{P}\Omega\mathcal{T}_{g,n}(\underline{\kappa})$ then quotienting out by the action of the entire mapping class group. More specifically, in [BCG$^+$19, §7] the authors construct an *augmented Teichmüller space of projective multi-scale differentials* $\mathbb{P}\Xi\mathcal{T}_{g,n}(\underline{\kappa})$ which is stratified into subspaces corresponding to enhanced multicurves Λ:

$$\mathbb{P}\Xi\mathcal{T}_{g,n}(\underline{\kappa}) = \bigsqcup_{\Lambda} \mathbb{P}\Omega\mathcal{B}_{\Lambda},$$

where $\mathbb{P}\Omega\mathcal{B}_{\Lambda}$ is the quotient of a so-called "prong-matched Teichmüller space" (a marked stratum with some extra combinatorial decoration) by $\mathbb{C}^{\#\text{levels}}$, representing the simultaneous rotation of the surfaces on a given level. In particular, if Λ is the empty multicurve then the corresponding stratum is just $\mathbb{P}\Omega\mathcal{T}_{g,n}(\underline{\kappa})$.

We now restrict to a single component $\mathcal{H}_\phi \subset \mathbb{P}\Omega\mathcal{T}_{g,n}(\underline{\kappa})$ and consider its closure $\mathbb{P}\Xi\mathcal{H}_\phi \subset \mathbb{P}\Xi\mathcal{T}_{g,n}(\underline{\kappa})$. The winding numbers of curves yield a necessary condition for how $\mathcal{H}_\phi$ meets the boundary:

Theorem 5.8. *If $\mathbb{P}\Xi\mathcal{H}_\phi \cap \mathbb{P}\Omega\mathcal{B}_{\Lambda}$ is nonempty, then Λ is ϕ-compatible.*

By the fact that the $\mathbb{P}\Xi\mathcal{H}_\phi$ surjects onto $\overline{\mathcal{H}_\phi}$ (via the IVC), we deduce a corresponding statement for unenhanced multicurves.

Corollary 5.9. *If $\overline{\mathcal{H}_\phi} \cap \mathcal{T}_{g,n}(\gamma)$ is nonempty, then γ is ϕ-compatible.*

The paper [MUW21] gives a sufficient condition for a topological type of unenhanced multicurve to appear in the boundary of an unmarked stratum. Using results from the literature, one can refine this result to stratum *components* $\overline{\mathcal{H}}$. In particular, using [Won24] one can determine exactly how the Arf invariants of subsurfaces plus prong matchings enforce the total Arf invariant, and the main result of [CF22] implies that the global residue condition does not impose any further restrictions.

To further upgrade this dicussion to marked strata (and their compactifications) one would also need to establish full transitivity results for the action of the framed mapping class group. We were unable to achieve this level of generality, and will instead focus our attention on the "largest" boundary strata (see Definition 5.12). For coarse-geometric questions, this distinction is irrelevant.

5.4. **Boundary complexes for marked strata.** In analogy with Section 5.2, we now build a variety of complexes capturing different facets of the boundary structure of $\mathbb{P}\Xi\mathcal{H}_\phi$.

The closest analogues of Tits buildings/standard curve complexes are based on the incidences of boundary strata. Suppose that $\mathcal{B}$ and $\mathcal{B}'$ are irreducible components of $\mathbb{P}\Xi\mathcal{H}_\phi \cap \mathbb{P}\Omega\mathcal{B}_{\Lambda}$ and $\mathbb{P}\Xi\mathcal{H}_\phi \cap \mathbb{P}\Omega\mathcal{B}_{\Lambda'}$, respectively. One says that $\mathcal{B}$ *degenerates* to $\mathcal{B}'$ if the closure of $\mathcal{B}$ in $\mathbb{P}\Xi\mathcal{T}_{g,n}(\underline{\kappa})$ contains $\mathcal{B}'$.

Definition 5.10. Let $\mathscr{B}(\mathbb{P}\Xi\mathcal{H}_\phi)$ be the order complex of the poset of such components $\mathcal{B}$, with order given by degenerations. Said another way, it is the flag complex whose vertices are components of $\mathbb{P}\Xi\mathcal{H}_\phi \cap \mathbb{P}\Omega\mathcal{B}_{\Lambda}$, as Λ ranges over all enhanced multicurves, and whose edges are given by degeneration.

The connectivity properties of $\mathscr{B}(\mathbb{P}\Xi\mathcal{H}_\phi)$ are still poorly understood.[2] Instead, we will focus on an intermediate complex that records not components $\mathcal{B}$ but rather their indexing enhanced multicurves.

[2]For example, given an enhanced multicurve Λ, the intersection $\mathbb{P}\Xi\mathcal{H}_\phi \cap \mathbb{P}\Omega\mathcal{B}_{\Lambda}$ is not necessarily connected. This is related to the difficulty of establishing transitivity properties for the action of $\mathrm{FMod}(S, \phi)$.

Definition 5.11. Let $\mathscr{C}(\mathbb{P}\Xi\mathcal{H}_\phi)$ be the flag complex whose vertices are enhanced multicurves Λ such that $\mathbb{P}\Xi\mathcal{H}_\phi \cap \mathbb{P}\Omega\mathcal{B}_\Lambda$ is nonempty. Two vertices Λ and Λ' are connected by an edge if and only if there are components $\mathcal{B}$ and $\mathcal{B}'$ of $\mathbb{P}\Xi\mathcal{H}_\phi \cap \mathbb{P}\Omega\mathcal{B}_\Lambda$ and $\mathbb{P}\Xi\mathcal{H}_\phi \cap \mathbb{P}\Omega\mathcal{B}_{\Lambda'}$ such that $\mathcal{B}$ degenerates to $\mathcal{B}'$.

Equivalently, $\mathscr{C}(\mathbb{P}\Xi\mathcal{H}_\phi)$ is obtained from $\mathscr{B}(\mathbb{P}\Xi\mathcal{H}_\phi)$ by collapsing the vertices corresponding to different components of $\mathbb{P}\Xi\mathcal{H}_\phi \cap \mathbb{P}\Omega\mathcal{B}_\Lambda$ to a single point corresponding to Λ (and then extending to simplices). Thus there is a 1-Lipschitz map $\mathscr{B}(\mathbb{P}\Xi\mathcal{H}_\phi) \to \mathscr{C}(\mathbb{P}\Xi\mathcal{H}_\phi)$, though the aforementioned connectivity issues mean that we do not know if it is a quasi-isometry.

The complex $\mathscr{C}(\mathbb{P}\Xi\mathcal{H}_\phi)$ is very similar to the graph $\mathscr{C}(\overline{\mathcal{H}_\phi})$ appearing in Theorem B, which we recall has a vertex for each *unenhanced* multicurve γ such that $\overline{\mathcal{H}_\phi} \cap \mathcal{T}_{g,n}(\gamma)$ is nonempty and with an edge between γ and γ' if the closure of $\overline{\mathcal{H}_\phi} \cap \mathcal{T}_{g,n}(\gamma)$ meets $\mathcal{T}_{g,n}(\gamma')$. It should come as no surprise, then, that they are quasi-isometric. We defer the formal statement and proof until the next section (Corollary 6.6), as it will be much easier once we have built other quasi-isometric models.

Comparing with the classical setting, we see that $\mathscr{C}(\mathbb{P}\Xi\mathcal{H}_\phi)$ is analogous to the complex of multicurves on S since its simplices are given by inclusion, not intersection. To get an analogue of the curve complex itself, we will want to restrict ourselves to considering only the "largest" boundary strata. Since its boundary is a normal-crossing divisor, the multi-scale compactification provides a good notion of size. Given an enhanced multicurve Λ, the (complex) codimension of $\mathbb{P}\Omega\mathcal{B}_\Lambda$ in $\mathbb{P}\Xi\mathcal{T}_{g,n}(\underline{\kappa})$ is equal to the number of admissible curves in Λ plus the number of levels minus 1 [BCG+19, p. 4]. We therefore make the following definition:

Definition 5.12. An enhanced multicurve Λ is *divisorial* for $\mathcal{H}_\phi$ if $\mathbb{P}\Xi\mathcal{H}_\phi \cap \mathbb{P}\Omega\mathcal{B}_\Lambda$ is nonempty (hence Λ is ϕ-compatible) and Λ is either a single horizontal curve or a BIC.

We now define a complex recording only the intersection pattern of boundary divisors.

Definition 5.13. Set $\mathscr{D}(\mathbb{P}\Xi\mathcal{H}_\phi)$ to be the complex whose vertices are given by divisorial enhanced multicurves, and such that $\{\Lambda_1, \ldots, \Lambda_n\}$ span a simplex if

$$\bigcap_{i=1}^{n} \overline{\mathbb{P}\Xi\mathcal{H}_\phi \cap \mathbb{P}\Omega\mathcal{B}_{\Lambda_i}} \neq \emptyset,$$

where the closures are taken in $\mathbb{P}\Xi\mathcal{H}_\phi$.

We caution the reader that despite first appearances, the barycentric subdivision of $\mathscr{D}(\mathbb{P}\Xi\mathcal{H}_\phi)$ is *not necessarily* equal to $\mathscr{C}(\mathbb{P}\Xi\mathcal{H}_\phi)$. This is because the intersection of two boundary divisors may not be irreducible, and different components may correspond to different enhanced multicurves. See [CMZ23, Figure 6] for an example of this phenomenon and [Dev26, Section 2] for a discussion of the genus 0 case. Since these subtleties appear only when considering positive-dimensional simplices, we have the following:

Lemma 5.14. *$\mathscr{C}(\mathbb{P}\Xi\mathcal{H}_\phi)$ and $\mathscr{D}(\mathbb{P}\Xi\mathcal{H}_\phi)$ are quasi-isometric.*

Proof. We note first that since $\mathscr{D}(\mathbb{P}\Xi\mathcal{H}_\phi)$ is finite-dimensional, it is quasi-isometric to its 1-skeleton $\mathscr{D}^{(1)}$. There is an obvious map f from the vertices of $\mathscr{D}(\mathbb{P}\Xi\mathcal{H}_\phi)$ into $\mathscr{C}(\mathbb{P}\Xi\mathcal{H}_\phi)$.

It is not hard to see f extends to a 2-Lipschitz map on the $\mathscr{D}^{(1)} \to \mathscr{C}(\mathbb{P}\Xi\mathcal{H}_\phi)$, as if Λ_1, Λ_2 are adjacent in $\mathscr{D}(\mathbb{P}\Xi\mathcal{H}_\phi)$ then there are boundary components $\mathcal{B}_i$ of $\mathbb{P}\Xi\mathcal{H}_\phi \cap \mathbb{P}\Omega\mathcal{B}_{\Lambda_i}$ whose closures intersect. Pick any component of $\overline{\mathcal{B}_1} \cap \overline{\mathcal{B}_2}$; then it lives in some $\mathbb{P}\Omega\mathcal{B}_{\Lambda_3}$, and hence Λ_1 is connected to Λ_3 is connected to Λ_2 in $\mathscr{C}(\mathbb{P}\Xi\mathcal{H}_\phi)$. Composing this with the natural quasi-isometry $\mathscr{D}(\mathbb{P}\Xi\mathcal{H}_\phi) \cong \mathscr{D}^{(1)}$, we get a coarsely Lipschitz map (also called f) from $\mathscr{D}(\mathbb{P}\Xi\mathcal{H}_\phi)$ to $\mathscr{C}(\mathbb{P}\Xi\mathcal{H}_\phi)$.

On the other hand, there is a cellular (hence 1-Lipschitz) map $\bar{f}$ from $\mathscr{C}(\mathbb{P}\Xi\mathcal{H}_\phi)$ to the barycentric subdivision of $\mathscr{D}(\mathbb{P}\Xi\mathcal{H}_\phi)$, obtained by sending each Λ to the vertex corresponding to the simplex spanned by its divisorial undegenerations.[3] To check that f and $\bar{f}$ are coarse inverses, we observe that $\bar{f} \circ f$ is the identity on the vertices of $\mathscr{D}(\mathbb{P}\Xi\mathcal{H}_\phi)$, while $f \circ \bar{f}$ sends Λ to its divisorial undegenerations, which are by definition distance 1 away in $\mathscr{C}(\mathbb{P}\Xi\mathcal{H}_\phi)$. □

Finally, we define one more complex which is amenable to the HHS techniques from the previous sections. First, we note that when one restricts to boundary divisors, the distinction between enhanced and unenhanced multicurves fades away:

Lemma 5.15. *Given a framing ϕ, every (unenhanced) multicurve γ admits at most one ϕ-compatible, divisorial enhancement.*

Proof. Since fixing a framing ϕ stipulates the winding numbers of curves, an unenhanced multicurve γ can be enhanced to be horizontal and ϕ-compatible if and only if it is a single admissible curve. Similarly, if γ is a BIC then for each constituent curve $b \subset \gamma$, the winding number of b stipulates if Y_0 is to its left or right. Thus γ can be enhanced to a ϕ-compatible BIC if and only if these comparisons are all consistent. □

In light of this Lemma, we say that an unenhanced multicurve γ is *divisorial* for ϕ if it admits an enhancement which is divisorial for ϕ.

Definition 5.16. Set $\mathscr{E}(\overline{\mathcal{H}_\phi})$ to be the graph whose vertices are divisorial multicurves, with an edge between γ and γ' if and only if they have disjoint representatives.

We note that with this definition, adjacent γ and γ' may share curves. One could define yet another graph in which the constituent curves of any two adjacent γ and γ' must also be distinct, but the result will be quasi-isometric to $\mathscr{E}(\overline{\mathcal{H}_\phi})$. (We leave a proof to the interested reader — it follows quickly from Corollary 6.5.)

By Lemma 5.15, $\mathscr{D}(\mathbb{P}\Xi\mathcal{H}_\phi)$ and $\mathscr{E}(\overline{\mathcal{H}_\phi})$ have the same vertices. Since any two boundary divisors intersect in a component in which both are pinched, the collar lemma implies that the edges of $\mathscr{D}(\mathbb{P}\Xi\mathcal{H}_\phi)$ are all edges of $\mathscr{E}(\overline{\mathcal{H}_\phi})$. However, $\mathscr{E}(\overline{\mathcal{H}_\phi})$ may have extra edges which do not correspond to intersections. We thank Martin Möller for first bringing this phenomenon to our attention.

Example 5.17 (Pinching curves but not their union)**.** Suppose that S has a single puncture and let Δ denote a curve encircling that puncture. Let c and d be separating curves on S such that (Δ, c, d) bounds a pair of pants P. Then $Y_i(c)$ and $Y_i(d)$ both have genus for $i = 0, -1$, and by the main Theorem of [MUW21] (or explicit construction), one sees that both c and d are divisorial.

[3]Here we are using the fact that the boundary of $\mathbb{P}\Xi\mathcal{H}$ is a normal crossing divisor.

However, if $\overline{\mathcal{H}_\phi}$ were to meet $\mathcal{T}_{g,n}(c \cup d)$, then on any multiscale differential corresponding to this boundary stratum the pair of pants P would be equipped with a meromorphic differential with a single zero and two poles. Stokes' theorem would then imply that the residues at each pole would be 0, but there is no meromorphic differential on $\widehat{\mathbb{C}}$ with a single zero and two poles of zero residue.

6. From transitivity to geometry

All of the graphs defined in the previous section carry a natural action of the relevant framed mapping class group. Throughout the section, fix a non-hyperelliptic stratum component $\mathcal{H} \subset \mathcal{M}_{g,n}$ of holomorphic type with $g \geq 3$. Let $\mathcal{H}_\phi$ be as in Definition 5.2. Set $S = S_{g,n}$. Then if $\overline{\mathcal{H}_\phi} \cap \mathcal{T}_{g,n}(\gamma) \neq \emptyset$, we have for any $f \in \mathrm{FMod}(S, \phi)$,

$$f(\overline{\mathcal{H}_\phi} \cap \mathcal{T}_{g,n}(\gamma)) = \overline{\mathcal{H}_\phi} \cap \mathcal{T}_{g,n}(f(\gamma)) \neq \emptyset.$$

Similar statements are true for the augmented Teichmüller space of multiscale differentials and enhanced multicurves. In this section, we use similar transitivity properties to relate the geometries of the complexes from the previous section, a hierarchically hyperbolic model, and the admissible curve graph $\mathscr{C}_{\mathrm{adm}}(S, \phi)$. This will complete the proof of our main Theorem B.

As a first example of this technique, let us prove the following:

Lemma 6.1. *Both $\mathscr{D}(\mathbb{P}\Xi\mathcal{H}_\phi)$ and $\mathscr{E}(\overline{\mathcal{H}_\phi})$ contain $\mathscr{C}_{\mathrm{adm}}(S, \phi)$.*

Proof. It suffices to prove the statement for $\mathscr{D}(\mathbb{P}\Xi\mathcal{H}_\phi)$ as its 1-skeleton is a subgraph of $\mathscr{E}(\overline{\mathcal{H}_\phi})$. Since every admissible curve is divisorial (compare [CS22, Corollary 1.2] and the discussion at the end of Section 5.1), $\mathscr{D}(\mathbb{P}\Xi\mathcal{H}_\phi)$ contains the vertices of $\mathscr{C}_{\mathrm{adm}}(S, \phi)$, so it remains to show that it also contains the edges.

We begin by noting that every multicurve α, all of whose curves are ϕ-admissible, admits a single ϕ-compatible enhancement. Proposition 2.10 states that the framed mapping class group $\mathrm{FMod}(S, \phi)$ acts transitively on pairs of admissible curves of the same topological type. Thus, it suffices to show that $\mathbb{P}\Xi\overline{\mathcal{H}_\phi}$ meets a boundary stratum for each *topological type* of pair α of admissible curves.

This can be done by explicit construction, one possibility of which we sketch below. The restriction of ϕ to $S \setminus \alpha$ is a framing with four boundary components of winding number 0. By holomorphicity of ϕ and homological coherence, each component of $S \setminus \alpha$ either has positive genus or each peripheral curve on that component has winding number -1. Pick meromorphic differentials on the components of $S \setminus \alpha$ inducing the same framing and with simple poles corresponding to α, all of the same residue (this can be done because strata of meromorphic differentials on surfaces of genus ≥ 1 with simple poles are always nonempty [Boi15], and the genus 0 case corresponds to adding free marked points on a cylinder). Cutting the infinite cylinders and gluing them together along α yields a holomorphic differential in the correct stratum; applying the (unframed) mapping class group then allows us to ensure that it actually lies in $\mathcal{H}_\phi$. Degenerating these cylinders by letting their heights go to ∞ then produces a path in $\mathcal{H}_\phi$ to $\mathbb{P}\Omega\mathcal{B}_\Lambda$, where Λ is the unique ϕ-compatible enhancement of α.

The only thing one might worry about is matching the Arf invariants of the subsurfaces to ensure that the plumbed surface has the correct Arf invariant: this turns out not to be an issue for the following reason. If $g \geq 4$ then each stratum of meromorphic differentials in genus ≥ 2 has components of both spin parities [Boi15, Theorem 1.2], so by choosing the appropriate Arf invariants on pieces we can ensure that the plumbed surface has the appropriate Arf invariant. In the special case that $g = 3$, there is a unique component of meromorphic differentials on a genus 1 surface with two simple poles and a single zero of order 2, and so the plumbed surface is forced to have odd Arf invariant. Fortunately, this only happens in the stratum $\mathbb{P}\Omega\mathcal{M}_3(2,2)$, which has a unique non-hyperelliptic component of odd Arf invariant [KZ03, Theorem 2]. □

6.1. **The action on BICs.** Recall from Definition 5.5 that an unenhanced multicurve γ is called a BIC if it admits a (necessarily unique) 2-level vertical enhancement. BICs can have many different topological types, so $\mathrm{FMod}(S, \phi)$ will not act transitively on them. However, even controlling for topological type and winding numbers, the Arf invariants of complementary subsurfaces present additional invariants of the $\mathrm{FMod}(S, \phi)$ orbit. We show below that these are the only obstructions to transitivity.

While we will not use this in the sequel, we note that the following statement is true for *all* BICs compatible with a framing of holomorphic type, not just divisorial ones.

Proposition 6.2. *Let ϕ be a framing of holomorphic type on a surface S of genus at least 3 and let β be any BIC. Then a multicurve β' is in the* $\mathrm{FMod}(S, \phi)$ *orbit of β if and only if there exists an $h \in \mathrm{Mod}(S)$ such that:*

(1) $h(\beta) = \beta'$.
(2) $\phi(b) = \phi(h(b))$ for every curve $b \in \beta$.
(3) For each component U of $S \setminus \beta$ of genus at least 2 such that $\phi|_U$ is of spin type,

$$\mathrm{Arf}(\phi|_U) = \mathrm{Arf}(\phi|_{h(U)}),$$

and similarly, for any complementary component U of genus 1,

$$\mathrm{Arf}_1(\phi|_U) = \mathrm{Arf}_1(\phi|_{h(U)}).$$

Proof. We are given $h \in \mathrm{Mod}(S)$ taking β to β'; our goal is to find an element in the (orientation-preserving, component-wise) stabilizer of β' such that its composition with h preserves the winding numbers of a GSB for S. We will construct this element and the associated GSB in steps, starting from the bottom and working up. The reader is invited to compare with the discussion of "perturbed period coordinates," especially Figure 5, in [BCG$^+$19]. Throughout the proof, given any curve of β or component U of $S \setminus \beta$, we will add a prime to denote its image under h, i.e., $U' := h(U)$.

Bottom level: Choose a GSB $\mathcal{B}_U$ on each component U of $Y_{-1}(\beta)$. Hypothesis (3) allows us to apply Lemma 2.8 to choose a GSB $\mathcal{B}_{U'}$ for U' with the same set of winding numbers as appear in $\mathcal{B}_U$. Using the classical change-of-coordinates principle, we can find some element $f_{U'} \in \mathrm{Mod}(U')$ (which we can then think of as living in $\mathrm{Mod}(S)$ by extending by the identity outside of U') that takes $h(\mathcal{B}_U)$ to $\mathcal{B}_{U'}$. Set

$$f_{bot} = \prod_{U' \subset Y_{-1}(\beta')} f_{U'} \circ h;$$

by construction it takes β to β' and preserves the winding numbers of a GSB for $Y_{-1}(\beta)$.

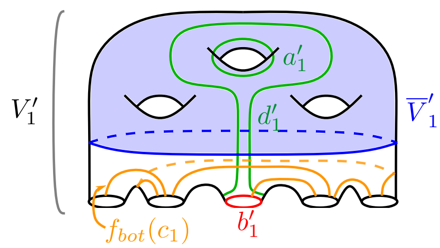

FIGURE 8. The multitwist needed to fix the winding number of $f_{bot}(c_1)$. Note that it may not always be possible to choose the curve d_1' to be disjoint from the other curves of $f_{bot}(\mathsf{c})$.

Level passage: For this and the next step, for each component V of the top level $Y_0(\beta)$, pick a subsurface $\overline{V} \subset V$ with full genus and a single boundary component.

Pick a maximally homologically independent subset $\mathsf{b} = (b_1, \ldots, b_h)$ of β and extend

$$\bigcup_{U \subset Y_{-1}(\beta)} \mathcal{B}_U \cup \mathsf{b}$$

to a GSB for the complement of all of the $\overline{V}$. Let $\mathsf{c} = (c_1, \ldots, c_h)$ denote the resulting set of curves symplectically dual to those of b. The element f_{bot} will most likely not preserve the winding numbers of c, but we can rectify this using elements supported entirely on $Y_0(\beta')$.

Consider first c_1; the dual curve b_1 is a curve of the BIC β, and we use V_1 to denote the component of $Y_0(\beta)$ adjacent to b_1. Since ϕ is of holomorphic type, V_1 has genus, hence so does the full-genus subsurface $\overline{V}_1 \subset V_1$. Set $\mathsf{c}' = f_{bot}(\mathsf{c})$ and likewise for its components. Pick some admissible curve $a_1' \subset \overline{V}_1' := h(\overline{V}_1) = f_{bot}(\overline{V}_1)$ and let d_1' denote the connect sum of b_1' with a_1' along some arc contained in V_1'. Since a_1' is disjoint from c' and b_1' only meets c_1', we see that the algebraic intersection number of d_1' with each c_j' is 0 unless $j = 1$, in which case it is exactly 1. See Figure 8.

By homological coherence (Lemma 2.2.2),

$$\phi(b_1') - \phi(d_1') = -1$$

when appropriately oriented. Thus, if we set

$$f_1 := \left(T_{b_1'}^{\pm} T_{d_1'}^{\pm}\right)^{\phi(c_1') - \phi(c_1)}$$

for an appropriate choice of signs, then by twist-linearity (Lemma 2.2.1) we see that $f_1(c_1')$ has the same winding number as c_1 and that f_1 preserves the winding numbers of all other c_j'.

We now repeat the above procedure but with $f_1 \circ f_{bot}$ instead of f_{bot}. More precisely, set $\overline{V}_2 \subset V_2$ to be the full-genus subsurface of the top-level component V_2 of $S \setminus \beta$ adjacent to b_2.[4] There is an admissible $a_2' \subset \overline{V}_2' := h(\overline{V}_2) = f_1 f_{bot}(\overline{V}_2)$, and taking the connect sum of b_2' with this curve yields some d_2' whose algebraic intersection with each curve of $f_1 f_{bot}(\mathsf{c})$ is 0 except for $f_1 f_{bot}(c_2)$. Taking an appropriate multitwist in b_2' and d_2' yields some f_2 supported on V_2' such that $f_2 f_1 f_{bot}$ preserves the winding numbers of both c_1 and c_2.

[4]The component V_2 (and subsurface $\overline{V}_2$) may be the same as V_1 (and $\overline{V}_1$).

Iterating, we get a sequence of mapping classes $f_1, \dots, f_h$ all supported on $Y_0(\beta')$ such that the composite

$$f_{mid} := f_h \circ \dots \circ f_1 \circ f_{bot}$$

takes β to β' and preserves the winding numbers of the curves of a GSB for the complement of a full-genus subsurface of the top-level subsurface $Y_0(\beta)$.

Top level: To finish, we can simply use the action of the mapping class group of top-level subsurfaces to amend the winding numbers of the remaining curves.

Consider one component V of $Y_0(\beta)$ and let $\overline{V} \subset V$ denote the full-genus subsurface coming from the previous step. Pick a GSB $\mathcal{B}_{\overline{V}}$ for $\overline{V}$. Then by Lemma 2.6 and hypothesis (3), we have that

$$\mathrm{Arf}(\phi|_{\overline{V}}) = \mathrm{Arf}(\phi|_V) = \mathrm{Arf}(\phi|_{V'}) = \mathrm{Arf}(\phi|_{f_{mid}(\overline{V})})$$

when defined. If V is of genus 1, the same thing holds for the genus 1 Arf invariant (note that this requires the fact that $\phi|_V$ is of holomorphic type!). Lemma 2.8 then implies that $f_{mid}(\overline{V})$ admits a GSB with the same winding numbers as $\mathcal{B}_{\overline{V}}$, and we pick an element $f_V \in \mathrm{Mod}(f_{mid}(\overline{V}))$ taking $f_{mid}(\mathcal{B}_{\overline{V}})$ to this GSB.

Finally, we observe that the mapping class $\prod_V f_V \circ f_{mid}$ takes β to β' and preserves the winding numbers of the following GSB for S:

$$\bigcup_{U \subset Y_{-1}(\beta)} \mathcal{B}_U \cup \mathsf{b} \cup \mathsf{c} \cup \bigcup_{V \subset Y_0(\beta)} \mathcal{B}_V.$$

We have therefore constructed the desired framed mapping class. □

6.2. **Pinching admissible curves.** Using the same ideas as Proposition 6.2, we show that every BIC is connected to some admissible curve in $\mathscr{D}(\mathbb{P}\Xi\mathcal{H}_\phi)$.

Proposition 6.3. *Suppose β is a divisorial BIC for ϕ. For any admissible $a \subset Y_0(\beta)$,*

$$\mathbb{P}\Xi\mathcal{H}_\phi \cap \mathbb{P}\Omega\mathcal{B}_{\Lambda \cup a} \neq \emptyset,$$

where Λ is the unique enhancement of β to a ϕ-compatible BIC and a is horizontal in $\Lambda \cup a$.

Proof. We first show that one can pinch *some* admissible curve in $Y_0(\beta)$. The restriction of any multiscale differential in $\mathbb{P}\Omega\mathcal{B}_\Lambda$ is holomorphic on $Y_0(\beta)$ (Fact 5.7). Every holomorphic differential contains an embedded cylinder [Mas86] whose core curve a' is necessarily admissible, and one can degenerate into $\mathbb{P}\Omega\mathcal{B}_{\Lambda \cup a'}$ by sending the height of this cylinder to ∞.

Thus it suffices to show that the stabilizer of β in $\mathrm{FMod}(S, \phi)$ acts transitively on the set of admissible curves contained in each component of $Y_0(\beta)$. Let a and a' be different admissible curves contained in the same component V of $Y_0(\beta)$. Note that since $\phi|_V$ has holomorphic type, this implies that a and a' must both be nonseparating in V, so the multicurves $\beta \cup a$ and $\beta \cup a'$ are of the same topological type. Postcomposing by an element of $\mathrm{Mod}(V)$ as necessary, we may therefore assume that there is an element $h \in \mathrm{Mod}(S)$ taking β to β' and a to a'. This is the starting point for the proof of Proposition 6.2.

The proof now proceeds by upgrading h into a framed mapping class: we show that each step, this can done be done in a way that preserves a' and so the composite element still takes a to a'.

Bottom level: The element f_{bot} differs from h by an element supported entirely on the bottom level $Y_{-1}(\beta)$, so still takes a to a'.

Level passage: Pick the full-genus subsurface $\overline{V} \subset V$ to contain a', so $f_{bot}(\overline{V})$ contains a'. Each element f_j is constructed by taking a multitwist disjoint from some choice of admissible curve. So long as we pick a' to be this admissible base curve each time that V is the relevant subsurface in the iteration, then the resulting multitwist f_j will preserve a' and the new subsurface $f_j \cdots f_1 f_{bot}(\overline{V})$ will still contain a'. Thus f_{mid} must also take a to a'.

Top level: The last step of the construction takes a fixed GSB of V to a GSB of $f_{mid}(V)$ with the same winding numbers. We now just make sure to take a as an element of the GSB of V and take a' to be the corresponding element of the GSB of $f_{mid}(V)$. The fact that we can extend a' to a GSB with the appropriate winding numbers is immediate when $g(V) = 1$, as divisoriality implies that $\mathrm{Arf}_1(\phi_V) = 0$, so every nonseparating curve is admissible.

Otherwise, suppose $g(V) \geq 2$. Let b denote the curve of the GSB of V dual to a, and let W denote the complement in V of the subsurface filled by $a \cup b$. Note $g(W) = g(V) - 1$. Using Corollary 2.12, we can choose a curve b' in $f_{mid}(V)$ dual to a' with $\phi(b') = \phi(b)$. Let W' denote the complement in $f_{mid}(V)$ of the subsurface filled by $a' \cup b'$. By additivity of the Arf invariant [RW14, Lemma 2.11], we get that $\mathrm{Arf}(\phi|_W) = \mathrm{Arf}(\phi|_{W'})$. If $g(W) \geq 2$, then Lemma 2.8 implies there exist GSBs on W and W' with the same winding numbers, completing the proof.

In the special case when $g(W) = 1$, we observe that $\phi|_W$ and $\phi|_{W'}$ both have signature $(-3, 1)$. In particular, $\mathrm{Arf}_1(\phi|_W) = 1$ or 2, and inspection of definitions gives

$$\mathrm{Arf}_1(\phi|_W) = \mathrm{Arf}(\phi|_W) + 1 \pmod 2$$

(and likewise for W'). In particular, equality of the (usual) Arf invariants implies equalities of the genus 1 Arf invariants, so we can apply Lemma 2.8 to conclude. ☐

Combining this with Lemma 6.1 allows us to quickly conclude connectivity.

Corollary 6.4. *$\mathscr{D}(\mathbb{P}\Xi\mathcal{H}_\phi)$ and $\mathscr{E}(\overline{\mathcal{H}_\phi})$ are connected.*

Proof. Since $\mathscr{D}(\mathbb{P}\Xi\mathcal{H}_\phi)$ and $\mathscr{E}(\overline{\mathcal{H}_\phi})$ have the same vertices and every edge of $\mathscr{D}(\mathbb{P}\Xi\mathcal{H}_\phi)$ is an edge of $\mathscr{E}(\overline{\mathcal{H}_\phi})$, it suffices to prove this for $\mathscr{D}(\mathbb{P}\Xi\mathcal{H}_\phi)$.

By Lemma 6.1, $\mathscr{D}(\mathbb{P}\Xi\mathcal{H}_\phi)$ contains the admissible curve graph $\mathscr{C}_{\mathrm{adm}}(S, \phi)$, and as shown in Lemma 3.1, $\mathscr{C}_{\mathrm{adm}}(S, \phi)$ is connected. Every vertex of $\mathscr{D}(\mathbb{P}\Xi\mathcal{H}_\phi)$ that is not an admissible curve is a 2-level multicurve β. There is some admissible curve a contained inside of each component of $Y_0(\beta)$ (Fact 5.7 and Lemma 2.7) and so Proposition 6.3 implies that $\overline{\mathcal{H}_\phi}$ meets $\mathcal{T}_{g,n}(\beta \cup a)$. Thus β and a are connected in $\mathscr{D}(\mathbb{P}\Xi\mathcal{H}_\phi)$. Since we have connected every vertex of $\mathscr{D}(\mathbb{P}\Xi\mathcal{H}_\phi)$ to the connected graph $\mathscr{C}_{\mathrm{adm}}(S, \phi)$, we conclude that $\mathscr{D}(\mathbb{P}\Xi\mathcal{H}_\phi)$ is connected. ☐

Pushing this line of reasoning slightly further, we also get the following:

Corollary 6.5. *$\mathscr{D}(\mathbb{P}\Xi\mathcal{H}_\phi)$ and $\mathscr{E}(\overline{\mathcal{H}_\phi})$ are quasi-isometric.*

Proof. It is enough to prove that any two adjacent vertices in $\mathscr{E}(\overline{\mathcal{H}_\phi})$ can be connected by a path of length 4 in $\mathscr{D}(\mathbb{P}\Xi\mathcal{H}_\phi)$.

To that end, consider any edge of $\mathscr{E}(\overline{\mathcal{H}_\phi})$ connecting disjoint, divisorial multicurves γ and δ. If both γ and δ are single admissible curves, then Lemma 6.1 implies they are adjacent in $\mathscr{D}(\mathbb{P}\Xi\mathcal{H}_\phi)$. Now suppose γ is admissible and δ is a BIC. As in the proof of Corollary 6.4, each component of $Y_0(\delta)$ has genus, and if $\gamma \subset Y_0(\delta)$ then Proposition 6.3 implies that γ and δ are adjacent in $\mathscr{D}(\mathbb{P}\Xi\mathcal{H}_\phi)$. Otherwise, $\gamma \subset Y_{-1}(\delta)$ and in particular it is disjoint from $Y_0(\delta)$. By Lemma 2.7, there is an admissible curve a on $Y_0(\delta)$. Applying Proposition 6.3 again we see that γ is adjacent to a which is adjacent to δ (by Lemma 6.1).

Finally, suppose that both γ and δ are (divisorial) BICs. If any components of $Y_0(\gamma)$ and $Y_0(\delta)$ are nested, then since they both have genus their intersection does. In particular by Lemma 2.7 there is some admissible curve a disjoint from both γ and δ. We may then invoke Proposition 6.3 again to connect γ to δ in $\mathscr{D}(\mathbb{P}\Xi\mathcal{H}_\phi)$ through a. Otherwise, $Y_0(\gamma)$ and $Y_0(\delta)$ are disjoint. In this case we choose admissible curves a_γ and a_δ inside $Y_0(\gamma)$ and $Y_0(\delta)$, respectively. The previous two paragraphs then imply that $(\gamma, a_\gamma, a_\delta, \delta)$ is an edge path in $\mathscr{D}(\mathbb{P}\Xi\mathcal{H}_\phi)$. □

Finally, we relate all of these complexes to the graph introduced in the introduction.

Corollary 6.6. *$\mathscr{C}(\overline{\mathcal{H}_\phi})$ and $\mathscr{D}(\mathbb{P}\Xi\mathcal{H}_\phi)$ are quasi-isometric.*

Proof. Define a map f on the vertices of $\mathscr{D}(\mathbb{P}\Xi\mathcal{H}_\phi)$ into $\mathscr{C}(\overline{\mathcal{H}_\phi})$ taking a given divisorial enhanced multicurve Λ to its underlying unenhanced multicurve. By definition of their edges, any two adjacent vertices of $\mathscr{D}(\mathbb{P}\Xi\mathcal{H}_\phi)$ are sent to points distance 2 away in $\mathscr{C}(\overline{\mathcal{H}_\phi})$, so f is Lipschitz. Since the forgetful map is a surjection $\mathbb{P}\Xi\mathcal{H}_\phi \twoheadrightarrow \overline{\mathcal{H}_\phi}$ and every boundary stratum in $\mathbb{P}\Xi\mathcal{H}_\phi$ is an intersection of divisorial ones, we see that f is coarsely onto.

We now build a coarse inverse $\bar{f}$ to conclude that f is a quasi-isometry, as follows. Let δ_i denote the family of all divisorial (unenhanced) multicurves adjacent to γ in $\mathscr{C}(\overline{\mathcal{H}_\phi})$. We then set $\bar{f}(\gamma)$ to be the set of the (unique) divisorial enhancements of the δ_i; see Lemma 5.15. The δ_i are all disjoint, so the proof of Corollary 6.5 shows they are distance at most 4 from each other in $\mathscr{D}(\mathbb{P}\Xi\mathcal{H}_\phi)$, hence $\bar{f}$ is coarsely well-defined. By definition, $\bar{f} \circ f$ is the identity and $f \circ \bar{f}$ moves points distance at most 1 in $\mathscr{C}(\overline{\mathcal{H}_\phi})$, so this is a coarse inverse. □

6.3. **A quasi-isometry with the model.** We now build off our work showing $\mathscr{C}_{\mathrm{adm}}(S, \phi)$ is hierarchically hyperbolic to prove that $\mathscr{E}(\overline{\mathcal{H}_\phi})$ (hence $\mathscr{D}(\mathbb{P}\Xi\mathcal{H}_\phi)$, $\mathscr{C}(\mathbb{P}\Xi\mathcal{H}_\phi)$, and $\mathscr{C}(\overline{\mathcal{H}_\phi})$) are as well.

As in the case of the admissible curve graph, we establish the hierarchical hyperbolicity of $\mathscr{E}(\overline{\mathcal{H}_\phi})$ by showing it is quasi-isometric to a model graph constructed from its witnesses. Because $\mathrm{Wit}(\mathscr{E}(\overline{\mathcal{H}_\phi}))$ is a proper subset of $\mathrm{Wit}(\mathscr{C}_{\mathrm{adm}}(S, \phi))$, our proof for $\mathscr{E}(\overline{\mathcal{H}_\phi})$ will actually rely on the proof for $\mathscr{C}_{\mathrm{adm}}(S, \phi)$. To describe this setup, we need the following notation:

- Set $\xi = \xi(S)$, the cardinality of the largest set of disjoint curves on $S = S_{g,n}$.
- Let $\mathcal{D}$ be the set of divisorial BICs for $\overline{\mathcal{H}_\phi}$ (these are exactly the vertices of $\mathscr{E}(\overline{\mathcal{H}_\phi})$ that are not in $\mathscr{C}_{\mathrm{adm}}(S, \phi)$).
- Let $\mathfrak{S} = \mathrm{Wit}(\mathscr{C}_{\mathrm{adm}}(S, \phi))$ and $\overline{\mathfrak{S}} = \mathrm{Wit}(\mathscr{E}(\overline{\mathcal{H}_\phi}))$. By definition,
$$\overline{\mathfrak{S}} = \mathfrak{S} \setminus \{W \in \mathfrak{S} : \exists \delta \in \mathcal{D} \text{ with } \delta \cap W = \emptyset\}.$$
- Let $\mathcal{K} = \mathcal{K}_{\mathfrak{S}}$ denote the quasi-isometric model for $\mathscr{C}_{\mathrm{adm}}(S, \phi)$ (Definition 3.4) and let $\overline{\mathcal{K}}$ denote $\mathcal{K}_{\overline{\mathfrak{S}}}$.

By construction, there are 1-Lipschitz inclusion maps

$$i\colon \mathcal{K} \to \overline{\mathcal{K}} \text{ and } \iota\colon \mathscr{C}_{\mathrm{adm}}(S,\phi) \to \mathscr{E}(\overline{\mathcal{H}_\phi}).$$

The idea behind our proof that $\overline{\mathcal{K}}$ is quasi-isometric to $\mathscr{E}(\overline{\mathcal{H}_\phi})$ is to show that the decreases in distances that happen under $i\colon \mathcal{K} \to \overline{\mathcal{K}}$ coarsely match the decreases that happen under $\iota\colon \mathscr{C}_{\mathrm{adm}}(S,\phi) \to \mathscr{E}(\overline{\mathcal{H}_\phi})$.

To formalize this idea, we define

$$P(\mu) = \{\alpha \in \mathcal{K} : \mu \subseteq \alpha\}$$

for any multicurve μ on S. If μ is a multicurve such that $S \setminus \mu$ does not contain a subsurface in $\mathfrak{S}$, then μ is a vertex of $\mathcal{K}$ and every vertex of $P(\mu)$ is connected to μ by a path with at most ξ edges (corresponding to removing curves until only μ is left). On the other hand, if $S \setminus \mu$ does contain a subsurface in $\mathfrak{S}$, then $P(\mu)$ has infinite diameter; see [RV19, Corollary 4.10]. Hence, if μ is a vertex of $\overline{\mathcal{K}}$, but not $\mathcal{K}$, then $P(\mu)$ is an infinite diameter subset of $\mathcal{K}$ that becomes finite diameter under $i\colon \mathcal{K} \to \overline{\mathcal{K}}$. If $\overline{\mathcal{K}}$ is to be quasi-isometric to $\mathscr{E}(\overline{\mathcal{H}_\phi})$, we would like the image of $P(\mu)$ under the quasi-isometry $\mathcal{K} \to \mathscr{C}_{\mathrm{adm}}(S,\phi)$ to have uniformly bounded diameter under $\iota\colon \mathscr{C}_{\mathrm{adm}}(S,\phi) \to \mathscr{E}(\overline{\mathcal{H}_\phi})$.

Our candidate quasi-isometry $\Theta\colon \overline{\mathcal{K}} \to 2^{\mathscr{E}(\overline{\mathcal{H}_\phi})}$ is therefore

$$\Theta(\mu) = \iota \circ \Pi \circ \Psi(P(\mu))$$

where $\Pi \circ \Psi$ is the quasi-isometry from $\mathcal{K}$ to $\mathscr{C}_{\mathrm{adm}}(S,\phi)$ constructed in Section 4.

As suggested above, the main work required to prove Θ is a quasi-isometry is to show that $\iota \circ \Pi \circ \Psi(P(\mu))$ is a bounded diameter subset of $\mathscr{E}(\overline{\mathcal{H}_\phi})$. To achieve this, we need to know that $\mathscr{E}(\overline{\mathcal{H}_\phi})$ is obtained from $\mathscr{C}_{\mathrm{adm}}(S,\phi)$ by adding enough divisorial BICs to collapse the image of $P(\mu)$ for each μ that is in $\overline{\mathcal{K}}$ but not $\mathcal{K}$. The abundance of BICs comes from the following lemma.

Lemma 6.7. *For any $\mathcal{H}_\phi$, there is an N, depending only on S, such that the following holds. Let W be a genus 0 witness of $\mathscr{C}_{\mathrm{adm}}(S,\phi)$ and let $\beta \in \mathcal{D}$ be a divisorial BIC disjoint from W. Then for any multicurve α on $S \setminus W$, there is an $f \in \mathrm{FMod}(S,\phi)$ such that $f(\beta)$ remains disjoint from W and $i(\alpha, f(\beta)) \leq N$.*

Note that in particular, $f(\beta) \in \mathcal{D}$ since $\mathcal{D}$ is a union of $\mathrm{FMod}(S,\phi)$ orbits.

This a weaker, framed version of the following standard "change of coordinates" lemma (compare Claim 4.12).

Lemma 6.8. *For any surface Z, there is an N_Z such that for any multicurves or multiarcs α and β, there is a $g \in \mathrm{Mod}(Z)$ such that $i(\alpha, g(\beta)) \leq N_Z$.*

As in Section 4, the reason that we cannot use a similar "change of coordinates" argument to deduce Lemma 6.7 (even though we have shown the set of divisorial BICs for $\overline{\mathcal{H}_\phi}$ is a finite union of $\mathrm{FMod}(S,\phi)$ orbits) is that there are infinitely many $\mathrm{FMod}(S,\phi)$ orbits of witnesses. Instead of making a finite number of arbitrary choices, we will instead need to be more clever and make a infinite number of good ones.

Proof of Lemma 6.7. We first record a number of topological consequences of our hypotheses. Let Y_0 and Y_{-1} denote the two levels of the unique ϕ-compatible enhancement of β. Then since W is a witness and (S,ϕ) is of holomorphic type, we see that

- $W \subset Y_0$,
- each component of Y_{-1} has genus 0,
- Y_0 is connected, and
- Y_0 has genus at least 2.

Indeed, contradicting any of the first three statements immediately implies that there is an admissible curve disjoint from W (Lemma 2.7). The last assertion follows similarly: Y_0 always has positive genus, and some curve of ∂W is always non-separating on Y_0. If the genus of Y_0 were equal to 1, then since β is divisorial and the differential on Y_0 must be holomorphic (Fact 5.7), we see $\mathrm{Arf}_1(Y_0) = 0$. Thus every nonseparating curve on Y_0, and in particular some peripheral curve of our witness W, is admissible. This is a contradiction.

We observe that since $S \setminus W$ has genus 0 (again since W is a witness plus Lemma 2.7), the winding number of any curve on $S \setminus W$ is determined by how it partitions the curves of ∂W and the punctures of S. Thus, any mapping class supported entirely on $S \setminus W$ necessarily preserves the winding numbers of the curves of β. Applying Lemma 6.8 for $Z = S \setminus W$ (and using the inclusion homomorphism for subsurfaces [FM12, Theorem 3.18]), we can therefore find some $h(\beta) \in \mathrm{Mod}(S) \cdot \beta$ that has bounded intersection with α and with the same winding numbers as β. Note that $h(\beta)$ is a ϕ-compatible BIC, though it may not be divisorial.

In the case that Y_0 is not of spin type, or if $\mathrm{Arf}(\phi|_{hY_0})$ equals $\mathrm{Arf}(\phi|_{Y_0})$, then Proposition 6.2 ensures that β and $h(\beta)$ are in the same $\mathrm{FMod}(S, \phi)$ orbit, completing the proof.

Otherwise, $\mathrm{Arf}(\phi|_{hY_0}) = \mathrm{Arf}(\phi|_{Y_0}) + 1$. Our goal is now to amend the Arf invariant of hY_0 while introducing a uniformly bounded number of intersections with α. We note first that the topology of the situation at hand forces there to be a curve w of ∂W that is nonseparating on Y_0 (hence on hY_0) and such that $\phi(w)$ is even. Indeed, if the winding numbers of all such w were odd, then since $Y_0 \setminus W$ has genus 0 this would imply that ϕ is odd on curves spanning a Lagrangian subspace of $H_1(Y_0; \mathbb{Z})$. In particular, this would imply that each term in the formula for the Arf invariant (1) would be 0, hence $\mathrm{Arf}(\phi|_{Y_0}) = 0$. The same would therefore also be true for $\mathrm{Arf}(\phi|_{hY_0})$, but this is a contradiction.

Now every mapping class supported entirely on $S \setminus W$ can be written as a product of Dehn twists, and every curve on $S \setminus W$ is separating. Thus, for any curve c on S and any curve $d \subset S \setminus W$,

$$\langle hc, d \rangle = \langle c, d \rangle.$$

Combining this with twist-linearity, we see that if we factorize $h = T_{d_1}^{k_1} \cdots T_{d_n}^{k_n}$ where d_i are curves on $S \setminus W$, then

$$\phi(hc) = \phi(c) + k_1 \langle c, d_1 \rangle \phi(d_1) + \ldots + k_n \langle c, d_n \rangle \phi(d_n)$$

for any curve c on S.

Returning to the situation at hand, since $\mathrm{Arf}(\phi|_{hY_0}) = \mathrm{Arf}(\phi|_{Y_0}) + 1$, the discussion above implies there must be some curve $c \subset Y_0$, part of a GSB for Y_0 and symplectically dual to a curve $w \subset \partial W$ of even winding number, and some curve $d \subset S \setminus W$ such that $\langle c, d \rangle$ and $\phi(d)$ are both odd. Moreover, $\langle hc, d \rangle$ is also odd, and since algebraic intersection number and winding number properties on a genus 0 surface depend only on how a curve separates the surface, Lemma 6.8 ensures there is a d' on $S \setminus W$ such that

(1) $\langle hc, d' \rangle$ is odd
(2) $\phi(d')$ is odd
(3) The geometric intersection number of d' with α is uniformly bounded.

Comparing with Formula (1), items (1) and (2) ensure that $T_{d'}hY_0$ has the same Arf invariant (and boundary winding numbers) as Y_0, hence Proposition 6.2 implies that β and $T_{d'}h\beta$ are in the same $\mathrm{FMod}(S, \phi)$ orbit. Since the geometric intersection of $h\beta$ and α was uniformly bounded, item (3) ensures that the geometric intersection of $T_{d'}h\beta$ and α is as well. □

We can now prove $\overline{\mathcal{K}}$ is quasi-isometric to $\mathscr{E}(\overline{\mathcal{H}_\phi})$; thus $\mathscr{E}(\overline{\mathcal{H}_\phi})$ is hierarchically hyperbolic.

Theorem 6.9. *The map* $\Theta\colon \overline{\mathcal{K}} \to 2^{\mathscr{E}(\overline{\mathcal{H}_\phi})}$ *is a quasi-isometry.*

Proof. Throughout the proof, we say a quantity is uniform if it depends only on the surface S. Set $\mathscr{E} := \mathscr{E}(\overline{\mathcal{H}_\phi})$ and let $\theta = \iota \circ \Pi \circ \Psi$, so $\Theta(\mu) = \theta(P(\mu))$.

Our proof has three steps. First we prove that $\mathrm{diam}(\Theta(\mu))$ is uniformly bounded for each vertex $\mu \in \overline{\mathcal{K}}$. Then we show that if μ and ν are joined by an edge of $\overline{\mathcal{K}}$ then $\mathrm{diam}(\Theta(\mu) \cup \Theta(\nu))$ is also uniformly bounded. Together these show that Θ is coarsely Lipschitz. Finally, we check that Θ is a coarse inverse to the 2-Lipschitz inclusion map $\mathscr{E} \to \overline{\mathcal{K}}$.

Step 1: vertices have uniform diameter. If no component of $S \setminus \mu$ is an element of $\mathfrak{S}$, then $\mu \in \mathcal{K}$ and every vertex of $P(\mu)$ is obtained by adding fewer than ξ curves to μ. Hence $\mathrm{diam}(P(\mu)) \leq 2\xi$. Since $\theta = \iota \circ \Pi \circ \Psi$ is coarsely Lipschitz, this implies $\theta(P(\mu)) = \Theta(\mu)$ is uniformly bounded. Hence, we can assume there exists a component W of $S \setminus \mu$ that is in $\mathfrak{S}$, i.e., is a witness for $\mathscr{C}_{\mathrm{adm}}$.

For our fixed $\mu \in \overline{\mathcal{K}}$, let $\alpha \in P(\mu)$. The multicurve α is the union of three distinct sets: μ, $\alpha_W = \alpha \cap W$, and $\alpha \setminus (\mu \cup \alpha_W)$; see Figure 9. Note that α_W must always be non-empty since $S \setminus \alpha$ cannot contain any witnesses of $\mathscr{C}_{\mathrm{adm}}$. Set $\alpha' := \mu \cup \alpha_W$. We divide the remainder of our proof into three cases based on the subsurface W.

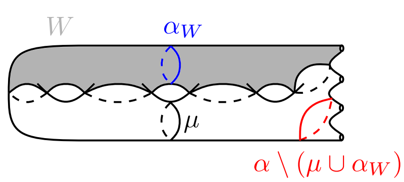

FIGURE 9. The partition of $\alpha \in P(\mu)$ into α_W, μ and $\alpha \setminus (\mu \cup \alpha_W)$.

Case 1: $g(W) \geq 1$. Since $W \in \mathfrak{S}$ but not in $\overline{\mathfrak{S}}$, there must exist some divisorial BIC $\delta \in \mathcal{D}$ disjoint from W. Since W has genus, no component of $S \setminus W$ can be in $\mathfrak{S}$, thus α' is a vertex of $\mathcal{K}$ that is joined by a path of length at most ξ to α. By Lemma 4.8, there exists a curve $c \subset W$ that cuts off a genus 1 subsurface and such that $c \in \Psi(\alpha')$. Thus, there is an admissible curve $a \in \Pi \circ \Psi(\alpha')$ contained in the subsurface W. Since a is disjoint from δ, there is an edge of $\mathscr{E}$ from a vertex of $\theta(\alpha')$ to δ. Since θ is coarsely Lipschitz, this implies $\theta(\alpha)$ is uniformly close to δ for all $\alpha \in P(\mu)$. This shows $\Theta(\mu)$ is uniformly bounded in this case.

Case 2: $g(W) = 0$, and none of the components of $S \setminus W$ are in $\mathfrak{S}$. This implies that α' is a vertex of $P(\mu)$. Moreover, α can be connected to α' with at most ξ edges of $\mathcal{K}$ (one for each curve removed to go from α to α').

Since $W \in \mathfrak{S}$ but not in $\overline{\mathfrak{S}}$, there exists a multicurve in $\mathcal{D}$ that is disjoint from W. Each component of $S \setminus W$ has genus zero by Lemma 3.2. Thus, we can apply Lemma 6.7 to find some $\delta \in \mathcal{D}$ and $N > 0$ depending only on S such that $i(\mu, \delta) \leq N$ and δ is disjoint from W; see Figure 10. Note that this choice depends only on μ, not on $\alpha \in P(\mu)$.

Since δ is a BIC, $S \setminus \delta$ has a component $Z \subset Y_0(\delta)$ with $g(Z) \geq 1$. Since Z contains an admissible curve and $W \in \mathfrak{S}$, we must have $W \subseteq Z$. By Lemma 6.8, there is a uniform $N' > 0$ and a (possibly empty) multicurve m on $Z \setminus W$ such that m cuts $Z \setminus W$ into three-holed spheres and $i(m, \mu) \leq N'$; see Figure 10. Since m cuts $Z \setminus W$ into three-holed spheres, $\alpha_W \cup m \cup \partial W \cup \delta$ is a vertex of $\mathcal{K}$ that intersects α' at most $N + N'$ times; see Figure 10 for a schematic of the situation. Let $\delta' = m \cup \partial W \cup \delta$, so $i(\alpha', \alpha_W \cup \delta') \leq N + N'$. Thus $d_{\mathcal{K}}(\alpha', \alpha_W \cup \delta')$ is bounded uniformly by some number determined by $N + N'$.

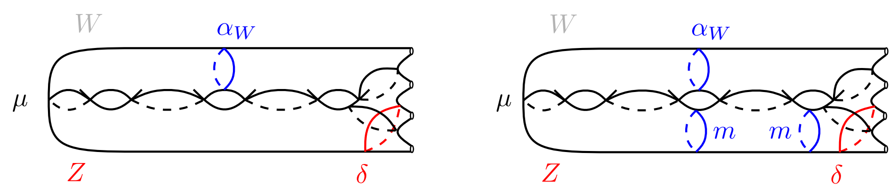

FIGURE 10. A schematic of the curves δ and m relative to μ and W. The actual intersection number between $\delta \cup m$ and μ is possibly higher, but still uniformly bounded.

As in the previous case, Lemma 4.8 says $\Psi(\alpha_W \cup \delta')$ will contain a curve $c \subset Z$ that cuts off a genus 1 subsurface of Z. Hence $\Psi(\alpha_W \cup \delta')$ will contain an admissible curve that is disjoint from δ. This means $\theta(\alpha_W \cup \delta')$ is a bounded diameter set that contains a vertex that is adjacent to δ in $\mathscr{E}$. Thus, since α is uniformly close to α' which is in turn uniformly close to $\alpha_W \cup \delta'$ and θ is coarsely Lipschitz, we conclude that $\theta(\alpha)$ is uniformly close to δ. Since δ depended only on μ, this implies $\mathrm{diam}(\Theta(\mu))$ is uniformly bounded.

Case 3: $g(W) = 0$ and $S \setminus W$ has a component V that is in $\mathfrak{S}$. By Lemma 3.2, there is only one such component V. Let β be a second vertex of $P(\mu)$ alongside α. Recall $\alpha_W = \alpha \cap W$ and similarly define α_V, β_W, and β_V. Note that $\mu \cup \alpha_W$ and $\mu \cup \beta_V$ are vertices of $\overline{\mathcal{K}}$ that we have already shown in the previous cases to have bounded diameter image under Θ. Now observe

$$\alpha \in P(\mu \cup \alpha_W),\ \alpha_W \cup \mu \cup \beta_V \in P(\mu \cup \alpha_W) \cap P(\mu \cup \beta_V),\ \text{and}\ \beta \in P(\mu \cup \beta_V).$$

Hence $\theta(\alpha)$ is uniformly close to $\theta(\alpha_W \cup \mu \cup \beta_V)$, which is in turn uniformly close to $\theta(\beta)$. Thus $\Theta(\mu)$ is uniformly bounded.

Step 2: edges have uniform diameter. If ν is obtained from μ by adding a curve, then $P(\nu) \subseteq P(\mu)$. Hence $\Theta(\mu) \cup \Theta(\nu) = \Theta(\mu)$ which has bounded diameter by Step 1.

Otherwise, the edge from μ to ν corresponds to a flip move. Let $x \in \mu$ and $x' \in \nu$ be such that x is flipped to x' and let Y be the component of $S \setminus (\mu \setminus x)$ containing x and x'. If

$\xi(Y) > 1$, then $i(x,y) = 0$ and $\mu \cup x'$ is a vertex of $\overline{\mathcal{K}}$. Now $\mu \cup x'$ is joined by a "remove" edge to both μ and ν (removing x' gives μ and removing x gives ν), so the diameter bound follows from the add/remove edge case. If $\xi(Y) = 1$, then x and x' intersect minimally on Y. We can therefore find two pants decompositions $\alpha \in P(\mu)$ and $\alpha' \in P(\nu)$ such that α differs from α' by flipping x to x'. Since α and α' are joined by an edge in $\mathcal{K}$, the sets $\theta(\alpha)$ and $\theta(\alpha')$ are uniformly close in $\mathscr{E}$. Since $\alpha \in P(\mu)$ and $\alpha' \in P(\nu)$, this implies that $\Theta(\mu) \cup \Theta(\nu)$ has uniformly bounded diameter.

Step 3: Θ is a coarse inverse of the inclusion. Let $j\colon \mathscr{E} \to \overline{\mathcal{K}}$ be the 2-Lipschitz inclusion map. Since $\mathscr{E}$ contains the admissible curve graph and $\overline{\mathcal{K}}$ contains the pants graph and the genus-separting curve graph, the argument from Lemma 4.7 shows that every vertex of $\overline{\mathcal{K}}$ is uniformly close to a vertex in $\mathscr{E}$, so j is coarsely surjective.

Let $\mu \in j(\mathscr{E})$, that is, μ is either an admissible curve or $\mu \in \mathcal{D}$. If μ is admissible, then μ is a vertex of $P(\mu)$ and $\Pi \circ \Psi(\mu)$ contains the admissible curve μ (as in the proof of Proposition 4.13), so $\mu \in \Theta(\mu)$. If $\mu \in \mathcal{D}$, then there is some admissible curve a contained in $Y_0(\mu)$ that is in particular disjoint from μ. Since $a \cup \mu$ and μ are joined by an edge of $\overline{\mathcal{K}}$, this means $\Theta(a \cup \mu)$ and $\Theta(\mu)$ are uniformly close in $\mathscr{E}$ by Step 2. Now $a \in \Theta(a \cup \mu)$ because it is an admissible curve. Thus $\Theta(\mu)$ is uniformly close to a, which is joined to μ by an edge in $\mathscr{E}$. In summary, $\Theta \circ j$ moves points a finite distance in $\mathscr{E}$, hence it is a coarse inverse to j. □

Proof of Theorem B. By Theorem 6.9, $\mathscr{E}(\overline{\mathcal{H}_\phi})$ is quasi-isometric to the hierarchically hyperbolic model $\overline{\mathcal{K}}$, and by Corollaries 6.5 and 6.6 it is quasi-isometric to $\mathscr{D}(\mathbb{P}\Xi\mathcal{H}_\phi)$ and $\mathscr{C}(\overline{\mathcal{H}_\phi})$. Thus all of these graphs are hierarchically hyperbolic.

The proof that $\mathscr{E}(\overline{\mathcal{H}_\phi})$ has a pair of disjoint witnesses is a slight variation of the one appearing in the proof of Theorem A. Since the restriction of ϕ to the top level of any ϕ-compatible BIC must be of holomorphic type, the winding numbers of each curve of β must be negative, and by homological coherence are all bounded by $\chi(S)$.

Now consider a multicurve α separating S into a pair of connected genus 0 subsurfaces $W^\pm$ such that W^+ contains no punctures and lies to the left of each curve of α. In order for W^+ to contain either an admissible curve or a curve of a BIC then there must be some subset $\mathcal{C}$ of the curves of α such that

$$1 - |\mathcal{C}| \le \sum_{c \in \mathcal{C}} \phi(c) \le 1 - |\mathcal{C}| - \chi(S).$$

(Lemma 3.2 gives the statement for admissibles, and the same logic gives the statement for curves with a given winding number.) Thus, by choosing an α with large enough winding numbers such that no subset sums of the winding numbers of its curves are in this small range, we see that there can be no admissible curves *or* BICs contained in W^+ and hence W^- is a witness. A similar argument implies that for sufficiently large choices of the winding numbers of α, the subsurface W^+ will also be a witness, and so $\mathscr{E}(\overline{\mathcal{H}_\phi})$ (hence $\mathscr{D}(\mathbb{P}\Xi\mathcal{H}_\phi)$, $\mathscr{C}(\mathbb{P}\Xi\mathcal{H}_\phi)$, and $\mathscr{C}(\overline{\mathcal{H}_\phi})$) cannot be Gromov hyperbolic. □

References

[ABW23] Paul Apisa, Matt Bainbridge, and Jane Wang, *Moduli spaces of complex affine and dilation surfaces*, J. Reine Angew. Math. **796** (2023), 229–243.

[AHW24] F. Arana-Herrera and A. Wright, *The geometry of totally geodesic subvarieties of moduli spaces of riemann surfaces*, Preprint, arXiv:2412.16330, 2024.

[BCG+18] M. Bainbridge, D. Chen, Q. Gendron, S. Grushevsky, and M. Möller, *Compactification of strata of Abelian differentials*, Duke Math. J. **167** (2018), no. 12, 2347–2416.

[BCG+19] ———, *The moduli space of multi-scale differentials*, Preprint, arXiv:1910.13492, 2019.

[BCM12] J. F. Brock, R. D. Canary, and Y. N. Minsky, *The classification of Kleinian surface groups, II: The ending lamination conjecture*, Ann. of Math. (2) **176** (2012), no. 1, 1–149. MR 2925381

[BF06] J. Brock and B. Farb, *Curvature and rank of Teichmüller space*, Amer. J. Math. **128** (2006), no. 1, 1–22. MR 2197066

[BHS17a] J. Behrstock, M. F. Hagen, and A. Sisto, *Asymptotic dimension and small-cancellation for hierarchically hyperbolic spaces and groups*, Proc. Lond. Math. Soc. (3) **114** (2017), no. 5, 890–926. MR 3653249

[BHS17b] ———, *Hierarchically hyperbolic spaces I: Curve complexes for cubical groups*, Geom. Topol. **21** (2017), no. 3, 1731–1804. MR 3650081

[BHS19] ———, *Hierarchically hyperbolic spaces II: Combination theorems and the distance formula*, Pacific J. Math. **299** (2019), 257–338.

[BM04] T. E. Brendle and D. Margalit, *Commensurations of the Johnson kernel*, Geom. Topol. **8** (2004), 1361–1384. MR 2119299

[Boi15] Corentin Boissy, *Connected components of the strata of the moduli space of meromorphic differentials*, Comment. Math. Helv. **90** (2015), no. 2, 255–286.

[Bro03] J. F. Brock, *The Weil-Petersson metric and volumes of 3-dimensional hyperbolic convex cores*, J. Amer. Math. Soc. **16** (2003), no. 3, 495–535. MR 1969203

[BSW22] M. Bainbridge, J. Smillie, and B. Weiss, *Horocycle dynamics: new invariants and eigenform loci in the stratum H(1,1)*, Mem. Amer. Math. Soc. **280** (2022), no. 1384, v+100.

[Cal20] A. Calderon, *Connected components of strata of Abelian differentials over Teichmüller space*, Comment. Math. Helv. **95** (2020), no. 2, 361–420. MR 4115287

[CF22] D. Chen and G. Faraco, *Period realization of meromorphic differentials with prescribed invariants*, Preprint, arXiv:2212.05754, 2022.

[CGH+25] D. Chen, S. Grushevsky, D. Holmes, M. Möller, and J. Schmitt, *A tale of two moduli spaces: logarithmic and multi-scale differentials*, Épijournal Géom. Algébrique **9** (2025), Art. 21, 51.

[CMSZ20] D. Chen, M. Möller, A. Sauvaget, and D. Zagier, *Masur-Veech volumes and intersection theory on moduli spaces of Abelian differentials*, Invent. Math. **222** (2020), no. 1, 283–373.

[CMZ22] M. Costantini, M. Möller, and J. Zachhuber, *The Chern classes and Euler characteristic of the moduli spaces of Abelian differentials*, Forum Math. Pi **10** (2022), Paper No. e16, 55.

[CMZ23] ———, `diffstrata`—*a Sage package for calculations in the tautological ring of the moduli space of abelian differentials*, Exp. Math. **32** (2023), no. 3, 545–565.

[CS22] A. Calderon and N. Salter, *Framed mapping class groups and the monodromy of strata of abelian differentials*, J. Eur. Math. Soc. (2022), (to appear).

[Dev26] P. Devkota, *Cohomology of Moduli Space of Multiscale Differentials in Genus 0*, Int. Math. Res. Not. IMRN (2026), no. 3, Paper No. rnag006.

[EMZ03] A. Eskin, H. Masur, and A. Zorich, *Moduli spaces of abelian differentials: the principal boundary, counting problems, and the Siegel-Veech constants*, Publ. Math. Inst. Hautes Études Sci. (2003), no. 97, 61–179.

[FI05] B. Farb and N. V. Ivanov, *The Torelli geometry and its applications: research announcement*, Math. Res. Lett. **12** (2005), no. 2-3, 293–301. MR 2150885

[FM12] B. Farb and D. Margalit, *A primer on mapping class groups*, Princeton Mathematical Series, vol. 49, Princeton University Press, Princeton, NJ, 2012.

[Har81] W. J. Harvey, *Boundary structure of the modular group*, pp. 245–252, Princeton University Press, 1981.

[Hen20] S. Hensel, *A primer on handlebody groups*, Handbook of group actions. V, Adv. Lect. Math. (ALM), vol. 48, Int. Press, Somerville, MA, [2020] ©2020, pp. 143–177. MR 4237892

[HJ89] S. Humphries and D. Johnson, *A generalization of winding number functions on surfaces*, Proc. London Math. Soc. **58** (1989), no. 2, 366–386.

[HK14] J. H. Hubbard and S. Koch, *An analytic construction of the Deligne-Mumford compactification of the moduli space of curves*, J. Differential Geom. **98** (2014), no. 2, 261–313.

[Iva97] N. V. Ivanov, *Automorphisms of complexes of curves and of Teichmüller spaces*, Progress in knot theory and related topics, Travaux en Cours, vol. 56, Hermann, Paris, 1997, pp. 113–120. MR 1603146

[Kaw18] N. Kawazumi, *The mapping class group orbits in the framings of compact surfaces*, Q. J. Math. **69** (2018), no. 4, 1287–1302.

[KZ] M. Kontsevich and A. Zorich, *Lyapunov exponents and Hodge theory*, Preprint, 1–16.

[KZ03] ———, *Connected components of the moduli spaces of Abelian differentials with prescribed singularities*, Invent. Math. **153** (2003), no. 3, 631–678.

[LM14] E. Looijenga and G. Mondello, *The fine structure of the moduli space of abelian differentials in genus 3*, Geom. Dedicata **169** (2014), no. 1, 109–128.

[Mas86] H. Masur, *Closed trajectories for quadratic differentials with an application to billiards*, Duke Math. J. **53** (1986), no. 2, 307–314.

[Min10] Y. N. Minsky, *The classification of Kleinian surface groups. I. Models and bounds*, Ann. of Math. (2) **171** (2010), no. 1, 1–107. MR 2630036

[MM98] H. Masur and Y. N. Minsky, *Geometry of the complex of curves I: Hyperbolicity*, Invent. Math. **138** (1998), 103–149.

[MM00] ———, *Geometry of the complex of curves. II. Hierarchical structure*, Geom. Funct. Anal. **10** (2000), no. 4, 902–974. MR 1791145

[MS13] H. Masur and S. Schleimer, *The geometry of the disk complex*, J. Amer. Math. Soc. **26** (2013), no. 1, 1–62. MR 2983005

[MUW21] M. Möller, M. Ulirsch, and A. Werner, *Realizability of tropical canonical divisors*, J. Eur. Math. Soc. **23** (2021), no. 1, 185–217.

[Put08] A. Putman, *A note on the connectivity of certain complexes associated to surfaces*, Enseign. Math. (2) **54** (2008), no. 3-4, 287–301. MR 2478089

[Raf05] K. Rafi, *A characterization of short curves of a Teichmüller geodesic*, Geom. Topol. **9** (2005), 179–202. MR 2115672

[RS09] Kasra Rafi and Saul Schleimer, *Covers and the curve complex*, Geom. Topol. **13** (2009), no. 4, 2141–2162.

[RV19] Jacob Russell and Kate M. Vokes, *The (non)-relative hyperbolicity of the separating curve graph*, arXiv:1910.01051 (2019).

[RW14] O. Randal-Williams, *Homology of the moduli spaces and mapping class groups of framed, r-Spin and Pin surfaces*, J. Topol. **7** (2014), no. 1, 155–186.

[Sal] N. Salter, *Higher spin mapping class groups in algebraic and flat geometry*, Notes available at `https://nsalter.science.nd.edu/expository-notes/cuernavacalectures.pdf`.

[Sal25] N. Salter, *Fundamental groups of strata of abelian differentials of low codimension*, Preprint, arXiv:2509.25357, 2025.

[Tan21] Robert Tang, *Affine diffeomorphism groups are undistorted*, J. Lond. Math. Soc. (2) **104** (2021), no. 2, 747–769.

[Vok22] K. Vokes, *Hierarchical hyperbolicity of graphs of multicurves*, Algebr. Geom. Topol. **22** (2022), no. 1, 113–151.

[Won24] Y. M. Wong, *An algorithm to compute the fundamental classes of spin components of strata of differentials*, Int. Math. Res. Not. IMRN (2024), no. 6, 4893–4962.

[Zor06] A. Zorich, *Flat surfaces*, Frontiers in number theory, physics, and geometry I, Springer, 2006, pp. 439–585.