HIERARCHICAL HYPERBOLICITY OF ADMISSIBLE CURVE GRAPHS AND THE BOUNDARY 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 $\Omega\mathcal{M}_g$ of genus g Abelian differentials forms a bundle over the usual moduli space $\mathcal{M}_g$ of genus g Riemann surfaces. This bundle 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 are closely related to *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]. 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 2.9), 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 $\mathcal{H} \subset \Omega^1\mathcal{M}_g$ is an (orbifold) $\mathbb{C}^*$-bundle over 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 $\mathcal{H} \subset \mathcal{M}_{g,n}$, let $\mathcal{H}_\phi$ be any component of the preimage of $\mathcal{H}$ in $\mathcal{T}_{g,n}$, and consider its closure $\overline{\mathcal{H}_\phi}$ in the augmented Teichmüller space $\overline{\mathcal{T}_{g,n}}$. 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, and whose edges are given by inclusion. The intricate structure of the boundary of $\overline{\mathcal{H}_\phi}$ means there are other natural ways to define this graph (see Sections 5.3 and 5.4), but they 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 $\mathscr{C}_{\mathrm{adm}}(S,\phi)$, 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 $\mathcal{H} \subset \Omega^1\mathcal{M}_g$ 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 $\mathscr{C}_{\mathrm{adm}}(S,\phi)$. 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 $\mathrm{FMod}(S,\phi)$ 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 $g \geq 3$.

Curve graph techniques have been used successfully to study certain $\mathsf{GL}_2\mathbb{R}$–invariant subvarieties of $\Omega^1\mathcal{M}_g$: [Tan21] proved that Veech groups are undistorted in $\mathrm{Mod}(S)$, [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 1.6. *Is $\mathrm{FMod}(S,\phi)$ distorted in $\mathrm{Mod}(S)$? 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 cobounded 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 boundary, as developed in [BCG$^+$19], and giving topological interpretations to many of the objects involved. These steps are accomplished in Section 5. In this section, we also build a trio of graphs whose definitions interpolate between the structure of $\partial\overline{\mathcal{H}_\phi}$ and framing data.

In the final Section 6, we show that the three 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 2.9, 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.

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, (relatively) closed subsurface of S. 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. We say 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 isotopy class (relative to ∂W) 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 isotopic by isotopies fixing ∂W setwise but not pointwise.

If α is a multicurve on S, then $S \setminus \alpha$ will denote the closed subsurface obtained by removing a small open neighborhood of each curve in α from S. Similarly, if W is a subsurface of S, then $S \setminus W$ is the closed subsurface obtained by removing a small open neighborhood 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. This is 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 $c_1, \dots, c_k$ denote its boundary components and the peripheral loops about its punctures, 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 $\Delta_1, \dots, \Delta_k$ denote small loops about the punctures of S (oriented with the surface on their left); then the *signature* of a framing ϕ is the tuple

$$\mathrm{sig}(\phi) := (\phi(\Delta_1), \dots, \phi(\Delta_k)) \in \mathbb{Z}^k.$$

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 a $k-1$ dimensional subspace of $H_1(S)$, so we can construct all framings with a given signature by specifying the values on $2g$ homologically independent curves [CS22, Remark 2.7]. One particularly nice configuration is as follows:

Definition 2.4. A collection of simple closed curves $\mathcal{B} = \{a_1, b_1, \dots, a_g, b_g\}$ 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.

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 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

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

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

- $(g = 0)$ *$sig(\phi) = sig(\phi')$*
- $(g = 1)$ *$sig(\phi) = sig(\phi')$ 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.

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 2.9 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.

[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.

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 nonseparating curves of arbitrary winding number.

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 2.9. *Let ϕ be a framing of a surface S of genus at least 3. Then $\mathrm{FMod}(S, \phi)$ acts transitively on the set of pairs of admissible curves of the same 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 2.9, 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 2.10. *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 the listed curves, such that either*

- *$g(T) \geq 2$, or*
- *$g(T) = 1$ and $\mathrm{Arf}_1(\phi|_T) = 1$.*

Suppose also that there is some arc ε connecting d to T that is disjoint from all c_i. 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. Orient d such that the arc from d to T exits d from its left-hand side.

Suppose first that $g(T) = 2$. Then by Lemma 2.8 there is a nonseparating curve e on T with winding number $-z - \phi(d) - 1$. Since d is not separated from T, we may concatenate ε with an arc connecting ∂T to the left side of e and take the connect sum of d and e along this composite arc. Let d_z be the resulting curve; then by homological coherence (Lemma 2.2.2) we have that

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

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 (see also the proof of [Sal, Corollary 4.4]).

In the case that $g(T) = 1$, our assumption on $\mathrm{Arf}_1(\phi|_T)$ implies (via Lemma 2.8) that there is some GSB (a, b) on T with $\phi(a) = 1$. Choose an arc from ∂T to b disjoint from a, then take the connected sum of d with b along the concatenation of ε with this arc. This results in a new curve d' that has the same intersection pattern as d with each c_i and meets a exactly once. Twist-linearity (Lemma 2.2.1) now implies that by twisting around a we can alter the winding number of d' by an arbitrary amount to find our desired d_z. □

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 2.11. *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 2.10 to adjust $\phi(b')$ at will. □

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

Proof of Proposition 2.9. 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 2.11, 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 2.11 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 2.11 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.

Pick an arbitrary curve b_1 meeting a_1 and a_2 each exactly once. Since at least one of T or U has genus at least 2 or genus 1 with $\mathrm{Arf}_1 = 1$, we can use Lemma 2.10 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 disjoint from $\mathcal{B}_{T'} \cup \mathcal{B}_{U'}$ that meets a_1' and a_2' exactly once.

Suppose ϕ is of spin type. Then we see that for any choice of b_1' meeting a_1' exactly once and disjoint from $\mathcal{B}_T \cup \mathcal{B}_U$, we have

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

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

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.

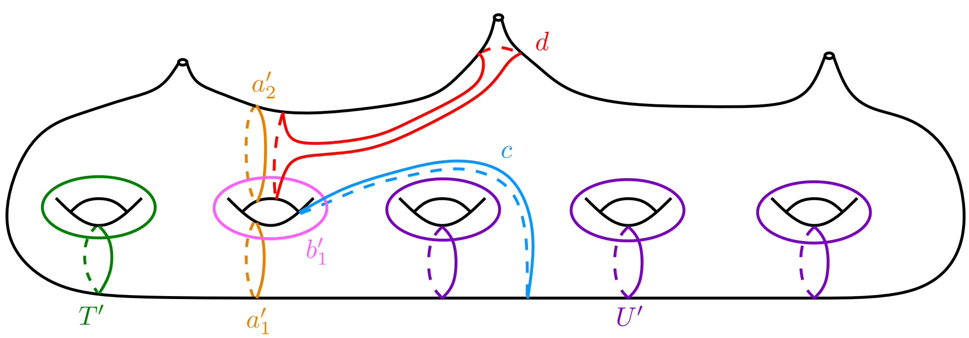

FIGURE 1. GSBs and auxiliary curves as in the proof of Proposition 2.9.

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. $\square$

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 2.9 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 for any framing on S, $\mathscr{C}_{\mathrm{adm}}(S,\phi)$ is connected.*

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 subsurface $W\subseteq S$ is a *witness* for $\mathcal{X}$ if every vertex of $\mathcal{X}$ intersects W and $\xi(W)<0$. 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 where the winding numbers of the boundary curves 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) If $Z\subseteq S$ is a genus 0 subsurface and $z_1,\dots,z_k$ are the boundary components of Z and peripheral loops about its punctures, oriented such that Z is to the left of each z_i, then Z contains a nonperipheral curve of winding number 0 if and only if there is no $I\subsetneq\{z_1,\dots,z_k\}$ 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 boundary components and peripheral loops of Z as in the previous item. Then there is no I such that (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 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; 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* 2*-Lipschitz*

Proof. If a, b are a pair of disjoint admissible curves, then $a \cup b$ is also a vertex of $\mathcal{K}_{\mathfrak{S}}(S)$, hence $a, a \cup b, b$ is a path of length 2 connecting a and b in $\mathcal{K}_{\mathfrak{S}}(S)$. □

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 cobounded 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 quasi-isometric model is that she showed that $\mathcal{K}_{\mathfrak{S}}(S)$ is a *hierarchically hyperbolic space* in a natural way. This means 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 cobounded 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 4.11).

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 $g \geq 3$ and $n \geq 1$ 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

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

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 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 $\mathrm{diam}(f(x) \cup f(x'))$ is uniformly bounded.

Two spaces are *quasi-isometric* if there exist two coarsely Lipschitz map $f \colon X \to 2^Y$ and $\overline{f} \colon Y \to 2^Y$ such that $d_X(x, \overline{f} \circ f(x))$ is uniformly bounded for all $x \in X$. In this case, f is a *quasi-isometry* from X to Y and $\overline{f}$ is the *quasi-inverse* of f.

4.2. **The genus-separating curve graph.** We begin building our quasi-isometry from $\mathcal{K}$ to $\mathscr{C}_{\mathrm{adm}}$ 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 are 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 $\mathscr{C}_{\mathrm{adm}}$. 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 $\mathcal{K}_{\mathrm{Wit}(\mathcal{G})}(S)$ 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 $\mathcal{G}$. $\square$

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 $i(a, c)$.

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 $d_{\mathscr{C}_{\mathrm{adm}}}(a, \Pi(c)) \leq M$.*

We will only ever apply this lemma with $m = 2$, but since the proof for general m is not much harder we choose to include it here.

Proof of Lemma 4.6. 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 homotopic to a boundary curve of S (they cannot both be homotopic to a boundary curve 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 both 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. □

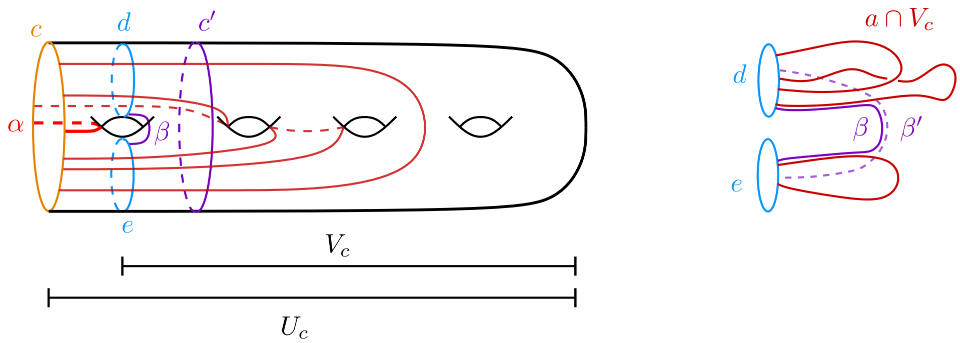

FIGURE 2. On the left, the subsurfaces involved in the proof of Lemma 4.6. On the right, surgering an arbitrary arc β' from d to e along $a \cap V_c$ to obtain a disjoint arc β.

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 $\mathscr{C}_{\text{adm}} \to \mathcal{K}$. 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 4.7. *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 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. $\square$

In light of Lemma 4.7, 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 4.7.

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 point 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 $\mathscr{C}_{\mathrm{adm}} \to \mathcal{K}$.

The key technical step is the next lemma, which takes a component Y of $S \setminus \alpha$ and a genus-separating curve $c \in \Psi(\alpha)$ and produces an admissible curve a that intersects c at most 4 times and is disjoint from Y. This admissible curve provides an "anchor" that allows us to modify c inside the component Y 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 curve in $\Psi(\alpha)$ ensured by Lemma 4.7 as opposed to defining $\Psi(\alpha)$ to be all genus-separating curves that intersect each curve of α some fixed number of times.

Lemma 4.8. *Let α be a multicurve in $\mathcal{K}$ and $c \in \Psi(\alpha)$. For each component Y of $S \setminus \alpha$ that c intersects, there exists an admissible curve a_Y that is disjoint from Y and has $i(c, a_Y) \leq 4$.*

Proof. Let Y be a component of $S \setminus \alpha$ that c intersects. If any curve of α is admissible, then c intersects that curve at most twice and we are done. This also allows us to proceed by assuming that $S \setminus \alpha$ is disconnected: because each component of $S \setminus \alpha$ is not a witness, if $S \setminus \alpha$ is connected then α must contain an admissible curve.

Since Y is not a witness for $\mathscr{C}_{\mathrm{adm}}$ by the definition of $\mathcal{K}$, some component Z of $S \setminus Y$ contains an admissible curve. If c is disjoint from Z, then c is disjoint from the admissible curve on Z and again we are done. So suppose that c intersects Z; then $c \cap Z$ separates Z since c is separating. Since c is genus-separating, if Z has positive genus then at least one of the components of $Z - (c \cap Z)$ must also have genus. Applying Lemma 2.7, this implies there is an admissible curve in Z that is disjoint from c whenever Z contains genus.

We can therefore concentrate on the case where Z has no genus. In this case, every curve on Z is separating, and which curves of Z are admissible are determined by how they

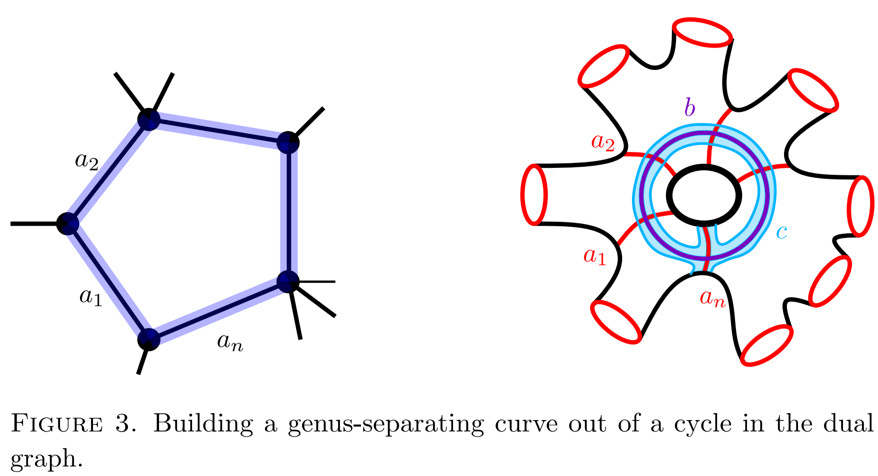

FIGURE 3. Building a genus-separating curve out of a cycle in the dual graph.

separate the boundary components and peripheral curves of Z (Lemma 2.2.2). Let A be a set of boundary and peripheral curves of Z such that any curve partitioning the boundaries and peripheral curves into A and its complement must be admissible. We argue below that one can always draw a curve a that cuts off the curves of A and intersects c at most 4 times.

To facilitate this, we first show that $c \cap Z$ cuts Z into at most 3 components. Since c intersects at most 2 components of ∂Y, it also intersects at most 2 components of ∂Z (and intersects each component at most twice) and must be disjoint from all peripheral curves. If c intersects exactly one component of ∂Z, then we are in case 3 of Lemma 4.7 and so $c \cap Z$ must be a single arc with both endpoints on the same boundary component of Z; in this case $Z - (c \cap Z)$ has two components. When c intersects two distinct components z_1, z_2 of ∂Z, then we are in case 4 of Lemma 4.7 and so $c \cap Z$ is a pair of arcs c_1, c_2 such that either

- both endpoints of c_i are on z_i for each $i \in \{1, 2\}$, or
- c_1, c_2 are parallel arcs each running from z_1 to z_2.

In the first case, $Z - (c \cap Z)$ has either two or three components and in the second it has two.

To find an admissible curve on Z that intersects c at most 4 times, let Z_1, Z_2, Z_3 be the components of $Z - (c \cap Z)$, with Z_3 being omitted in the case of two components. Without loss of generality, assume ∂Z_2 contains an arc of $c \cap Z$ in common with both ∂Z_1 and ∂Z_3 when there are three components. Partition the curves in A into 5 (possibly empty) sets: A_1, A_2, A_3 and B_1, B_2. The A_i are the subsets of curves in A that are contained in Z_i for each i, while B_1 are the curve(s) that contains the endpoints of the arc in $c \cap Z$ shared by ∂Z_1 and ∂Z_2 and B_2 is the same for ∂Z_2 and ∂Z_3 (when Z_3 exists). Note that the A_i may contain curves peripheral to the punctures, but the B_i must always consist of essential curves on S.

Order the curves in each A_i and B_i in any sequence, then join successive curves by disjoint arcs in the following order, skipping any empty sets: A_1, B_1, A_2, B_2, A_3. We further stipulate that the arcs must be disjoint from $c \cap Z$ unless some set is empty, in which case their intersection with $c \cap Z$ is allowed to be the difference of the indices of the Z_i that the two sets border. For example, if only A_2 is empty then the arc from B_1 to B_2 must still be disjoint from c, since both B_1 and B_2 border Z_2, but if B_1, A_2, and B_2 are empty then the arc from A_1 to A_3 is allowed to meet $c \cap Z$ twice. Compare Figure 4.

A regular neighborhood of A together with these arcs produces a curve a that cuts off all of the curves in A, and hence must be admissible. It remains to note that the arcs and curves in the construction of a are all disjoint from $c \cap Z$ except for the B_i's and arcs that travel between different Z_i's (which exist only when one of the B_i's is empty). In particular, this means that a intersects c only in a neighborhood of the B_i or the above-mentioned arcs, and only does so at most twice for each component of the construction. □

We now prove that $\Pi \circ \Psi(\alpha)$ has uniformly bounded diameter for each $\alpha \in \mathcal{K}$. The proof will use Lemma 4.8 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 4.9. *There is an $N \geq 0$ such that for any $\alpha \in \mathcal{K}$ and $c, d \in \Psi(\alpha)$, there is $c' \in \Psi(\alpha)$ with*

(1) $i(c', d) \leq 2|\chi(S)|$ and

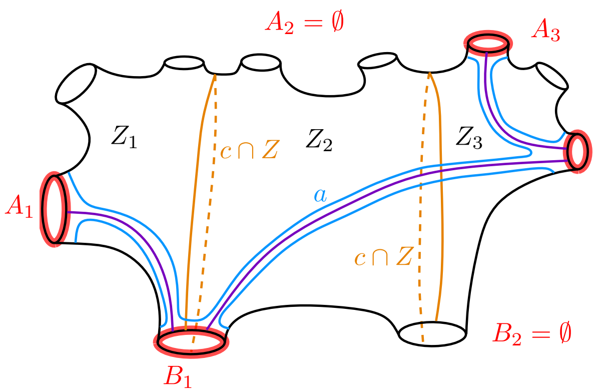

FIGURE 4. Building a curve that cuts off A, and is hence admissible. The highlighted curves are in A. In this example, A_2 and B_2 are empty, so the arc from B_1 to A_3 meets $c \cap Z$ exactly once.

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

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 4.10. *If* Y *is a component 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 twice on* Y *and* c_Y *agrees with* c *on* $S \setminus Y$.

Proof. We will show that c_Y can be obtained by replacing $c \cap Y$ with some well chosen arcs that intersect $d \cap Y$ at most twice. 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 separate Y).

We first handle the case where $c \cap Y$ is a pair of parallel arcs. Let $c_1^1, c_1^2, c_2^1, c_2^2$ be the four endpoints of $c \cap Y$ in Y such that c_i^1 is joined by an arc of $c \cap Y$ to c_i^2. If $d \cap Y$ is a single arc, then c_i^1 and c_i^2 are either on the same or different sides of $d \cap Y$. In either case, we can connect each c_i^1 to its corresponding c_i^2 with an arc γ_i such that γ_1 and γ_2 are parallel arcs and $i(\gamma_i, d) \leq 1$. If $d \cap Y$ is instead a pair of parallel arcs, let δ_1, δ_2 be the arcs of $d \cap Y$. Now $Y \setminus \delta_1$ is connected, but $(Y \setminus \delta_1) \setminus \delta_2$ has two components. Thus c_i^1 and c_i^2 are either on the same or different sides of of δ_2 in $Y \setminus \delta_1$. As before, this means we can connect each pair c_i^1 and c_i^2 with an arc γ_i such that γ_1 and γ_2 are parallel, $i(\gamma_i, \delta_2) \leq 1$, and $i(\delta_1, \gamma_i) = 0$. In either case, let c_Y be the curve obtained from c be replacing $c \cap Y$ with $\gamma_1 \cup \gamma_2$. Since $c \cap Y$ and $c_Y \cap Y$ are both parallel arcs between the same boundary components of Y, we see that

$S \setminus c$ is homeomorphic to $S \setminus c_Y$, and in particular c_Y is genus-separating. By construction, it is also clear that $c_Y \in \Psi(\alpha)$, so we are done.

Now consider the case where $c \cap Y$ is a single arc. Since $c \cap Y$ separates Y, we orient c and then label each boundary component and peripheral curve of Y by "left" or "right" depending on which side of $c \cap Y$ it lies on. Let g_l and g_r be the genus of the left and right sides of $Y \setminus (c \cap Y)$ respectively. We will find c_Y by replacing $c \cap Y$ with an arc γ that separates Y into two components, one with genus g_l and all the left curves of Y and the other with genus g_r and all the right curves of Y (any such arc is essential on Y since $c \cap Y$ is an essential arc and γ will separate Y in the same way as c). This ensures $S \setminus c$ is homeomorphic to $S \setminus c_Y$, which makes c_Y a genus-separating curve which is in $\Psi(\alpha)$ by construction. Let c_1, c_2 be the end points of $c \cap Y$ in ∂Y.

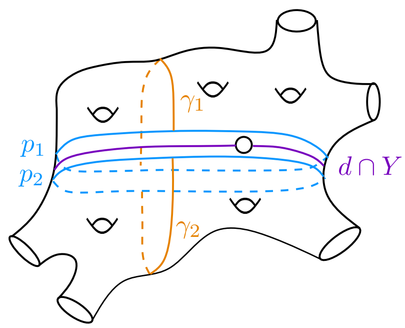

FIGURE 5. The curves p_1, p_2 cobounding the pair of pants P. The arcs γ_1 and γ_2 cut $S \setminus P$ into "left" and "right" sides.

If $d \cap Y$ is a single arc, let y be the curve of ∂Y that d intersects. The boundary of a neighborhood of $(d \cap Y) \cup y$ is a pair of curves p_1, p_2 that cobound a pair of pants P with the boundary curve y. The complement $Y \setminus P$ has two components Z_1, Z_2 where Z_i contains p_i as a boundary curve; see Figure 5.

Suppose that c also intersects the boundary curve y. On each Z_i, we can draw an arc γ_i with both endpoints on p_i such that γ_i separates Z_i into two components, one that contains the left boundary components of Y that also live on Z_i and the other that contains the right boundary components. Moreover, we can choose the γ_i such that the sum of the genera on the "left" sides of $Z_i \setminus \gamma_i$ is g_l and the sum of the genera on the "right" sides is g_r. The γ_i also separate p_i into "left" and "right" arcs.

We can now complete $\gamma_1 \cup \gamma_2$ to an arc on all of Y by adding arcs in the pair of pants P. Select three disjoint arcs a, b_1, b_2 such that a joins one endpoint of γ_1 to one endpoint of γ_2 and each b_i joins the other endpoint of γ_i to c_i by an arc in P. These arcs can be chosen such that a intersects $d \cap Y$ once, b_1 is disjoint from $d \cap Y$, and b_2 intersects $d \cap Y$ at most once. Moreover, we can choose these arcs such that the left arcs of p_i are in one component of $P \setminus (a \cup b_1 \cup b_2)$ and the right arcs are in the other; see Figure 6. The desired arc γ is the concatenation of γ_1, γ_2 and these arcs in P.

The case when c does not intersect the boundary curve y is similar. In this case c intersects a different boundary curve $y' \in \partial Y$ and without loss of generality, $y' \subset Z_2$. We draw γ_1 as

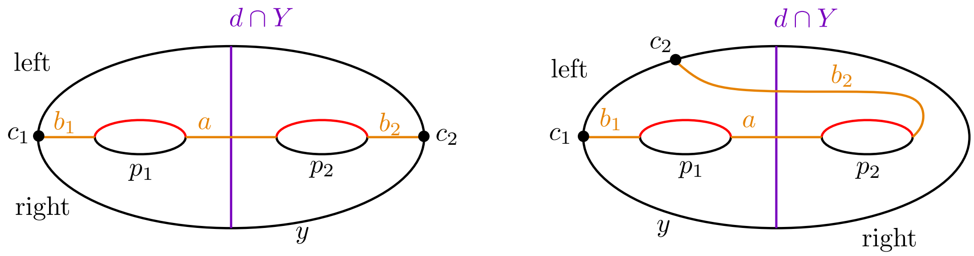

FIGURE 6. The arcs a, b_1, b_2 one must add in the pair of pants P to complete $\gamma_1 \cup \gamma_2$ to γ.

we did in the previous case, but instead of γ_2, we draw two arcs γ_2^1, γ_2^2 where γ_2^1 connects c_1 to p_2 and γ_2^2 connects c_2 to p_2 such that $\gamma_2^1 \cup \gamma_2^2$ cuts Z_2 into two pieces with the appropriate boundary components and number of genus on the "left" and "right" sides. We now finish γ by joining each end point of γ_2^i on p_2 to one of the endpoint of γ_1 on p_1 by arcs in P that intersect $d \cap Y$ exactly once and separate the left and right arc of p_1, p_2 to the correct sides.

Now suppose $d \cap Y$ is a pair of parallel arcs between two boundary component $y_1, y_2 \in \partial Y$. There is a unique curve $p \subset Y$ that forms a pair of pants P with y_1 and y_2 such that P contains $d \cap Y$; this curve p is found by taking the boundary of a neighborhood of $(d \cap Y) \cup y_1 \cup y_2$. Note that $Y \setminus P$ is a connected subsurface with the same genus as Y but one fewer boundary.

Assume first that both y_1 and y_2 are on the same side of $c \cap Y$; this implies c is disjoint from y_1 and y_2. Since $g(Y) = g(Y \setminus P)$ and y_1, y_2 are on the same side of $c \cap Y$, we can draw an arc γ on $Y \setminus P$ with connects c_1 to c_2 and cuts Y into two components, one with g_l genus and all the "left" components of ∂Y and one with g_r genus and all the "right" components.

Now assume that both y_1 and y_2 are on different sides of $c \cap Y$ (again this implies c is disjoint from y_1 and y_2). Without loss of generality let y_1 be on the left side of c and y_2 on the right. In this case we draw two arcs γ_1, γ_2 on $Y \setminus P$ such that γ_1 connects c_1 to p, γ_2 connects c_2 to p, and $\gamma_1 \cup \gamma_2$ separates $Y \setminus P$ into "left" and "right" components where the left component has g_l genus and all the left curves of Y except y_1 and the right component has g_r genus and all the right curves except y_2. We complete $\gamma_1 \cup \gamma_2$ to the arc γ on Y by joining γ_1 to γ_2 by an arc in P that separates y_1 and y_2 to the correct side of $Y \setminus \gamma$; this can be done such that the final arc has $i(\gamma, d \cap Y) \leq 2$; see Figure 7.

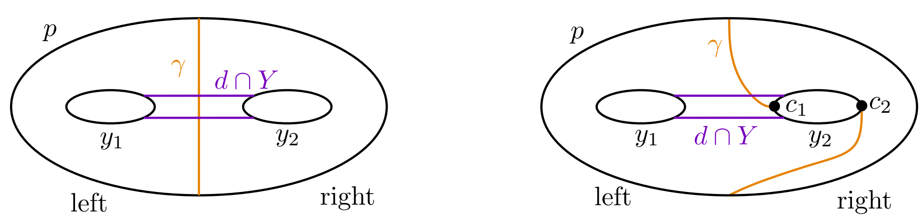

FIGURE 7. The arc drawn in P to complete the arc γ. One the left, the case where y_1 and y_2 are on different sides of $c \cap Y$. On the right, the case where c intersects y_2.

Finally, assume that c intersects exactly one of y_1 or y_2. Without loss of generality, assume c intersects y_2 and y_1 is on the left side of c. As in the previous cases, pick an arc γ_0 on $Y \setminus P$ that has both endpoints on p and separates $Y \setminus P$ into two components where the "left" component has g_l genus and contains all left curves of Y except y_1 and the "right" component has g_r genus and contains all right curves. We complete γ_0 to an arc γ on Y by joining the endpoints of γ_0 to c_1 and c_2 by arcs in P that separate y_1 to the "left" side of $Y \setminus \gamma$; this can be done such that the final arc has $i(\gamma, d \cap Y) \leq 2$; see Figure 7.

We conclude by observing that in any of the three above cases, we have produced an arc γ on Y with the same topological type as $c \cap Y$ but that intersects d at most twice on Y. Surgering c along γ as before we produce the desired curve c_Y. $\square$

To prove Proposition 4.9, let $Y_1, \ldots Y_k$ be the components of $S \setminus \alpha$ on which c and d intersect. Applying Claim 4.10 to Y_1, we get a genus-separating curve $c_1 \in \Psi(\alpha)$ that intersects d at most 4 times in Y_1 and agrees with c outside of Y_1. By Lemma 4.8, there is an admissible curve a_1 on $S \setminus Y_1$ that intersects c, and hence c_1, at most twice. Applying Lemma 4.6, this implies that a_1 is M-close to both $\Pi(c)$ and $\Pi(c_1)$ in $\mathscr{C}_{\mathrm{adm}}$ for some universal M. Hence, $\Pi(c)$ and $\Pi(c_1)$ are $2M$-close to each other. Repeating this argument, we produce a sequence of genus-separating curves $c = c_0, c_1, \ldots, c_k$ in $\Psi(\alpha)$ such that $\Pi(c_i)$ and $\Pi(c_{i+1})$ are $2M$-close in $\mathscr{C}_{\mathrm{adm}}$ and $i(c_k, d)$ is at most 2 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'.

We now establish the requisite diameter bounds. Since the length of the sequence from c to c' is bounded by $|\chi(S)|$, each $\Pi(c_i)$ has uniformly bounded diameter in $\mathscr{C}_{\mathrm{adm}}$, and each $\Pi(c_i)$ and $\Pi(c_{i+1})$ are $2M$-close, we conclude that $\Pi(c) \cup \Pi(c')$ has uniformly bounded diameter. This gives (2).

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 4.9 now follows by the triangle inequality. $\square$

We now show that the admissible curve graph $\mathscr{C}_{\mathrm{adm}}$ is quasi-isometric to the model $\mathcal{K}$. Since the inclusion $\mathscr{C}_{\mathrm{adm}} \to \mathcal{K}$ is simplicial and hence 1-Lipschitz, this statement is implied by the following:

Proposition 4.11. *The map $\Pi \circ \Psi \colon \mathcal{K} \to \mathscr{C}_{\mathrm{adm}}$ is a quasi-inverse to the inclusion $\mathscr{C}_{\mathrm{adm}} \to \mathcal{K}$.*

Proof. We first check that for all $a \in \mathscr{C}_{\mathrm{adm}}$, the image $\Pi \circ \Psi(a)$ is uniformly close to a in $\mathscr{C}_{\mathrm{adm}}$. 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 4.11. We have already shown in Proposition 4.9 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 4.9.

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 $g(S) \geq 3$, this is only possible if $g(S) = 3$ and Y is a 4-holed sphere where every curve in ∂Y is non-peripheral and non-separating on 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 8. Since no curves in ∂Y are separating or peripheral on S, 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 $\operatorname{diam}(\Pi \circ \Psi(\alpha) \cup \Pi \circ \Psi(\alpha'))$ is uniformly bounded by Proposition 4.9. $\square$

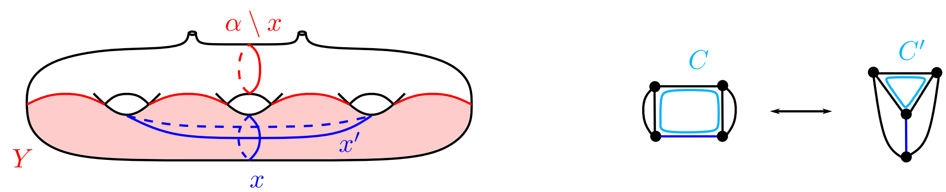

FIGURE 8. One the left, the subsurface Y where x is flipped to x'. One the right, the Whitehead move on the dual graph corresponding to flipping x to x'. The cycle C is sent to the cycle C' under this move.

Proof of Theorem A. Lemma 3.5 and Proposition 4.11 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 $\Delta_1, \ldots, \Delta_k$ and W^- contains $\Delta_{k+1}, \ldots, \Delta_b$; see Figure 9. Orient each curve of α such that W^+ is to the left.

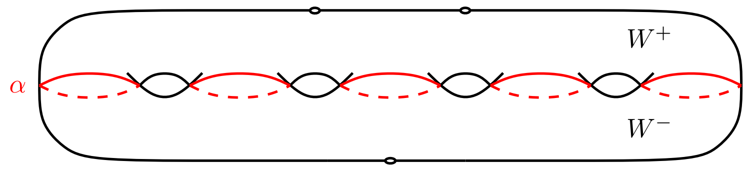

Figure 9. The multicurve α whose complement is a pair of witnesses for $\mathscr{C}_{\mathrm{adm}}$.

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 (non-peripheral) ψ-admissible curve if and only if there is some subset $\mathcal{C}$ of $\alpha \cup \Delta_1 \cup \ldots \cup \Delta_k$ such that

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

A similar condition tells us if W^- contains any non-peripheral admissible curves.

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 (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]). Moreover, we can choose x_i not to satisfy (4) for any subset $\mathcal{C}$ of ∂W^+ or the corresponding equations for W^- since these all linearly independent from (3). Thus W^+ and W^- are a pair of disjoint witnesses for $\mathscr{C}_{\mathrm{adm}}(S, \psi)$.

Set $K = \sum |\phi(\Delta_j)|$. The choices in the previous paragraph can all be made explicitly by choosing $x_1, \ldots, x_g$ 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 $\alpha \cup \Delta_1, \ldots, \Delta_k$, the left-hand side of (4) has magnitude larger than K unless it contains all of α. In this case, any curve separating off (a subset of) the Δ_j appearing in W^+ must have negative winding number, which is in particular not zero. Thus W^+ contains no admissible curves, so W^- is a witness. The argument to show W^+ is a witness 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 $f \in \mathrm{Mod}(S)$ such that $\phi = f(\psi)$, and thus $f(W^+)$ and $f(W^-)$ are the desired pair of disjoint witnesses for $\mathscr{C}_{\mathrm{adm}}(S, \phi)$. □

5. Curve graphs for strata

In this section we define a number of analogous graphs for (bordifications of) strata. We start by recalling some of the results of [CS22] on the relationship between strata, markings, and framed mapping class groups and discussing how the curve complex captures the intersection pattern of the boundary of a bordification of Teichmüller space. 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 all quasi-isometric (Corollary 6.5).

5.1. **Framings and strata.** 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 $\underline{\kappa} = (k_1, \dots, k_n)$ 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 $k_1, \dots, k_n$. Since a holomorphic 1-form is entirely determined (up to global scaling by $\mathbb{C}^*$) by the order and position of its zeros, any stratum can be thought of as a $\mathbb{C}^*$ bundle over a subvariety of $\mathcal{M}_{g,n}$ (after taking a manifold cover). In the sequel, we will freely conflate a stratum and its image in $\mathcal{M}_{g,n}$; we trust this will not cause any confusion.

Let $\Omega\mathcal{T}_{g,n}(\underline{\kappa})$ denote the full preimage of the stratum $\Omega\mathcal{M}_g(\underline{\kappa})$ inside of $\mathcal{T}_{g,n}$. In order to understand its connected components, one needs to understand which mapping classes can be realized inside a stratum, that is, 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, where $\mathcal{H}$ is any stratum component. When $\mathcal{H}$ is hyperelliptic, 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 $1/f^*\omega$. Then we use $\mathcal{H}_\phi$ to denote the subset of $\Omega\mathcal{T}_{g,n}(\underline{\kappa})$ parametrizing those marked differentials (X', ω', f') such $(X', \omega') \in \mathcal{H}$ and $1/(f')^*(\omega')$ is isotopic to ϕ.

By Theorem 5.1, if $g \geq 5$ then $\mathcal{H}_\phi$ is just a specified connected component of $\Omega\mathcal{T}_{g,n}(\underline{\kappa})$. 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 2.9) 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 boundary strata 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 of non-degenerate structures: 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 1-skeleton of the barycentric subdivision of the curve complex is the *multicurve graph*, which has a vertex for each (simple) multicurve on S and whose edges are given by inclusion: γ is connected to δ if and only if $\gamma \subset \delta$ or $\delta \subset \gamma$. Equivalently, the multicurve graph is the nerve of the (closures of) all boundary strata of $\overline{\mathcal{T}_{g,n}}$.

5.3. **Multi-scale differentials and level splittings.** We now perform a similar construction for (marked) strata of abelian differentials. Our discussion will be made more 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 10. In this example, curves α and β are homologous and so their periods must

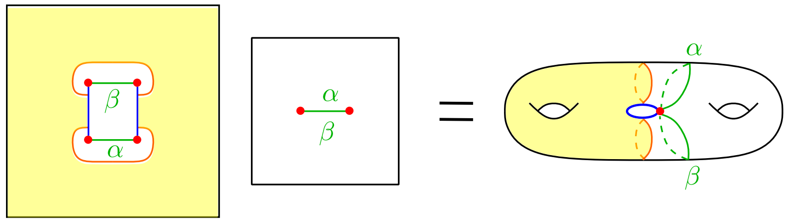

FIGURE 10. A 2-level splitting consisting of two homologous curves. Neither α nor β is admissible, and neither defines a level splitting by itself, so by Theorem 5.9 below, $\overline{\mathcal{H}_\phi}$ does not meet $\mathcal{T}_{g,n}(\alpha)$ or $\mathcal{T}_{g,n}(\beta)$.

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). Analogous to the multicurve graph, we now define a graph capturing the pattern of intersections of $\overline{\mathcal{H}_\phi}$ with the boundary strata of $\overline{\mathcal{T}_{g,n}}$:

Definition 5.5. Let $\mathscr{C}(\overline{\mathcal{H}_\phi})$ be the graph whose vertices are multicurves γ such that $\overline{\mathcal{H}_\phi} \cap \mathcal{T}_{g,n}(\gamma)$ is nonempty and whose edges are given by inclusion.

Exactly which multicurves appear as vertices of $\mathscr{C}(\overline{\mathcal{H}_\phi})$ is a very intricate question, and is related to subtle properties of a certain compactification 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" $\Xi\mathcal{H}$ which has nicer geometric properties (e.g., its boundary is a normal crossing divisor). A multi-scale differential is encoded by three pieces of data: an "enhanced level graph," a twisted differential compatible with the level graph and the enhancement, and a "prong matching."

We will not give precise definitions of these compactifications here, and direct the reader to the original papers (especially Section 5.1 of [BCG+19] and Section 3 of [CMZ22]). Instead, we record some of the relevant combinatorial data using our terminology of multicurves and winding numbers. We keep the numbering conventions of [BCG+19].

Definition 5.6. Let $S = S_{g,n}$ and let ϕ be a framing of S. An N-*level splitting* is an oriented multicurve $\vec{\beta}$ together with a partition of $S \setminus \beta$ into (nonempty, but possibly disconnected) subsurfaces $Y_0, \ldots, Y_{-N+1}$ such that:

- The winding number $\phi(b)$ is negative for every curve $b \subset \vec{\beta}$.
- Let b be a curve of $\vec{\beta}$. Then if the subsurfaces it sees on its left and right are Y_i and Y_j, respectively, then $i > j$.

A multicurve β is an N*-level multicurve* if β can be oriented and $S \setminus \beta$ can be partitioned to yield an level splitting with N levels.

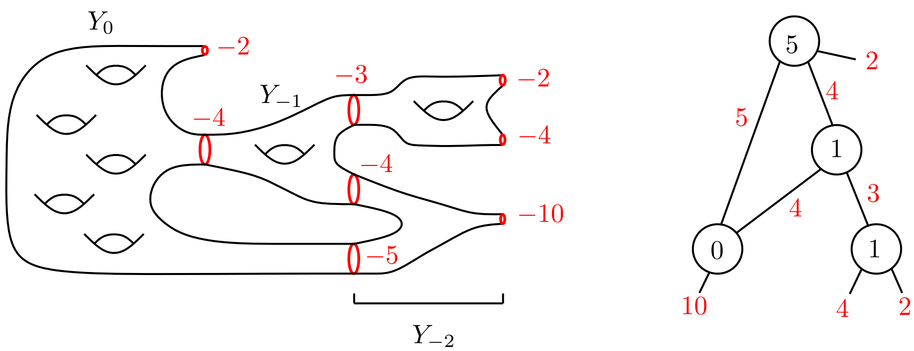

FIGURE 11. A 3-level multicurve and corresponding enhanced level graph.

Remark 5.7. Homological coherence implies that a 2-level multicurve determines a unique level splitting (because there are only two options for the partition of $S \setminus \beta$, and one option does not satisfy homological coherence). However, as shown in [BCG+18, Examples 3.3 and 3.4], this is not true for general N-level multicurves as soon as $N \geq 3$. Moreover, it is possible that an N-level multicurve is compatible with N'-level splittings.

Comparing this to [BCG+19], a level splitting records slightly different information than an enhanced level graph/enhanced multicurve without horizontal edges. The dual graph to a splitting (together with the partition of $S \setminus \beta$) is a level graph, and the winding numbers of the curves correspond to the enhancement, i.e., the orders of the zero and pole on each side of the node. In particular, for each oriented curve $b \subset \vec{\beta}$ corresponding to an edge of the dual graph with enhancement κ (so a zero z of order $o(z) = \kappa - 1$ and a pole p of order $o(p) = -\kappa - 1$), we have

$$\phi(b) = -\kappa = -1 - o(z) = 1 + o(p).$$

Thus every N-level splitting gives rise to an enhanced level graph.

However, a single enhanced level graph may be compatible with multiple level splittings, as the splitting enforces the Arf invariants of the components of $\phi|_{S\setminus\beta}$ (even up to the action of $\mathrm{Stab}_{\mathrm{Mod}(S)}(\beta)$) and the level graph does not. This is related to the fact that the $\mathrm{Mod}(S)$ orbit of multicurves are generally larger than the $\mathrm{FMod}(S, \phi)$ orbits (even controlling for winding numbers). Compare Proposition 6.2 below.

When ϕ has holomorphic type, every boundary component and peripheral curve of the top (that is, 0^{th}) level of a level splitting has negative winding number. Homological coherence (Lemma 2.2.2) then implies that each of these top component 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, hence is supported on a surface with genus. We record this for later use:

Fact 5.8. *Let $\vec{\beta}, (Y_0, \dots, Y_{-N+1})$ be a level splitting of a framed surface of holomorphic type. Then each component of Y_0 has positive genus.*

A corollary of the description of the moduli space of multi-scale differentials in [BCG+19] is the following statement, which gives an important necessary condition for which multicurves can be pinched in $\overline{\mathcal{H}_\phi}$:

Theorem 5.9. *If $\overline{\mathcal{H}_\phi} \cap \mathcal{T}_{g,n}(\gamma)$ is nonempty, then γ is the union of an admissible multicurve α and a disjoint N-level multicurve β.*

[MUW21] gives a sufficient condition for a topological type of multicurve to appear in the boundary of a 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 when the Arf invariants of subsurfaces 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 to a result for $\overline{\mathcal{H}_\phi}$, one would also need to establish very strong transitivity results for the action of the framed mapping class group (for example, one needs transitivity on all admissible multicurves of the same topological type, not just pairs). We were unable to achieve this level of generality, and so instead focus our attention on the "largest" boundary strata. For coarse-geometric questions, this distinction will be irrelevant.

5.4. **Divisorial multicurves.** As mentioned above, the map from the space of multiscale differentials $\Xi\mathcal{H}$ to the closure $\overline{\mathcal{H}} \subset \overline{\mathcal{M}_{g,n}}$ is highly singular. All the same, because the boundary of $\Xi\mathcal{H}$ is a normal crossing divisor, it gives us a good notion of what the largest boundary strata are.

Irreducible components of the boundary divisor of the moduli space $\Xi\mathcal{H}$ of multiscale differentials correspond to 1-level graphs with a single horizontal edge (i.e., admissible curves) and certain 2-level graphs with only vertical edges (i.e., 2-level multicurves) [BCG+19]. We therefore make the following definition:

Definition 5.10. A multicurve γ is called *divisorial* for $\mathcal{H}_\phi$ if $\overline{\mathcal{H}_\phi} \cap \mathcal{T}_{g,n}(\gamma)$ is nonempty and γ is either a single admissible curve or a 2-level multicurve.

It is still fairly complicated to identify exactly which 2-level multicurves are divisorial (compare the discussion at the end of the previous subsection as well as Proposition 6.2 below). As a first example, if β is divisorial and $Y_0(\beta)$ has a genus 1 component U, then in the associated boundary component of $\Xi\mathcal{H}$ the surface U is equipped with a holomorphic abelian differential. Thus $\mathrm{Arf}_1(\phi|_U)$ must be 0 (Remark 2.3).

All the same, we can build a graph that records only the intersections of (multicurves corresponding to) strata of $\overline{\mathcal{H}_\phi}$ coming from boundary divisors of $\Xi\mathcal{H}$.

Definition 5.11. Set $\mathscr{D}(\overline{\mathcal{H}_\phi})$ to be the graph with vertices given by divisorial multicurves γ, and with an edge between γ and δ if and only if $\overline{\mathcal{H}_\phi} \cap \mathcal{T}_{g,n}(\gamma \cup \delta)$ is nonempty.

Observe that $\mathscr{D}(\overline{\mathcal{H}_\phi})$ is to the curve graph as $\mathscr{C}(\overline{\mathcal{H}_\phi})$ is to the multicurve graph. In particular, its subdivision is a subgraph of $\mathscr{C}(\overline{\mathcal{H}_\phi})$ by definition, and since every boundary stratum of $\Xi\mathcal{H}$ is an intersection of boundary divisors, this subgraph is coarsely dense.

Unfortunately, while it is simpler than $\mathscr{C}(\overline{\mathcal{H}_\phi})$, the edges of $\mathscr{D}(\overline{\mathcal{H}_\phi})$ are still defined in terms of intersections of boundary strata. This is a subtle question even for disjoint 2-level splittings, so we define one further, simpler curve graph that is more amenable to the HHS techniques from the previous sections. Eventually, we will show that all of the graphs we have defined are quasi-isometric (Corollary 6.5).

Definition 5.12. Set $\mathscr{E}(\overline{\mathcal{H}_\phi})$ to have the same vertex set as $\mathscr{D}(\overline{\mathcal{H}_\phi})$, with an edge between γ and δ if and only if the two multicurves are disjoint (but are allowed to share components).

The graphs $\mathscr{D}(\overline{\mathcal{H}_\phi})$ and $\mathscr{E}(\overline{\mathcal{H}_\phi})$ are indeed different. We thank Martin Möller for first bringing this phenomenon to our attention.

Example 5.13 (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 $\overline{\mathcal{H}_\phi}$ meets both $\mathcal{T}_{g,n}(c)$ and $\mathcal{T}_{g,n}(d)$.

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 (more generally, the global residue condition [BCG$^+$19, §2.4 item (4)]) 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.

The inclusion gives a 1-Lipschitz map from $\mathscr{D}(\overline{\mathcal{H}_\phi})$ to $\mathscr{E}(\overline{\mathcal{H}_\phi})$; below, we show that this map actually extends to $\mathscr{C}(\overline{\mathcal{H}_\phi})$.

Lemma 5.14. *There is a coarsely Lipschitz map*

$$\xi : \mathscr{C}(\overline{\mathcal{H}_\phi}) \to \mathscr{E}(\overline{\mathcal{H}_\phi})$$

that coarsely agrees with the inclusion $\mathscr{D}(\overline{\mathcal{H}_\phi}) \hookrightarrow \mathscr{E}(\overline{\mathcal{H}_\phi})$*.*

Proof. Any boundary stratum of $\Xi\mathcal{H}$ is an intersection of boundary divisors; exactly which divisors can be recovered from the "undegenerations" of the associated enhanced level graph [BCG$^+$19, Definition 5.1]. The precise details of the situation will not be important to us; all we need is the following:

Fact 5.15. *Given any boundary stratum of* $\Xi\mathcal{H}$ *in which* γ *is pinched, all of its undegenerations correspond to pinching sub-multicurves of* γ*.*

We now define the desired coarsely Lipschitz map

$$\xi : \mathscr{C}(\overline{\mathcal{H}_\phi}) \to 2^{\mathscr{E}(\overline{\mathcal{H}_\phi})}$$

by sending a multicurve γ to the set of multicurves pinched in any divisorial undegeneration of any boundary stratum corresponding to γ.[2] The image of any vertex of $\mathscr{C}(\overline{\mathcal{H}_\phi})$ lies in a

[2]This is one place where our viewpoint of taking level multicurves, not level splittings, makes the discussion more complicated. Which undegenerations occur, and which boundary divisors intersect, depend not just on the multicurve but also on the level structure.

clique since all of the undegenerations are disjoint, and since edges of $\mathscr{C}(\overline{\mathcal{H}_\phi})$ are given by inclusion, it follows that if $\gamma \subset \delta$ then $\xi(\gamma) \subset \xi(\delta)$. Thus ξ is coarsely (2-)Lipschitz.

To see that ξ coarsely agrees with the inclusion, we simply observe that it agrees with the inclusion on vertices of $\mathscr{D}(\overline{\mathcal{H}_\phi})$, and that the edges of $\mathscr{D}(\overline{\mathcal{H}_\phi})$ get mapped to sets containing both endpoints. □

6. From transitivity to geometry

All of the graphs defined in the previous section carry a natural action of the framed mapping class group. Throughout the section, fix a non-hyperelliptic stratum component $\mathcal{H} \subset \mathcal{M}_{g,n}$ (necessarily with $g \geq 3$) and 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 $f \in \mathrm{FMod}(S, \phi)$,

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

In this section, we analyze this action in more detail and use certain transitivity properties to relate the geometries of $\mathscr{C}(\overline{\mathcal{H}_\phi})$, $\mathscr{D}(\overline{\mathcal{H}_\phi})$, and $\mathscr{E}(\overline{\mathcal{H}_\phi})$ to each other, to a hierarchically hyperbolic model, and to 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 $\mathscr{D}(\overline{\mathcal{H}_\phi})$ and $\mathscr{E}(\overline{\mathcal{H}_\phi})$ contain $\mathscr{C}_{\mathrm{adm}}(S, \phi)$.*

Proof. It suffices to prove the statement for $\mathscr{D}(\overline{\mathcal{H}_\phi})$ as it is a subgraph of $\mathscr{E}(\overline{\mathcal{H}_\phi})$. Since every admissible curve is divisorial, $\mathscr{D}(\overline{\mathcal{H}_\phi})$ contains the vertices of $\mathscr{C}_{\mathrm{adm}}(S, \phi)$, so it remains to show that it also contains the edges. By Proposition 2.9, the framed mapping class group $\mathrm{FMod}(S, \phi)$ acts transitively on pairs of admissible curves of the same topological type. Thus, it only remains to show that $\overline{\mathcal{H}_\phi}$ meets *some* boundary stratum $\mathcal{T}_{g,n}(\alpha)$ for each topological type α of pair of admissible curves.

One can do this 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, 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 $\mathcal{T}_{g,n}(\alpha)$.

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 $\Omega\mathcal{M}_3(2,2)$, which has a unique non-hyperelliptic component of odd Arf invariant [KZ03, Theorem 2]. □

6.1. **The action on 2-level multicurves.** 2-level multicurves can have many different topological types, so $\mathrm{FMod}(S,\phi)$ will certainly not act transitively on them. However, even controlling for topological type and winding numbers, the Arf invariants of 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, note that the following statement is true for *all* 2-level splittings of a surface 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. Let β be any 2-level multicurve. 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 $\mathrm{Mod}(S)$ via inclusion) 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 $Y_{-1}(\beta)$.

Level passage: For this and the next step, for each component U of the top level $Y_0(\beta)$, pick a subsurface $V_U \subset U$ 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

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

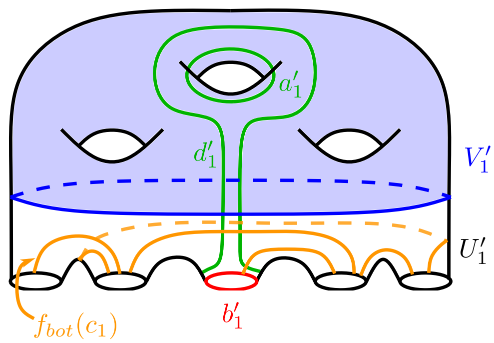

FIGURE 12. The multitwist needed to fix the winding number of $f_{bot}(c_1)$. Note that it may not always be possible to choose the curve d_1' to be disjoint from the other curves of $f_{bot}(\mathsf{c})$.

to a GSB for the complement of all of the V_U. 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 2-level splitting β, and we use U_1 to denote the component of $Y_0(\beta)$ adjacent to b_1. Since ϕ is of holomorphic type, U_1 has genus, hence so does the full-genus subsurface $V_1 := V_{U_1}$. Set $\mathsf{c}' = f_{bot}(\mathsf{c})$ and likewise for its components. Pick some admissible curve $a_1' \subset f_{bot}(V_1)$ and let d_1' denote the connect sum of b_1' with a_1' along some arc contained in U_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 12.

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 $V_2 \subset U_2$ to be the full-genus subsurface of the top-level component of $S \setminus \beta$ adjacent to b_2.[3] There is an admissible $a_2' \subset f_1 f_{bot}(V_2)$, 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 U_2' such that $f_2 f_1 f_{bot}$ preserves the winding numbers of both c_1 and c_2.

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

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

[3]The component U_2 (and subsurface V_2) may be the same as U_1 (and V_1).

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

Top level: To finish, we can simply use the action of $\mathrm{Mod}(f_{mid}(V))$ to amend the winding numbers of the remaining curves as we did for the bottom level.

Pick a GSB $\mathcal{B}_V$ for each such V. Then by Lemma 2.6 and hypothesis (3), we have that

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

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

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

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 2-level splitting is connected to some admissible curve in $\mathscr{D}(\overline{\mathcal{H}_\phi})$.

Proposition 6.3. *Suppose that β is a divisorial 2-level multicurve for $\overline{\mathcal{H}_\phi}$, i.e., $\overline{\mathcal{H}_\phi}$ meets $\mathcal{T}_{g,n}(\beta)$. Then for any admissible $a \subset Y_0(\beta)$, we have that $\overline{\mathcal{H}_\phi}$ meets $\mathcal{T}_{g,n}(\beta \cup a)$.*

Proof. We first show that one can further pinch *some* admissible curve in $Y_0(\beta)$. The restriction of a multiscale differential lying over $\overline{\mathcal{H}_\phi} \cap \mathcal{T}_{g,n}(\beta)$ is holomorphic on $Y_0(\beta)$. Every holomorphic differential contains an embedded nonsingular cylinder [Mas86] whose core curve a' is necessarily admissible, and one can degenerate into $\overline{\mathcal{H}_\phi} \cap \mathcal{T}_{g,n}(\beta \cup a')$ 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 U of $Y_0(\beta)$ and suppose (postcomposing by an element of $\mathrm{Mod}(U)$ as necessary) that the element h taking β to β' also takes a to a'. The proof of Proposition 6.2 then 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 $V \subset U$ to contain a', so $f_{bot}(V)$ 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 U is the relevant subsurface in the iteration, then the resulting multitwist f_j will preserve a' and the new subsurface $f_j \cdots f_1 f_{bot}(V)$ 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 in the genus 1 case, as divisoriality implies that $\mathrm{Arf}_1(\phi_V) = 0$, so every nonseparating curve is admissible. If $g(V) \geq 2$, it follows from Corollary 2.11 that we can choose a curve transverse to a' with the desired winding number, then apply Lemma 2.8 to the complement in V of the subsurface filled by a and this curve (note that if $g(V) = 2$, then the boundaries of this subsurface have winding numbers -3 and $+1$, so its genus 1 Arf invariant is either 1 or 2 and there are no more restrictions than already appear from fixing the Arf invariant of V). □

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

Corollary 6.4. *The graphs $\mathscr{C}(\overline{\mathcal{H}_\phi})$, $\mathscr{D}(\overline{\mathcal{H}_\phi})$, and $\mathscr{E}(\overline{\mathcal{H}_\phi})$ are connected.*

Proof. Since $\mathscr{D}(\overline{\mathcal{H}_\phi})$ is a subgraph of $\mathscr{E}(\overline{\mathcal{H}_\phi})$ (and its subdivision is a subgraph of $\mathscr{C}(\overline{\mathcal{H}_\phi})$), it suffices to prove this for $\mathscr{D}(\overline{\mathcal{H}_\phi})$. By Lemma 6.1, $\mathscr{D}(\overline{\mathcal{H}_\phi})$ 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 $\mathscr{D}(\overline{\mathcal{H}_\phi})$ 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 5.8 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 $\mathscr{D}(\overline{\mathcal{H}_\phi})$. Since we have connected every vertex of $\mathscr{D}(\overline{\mathcal{H}_\phi})$ to the connected graph $\mathscr{C}_{\mathrm{adm}}(S, \phi)$, we conclude that $\mathscr{D}(\overline{\mathcal{H}_\phi})$ is connected. □

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

Corollary 6.5. *The graphs $\mathscr{C}(\overline{\mathcal{H}_\phi})$, $\mathscr{D}(\overline{\mathcal{H}_\phi})$, and $\mathscr{E}(\overline{\mathcal{H}_\phi})$ are quasi-isometric.*

Proof. Observe that via the inclusion (of its subdivision) and Lemma 5.14, we have already built coarsely Lipschitz maps

$$\mathscr{D}(\overline{\mathcal{H}_\phi}) \to \mathscr{C}(\overline{\mathcal{H}_\phi}) \to \mathscr{E}(\overline{\mathcal{H}_\phi})$$

such that the final map $\mathscr{D}(\overline{\mathcal{H}_\phi}) \to \mathscr{E}(\overline{\mathcal{H}_\phi})$ coarsely agrees with the inclusion. Thus, it remains to show that the inclusion $\mathscr{D}(\overline{\mathcal{H}_\phi}) \hookrightarrow \mathscr{E}(\overline{\mathcal{H}_\phi})$ has a coarse inverse.

Define $\zeta \colon \mathscr{E}(\overline{\mathcal{H}_\phi}) \to 2^{\mathscr{D}(\overline{\mathcal{H}_\phi})}$ to be the identity on the vertices of $\mathscr{E}(\overline{\mathcal{H}_\phi})$ and to send each edge of $\mathscr{E}(\overline{\mathcal{H}_\phi})$ to the pair of vertices in $\mathscr{D}(\overline{\mathcal{H}_\phi})$ that are its endpoints. If ζ is coarsely Lipschitz, then it is necessarily a coarse inverse of the inclusion $\mathscr{D}(\overline{\mathcal{H}_\phi}) \hookrightarrow \mathscr{E}(\overline{\mathcal{H}_\phi})$.

To that end, consider any edge of $\mathscr{E}(\overline{\mathcal{H}_\phi})$ connecting disjoint multicurves γ and δ that correspond to divisorial boundary strata. We just need to show a bound on the length of a path from γ to δ in $\mathscr{D}(\overline{\mathcal{H}_\phi})$. If both γ and δ are single admissible curves, then Lemma 6.1 implies they are connected in $\mathscr{D}(\overline{\mathcal{H}_\phi})$. Now suppose γ is admissible and δ is a 2-level curve. 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 connected by an edge of $\mathscr{D}(\overline{\mathcal{H}_\phi})$. 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 connected to a which is connected to δ (by Lemma 6.1).

Finally, suppose that both γ and δ are 2-level curves. 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 $\mathscr{D}(\overline{\mathcal{H}_\phi})$ 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 $\mathscr{D}(\overline{\mathcal{H}_\phi})$. Thus, collecting cases, we have shown that each edge of $\mathscr{E}(\overline{\mathcal{H}_\phi})$ is sent to a set of diameter at most 4, hence ζ is coarsely Lipschitz. □

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 $\mathscr{D}(\overline{\mathcal{H}_\phi})$ 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 2-level splittings 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 2-level multicurves to collapse the image of $P(\mu)$ for each μ that is in $\overline{\mathcal{K}}$ but not $\mathcal{K}$. The abundance of 2-level multicurves comes from the following lemma.

Lemma 6.6. *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 2-level multicurve 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 is a weaker, framed version of the following standard "change of coordinates" lemma.

Lemma 6.7. *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 (even though we have shown the set of divisorial 2-level splittings 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 6.6. We first record a number of topological consequences of our hypotheses. Let Y_0 and Y_{-1} denote the two levels of the 2-level splitting associated to 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 $\mathrm{Arf}_1(Y_0) = 0$ must be 0 by divisoriality (see the discussion right after Definition 5.10) and hence some boundary curve of our witness W would be admissible, a contradiction.

We observe that since $S \setminus W$ has genus 0, 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 W necessarily preserves the winding numbers of the curves of β. Applying Lemma 6.7 (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 2-level splitting, as the definition of 2-level splitting depends only on winding numbers. Note also that $h(\beta)$ may not be divisorial.

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

Otherwise, $\mathrm{Arf}(hY_0) = \mathrm{Arf}(Y_0) + 1$. 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 $\mathrm{Arf}(Y_0) = 0$. The same would therefore also be true for $\mathrm{Arf}(hY_0)$, but this is a contradiction.

Now every mapping class supported entirely on 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 $d \subset S \setminus W$, we have that

$$\hat{\imath}(hc, d) = \hat{\imath}(c, d)$$

where $\hat{\imath}(\cdot, \cdot)$ denotes the algebraic intersection number. 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

$$\phi(hc) = \phi(c) + k_1 \hat{\imath}(c, d_1)\phi(d_1) + \ldots + k_n \hat{\imath}(c, d_n)\phi(d_n)$$

for any curve c on S.

Returning to the situation at hand, since $\mathrm{Arf}(hY_0) = \mathrm{Arf}(Y_0) + 1$, 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 $\hat{\imath}(c, d)$ and $\phi(d)$ are both odd. Moreover, $\hat{\imath}(hc, d)$ 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 6.7 ensures there is a d' on $S \setminus W$ such that

(1) $\hat{\imath}(gc, d')$ 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. $\square$

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 6.8. *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 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 exist 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 13. Set $\alpha' := \mu \cup \alpha_W$. We divide the remainder of our proof into three cases based on the subsurface W.

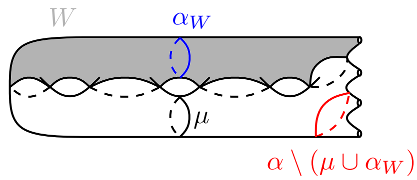

FIGURE 13. The partition of $\alpha \in P(\mu)$ into α_W, μ and $\alpha \setminus (\mu \cup \alpha_W)$.

Case 1: $g(W) \geq 1$. Since $\mu \in \overline{\mathcal{K}}$, then by definition of $\overline{\mathcal{K}}$ there must exist some 2-level splitting $\delta \in \mathcal{D}$ such that W is disjoint from δ. 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 4.7, 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 6.6 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 14. Note that this choice depends only on μ, not on $\alpha \in P(\mu)$.

Since δ is a 2-level splitting, $S \setminus \delta$ 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 6.7, 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 14. 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 14 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 $N + N'$.

As in the previous case, Lemma 4.7 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 $\alpha_W = \alpha \cap W$

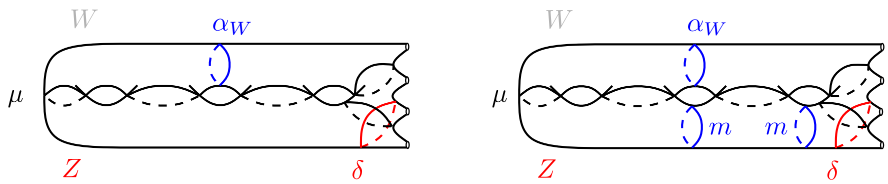

FIGURE 14. A schematic of the curves δ and m relative to μ and W. The actual intersection number between $\delta \cup m$ and μ is possibly higher, but still uniformly bounded.

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 $\theta(\beta_W \cup \mu \cup \beta_V)$. 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. 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 4.11), 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 $\mathscr{E}(\overline{\mathcal{H}_\phi})$. $\square$

Proof of Theorem B. By Corollary 6.5, $\mathscr{E}(\overline{\mathcal{H}_\phi})$ is quasi-isometric to $\mathscr{D}(\overline{\mathcal{H}_\phi})$ and $\mathscr{C}(\overline{\mathcal{H}_\phi})$, and by Theorem 6.8, it is quasi-isometric to the hierarchically hyperbolic model $\overline{\mathcal{K}}$. Thus all of these graphs are hierarchically hyperbolic.

The proof that $\mathscr{C}(\overline{\mathcal{H}_\phi})$ 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 2-level splitting 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 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 2-level splitting then there must be some subset $\mathcal{C}$ of the curves of α such that

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

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* 2-level splittings 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 $\overline{\mathcal{K}}$ 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

[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.

[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.

[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).

[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`.

[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.