HIERARCHICAL HYPERBOLICITY OF ADMISSIBLE CURVE GRAPHS

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 curves with winding number 0 with respect to the framing is hierarchically hyperbolic but not Gromov hyperbolic. We also describe how this graph can be viewed as encoding the combinatorics of a partial bordification of a marked stratum of abelian differentials.

1. Introduction

There is a storied history to using graphs built from curves on a surface S to understand the mapping class group, $\mathrm{Mod}(S)$, of the surface and related objects. The most famous and far reaching example is Harvey's curve graph, $\mathscr{C}(S)$, which 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]. The curve graph is a central object in low-dimensional topology, illuminating not only the mapping class group [Iva97, MM00], but also the geometry of Teichmüller space [MM98, Raf05] and the structure of hyperbolic 3-manifolds [Min10, BCM12]. Other examples of this paradigm include the pants graph for understanding the coarse geometry of the Weil–Petersson metric [Bro03, BF06], the Torelli and separating curve graphs for studying the Torelli subgroup and the Johnson kernel [FI05, BM04], and the disk graph for examining the handlebody group and Heegaard splittings [Hen20, MS13].

Recently, the first author and Salter have shown that the *framed mapping class group* plays a important role in understanding moduli spaces of abelian differentials [CS22]. A *framing* ϕ on a surface S is a trivialization of its tangent bundle, or equivalently (up to isotopy), a non-vanishing vector field. The framed mapping class group $\mathrm{FMod}(S, \phi)$ is the subgroup of $\mathrm{Mod}(S)$ that stabilizes the isotopy class of ϕ. A natural graph of curves on which $\mathrm{FMod}(S, \phi)$ acts is the *admissible curve graph*, $\mathscr{C}_{\mathrm{adm}}(S, \phi)$, the subgraph of $\mathscr{C}(S)$ spanned by curves that have winding number 0 with respect to ϕ. How closely the relationship between $\mathrm{FMod}(S, \phi)$ and $\mathscr{C}_{\mathrm{adm}}(S, \phi)$ mimics the relationship between $\mathrm{Mod}(S)$ and $\mathscr{C}(S)$ is an open question.

A marquee results on the curve graph $\mathscr{C}(S)$ is Masur and Minsky's proof of hyperbolicity [MM98]. This has had far-reaching implications for the coarse geometry of the mapping class group and is a central component of the resolution of the ending lamination conjecture [Min10]. A starting place to understand the relationship between $\mathrm{FMod}(S, \phi)$ and $\mathscr{C}_{\mathrm{adm}}(S, \phi)$ is thus to ask about the geometry of $\mathscr{C}_{\mathrm{adm}}(S, \phi)$. We show that the admissible curve graph is not hyperbolic, but does possess a generalized notion of hyperbolicity.

Theorem A. *For any surface S 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 hyperbolic).*

Hierarchical hyperbolicity was introduced by Behrstock, Hagen, and Sisto to unify the coarse geometry of the mapping class group and Teichmüller space with right-angled Artin groups [BHS17b]. This framework allows one to understand the geometry of a space by projecting it onto a collection of hyperbolic spaces. In the case of $\mathscr{C}_{\mathrm{adm}}(S, \phi)$, we use Masur and Minsky's subsurface projection maps to project $\mathscr{C}_{\mathrm{adm}}(S, \phi)$ on the curve graphs of *witnesses* — subsurfaces of S that intersect every admissible curve. The non-hyperbolicity of $\mathscr{C}_{\mathrm{adm}}(S, \phi)$ emerges from the fact that there exist pairs of disjoint witnesses for $\mathscr{C}_{\mathrm{adm}}(S, \phi)$. Using the hierarchically hyperbolic machinery, these disjoint witnesses induce undistorted product regions in $\mathscr{C}_{\mathrm{adm}}(S, \phi)$ which obstruct hyperbolicity.

This approach was inspired by work of Vokes, who showed that a wide variety of graphs of curves are similarly hierarchically hyperbolic using their subsurface projection maps to witnesses [Vok22]. Vokes first uses the set of witnesses to build a "model graph" $\mathcal{K}$ that she proves is hierarchically hyperbolic. She then shows that when the graph of curves admits a cobounded action of $\mathrm{Mod}(S)$ it is quasi-isometric to the model graph $\mathcal{K}$. While we can construct Vokes's hierarchically hyperbolic model graph $\mathcal{K}$ for $\mathscr{C}_{\mathrm{adm}}(S, \phi)$, 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 use her argument. Instead, we construct a quasi-isometry $\mathcal{K} \to \mathscr{C}_{\mathrm{adm}}(S, \phi)$ by hand, without relying on the change-of-coordinates principle for $\mathrm{Mod}(S)$.

The boundary of marked strata. In addition to its topological definition, $\mathscr{C}(S)$ can also be interpreted as encoding the intersection pattern of pieces of the thin part of Teichmüller space. It turns out that the admissible curve graph can also be viewed as capturing the combinatorics of a partial bordification of (marked) strata.

We recall that an *abelian differential* is a holomorphic 1-form on a Riemann surface. The moduli space of all abelian differentials of genus g forms a rank g (orbifold) vector bundle $\Omega\mathcal{M}_g$ over the usual moduli space of genus g closed Riemann surfaces $\mathcal{M}_g$. This bundle is broken into pieces called *strata*, which parametrize those differentials with a fixed number and order of zeros. Since strata parametrize Riemann surfaces with marked points, and differentials are determined up to scaling by the order and position of their zeros, we may as well think of strata as subvarieties of $\mathcal{M}_{g,n}$, the moduli space of genus g Riemann surfaces with n marked points. Strata are not always connected, but Kontsevich and Zorich classified their connected components [KZ03]: there are always at most 3, at most one of which is *hyperelliptic*, in that it consists entirely of hyperelliptic Riemann surfaces with point sets that are invariant under the hyperelliptic involution (preserving orders).

Given a (non-hyperelliptic) stratum component $\mathcal{H} \subset \mathcal{M}_{g,n}$, one can construct a partial bordification $\breve{\mathcal{H}} \subset \overline{\mathcal{M}_{g,n}}$ in which cylinders can be stretched into infinite poles but no other degenerations are allowed. We can then "lift" this bordification to any component $\mathcal{H}_\phi$ of the preimage of $\mathcal{H}$ in $\mathcal{T}_{g,n}$ to yield a bordification $\breve{\mathcal{H}}_\phi \subset \overline{\mathcal{T}_{g,n}}$. This should be thought of as an analogue of how the augmented Teichmüller space "lifts" the Deligne–Mumford compactification [HK14].

In Proposition 5.5 we show that the combinatorics of this bordification correspond to the admissible curve graph; equivalently, $\mathscr{C}_{\mathrm{adm}}(S, \phi)$ can also be thought of as the "graph of cylinders" for a given (marked) stratum component (compare [CS22, Corollary 1.2]). Thus Theorem A can also be viewed as a statement about the coarse geometry of (marked) strata.

One could also define a number of different graphs that capture the intersection pattern of the boundary of the entire closure of $\mathcal{H}_\phi$. The authors will consider the geometry of these graphs in a future version of this paper.

Remark 1.1. 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 the sphere, 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 $\mathcal{M}_{0,n}$, and the intersection pattern of the boundary of $\mathcal{T}_{0,n}$ is just the usual curve graph of an n-times punctured sphere.

Outline of paper. We begin in Section 2 by recalling some basic information on the framed mapping class group and proving fundamental "change-of-coordinate" style lemmas. After these preliminaries, we prove Theorem A in Sections 3 and 4. The former section records Vokes's construction of a model graph given a collection of witnesses, while in the latter we build a quasi-isometry between the model and $\mathscr{C}_{\mathrm{adm}}(S, \phi)$. The final Section 5 discusses the relevant background on strata and explains how to relate $\mathscr{C}_{\mathrm{adm}}(S, \phi)$ to their boundaries.

Acknowledgments. We thank Martin Möller and Frederik Benirschke for helping us to understand the structure of the boundary of strata.

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 $S = S_g^b$ denote an orientatable surface with genus g and b boundary components. We denote the boundary curves of S by ∂S. The *complexity* of $S = S_g^b$ is $\xi(S) = 3g - 3 + b$. By a *curve* on S we mean an isotopy class of an essential (i.e., non-nulhomotopic), non-peripheral (i.e., not homotopic into ∂S), simple closed curve on S. An *arc* on S is an isotopy class of essential, non-peripheral simple arcs with endpoints on ∂S and with isotopy classes taken relative to ∂S. Curves and arcs are unoriented unless we say otherwise. By a *subsurface* of S, we mean an isotopy class of an essential, non-peripheral, 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 mapping class group, $\mathrm{Mod}(S)$, is the group of homeomorphisms of S that fix ∂S pointwise modulo isotopies that leave ∂S fixed. The mapping class group is generated by *Dehn twists*: for any simple closed curve c, let 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 boundary; our results also apply equally well to surfaces with punctures after applying the "capping homomorphism" (see [CS22, Section 6.2] for a discussion in the context of framed mapping class groups).

We are interested in the set of framings up to isotopy, allowing ϕ to vary on ∂S: this corresponds to the notion of an "absolute framing" 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.

These 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 2.1 (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 with boundary components* $c_1, \ldots, c_k$*, oriented so 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.*

Suppose that S has boundary components $\Delta_1, \ldots, \Delta_k$ (oriented with the surface on their left); then the *signature* of a framing ϕ is the tuple

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

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).

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

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 2.1, 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 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 2.1.1), 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 2.1) 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*

[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.

invariant of ϕ is defined to be

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

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

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

Theorem 2.2. *Two framings ϕ and ϕ' of S are in the same* $\operatorname{Mod}(S)$ *orbit if and only if*

$(g = 0)$ *$sig(\phi) = sig(\phi')$*
$(g = 1)$ *$sig(\phi) = sig(\phi')$ and* $\operatorname{Arf}_1(\phi) = \operatorname{Arf}_1(\phi')$
$(g \geq 2)$ *$sig(\phi) = sig(\phi')$ and if ϕ and ϕ' are of spin type, then* $\operatorname{Arf}(\phi) = \operatorname{Arf}(\phi')$.

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

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 $f \in \operatorname{Mod}(S)$ 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 $\operatorname{Mod}(S)$ 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 2.2 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 2.3. *For any framing ϕ on a surface S of positive genus, there is some 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

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

Now by Theorem 2.2 there is some homeomorphism $f \in \operatorname{Mod}(S)$ 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 2.4. *Let ϕ be a framing of a surface S of genus $g \geq 1$ and fix any tuple of integers $(x_1, y_1, \ldots, x_g, y_g)$ so 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 so 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.

However, Theorem 2.2 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 that 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 boundary 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 2.5. *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 non-separating admissible curves of the same topological type. That is, if γ, γ' are pairs of non-separating 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'.*

In particular, if ϕ is of holomorphic type then $\mathrm{FMod}(S, \phi)$ acts transitively on the set of all admissible curves.

Remark 2.6. When ϕ does not have holomorphic type, $\mathrm{FMod}(S, \phi)$ does *not* necessarily transitively on the set of all admissible curves, even of the same topological type. If ϕ is of spin type and c is admissible, then the restriction of ϕ to each of the components of $S \setminus c$ is also of spin type and the Arf invariant of each piece provides an obstruction to transitivity of the $\mathrm{FMod}(S, \phi)$ action.

Before proving Proposition 2.5, 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 2.7. *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 so that $\phi(d_z) = z$ and $i(c_i, d_z) = i(c_i, d)$ for all i.

Proof. Orient d so that the arc from d to T exits d from its left-hand side.

Suppose first that $g(T) = 2$. Then by Lemma 2.4 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 2.1.2) 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 2.4) 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 2.1.1) 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 2.8. *For any surface of genus at least 2, any nonseparating 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 2.4 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 2.7 to adjust $\phi(b')$ at will. □

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

Proof of Proposition 2.5. 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 nonseparating and admissible. Complete a to a GSB $a = a_1, b_1, \ldots, a_g, b_g$ of S. Using Corollary 2.8, 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 2.4 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 2.1), 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 2.8 and 2.4 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 extending a_1 since the complement of $a_1 \cup a_2$ has genus 1 (and hence there are more possible $\mathrm{Mod}(S)$ orbits). Suppose first that ϕ is of spin type. Using Corollary 2.8 twice, we can choose disjoint curves b_1 and b_2, each meeting their respective a_i and disjoint from the other, so 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 so that $\phi|_{S\setminus Y}$ has fixed Arf_1, and so by Lemma 2.4 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 of nonseparating curves. 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 meeting a_1 and a_2 exactly once. Since at least one of T or U has genus at least 2, we can use Lemma 2.7 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 2.4 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
- bounds a surface homeomorphic to $S_{1,2}$ together with a_1'.

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 2.1.2), it must be that $\phi(c) = \pm 2$ (where sign depends on orientation). Twist-linearity (Lemma 2.1.1) then implies that some twist of b_1' about c will be admissible. Thus the configurations of curves

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

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

If ϕ is not of spin type, then we can conclude by picking an arbitrary b_1' disjoint from $\mathcal{B}_{T'} \cup \mathcal{B}_{U'}$. We then note that since ϕ is not of spin type, then there is some 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

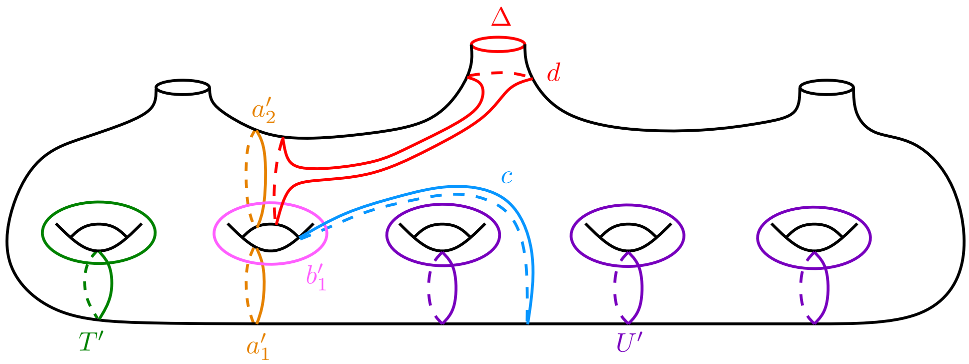

FIGURE 1. GSBs and auxiliary curves as in the proof of Proposition 2.5.

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* on 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 curves that are admissible with respect to ϕ.

Proposition 2.5 implies that the framed mapping class group $\mathrm{FMod}(S, \phi)$ acts with finitely many orbits on its vertices and edges (when ϕ is of holomorphic type, it acts with a single orbit on vertices). As a consequence of Lemma 2.3, 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 2.3 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 $\mathscr{C}_{\mathrm{adm}}(S, \phi)$. $\square$

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 $S = S_g^b$ with $g \geq 3$ and $b \geq 1$. Fix a framing ϕ of S.*

(1) *If $Z \subseteq S$ is a genus 0 subsurface and $z_1, \ldots, z_k$ are the boundary components of Z, oriented so that Z is to the left of each z_i, then Z contains a nonperipheral admissible curve if and only if there is no $I \subsetneq \{1, \ldots, k\}$ such that*

$$\sum_{i \in I} \phi(z_i) = 1 - |I|.$$

(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 of $S \setminus W$ is a genus 0 subsurface that does not contain any admissible curves.*

(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 is separating. The second item follows from the first plus Lemma 2.3's guarantee that every subsurface with genus contains an admissible curve. The third item is an immediate consequence of the second item. □

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 so 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 $\mathcal{K}_{\mathfrak{S}}(S)$ as a quasi-isometric model 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 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 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, $S = S_g^b$ will be a surface with $g \geq 3$ and $b \geq 1$ 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 $C \geq 0$ so that

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

for every $x, x' \in X$. In particular, note that coarsely Lipschitz maps are in particular coarsely well-defined. Prototypical examples are the inclusion of a connected subgraph into a connected graph, the subsurface projection map from the 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 end points of that edge. In this case, to show f is coarsely Lipschitz, it suffices to show that

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

Two =spaces are *quasi-isometric*, if there exists two coarsely Lipschitz map $f \colon X \to 2^Y$ and $\overline{f} \colon Y \to 2^Y$ so 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 full separating curve graph is connected also shows that $\mathcal{G}$ is connected [Put08].

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 $N = N(n) \geq 0$ so 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 so that every pair of genus-separating curves that intersect at most n times is in the $\mathrm{Mod}(S)$–orbit of some (c_i, d_i). Setting $N = \max\{d_{\mathcal{G}}(c_i, d_i) : 1 \leq i \leq k\}$, the fact that $\mathrm{Mod}(S)$ acts by isometries on $\mathcal{G}$ implies any two genus-separating curves that intersect at most n times are at most N far apart in $\mathcal{G}$. □

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

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

by sending a genus-separating curve to the collection of admissible curves disjoint from it. This set is always non-empty by Lemma 2.3.

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

Proof. 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 so 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 $M = M(m) \geq 0$ so 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 exist 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. $\square$

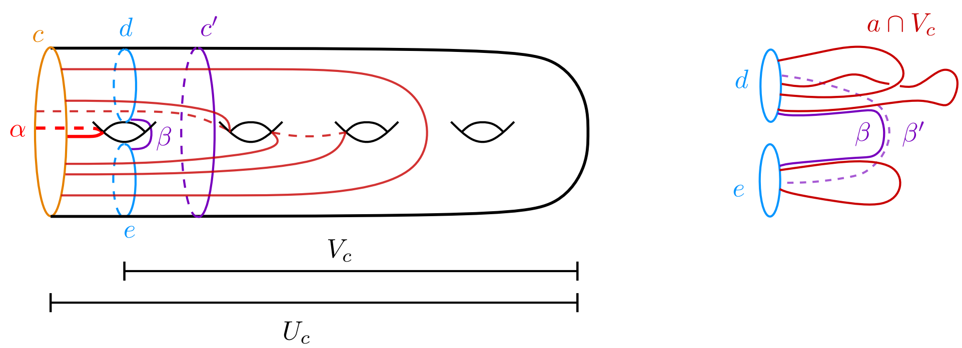

FIGURE 2. On the left, the subsurfaces involved in the proof of Lemma 4.6. On the right, surgering an arbitrary arc β' from d to e along $a \cap V_c$ to obtain a disjoint arc β.

4.4. **A quasi-inverse.** We now construct a map Ψ that assigns 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. $\square$

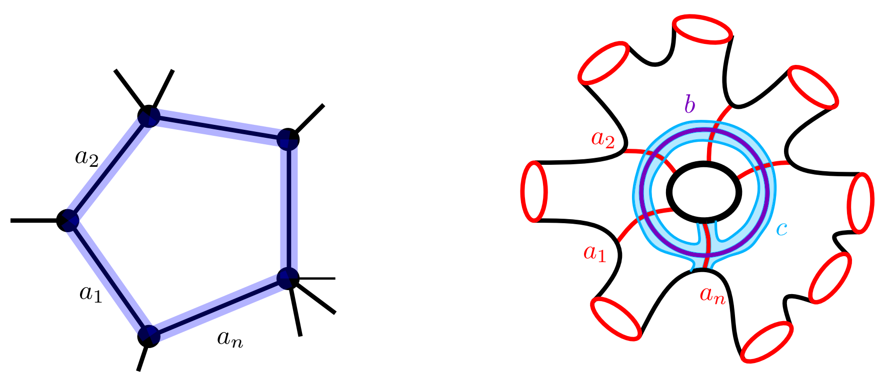

FIGURE 3. Building a genus-separating curve out of a cycle in the dual graph.

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 $\mathscr{C}_{\text{adm}}$ 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 2.3, 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 separate the boundary components of Z (Lemma 2.1.2). Let A be a set of curves in ∂Z so that any curve in Z partitioning ∂Z into A and $\partial Z \setminus A$ must be admissible. We argue below that one can always draw a curve a that cuts off the boundary components in 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). 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 so 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).

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

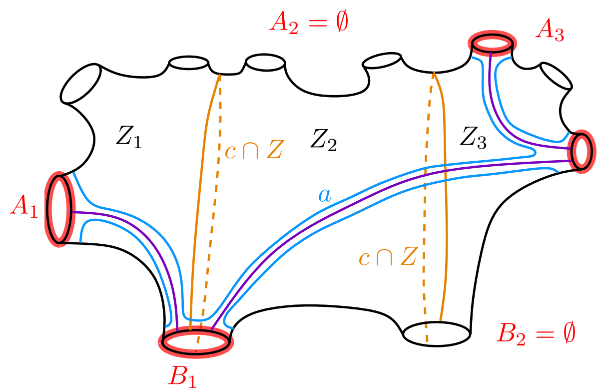

FIGURE 4. Building a curve that cuts off A, and is hence admissible. The highlighted curves are in A. In this example, A_2 and B_2 are empty, so the arc from B_1 to A_3 meets $c \cap Z$ exactly once.

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. This proves Lemma 4.8. □

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 $N \geq 0$ so 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

(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 so that c and d are each in minimal position with respect to α, and so 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 $c_Y \in \Psi(\alpha)$ so 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 so 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 so 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 so 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 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 boundary components of Y and the other with genus g_r and all the right boundary components 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.

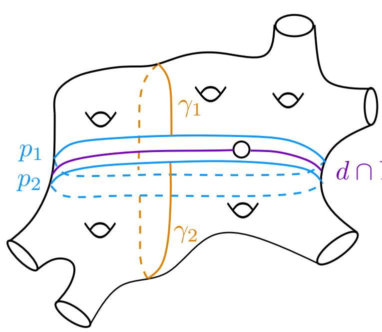

FIGURE 5. The curves p_1, p_2 cobounding the pair of pants P. The arcs γ_1 and γ_2 cut $S \setminus P$ into "left" and "right" sides.

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 so 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 so 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 so 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 so 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 so 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.

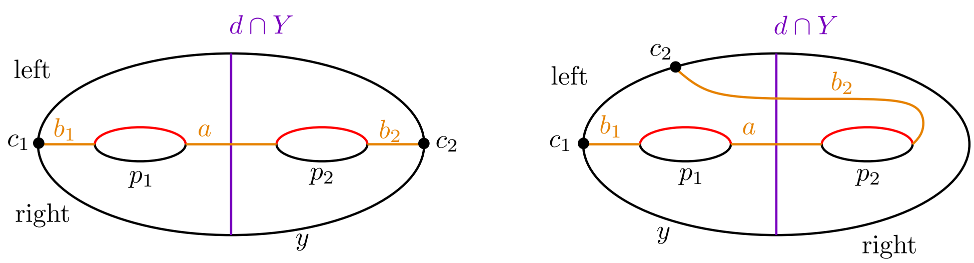

FIGURE 6. The arcs a, b_1, b_2 one must add in the pair of pants P to complete $\gamma_1 \cup \gamma_2$ to γ.

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 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 so 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 so 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 $Y \setminus P$ so 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 boundary of Y except y_1 and the right component has g_r genus and all the right boundary 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 so that the final arc has $i(\gamma, d \cap Y) \leq 2$; see Figure 7.

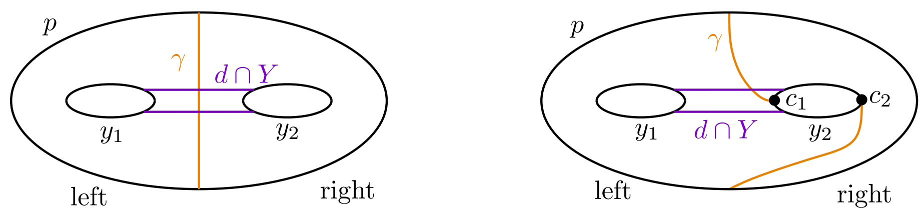

FIGURE 7. The arc drawn in P to complete the arc γ. One the left, the case where y_1 and y_2 are on different sides of $c \cap Y$. On the right, the case where c intersects y_2.

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 boundaries of Y except y_1 and the "right" component has g_r genus and contains all right boundaries. 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 so 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 $\mathscr{C}_{\text{adm}}$ 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 $\Psi(\alpha)$ so that $\Pi(c_i)$ and $\Pi(c_{i+1})$ are $2M$-close in $\mathscr{C}_{\text{adm}}$ 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 $\mathscr{C}_{\text{adm}}$, 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. □

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 $x' \in \alpha'$ so 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 you collapse the edge of D dual to x and then expand 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 so that when you perform 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. □

Corollary 4.12. *$\mathscr{C}_{\mathrm{adm}}$ is not Gromov hyperbolic.*

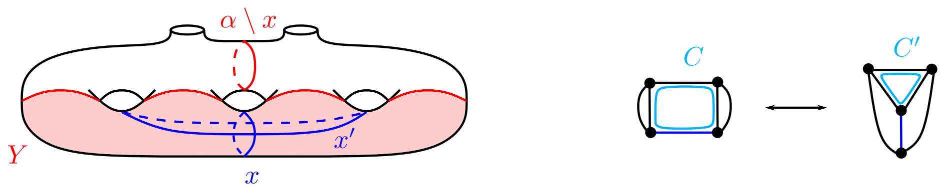

FIGURE 8. One the left, the subsurface Y where x is flipped to x'. One the right, the Whitehead move on the dual graph corresponding to flipping x to x'. The cycle C is sent to the cycle C' under this move.

Proof. Lemma 3.5 and Proposition 4.11 together show that $\mathscr{C}_{\text{adm}}$ is quasi-isometric to the hierarchically hyperbolic space $\mathcal{K}$. Since hierarchical hyperbolicity can be passed along quasi-isometries, $\mathscr{C}_{\text{adm}}$ 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 $\mathscr{C}_{\text{adm}}$ has a pair of disjoint witnesses.

Let $\Delta_1, \ldots, \Delta_b$ be the boundary curves of S. Without loss of generality, assume $\phi(\Delta_i) \geq 0$ for $i \in \{1, \ldots, k\}$ and $\phi(\Delta_i) < 0$ for $i \in \{k+1, \ldots, b\}$. Let α be a multicurve consisting of $g+1$ non-separating curves $a_1, \ldots, a_{g+1}$ so 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 so that W^+ is to the left.

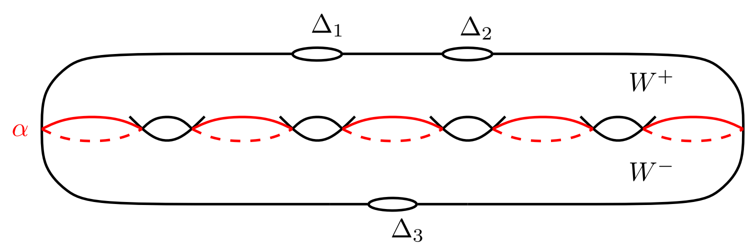

FIGURE 9. The multicurve α whose complement is a a pair of witnesses for $\mathscr{C}_{\text{adm}}$.

By homological coherence (Lemma 2.1.2), we have that for any framing ψ of S,

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

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 no subset $\mathcal{C}$ of its boundary $\alpha \cup \Delta_1 \cup \ldots \cup \Delta_k$ so that

$$\sum_{c \in \mathcal{C}} \psi(c) = 1 - |\mathcal{C}|. \tag{2}$$

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 $(x_1, \ldots, x_{g+1}) \in \mathbb{Z}^{g+1}$ so that (1) holds, there is a framing ψ of S so 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 (2) for any subset $\mathcal{C}$ of ∂W^+ or the corresponding equations for W^- since these all linearly independent from (1). 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 so that their differences are all larger than $2K$. Set x_{g+1} to satisfy (1), so it will necessarily be very negative. Then for any subset $\mathcal{C}$ of ∂W^+, the left-hand side of (2) 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 witnesses. The argument for W^- is completely analogous but with signs flipped.

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 2.2 now provides $f \in \mathrm{Mod}(S)$ so 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. A partial boundary complex for strata

In this section, we explain how the admissible curve graph can be viewed as capturing the combinatorics of a partial bordification of (marked) strata. For this section, we let $S_{g,n}$ donote the genus g surface with n marked points. $\mathcal{M}_{g,n}$ and $\mathcal{T}_{g,n}$ will denote the Moduli and Teichmüller spaces of $S_{g,n}$

5.1. Framings and strata.

Framings and strata. Let us first recall some of the results of [CS22] on the relationship between strata, markings, and framed mapping class groups.

A *stratum* of abelian differentials is a (quasi-projective) subvariety of the bundle of holomorphic abelian differentials $\Omega\mathcal{M}_g$ on genus g Riemann surfaces defined by conditioning the number and order of zeros. More explicitly, given any partition $\underline{\kappa} = (k_1, \ldots, k_n)$ of $2g-2$ into positive integers, we let $\Omega\mathcal{M}_g(\underline{\kappa}) \subset \Omega\mathcal{M}_g$ denote the stratum parametrizing pairs (X, ω) where X is a Riemann surface and ω is a holomorphic 1-form on X with n distinct zeros of orders $k_1, \ldots, k_n$. Since a holomorphic 1-form is entirely determined (up to global scaling by $\mathbb{C}^*$) by the order and position of its zeros, any stratum can be thought of as a $\mathbb{C}^*$ bundle over a subvariety of $\mathcal{M}_{g,n}$. 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.

In order to understand the connected components of preimages of strata in $\mathcal{T}_{g,n}$, 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 [LM14a, Cal20]. The main theorem of [CS22] characterizes the image of ρ for non-hyperelliptic components.

Before stating the theorem, we observe 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 (C.–Salter)**.** *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 set $\mathcal{H}_\phi$ to denote the subset of $\mathcal{T}_{g,n}$ parametrizing those marked differentials (X', ω', f') so that $1/(f')^*(\omega')$ is isotopic to ϕ.

Theorem 5.1 in particular implies that (for $g \geq 5$) any such $\mathcal{H}_\phi$ is a connected component of the preimage of $\mathcal{H}$ under the covering map $\mathcal{T}_{g,n} \to \mathcal{M}_{g,n}$.

Another consequence is the equivalence 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 is admissible with respect to the corresponding framing. Transitivity of the $\mathrm{FMod}(S, \phi)$ action on admissible curves (see Proposition 2.5) implies that every admissible curve is realized as a cylinder on some differential in $\mathcal{H}_\phi$.

Remark 5.3. In analogy with the fact that $\mathcal{M}_g$ is a $K(\pi, 1)$ for the mapping class group, Kontsevich conjectured that components of strata should be $K(\pi, 1)$'s for "some sort of mapping class group" [KZ]. This is true for hyperelliptic components as ρ is (essentially) injective, and components in genus 3 are indeed $K(\pi, 1)'s$ [LM14b], but an understanding of the fundamental groups of strata remains tantalizingly out of reach.

5.2. **Curve graphs as nerves.** 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 $\mathcal{T}_{g,n}$. This ensures, for example, that if X_i converges to a point X_∞ in $\mathcal{T}_{g,n}(\gamma)$ then the hyperbolic length of γ on X_i goes to 0, so X_i develops 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 5.4. 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 usual curve graph (with vertices for simple closed curves and edges for disjointness) is now the 1-skeleton of the nerve of the top-dimensional boundary strata of $\overline{\mathcal{T}_{g,n}}$. That is, it has a vertex for each $\mathcal{T}_{g,n}(c)$ where c is a simple closed curve, and two vertices are connected by an edge if the Weil–Petersson metric completions of $\mathcal{T}_{g,n}(c)$ and $\mathcal{T}_{g,n}(d)$ meet, which happens if and only if the curves are disjoint (this is a consequence of the collar lemma).

Degenerating cylinders. In order to define a boundary graph for (marked) strata analogous to the curve graph, we consider a (partial) bordification analogous to the augmented Teichmüller space and the Deligne–Mumford compactification. To do this, consider a stratum component $\mathcal{H}$ as a subvariety of $\mathcal{M}_{g,n}$ and then take its closure in $\overline{\mathcal{M}_{g,n}}$. Equivalently, one could take the completion $\overline{\mathcal{H}}$ of $\mathcal{H}$ with respect to the Weil–Petersson metric. This discussion can also be carried out with markings; let $\overline{\mathcal{H}_\phi}$ denote the closure of $\mathcal{H}_\phi$ in $\overline{\mathcal{T}_{g,n}}$ (equivalently, its Weil–Petersson completion).

The structure of $\partial\overline{\mathcal{H}}$ and $\partial\overline{\mathcal{H}_\phi}$ is determined by the so-called "incidence variety compactification" (IVC) [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 a 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 all of the relevant terms here, and instead direct the reader to the original papers as well as [CMZ22, Section 3]. The only fact that is relevant at the moment is that one of irreducible components of the boundary divisor of $\Xi\mathcal{H}$ corresponds to meromorphic differentials on a surface of genus $g-1$ with two glued simple poles (in the language of [BCG+19], these are 1-level graphs with a single horizontal edge). In flat-geometric terms, this degeneration is obtained by taking an embedded cylinder on a differential and then increasing its height while leaving its circumference and the rest of the surface fixed. This increases the modulus of the cylinder to ∞, so the underlying Riemann surfaces develop a node.

More generally, the boundary contains many (higher codimension) components corresponding to pinching the core curves of multiple disjoint cylinders (equivalently, 1-level graphs with multiple horizontal edges). Adjoining only these parts of $\partial\Xi\mathcal{H}$ to $\mathcal{H}$, we obtain a partial bordification $\breve{\mathcal{H}} \subset \overline{\mathcal{M}_{g,n}}$ in which only cylinders are allowed to degenerate. Lifting this to the marked stratum $\mathcal{H}_\phi$ we likewise get a partial bordification $\breve{\mathcal{H}}_\phi \subset \overline{\mathcal{T}_{g,n}}$.

We now define a graph $\mathscr{C}(\breve{\mathcal{H}}_\phi)$ that captures the intersection pattern of the boundary components of $\overline{\mathcal{T}_{g,n}}$ induced by $\breve{\mathcal{H}}_\phi$ as follows: it has a vertex for each curve a so that $\breve{\mathcal{H}}_\phi$ meets $\mathcal{T}_{g,n}(a)$ and an edge between a and b if and only if $\breve{\mathcal{H}}_\phi$ meets $\mathcal{T}_{g,n}(a \cup b)$. Equivalently, $\mathscr{C}(\breve{\mathcal{H}}_\phi)$ has a vertex for every simple closed curve c that is a cylinder on some differential in $\mathcal{H}_\phi$ and an edge between c and d if they can be simultaneously realized as disjoint cylinders on a differential in $\mathcal{H}_\phi$.

Since $\breve{\mathcal{H}}_\phi$ is defined in terms of cylinders, it follows that every vertex of $\mathscr{C}(\breve{\mathcal{H}}_\phi)$ corresponds to an admissible curve. In fact, it turns out this graph is just the admissible curve graph.

Proposition 5.5. *Let $\mathcal{H}_\phi$ be a marked non-hyperelliptic stratum component of genus ≥ 3 abelian differentials. Then $\mathscr{C}(\breve{\mathcal{H}}_\phi) = \mathscr{C}_{\mathrm{adm}}(S, \phi)$.*

Proof. We just need to prove that if a and b are disjoint admissible curves, then $\breve{\mathcal{H}}_\phi$ meets both $\mathcal{T}_{g,n}(a)$ and $\mathcal{T}_{g,n}(a \cup b)$.

Let us first prove that $\breve{\mathcal{H}}_\phi$ meets a boundary stratum corresponding to *some* admissible curve. Let $(X, \omega) \in \mathcal{H}_\phi$; then since ω is holomorphic the corresponding flat surface contains an embedded cylinder with core curve a'. By increasing the height of this cylinder while leaving the rest of the surface fixed, we can degenerate our surface, pushing it into $\mathcal{T}_{g,n}(a')$. We now use the fact that $\mathrm{FMod}(S, \phi)$ acts transitively on admissible curves (Proposition 2.5): there is some f taking a' to a, and so since f stabilizes $\breve{\mathcal{H}}_\phi$ we have that $f(\breve{\mathcal{H}}_\phi) = \breve{\mathcal{H}}_\phi$ meets $f(\mathcal{T}_{g,n}(a')) = \mathcal{T}_{g,n}(a)$.

The proof of the second statement is similar: the main difficulty is to find some pair of admissible curves a' and b' with the same topological type as a and b together with some $(X, \omega) \in \mathcal{H}_\phi$ on which a' and b' are cylinders: thus $\breve{\mathcal{H}}_\phi$ meets $\mathcal{T}_{g,n}(a' \cup b')$. This can be done by explicit construction (e.g., via the prototypes from [CS21, Section 6.3]), by plumbing a meromorphic differential with 4 simple poles, glued together in 2 pairs [BCG+18, Proposition 4.4], or by applying the main result of [MUW21]. We can then conclude using Proposition 2.5 as above. ☐

The entire completion $\overline{\mathcal{H}_\phi}$ of the marked stratum meets more than just admissible curves. In order to study the intersection pattern of its boundary, one needs to also include curves corresponding to "2-level graphs;" this has the effect of coning off regions of $\mathscr{C}_{\mathrm{adm}}(S, \phi)$. The authors will address the geometry of this graph in a future version of this paper.

References

[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

[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

[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.

[CS21] A. Calderon and N. Salter, *Higher spin mapping class groups and strata of Abelian differentials over Teichmüller space*, Adv. Math. **389** (2021), Paper no. 107926.

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

[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

[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.

[LM14a] 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.

[LM14b] ———, *The fine structure of the moduli space of abelian differentials in genus 3*, Geom. Dedicata **169** (2014), no. 1, 109–128.

[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

[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`.

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