J. Eur. Math. Soc. 25, 4719–4790 (2023) DOI 10.4171/JEMS/1290

Aaron Calderon · Nick Salter

Framed mapping class groups and the monodromy of strata of abelian differentials

Received August 12, 2020

Abstract. This paper investigates the relationship between strata of abelian differentials and various mapping class groups afforded by means of the topological monodromy representation. Building off of prior work of the authors, we show that the fundamental group of a stratum surjects onto the subgroup of the mapping class group which preserves a fixed framing of the underlying Riemann surface, thereby giving a complete characterization of the monodromy group. In the course of our proof we also show that these "framed mapping class groups" are finitely generated (even though they are of infinite index) and give explicit generating sets.

Keywords. Mapping class groups, framed surfaces, abelian differentials, strata

1. Introduction

The moduli space $\Omega\mathcal{M}_g$ of holomorphic 1-forms (*abelian differentials*) of genus g is a complex g-dimensional vector bundle over the moduli space $\mathcal{M}_g$. The complement of its zero section is naturally partitioned into *strata*, suborbifolds with fixed number and degree of zeros. Fixing a partition $\underline{\kappa} := (\kappa_1, \dots, \kappa_n)$ of $2g - 2$, we let $\Omega\mathcal{M}_g(\underline{\kappa})$ denote the stratum consisting of those pairs (X, ω) where ω is an abelian differential on $X \in \mathcal{M}_g$ with zeros of orders $\underline{\kappa}$.

As strata are quasi-projective varieties their (orbifold) fundamental groups are finitely presented. Kontsevich and Zorich famously conjectured that strata should be $K(G, 1)$'s for "some mapping class group" [20], but little progress has been made in this direction. This paper continues the work begun in [4,5], where the authors investigate these orbifold fundamental groups by means of a "topological monodromy representation."

Any (homotopy class of) loop in $\Omega\mathcal{M}_g(\underline{\kappa})$ based at (X, ω) gives rise to (the isotopy class of) a self-homeomorphism $X \to X$ which preserves $Z = \mathrm{Zeros}(\omega)$. This gives rise

Aaron Calderon: Department of Mathematics, University of Chicago, Chicago, IL 60637, USA; aaroncalderon@uchicago.edu

Nick Salter: Department of Mathematics, University of Notre Dame, Notre Dame, IN 46556, USA; nsalter@nd.edu

Mathematics Subject Classification (2020): Primary 57K20; Secondary 30F30

to the *topological monodromy representation*

$$\rho : \pi_1^{\mathrm{orb}}(\mathcal{H}) \to \mathrm{Mod}_g^n$$

where $\mathcal{H}$ is the component of $\Omega\mathcal{M}_g(\underline{\kappa})$ containing (X, ω) and Mod_g^n is the mapping class group of X relative to the set Z of n marked points.

The fundamental invariant: framings. The horizontal vector field $1/\omega$ of any $(X, \omega) \in \Omega\mathcal{M}_g(\underline{\kappa})$ defines a trivialization $\bar{\phi}$ of the tangent bundle of $X \setminus Z$, or an "absolute framing" (see §§2.1 and 6.2; the terminology reflects the finer notion of a "relative framing" to be discussed below). The mapping class group Mod_g^n generally does not preserve (the isotopy class of) this absolute framing, and its stabilizer $\mathrm{Mod}_g^n[\bar{\phi}]$ is of infinite index. On the other hand, the canonical nature of $1/\omega$ means that the image of ρ does leave some such absolute framing fixed. Our first main theorem identifies the image of the monodromy representation as the stabilizer of an absolute framing.

Theorem A. *Suppose that $g \geq 5$ and $\underline{\kappa}$ is a partition of $2g - 2$. Let $\mathcal{H}$ be a nonhyperelliptic component of $\Omega\mathcal{M}_g(\underline{\kappa})$. Then*

$$\rho(\pi_1^{\mathrm{orb}}(\mathcal{H})) \cong \mathrm{Mod}_g^n[\bar{\phi}]$$

where $\bar{\phi}$ is the absolute framing induced by the horizontal vector field of any surface in $\Omega\mathcal{M}_g(\underline{\kappa})$.

For an explanation as to why we restrict to the nonhyperelliptic components of strata, see the discussion after Theorem 7.2. The bound $g \geq 5$ is an artifact of the method of proof and can probably be relaxed to $g \geq 3$. We invoke $g \geq 5$ in Proposition 3.10, Proposition 4.1, and Lemma 5.5; it is not needed elsewhere.

Equivalently, the universal property of $\mathcal{M}_{g,n}$ implies there is a map

$$\ell : \mathcal{H} \to \mathcal{M}_{g,\underline{\kappa}}, \quad (X, \omega) \mapsto (X, \mathrm{Zeros}(\omega)).$$

where $\mathcal{M}_{g,\underline{\kappa}}$ denotes the moduli space of Riemann surfaces with n marked points labeled by $(\kappa_1, \ldots, \kappa_n)$, in which one may only permute marked points if they have the same label.[1] In this language, Theorem A characterizes the image of ℓ at the level of (orbifold) fundamental groups.

Remark 1.1. Via the mapping class group's action on relative homology, Theorem A also implicitly characterizes the image of the *homological* monodromy representation

$$\rho_H : \pi_1^{\mathrm{orb}}(\mathcal{H}) \to \mathrm{Aut}(H_1(X, Z; \mathbb{Z})).$$

In a companion paper [6], we give an explicit description of the image of ρ_H as the kernel of a certain crossed homomorphism on $\mathrm{Aut}(H_1(X, Z; \mathbb{Z}))$.

[1] Of course, $\mathcal{M}_{g,\underline{\kappa}}$ is a finite cover of $\mathcal{M}_{g,n}$, corresponding to the group $\mathrm{Sym}(\underline{\kappa})$ of label-preserving permutations.

Application: realizing curves and arcs geometrically. Using the identification of Theorem A, we can apply a framed version of the "change of coordinates" principle (see Proposition 2.15) to deduce the following characterization of which curves can be realized as the core curves of embedded cylinders. To formulate this, we observe that the data of an absolute framing $\bar{\phi}$ gives rise to a "winding number function" (also denoted $\bar{\phi}$) that sends an oriented simple closed curve to the $\mathbb{Z}$-valued holonomy of its forward-pointing tangent vector relative to the framing (§2.1).

Corollary 1.2 (cf. [5, Corollary 1.1]). *Fix $g \geq 5$ and a partition $\underline{\kappa}$ of $2g-2$. Pick some (X, ω) in a nonhyperelliptic component $\mathcal{H}$ of $\Omega\mathcal{M}_g(\underline{\kappa})$ and let $\bar{\phi}$ denote the induced (absolute) framing. Pick a nonseparating simple closed curve $c \subset X \setminus Z$. Then there is a loop $\gamma : [0,1] \to \mathcal{H}$ with $\gamma(0) = (X, \omega)$ and such that the parallel transport of c along γ is a cylinder on $\gamma(1) = (X, \omega)$ if and only if the winding number of c with respect to $\bar{\phi}$ is* 0.

Proof. The condition that $\bar{\phi}(c) = 0$ is necessary, as the core curve of a cylinder has constant slope.

To see that it is sufficient, we note that there is some cylinder on X with core curve d and the winding number of d with respect to $\bar{\phi}$ is 0. Therefore, by the framed change-of-coordinates principle (Proposition 2.15), there is some element $g \in \mathrm{Mod}_g^n[\bar{\phi}]$ taking d to c. By Theorem A, g lies in the monodromy group, so there is some $\gamma \in \pi_1^{\mathrm{orb}}(\mathcal{H})$ whose monodromy is g. This γ is the desired path. ∎

We can also deduce a complementary result for arcs using the same principle. Recall that a *saddle connection* on an abelian differential is a nonsingular geodesic segment connecting two zeros.

Corollary 1.3. *Fix $g \geq 5$ and a partition $\underline{\kappa}$ of $2g-2$. Pick some (X, ω) in a nonhyperelliptic component $\mathcal{H}$ of $\Omega\mathcal{M}_g(\underline{\kappa})$ and fix a nonseparating arc a connecting distinct zeros of Z. Then there is a path $\gamma : [0,1] \to \mathcal{H}$ with $\gamma(0) = (X, \omega)$ and such that the parallel transport of a along γ is realized as a saddle connection on $\gamma(1)$.*

The proof of this corollary uses machinery developed throughout the paper and is therefore deferred to §8.3. In §8 we collect other corollaries that we can obtain by the methods of the paper; we also give a classification of components of strata with marking data (§8.1) and we show in §8.2 that for $\mathcal{H}$ a sufficiently general stratum-component, the monodromy image $\rho(\pi_1^{\mathrm{orb}}(\mathcal{H})) \leq \mathrm{PMod}_g^n[\bar{\phi}]$ is *not* generated by shears about cylinders (here and throughout, PMod_g^n denotes the subgroup of the mapping class group of a genus g surface with n marked points each of which is fixed).

Other monodromy groups. Theorem A is a consequence of our characterization of the images of certain other monodromy representations: in Theorem 7.13, we compute the monodromy of a "prong-marked" stratum into the mapping class group $\mathrm{Mod}_{g,n}$ of a surface with boundary. Theorem 7.14 computes the monodromy of a stratum into the "pronged mapping class group," denoted $\mathrm{Mod}_{g,n}^*$, a refinement which captures the combinatorics of the zeros of the differential. In both Theorems 7.13 and 7.14 we find that the monodromy group is the stabilizer $\mathrm{Mod}_{g,n}[\phi]$, respectively $\mathrm{Mod}^*_{g,n}[\phi]$, of an appropriate "relative framing" ϕ.

A *relative framing* is an isotopy class of framing of $\Sigma_{g,n}$ where the isotopies are required to be trivial on $\partial\Sigma_{g,n}$ (see §2). To promote an absolute framing $\bar{\phi}$ to a relative framing ϕ we "blow up" the zeros of a differential (see §7.2); under this transformation, a zero of order k becomes a boundary component with winding number $-1-k$, so an element of $\Omega\mathcal{M}_g(\underline{\kappa})$ induces a relative framing ϕ on its blow-up with "signature" $-1-\underline{\kappa} := (-1-\kappa_1, \ldots, -1-\kappa_n)$ (see §2.1). Thus each boundary component has negative winding number; a framing with this property is said to be *of holomorphic type*.

Generating the framed mapping class group. The monodromy computations in Theorems A, 7.13, and 7.14 rest on a development in the theory of stabilizers of relative framings as subgroups of $\mathrm{Mod}_{g,n}$: we determine simple explicit finite generating sets.

We introduce some terminology used in the statement. Let $\mathcal{C} = \{c_1, \ldots, c_k\}$ be a collection of curves on a surface $\Sigma_{g,n}$, pairwise in minimal position, with the property that the geometric intersection number $i(c_i, c_j)$ is at most 1 for all pairs $c_i, c_j \in \mathcal{C}$. Associated to such a configuration is its *intersection graph* $\Lambda_{\mathcal{C}}$, whose vertices correspond to the elements of $\mathcal{C}$, with c_i and c_j joined by an edge whenever $i(c_i, c_j) = 1$. Such a configuration $\mathcal{C}$ *spans* $\Sigma_{g,n}$ if there is a deformation retraction of $\Sigma_{g,n}$ onto the union of the curves in $\mathcal{C}$. We say that $\mathcal{C}$ is *arboreal* if the intersection graph $\Lambda_{\mathcal{C}}$ is a tree, and *E-arboreal* if $\Lambda_{\mathcal{C}}$ moreover contains the E_6 Dynkin diagram as a subgraph. Note that an E-arboreal spanning configuration of $\Sigma_{g,n}$ necessarily contains $k = 2g + n - 1$ curves. See Figure 15 for the examples of spanning configurations we exploit in the pursuit of Theorem A.

When working with framings of meromorphic type we will need to consider sets of curves more general than spanning configurations (see the discussion in §5.6). To that end we define an *h-assemblage of type E* on $\Sigma_{g,n}$ as a set of curves $\mathcal{C} = \{c_1, \ldots, c_{2h+m-1}, c_{2h+m}, \ldots, c_{2g+n-1}\}$ such that (1) $\mathcal{C}_1 = \{c_1, \ldots, c_{2h+m-1}\}$ is an E-arboreal spanning configuration on a subsurface $S \subset \Sigma_{g,n}$ of genus $g(S) = h$ with m boundary components, (2) for $j \geq k$, let S_j denote a regular neighborhood of the curves $\{c_1, \ldots, c_j\}$; then for $j \geq 2h + m$, we require that $c_j \cap S_{j-1}$ be a single arc (possibly, but not necessarily, entering and exiting along the same boundary component of S_j), and (3) $S_{2g+n-1} = \Sigma_{g,n}$. In other words, an assemblage of type E is built from an E-arboreal spanning configuration on a subsurface by sequentially attaching (neighborhoods of) further curves, decreasing the Euler characteristic by exactly 1 at each stage but otherwise allowing the new curves to intersect individual old curves arbitrarily. Note then that an h-assemblage on $\Sigma_{g,n}$ again consists of exactly $2g + n - 1$ curves.

Theorem B. *Let $\Sigma_{g,n}$ be a surface of genus $g \geq 5$ with $n \geq 1$ boundary components.*

(I) *Suppose ϕ is a framing of $\Sigma_{g,n}$ of holomorphic type. Let $\mathcal{C} = \{c_1, \ldots, c_{2g+n-1}\}$ be an E-arboreal spanning configuration of curves on $\Sigma_{g,n}$ such that $\phi(c) = 0$ for all $c \in \mathcal{C}$. Then*

$$\mathrm{Mod}_{g,n}[\phi] = \langle T_c \mid c \in \mathcal{C} \rangle.$$

(II) *If* ϕ *is an arbitrary framing* (*of holomorphic or meromorphic type*) *and* $\mathcal{C} = \{c_1, \ldots, c_{2g+n-1}\}$ *is an* h*-assemblage of type* E *for* $h \geq 5$ *of curves such that* $\phi(c) = 0$ *for all* $c \in \mathcal{C}$, *then*

$$\mathrm{Mod}_{g,n}[\phi] = \langle T_c \mid c \in \mathcal{C} \rangle.$$

Theorem B also implies a finite generation result for stabilizers of absolute framings.

Corollary 1.4. *Let* $g, \underline{\kappa}$ *and* ϕ *be as above. Let* $\bar{\phi}$ *be the absolute framing on* Σ_g^n *obtained by shrinking the boundary components of* $\Sigma_{g,n}$ *to punctures. Then* $\mathrm{PMod}_g^n[\bar{\phi}]$ *is generated by finitely many Dehn twists.*

An explicit finite generating set for $\mathrm{PMod}_g^n[\bar{\phi}]$ is given in Corollary 6.12. In general, the set of Dehn twists described in Theorem B only generates a finite-index subgroup of $\mathrm{PMod}_g^n[\bar{\phi}]$ (Proposition 6.14).

Our methods of proof also yield a generalization of the main mapping class group-theoretic result of [5], allowing us to greatly expand our list of generating sets for "r-spin mapping class groups," the analogue of framed mapping class groups for closed surfaces (§2.1). See Corollary 3.11.

Remark 1.5. Both $\mathrm{Mod}_{g,n}[\phi]$ and $\mathrm{Mod}_g^n[\bar{\phi}]$ are of infinite index in their respective ambient mapping class groups, and so *a priori* could be infinitely generated. To the best of the authors' knowledge, Theorem B and Corollary 1.4 are the first recorded proofs that these groups are finitely generated. This is another instance of a surprising and poorly understood theme in the study of mapping class groups: *stabilizers of geometric structures often have unexpectedly strong finiteness properties*. The most famous instance of this principle is Johnson's proof that the Torelli group is finitely generated for all $g \geq 3$ [16]; this was recently and remarkably improved by Ershov–He and Church–Ershov–Putman to establish finite generation for each term in the Johnson filtration [8, 9].

Remark 1.6. As explained to the authors by Dick Hain [12], finite generation of framed mapping class groups can also be deduced from the perspective of the Torelli group. In particular, one can define a "generalized Johnson homomorphism" on the Torelli group together with a further contraction to $H_1(\Sigma_g; \mathbb{Z})$, the composition of which generalizes the Chillingworth map [7]. The intersection of the Torelli group with the framed mapping class group is then exactly the kernel of this generalized Chillingworth map (see [5, §5.4]) and so applying the results of Ershov–He and Church–Ershov–Putman [8, 9] one can deduce that the intersection of the Torelli group with $\mathrm{Mod}_{g,n}[\phi]$ is finitely generated. The framed mapping class group is therefore an extension of its action on homology (which is either $\mathrm{Sp}_{2g}(\mathbb{Z})$ or a finite-index subgroup) by another, and therefore is itself finitely generated.

We note that while this argument nicely parallels structural results for r-spin mapping class groups (see, e.g., [5, §5.4]), it alone cannot provide explicit generating sets due to the nonconstructive nature of the proofs of [8, 9].

Remark 1.7. In contemporary work [24], the second author and P. Portilla Cuadrado apply Theorem B to give a description in the spirit of Theorem A of the geometric monodromy group of an arbitrary isolated plane curve singularity as a framed mapping class group. The counterpart to Corollary 1.2 then yields an identification of the set of vanishing cycles for Morsifications of arbitrary plane curve singularities.

Context. As mentioned above, this paper serves as a sequel to [5]. The main result of that work considers a weaker version of the monodromy representation attached to a stratum of abelian differentials. In [5], we study the monodromy representation valued in the *closed* mapping class group Mod_g; here we enrich our monodromy representation so as to track the location of the zeros. There, we find that an object called an "r-spin structure" (see §2.1) governs the behavior of the monodromy representation. Here, the added structure of the locations of the zeros allows us to refine these r-spin structures to the more familiar notion of a globally invariant framing of the fibers. Where the technical core of [5] is an analysis of the group theory of the stabilizer in Mod_g of an r-spin structure, here the corresponding work is to understand these "framed mapping class groups" and to work out their basic theory, including the surprising fact that these infinite-index subgroups admit the remarkably simple finite generating sets described in Theorem B.

In recent preprints, Hamenstädt has also analyzed the monodromy representation into Mod_g^n. In [13] she gives generators for the image in terms of "completely periodic admissible configurations," which are analogous to the spanning configurations appearing in Theorem B. In [14], she identifies the image of the monodromy into the *closed* mapping class group Mod_g as the stabilizer of an "r-spin structure," recovering and extending work of the authors (see §2.1 as well as [5]). The paper [14] also contains a description of generators for the fundamental groups of certain strata.

Earlier work of Walker also bears mentioning, as it was a major source of inspiration for this result (though not for the proofs). In [29], she investigates the relationship between the monodromy image $\rho(\pi_1^{\mathrm{orb}}(\mathcal{H}))$ of a component $\mathcal{H}$ of a stratum of *quadratic* differentials and the surface braid group $SB_{g,n}$ (i.e., the kernel of the map $\mathrm{Mod}_g^n \to \mathrm{Mod}_g$). By building explicit deformations of differentials, she can in some cases identify the intersection $\rho(\pi_1^{\mathrm{orb}}(\mathcal{H})) \cap SB_{g,n}$ as the kernel of a certain "Abel–Jacobi map" (which is in turn related to the generalized Johnson homomorphism of Remark 1.6). Combining this with her previous partial results [28] on the monodromy of strata of quadratic differentials into Mod_g then yields an algebraic characterization of the monodromy image of the principal stratum of quadratic differentials (but no others).

Remark 1.8. It would be interesting to extend the methods of the current paper to the setting of quadratic differentials to recover and extend the results of Walker. Every quadratic differential defines an $\mathbb{R}P^1$-valued Gauss map, and so the general principles of this paper should be applicable to investigate the monodromy of strata of quadratic differentials. One basic obstruction to carrying this program out is our lack of knowledge of the $\mathrm{Mod}_{g,n}$ action on the associated winding number functions. A deeper and more serious obstruction is our lack of understanding of the structure of the Mod_g-stabilizer of a given root of the square of the canonical bundle (as compared with our quite detailed understanding of the stabilizers of r-spin structures; see §2.1 and [5]).

1.1. Structure of the paper

This paper is roughly divided into two parts: the first deals exclusively with relative framings on surfaces with boundary and their associated framed mapping class groups, while the second deals with variations on framed mapping class groups and their relationship with strata of abelian differentials. Readers interested only in Theorem B can read §§2–5 independently, while readers interested only in Theorem A need only read the introductory §2 together with §§6–8 (provided they are willing to accept Theorem B as a black box).

In Appendix A, we collect all of the various generating set statements obtained throughout the paper in one place. We hope that this will serve as an easy-to-use reference for future investigations of framed and r-spin mapping class groups.

Outline of the proof of Theorem B. The proof of Theorem B has two steps, which roughly parallel those of [5, Theorem B]. For the first step, we show in Proposition 3.1 that the Dehn twists on a spanning configuration of admissible curves as specified in the theorem generate the "admissible subgroup" $\mathcal{T}_\phi \leq \mathrm{Mod}_{g,n}[\phi]$ (see §2.3). The proof of this step relies on the theory of "subsurface push subgroups" from [27] and extends these results, establishing a general inductive procedure to build subsurface push subgroups from admissible twists and sub-subsurface push subgroups (Lemma 3.3).

The second step is to show that the admissible subgroup is the *entire* stabilizer of the relative framing; the proof of this step spans both §4 and §5. In [5, 27], the analogous step is accomplished using the "Johnson filtration" of the mapping class group, a strategy which does not work for surfaces with multiple boundary components. Instead, we prove that $\mathcal{T}_\phi = \mathrm{Mod}_{g,n}[\phi]$ by induction on the number of boundary components of $\Sigma_{g,n}$.

The base case of the induction (when there is a single boundary component of winding number $1-2g$) takes place in §4. Its proof relies heavily on the analysis of [5] and the relationship between framings and "r-spin structures," their analogues on closed surfaces (see the end of §2.1). Using a version of the Birman exact sequence adapted to framed mapping class groups (Lemma 4.6), we show that the equality $\mathcal{T}_\phi = \mathrm{Mod}_{g,n}[\phi]$ is equivalent to the statement that $\mathcal{T}_\phi$ contains "enough separating twists." We directly exhibit these twists in Proposition 4.1, refining [5, Lemma 6.4] and its counterpart in [27]; the reader is encouraged to think of Proposition 4.1 as the "canonical version" of this statement.

The inductive step of the proof that $\mathcal{T}_\phi = \mathrm{Mod}_{g,n}[\phi]$ is contained in §5. The overall strategy is to introduce a connected graph $\mathcal{A}^s$ on which $\mathrm{Mod}_{g,n}[\phi]$ acts vertex- and edge-transitively (see §§5.1–5.4). The heart of the argument is thus to establish these transitivity properties (Lemma 5.3) and the connectedness of $\mathcal{A}^s$ (Lemma 5.4); these both require a certain amount of care, and the arguments are lengthy. Standard techniques then imply that $\mathrm{Mod}_{g,n}[\phi]$ is generated by the stabilizer of a vertex (which we can identify with $\mathrm{Mod}_{g,n-1}[\phi']$ for some ϕ') together with an element that moves along an edge (Lemma 5.10). Applying the inductive hypothesis and explicitly understanding the action of certain Dehn twists on $\mathcal{A}^s$ together yield $\mathcal{T}_\phi = \mathrm{Mod}_{g,n}[\phi]$, completing the proof of Theorem B.

Variations on framed mapping class groups. §6 is an interlude into the theory of other framed mapping class groups. In §6.1 we introduce the theory of *pronged surfaces,* surfaces with extra tangential data which mimic the zero structure of an abelian differential. After discussing the relationship between the mapping class groups $\mathrm{Mod}^*_{g,n}$ of pronged surfaces and surfaces with boundaries or marked points, we introduce a theory of relative framings of pronged surfaces and hence a notion of framed, pronged mapping class group $\mathrm{Mod}^*_{g,n}[\phi]$. The main result of this subsection is Proposition 6.7, which exhibits $\mathrm{Mod}^*_{g,n}[\phi]$ as a certain finite extension of $\mathrm{Mod}_{g,n}[\phi]$.

We then proceed in §6.2 to a discussion of absolute framings of pointed surfaces, as at the beginning of this introduction. When a surface has marked points instead of boundary components, framings can only be considered up to absolute isotopy. Therefore, the applicable notion is not a relative but an absolute framing $\bar{\phi}$. In this section we prove Theorem 6.10, which states that the (pronged) relative framing stabilizer $\mathrm{Mod}^*_{g,n}[\phi]$ surjects onto the (pointed) absolute framing stabilizer $\mathrm{PMod}^n_g[\bar{\phi}]$. Combining this theorem with work of the previous subsection also gives explicit generating sets for $\mathrm{PMod}^n_g[\bar{\phi}]$ (see Corollary 6.12).

Outline of the proof of Theorem A. The proof of Theorem A is accomplished in §7. After recalling background material on abelian differentials (§7.1) and exploring the different sorts of framings a differential induces (§7.2), we record the definitions of certain marked strata, first introduced in [1] (§7.3). These spaces fit together in a tower of coverings (16) which evinces the structure of the pronged mapping class group, as discussed in §6.1. By a standard continuity argument, the monodromy of each covering must stabilize a framing (see Lemma 7.10 and Corollaries 7.11 and 7.12).

Using these marked strata, we can upgrade the Mod^n_g-valued monodromy of $\mathcal{H}$ into a $\mathrm{Mod}^*_{g,n}$-valued homomorphism, and passing to a certain finite cover of the stratum therefore results in a space $\mathcal{H}^{\mathrm{pr}}$ whose monodromy lies in $\mathrm{Mod}_{g,n}$. By realizing the generating set of Theorem B as cylinders on a prototype surface in $\mathcal{H}^{\mathrm{pr}}$, we can explicitly construct deformations whose monodromy is a Dehn twist, hence proving that the $\mathrm{Mod}_{g,n}$-valued monodromy group of $\mathcal{H}^{\mathrm{pr}}$ is the entire stabilizer of the appropriate framing (Theorem 7.13).

To deduce Theorem A from the monodromy result for $\mathcal{H}^{\mathrm{pr}}$ requires an understanding of the interactions between all three types of framed mapping class groups. Using the diagram of coverings (16) together with the structural results of §6, we conclude that the $\mathrm{Mod}^*_{g,n}$-valued monodromy of $\mathcal{H}$ is exactly the framing stabilizer $\mathrm{Mod}^*_{g,n}[\phi]$ (Theorem 7.14). An application of Theorem 6.10 together with a discussion of the permutation action of $\mathrm{Mod}^n_g[\phi]$ on Z finishes the proof of Theorem A.

The concluding §8 contains applications of our analysis to the classification of components of certain covers of strata (Corollaries 8.1 and 8.2) as well as to the relationship between cylinders and the fundamental groups of strata (§8.2). This section also contains the proof of Corollary 1.3.

2. Framings and framed mapping class groups

2.1. Framings

We begin by recalling the basics of framed surfaces. Our conventions ultimately follow those of Randal-Williams [26, §§1.1, 2.3], but we have made some convenient cosmetic alterations and use language compatible with our previous papers [5,27]. See Remark 2.1 below for an explanation of how to reconcile these two presentations.

Framings, (relative) isotopy. Let $\Sigma_{g,n}$ denote a compact oriented surface of genus g with $n \geq 1$ boundary components $\Delta_1, \dots, \Delta_n$. Through §5 we will work exclusively with boundary components, but in §6, we will also consider surfaces equipped with marked points. We formulate our discussion here for surfaces with boundary components; we will briefly comment on the changes necessary to work with marked points in §6.

Throughout this section we fix an orientation θ and a Riemannian metric μ of $\Sigma_{g,n}$, affording a reduction of the structure group of the tangent bundle $T\Sigma_{g,n}$ to SO(2). A *framing* of $\Sigma_{g,n}$ is an isomorphism of SO(2)-bundles

$$\phi : T\Sigma_{g,n} \to \Sigma_{g,n} \times \mathbb{R}^2.$$

With θ, μ fixed, framings are in one-to-one correspondence with nowhere-vanishing vector fields ξ; in what follows we will largely take this point of view. In this language, we say that two framings ϕ, ψ are *isotopic* if the associated vector fields ξ_ϕ and ξ_ψ are homotopic through nowhere-vanishing vector fields.

Suppose that ϕ and ψ restrict to the same framing δ of $\partial\Sigma_{g,n}$. In this case, we say that ϕ and ψ are *relatively isotopic* if they are isotopic through framings restricting to δ on $\partial\Sigma_{g,n}$. With a choice of δ fixed, we say that ϕ is a *relative framing* if ϕ is a framing restricting to δ on $\partial\Sigma_{g,n}$.

(Relative) winding number functions. Let $(\Sigma_{g,n}, \phi)$ be a (relatively) framed surface. We explain here how the data of the (relative) isotopy class of ϕ can be encoded in a topological structure known as a (*relative*) *winding number function*. Let $c : S^1 \to \Sigma_{g,n}$ be a C^1 immersion. Given two vectors $v, w \in T_x\Sigma_{g,n}$, we denote the angle (relative to the metric μ) between v, w by $\angle(v, w)$. We define the *winding number* $\phi(c)$ of c as the degree of the "Gauss map" restricted to c:

$$\phi(c) = \int_{S^1} d\angle(c'(t), \xi_\phi(c(t))).$$

The winding number $\phi(c)$ is clearly an invariant of the isotopy class of ϕ, and is furthermore an invariant of the isotopy class of c as an immersed curve in $\Sigma_{g,n}$.[2]

[2]It is not, however, an invariant of the *homotopy* class of the map c, since the winding number will change under the addition or removal of small self-intersecting loops. In this paper we will be exclusively concerned with winding numbers of *embedded* curves and arcs, so we will not comment further on this. See [15] for further discussion.

Possibly after altering ϕ by an isotopy, we can assume that each component Δ_i of $\partial\Sigma_{g,n}$ contains a point p_i such that $\xi_\phi(p_i)$ is orthogonally inward-pointing. We call such a point p_i a *legal basepoint* for Δ_i. We emphasize that even though Δ_i may contain several legal basepoints, we choose *exactly one* legal basepoint on each Δ_i, so that all arcs based at Δ_i are based at the same point.

Let $a : [0,1] \to \Sigma_{g,n}$ be a C^1 immersion with $a(0), a(1)$ equal to distinct legal basepoints p_i, p_j; assume further that $a'(0)$ is orthogonally inward-pointing and $a'(1)$ is orthogonally outward-pointing. We call such an arc *legal*. Then the winding number

$$\phi(a) := \int_0^1 d\angle(a'(t), \xi_\phi(a(t)))$$

is necessarily half-integral, and is invariant under the *relative* isotopy class of ϕ and under isotopies of a through legal arcs.

Thus a framing ϕ gives rise to an *absolute winding number function* which we denote by the same symbol. Let $\mathcal{S}$ denote the set of isotopy classes of oriented simple closed curves on $\Sigma_{g,n}$. Then the framing ϕ determines the winding number function

$$\phi : \mathcal{S} \to \mathbb{Z}, \quad x \mapsto \phi(x).$$

Likewise, let $\mathcal{S}^+$ be the set obtained from $\mathcal{S}$ by adding the set of isotopy classes of legal arcs. Then ϕ also determines a *relative winding number function*

$$\phi : \mathcal{S}^+ \to \tfrac{1}{2}\mathbb{Z}, \quad x \mapsto \phi(x).$$

Signature; holomorphic/meromorphic type. The *signature* of a framing δ of $\partial\Sigma_{g,n}$ (or of a framing ϕ of $\Sigma_{g,n}$ restricting to δ on $\partial\Sigma_{g,n}$) is the vector

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

where each Δ_i is oriented with $\Sigma_{g,n}$ lying to the left. A relative framing ϕ is said to be *of holomorphic type* if $\delta(\Delta_i) \le -1$ for all i and is *of meromorphic type* otherwise. In §7.2 we will see that if ω is an abelian differential on a Riemann surface X, then the relative framing induced by ω is indeed of holomorphic type. Given a partition $\underline{\kappa} = (\kappa_1, \ldots, \kappa_n)$ of $2g-2$, we say that a relative framing ϕ has signature $-1-\underline{\kappa}$ if the boundary components have signatures $(-1-\kappa_1, \ldots, -1-\kappa_n)$.

Remark 2.1. For the convenience of the reader interested in comparing the statements of this section with their counterparts in [26], we briefly comment on the places where the two expositions diverge. We have used the term "framing" where Randal-Williams uses "θ^r-structure" with $r = 0$, and we use the term "(relative) winding number function" where Randal-Williams uses an equivalent structure denoted "q_ξ." Randal-Williams also adopts some different normalization conventions. If x is a curve, then $q_\xi(x) = \phi(x) - 1$, and if x is an arc, $q_\xi(x) = \phi(x) - 1/2$ (in particular, q_ξ is integer-valued on arcs).

The lemma below allows us to pass between framings and winding number functions. Its proof is a straightforward exercise in differential topology (see [26, Proposition 2.4]).

Lemma 2.2. *Let $\Sigma_{g,n}$ be a surface with $n \geq 1$ boundary components and let ϕ and ψ be framings. Then ϕ and ψ are isotopic as framings if and only if the associated absolute winding number functions are equal. If $\phi|_{\partial\Sigma_{g,n}} = \psi|_{\partial\Sigma_{g,n}}$, then ϕ and ψ are relatively isotopic if and only if the associated relative winding number functions are equal.*

Moreover, if $\Sigma_{g,n}$ is endowed with the structure of a CW complex for which each 0*-cell is a legal basepoint and each* 1*-cell is either isotopic to a simple closed curve or a legal arc, then ϕ and ψ are* (*relatively*) *isotopic if and only if the* (*relative*) *winding numbers of each* 1*-cell are equal.*

Remark 2.3. Following Lemma 2.2, we will be somewhat lax in our terminology. Often we will use the term "(relative) framing" to refer to the entire (relative) isotopy class, or else conflate the (relative) framing with the associated (relative) winding number function.

Properties of (relative) winding number functions. The terminology of "winding number function" originates with the work of Humphries and Johnson [15] (although we are discussing what they call *generalized* winding number functions). We recall here some properties of winding number functions which they identified.[3]

Lemma 2.4. *Let ϕ be a relative winding number function on $\Sigma_{g,n}$ associated to a relative framing of the same name. Then ϕ satisfies the following properties.*

(1) (Twist-linearity) *Let $a \subset \Sigma_{g,n}$ be a simple closed curve, oriented arbitrarily. Then for any $x \in \mathcal{S}^+$,*
$$\phi(T_a(x)) = \phi(x) + \langle x, a\rangle \phi(a),$$
where
$$\langle \cdot, \cdot \rangle : H_1(\Sigma_{g,n}, \partial\Sigma_{g,n}; \mathbb{Z}) \times H_1(\Sigma_{g,n}; \mathbb{Z}) \to \mathbb{Z}$$
denotes the relative algebraic intersection pairing.

(2) (Homological coherence) *Let $S \subset \Sigma_{g,n}$ be a subsurface with boundary components $c_1, \ldots, c_k$, oriented so that S lies to the left of each c_i. Then*
$$\sum_{i=1}^{k} \phi(c_i) = \chi(S),$$
where $\chi(S)$ denotes the Euler characteristic.

Functoriality. In the body of the argument we will have occasion to consider maps between surfaces equipped with framings and related structures. We record here some relatively simple observations about this. Firstly, as we have already implicitly used, if $S \subset \Sigma_{g,n}$ is a subsurface, any (isotopy class of) framing ϕ of $\Sigma_{g,n}$ restricts to an (isotopy class of) framing of S. There is a converse as well; the proof is an elementary exercise in differential topology.

[3] In the nonrelative setting.

Lemma 2.5. *Let $S \subset \Sigma_{g,n}$ be a subsurface, and let ϕ be a framing of S. Enumerate the components of $\Sigma_{g,n} \setminus S$ as $S_1, \dots, S_k$. Call such a component S_i* relatively closed *if $\partial S_i \subset \partial S$. Then ϕ extends to a framing $\widetilde{\phi}$ of $\Sigma_{g,n}$ if and only if for each relatively closed component S_i, there is an equality*

$$\sum_{c \text{ a component of } \partial S_i} \phi(c) = \chi(S_i),$$

with each c oriented with S_i to the left.

In particular, suppose that $S = \Sigma_{g,n} \setminus D$, where D is an embedded disk, and let ϕ be a framing of S. Then ϕ extends over D to give a framing of $\Sigma_{g,n}$ if and only if $\phi(\partial D) = 1$ when ∂D is oriented with D to the left.

Closed surfaces; r-spin structures. For $g \geq 2$, the closed surface Σ_g does not admit any nonvanishing vector fields, but there is a "mod r analogue" of a framing called an *r-spin structure*. As r-spin structures will play only a passing role in the arguments of this paper (see §4.2), we present here only the bare bones of the theory. See [27, §3] for a much more complete discussion.

Definition 2.6. Let Σ_g be a closed surface, and as above, let $\mathcal{S}$ denote the set of oriented simple closed curves on Σ_g. An *r-spin structure* is a function

$$\widehat{\phi} : \mathcal{S} \to \mathbb{Z}/r\mathbb{Z}$$

that satisfies the twist-linearity and homological coherence properties of Lemma 2.4.

Above we saw how a nonvanishing vector field ξ_ϕ on $\Sigma_{g,n}$ gives rise to a winding number function ϕ on $\Sigma_{g,n}$. Suppose now that ξ is an arbitrary vector field on Σ_g with isolated zeros $p_1, \dots, p_k$. For $i = 1, \dots, k$, let γ_i be a small embedded curve encircling p_i (oriented with p_i to the left) and define

$$r = \gcd\{\mathrm{WN}_\xi(\gamma_i) - 1 \mid 1 \leq i \leq k\},$$

where $\mathrm{WN}_\xi(\gamma_i)$ means to take the $\mathbb{Z}$-valued winding number of γ_i viewed as a curve on $\Sigma_g \setminus \{p_1, \dots, p_n\}$ endowed with the framing given by ξ. Then it can be shown (see [15, §1]) that the map

$$\widehat{\phi} : \mathcal{S} \to \mathbb{Z}/r\mathbb{Z}, \quad \phi(x) = \mathrm{WN}_\xi(x) \pmod r,$$

determines an r-spin structure.

2.2. *The action of the mapping class group*

Recall that the mapping class group $\mathrm{Mod}_{g,n}$ of $\Sigma_{g,n}$ is the set of isotopy classes of self-homeomorphisms of $\Sigma_{g,n}$ which restrict to the identity on $\partial\Sigma_{g,n}$. Therefore $\mathrm{Mod}_{g,n}$ acts on the set of relative (isotopy classes of) framings, and hence the set of relative winding number functions, by pull-back. As we require $\mathrm{Mod}_{g,n}$ to act from the left, there is the formula

$$(f \cdot \phi)(x) = \phi(f^{-1}(x)).$$

We recall here the basic theory of this action, as developed by Randal-Williams [26, §2.4]. Throughout this section, we fix a framing δ of $\partial \Sigma_{g,n}$ and may therefore choose a legal basepoint on each boundary component once and for all.

The (generalized) Arf invariant. The $\mathrm{Mod}_{g,n}$ orbits of relative framings are classified by a generalization of the classical Arf invariant. To define this, we introduce the notion of a *distinguished geometric basis* for $\Sigma_{g,n}$. For $i = 1, \dots, n$, let p_i be a legal basepoint on the ith boundary component of $\Sigma_{g,n}$. A *distinguished geometric basis* is a collection

$$\mathcal{B} = \{x_1, y_1, \dots, x_g, y_g\} \cup \{a_2, \dots, a_n\}$$

of $2g$ oriented simple closed curves $x_1, \dots, y_g$ and $n-1$ legal arcs $a_2, \dots, a_n$ that satisfy the following intersection properties:

(1) $\langle x_i, y_i \rangle = i(x_i, y_i) = 1$ (here $i(\cdot,\cdot)$ denotes the geometric intersection number) and all other pairs of elements of $\{x_1, \dots, y_g\}$ are disjoint.

(2) Each arc a_i is a legal arc running from p_1 to p_i and is disjoint from all curves $x_1, \dots, y_g$.

(3) The arcs a_i, a_j are disjoint except at the common endpoint p_1.

Remark 2.7. A distinguished geometric basis can easily be used to determine a CW-structure on $\Sigma_{g,n}$ satisfying the hypotheses of Lemma 2.2. In particular, a (relative) winding number function (and hence the associated (relative) isotopy class of framing) is determined by its values on a distinguished geometric basis. Moreover, for any vector $(w_1, \dots, w_{2g+n-1}) \in \mathbb{Z}^{2g} \times (\mathbb{Z} + \frac{1}{2})^{n-1}$, there exists a framing ϕ of $\Sigma_{g,n}$ realizing the values $(w_1, \dots, w_{2g+n-1})$ on a chosen distinguished geometric basis (this is a straightforward construction).

Let $\mathcal{B}$ be a distinguished geometric basis for the framed surface $(\Sigma_{g,n}, \phi)$. The *Arf invariant* of ϕ relative to $\mathcal{B}$ is the element of $\mathbb{Z}/2\mathbb{Z}$ given by

$$\mathrm{Arf}(\phi, \mathcal{B}) = \sum_{i=1}^{g} (\phi(x_i) + 1)(\phi(y_i) + 1) + \sum_{j=2}^{n} \big(\phi(a_j) + \tfrac{1}{2}\big)(\phi(\Delta_j) + 1) \pmod 2 \tag{1}$$

(compare to [26, (2.4)]).

Lemma 2.8 (see [26, Proposition 2.8]). *Let $(\Sigma_{g,n}, \phi)$ be a framed surface. Then the Arf invariant* $\mathrm{Arf}(\phi, \mathcal{B})$ *is independent of the choice of distinguished geometric basis $\mathcal{B}$.*

Remark 2.9. We caution the reader that while the Arf invariant does not depend on the choice of basis, it *does* depend on the choice of legal basepoints on each boundary component. Since $\mathrm{Mod}_{g,n}$ fixes the boundary pointwise, it preserves our choice of legal basepoint on each boundary component.

Following Lemma 2.8, we write $\mathrm{Arf}(\phi)$ to indicate the Arf invariant $\mathrm{Arf}(\phi, \mathcal{B})$ computed on an arbitrary choice of distinguished geometric basis. The next result shows that for $g \geq 2$, the Arf invariant classifies $\mathrm{Mod}_{g,n}$ orbits of relative isotopy classes of framings.

Proposition 2.10 (cf. [26, Theorem 2.9]). *Let $g \geq 2$ and $n \geq 1$ be given, and let ϕ, ψ be two relative framings of $\Sigma_{g,n}$ which agree on $\partial \Sigma_{g,n}$. Then there is an element $f \in \mathrm{Mod}_{g,n}$ such that $f \cdot \psi = \phi$ if and only if* $\mathrm{Arf}(\phi) = \mathrm{Arf}(\psi)$.

For surfaces of genus 1, the action is more complicated. In this work we will only need to study the case of one boundary component; this was treated by Kawazumi [19]. Let $(\Sigma_{1,1}, \phi)$ be a framed surface. Consider the set

$$\mathfrak{a}(\phi) = \{\phi(x) \mid x \subset \Sigma_{1,1} \text{ an oriented s.c.c.}\}.$$

The twist-linearity formula (Lemma 2.4 (1)) implies that $\mathfrak{a}(\phi)$ is in fact an ideal of $\mathbb{Z}$. We define the *genus*-1 *Arf invariant* of ϕ to be the unique nonnegative integer $\mathrm{Arf}_1(\phi) \in \mathbb{Z}_{\geq 0}$ such that

$$\mathfrak{a}(\phi) = \mathrm{Arf}_1(\phi)\mathbb{Z}. \tag{2}$$

Remark 2.11. The normalization conventions for Arf as in (1) and Arf_1 as in (2) are different. In §4.1 we will reconcile them, but here we chose to present the "natural" definition of Arf_1.

Lemma 2.12 (cf. Theorem 0.3 of [19]). *Let ϕ and ψ be relative framings of $\Sigma_{1,1}$. Then there is $f \in \mathrm{Mod}_{1,1}$ such that $f \cdot \psi = \phi$ if and only if* $\mathrm{Arf}_1(\phi) = \mathrm{Arf}_1(\psi)$.

2.3. *Framed mapping class groups*

Having studied the orbits of the $\mathrm{Mod}_{g,n}$ action on the set of framings in the previous section, we turn now to the stabilizer of a framing.

Definition 2.13 (Framed mapping class group). Let $(\Sigma_{g,n}, \phi)$ be a (relatively) framed surface. The *framed mapping class group* $\mathrm{Mod}_{g,n}[\phi]$ is the stabilizer of the relative isotopy class of ϕ:

$$\mathrm{Mod}_{g,n}[\phi] = \{f \in \mathrm{Mod}_{g,n} \mid f \cdot \phi = \phi\}.$$

Remark 2.14. We pause here to note one somewhat counterintuitive property of *relatively* framed mapping class groups. Suppose that ϕ, ϕ' are distinct as relative isotopy classes of framings, but are equal as absolute framings (in terms of relative winding number functions, this means that ϕ and ϕ' agree when restricted to the set of simple closed curves but assign different values to arcs). Then the associated relatively framed mapping classes are *equal*: $\mathrm{Mod}_{g,n}[\phi] = \mathrm{Mod}_{g,n}[\phi']$. This is not hard to see: allowing the framing on the boundary to move under isotopy changes the winding numbers of all arcs in the same way, so that the ϕ'-winding number of an arc can be computed from the ϕ-winding number by adjusting by a universal constant. Necessarily then $\phi(f(\alpha)) = \phi(\alpha)$ for a mapping class f and an arc α if and only if $\phi'(f(\alpha)) = \phi(\alpha)$.

Admissible curves, admissible twists, and the admissible subgroup. In our study of $\mathrm{Mod}_{g,n}[\phi]$, a particularly prominent role will be played by the Dehn twists that preserve ϕ. An *admissible curve* on a framed surface $(\Sigma_{g,n}, \phi)$ is a nonseparating simple closed curve a such that $\phi(a) = 0$. It follows from the twist-linearity formula (Lemma 2.4 (1)) that the associated Dehn twist T_a preserves ϕ. We call the mapping class T_a an *admissible twist*. Finally, we define the *admissible subgroup* to be the group generated by all admissible twists:

$$\mathcal{T}_\phi := \langle T_a \mid a \text{ admissible for } \phi \rangle.$$

Change of coordinates for framed surfaces. The classical "change-of-coordinates principle" for surfaces is a body of techniques for constructing special configurations of curves and subsurfaces on a fixed surface (see [11, §1.3]). The underlying mechanism is the classification of surfaces, which provides a homeomorphism between a given surface and a "reference surface;" if a desired configuration exists on the reference surface, then the configuration can be pulled back along the classifying homeomorphism.

A similar principle exists for framed surfaces, governing when configurations of curves with prescribed winding numbers exist on framed surfaces. The classification results of Proposition 2.10 and Lemma 2.12 assert that the Arf invariant provides the only obstruction to constructing desired configurations of curves in the presence of a framing. We will make extensive and often tacit use of the "framed change-of-coordinates principle" throughout the body of the argument. Here we will illustrate some of the more frequent instances of which we avail ourselves. Recall that a *k-chain* is a sequence $c_1, \dots, c_k$ of curves such that $i(c_i, c_{i+1}) = 1$ for $i = 1, \dots, k-1$, and $i(c_i, c_j) = 0$ for $|i - j| \geq 2$.

Proposition 2.15 (Framed change of coordinates). *Let $(\Sigma_{g,n}, \phi)$ be a relatively framed surface with $g \geq 2$ and $n \geq 1$. A configuration $x_1, \dots, x_k$ of curves and/or arcs with prescribed intersection pattern and winding numbers $\phi(x_i) = s_i$ exists if and only if*

(a) *a configuration $\{x_1', \dots, x_k'\}$ of the prescribed topological type exists in the "unframed" setting where the values $\phi(x_i')$ are allowed to be arbitrary,*

(b) *there exists some framing ψ such that $\psi(x_i) = s_i$ for all i,*

(c) *if* $\mathrm{Arf}(\psi)$ *is determined by the constraints of* (b)*, then* $\mathrm{Arf}(\phi) = \mathrm{Arf}(\psi)$.

In particular:

(1) *For $s \in \mathbb{Z}$ arbitrary, there exists a nonseparating curve $c \subset \Sigma_{g,n}$ with $\phi(c) = s$.*

(2) *For $n = 1$, there exists a $2g$-chain of admissible curves on $\Sigma_{g,1}$ if and only if the pair $(g \pmod 4, \mathrm{Arf}(\phi)) \in (\mathbb{Z}/4\mathbb{Z}, \mathbb{Z}/2\mathbb{Z})$ is one of the four listed below:*

$$(0,0), (1,1), (2,1), (3,0). \tag{3}$$

Such a chain is called a maximal *chain of admissible curves.*

Proof. We will prove (1) and (2), from which it will be clear how the general argument works. We begin with (1). Let $\mathcal{B} = \{x_1 \dots, y_g, a_2, \dots, a_n\}$ be a distinguished geometric basis. Following Remark 2.7, a relative framing ψ of $\Sigma_{g,n}$ can be constructed by (freely) specifying the values $\psi(b)$ for each element $b \in \mathcal{B}$. Set $\psi(x_1) = s$ and let $\psi(y_1)$ be arbitrary. Since $g \geq 2$, it is possible to choose the values of $\psi(x_2), \psi(y_2)$ such that $\mathrm{Arf}(\psi) = \mathrm{Arf}(\phi)$. By Proposition 2.10, there exists a diffeomorphism $f : \Sigma_{g,n} \to \Sigma_{g,n}$ such that $f \cdot \psi = \phi$. We see that $f(x_1)$ is the required curve:

$$\phi(f(x_1)) = (f^{-1} \cdot \phi)(c_1) = \psi(x_1) = s$$

as required.

For (2), consider a maximal chain $a_1, \dots, a_{2g}$ on $\Sigma_{g,1}$. Define $b_1 := a_1$ and choose curves $b_2, b_3, \dots, b_g$, each disjoint from all a_j with j odd, such that

$$\mathcal{B} = \{b_1, a_2, b_2, a_4, \dots, b_g, a_{2g}\}$$

is a distinguished geometric basis. We now construct a framing ψ such that each a_i is admissible. By construction, the curves b_k, a_{2k+1}, b_{k+1} form pairs of pants for each $1 \leq k \leq g-1$. By the homological coherence property (Lemma 2.4 (2)), if each a_i is to be admissible, we must have $\psi(b_k) = 1 - k$ for $1 \leq k \leq g$ when b_k is oriented so that the pair of pants cobounded by b_{k-1} and a_{2k-1} lies to the left. $\mathrm{Arf}(\psi)$ is determined by these conditions, and is computed to be

$$\mathrm{Arf}(\psi) = \begin{cases} 0, & g \equiv 0, 3 \pmod 4, \\ 1, & g \equiv 1, 2 \pmod 4. \end{cases}$$

If the pair $(g, \mathrm{Arf}(\phi))$ is one of those listed in (3), then by Proposition 2.10, there exists $f : \Sigma_{g,1} \to \Sigma_{g,1}$ such that $f \cdot \psi = \phi$. As above, we find that $f(a_1), \dots, f(a_{2g})$ is the required maximal chain of admissible curves. Conversely, if $(g, \mathrm{Arf}(\phi))$ does not appear in (3), then the Arf invariant of ϕ obstructs the existence of a maximal chain of admissible curves. ∎

3. Finite generation of the admissible subgroup

Theorem B asserts that the framed mapping class group $\mathrm{Mod}_{g,n}[\phi]$ is generated by any spanning configuration $\mathcal{C}$ of admissible Dehn twists so long as the intersection graph $\Lambda_{\mathcal{C}}$ is a tree containing E_6 as a subgraph (recall the definition of "spanning configuration" prior to the statement of Theorem B). In this section, we take the first step to establishing this result. Proposition 3.1 establishes that such a configuration of twists generates the admissible subgroup. In the subsequent sections we will show that $\mathcal{T}_\phi = \mathrm{Mod}_{g,n}[\phi]$, establishing Theorem B.

Recall (see the discussion preceding Theorem B) that a collection $\mathcal{C}$ of curves is said to be an *E-arboreal spanning configuration* if each pair of curves intersects at most once, and the intersection graph is a tree containing E_6 as a subgraph.

Proposition 3.1 (Generating the admissible subgroup). *Let $\Sigma_{g,n}$ be a surface of genus $g \geq 5$ with $n \geq 1$ boundary components, and let ϕ be a framing of holomorphic type. Let $\mathcal{C}$ be an E-arboreal spanning configuration of admissible curves on $\Sigma_{g,n}$, and define*

$$\mathcal{T}_{\mathcal{C}} := \langle T_c \mid c \in \mathcal{C} \rangle.$$

Then $\mathcal{T}_{\mathcal{C}} = \mathcal{T}_\phi$.

The proof of Proposition 3.1 closely follows the approach developed in [27]. The heart of the argument (Lemma 3.6) is to show that our finite collection of twists generates a version of a point-pushing subgroup for a subsurface. This will allow us to express all admissible twists supported on this subsurface with our finite set of generators. Having shown this, we can import our method from [27] (appearing below as Proposition 3.10) which allows us to propagate this argument across the set of subsurfaces, proving the result.

3.1. *Framed subsurface push subgroups*

Let $S \subset \Sigma_{g,n}$ be a subsurface and suppose $\Delta \subset \partial S$ is a boundary component. Let $\overline{S}$ denote the surface obtained from S by capping Δ with a disk, and let $UT\overline{S}$ denote the associated unit tangent bundle. Recall the *disk–pushing homomorphism* $P : \pi_1(UT\overline{S}) \to \mathrm{Mod}(S)$ [11, Section 4.2.5]. The inclusion $S \hookrightarrow \Sigma_{g,n}$ induces a homomorphism $i : \mathrm{Mod}(S) \to \mathrm{Mod}(\Sigma_{g,n})$ which restricts to give a *subsurface push homomorphism* $\mathcal{P} := i \circ P$:

$$\mathcal{P} : \pi_1(UT\overline{S}) \to \mathrm{Mod}(\Sigma_{g,n}).$$

The *framed subsurface push subgroup* $\widetilde{\Pi}(S)$ is the intersection of this image with $\mathrm{Mod}_{g,n}[\phi]$:

$$\widetilde{\Pi}(S) := \mathrm{Im}(\mathcal{P}) \cap \mathrm{Mod}_{g,n}[\phi].$$

Note that $\widetilde{\Pi}(S)$ is defined relative to the boundary component Δ, suppressed in the notation. In practice, the choice of Δ will be clear from context.

There is an important special case of the above construction. Let $b \subset \Sigma_{g,n}$ be an oriented nonseparating curve satisfying $\phi(b) = -1$. The subsurface $\Sigma_{g,n} \setminus \{b\}$ has a distinguished boundary component Δ corresponding to the left-hand side of b. For this choice of (S, Δ), we streamline notation, defining

$$\widetilde{\Pi}(b) := \widetilde{\Pi}(\Sigma_{g,n} \setminus b)$$

(constructed relative to Δ). As $\phi(\Delta) = -1$ (oriented so that S lies to the left), it follows from Lemma 2.5 that the framing of S can be extended over the capping disk to $\overline{S}$. Such a framing of $\overline{S}$ gives rise to a section $s : \overline{S} \to UT\overline{S}$, and hence a splitting $s_* : \pi_1(\overline{S}) \to \pi_1(UT\overline{S})$.

Lemma 3.2. *If $\phi(\Delta) = -1$, then*

$$\widetilde{\Pi}(S) = \mathcal{P}(s_*(\pi_1(\overline{S}))).$$

Proof. Let $x_1, \dots, x_n \in \pi_1(\overline{S})$ be a system of generators such that each x_i is represented by a simple based loop on $\overline{S}$. Under $\mathcal{P} \circ s_*$, each such x_i is sent to a multitwist:

$$\mathcal{P}(s_*(x_i)) = T_{x_i^L} T_{x_i^R}^{-1} T_\Delta^{\phi(x_i^L)},$$

where the curves $x_i^L, x_i^R \subset S$ are characterized by the following two conditions:

(1) $x_i^L \cup x_i^R \cup \Delta$ form a pair of pants (necessarily lying to the left of Δ),

(2) x_i^L (resp. x_i^R) lies to the left (resp. right) of x_i as a based oriented curve.

Note that the assumption that $\phi(\Delta) = -1$ plus homological coherence (Lemma 2.4 (2)) implies that $\phi(x_i^L) = \phi(x_i^R)$. By the twist-linearity formula (Lemma 2.4 (1)), we then have, for any curve or arc y on $\Sigma_{g,n}$,

$$\begin{aligned}\phi(\mathcal{P}(s_*(x_i))(y)) &= \phi(y) + \langle y, x_i^L\rangle\phi(x_i^L) - \langle y, x_i^R\rangle\phi(x_i^R) + \phi(x_i^L)\langle y, \Delta\rangle\phi(\Delta)\\ &= \phi(y) + \phi(x_i^L)(\langle y, x_i^L\rangle - \langle y, x_i^R\rangle - \langle y, b\rangle)\\ &= \phi(y) + \phi(x_i^L)\cdot 0 = \phi(y)\end{aligned}$$

where the last line follows because $[b] + [x_i^R] - [x_i^L] = 0$ in the homology of $\Sigma_{g,n}$. Thus $\mathcal{P}(s_*(x_i))$ preserves ϕ. As the set of x_i generates $\pi_1(\overline{S})$, it follows that $\mathcal{P}(s_*(\pi_1(\overline{S}))) \le \widetilde{\Pi}(S)$.

To establish the opposite containment, we recall that s_* gives a splitting of the sequence

$$1 \to \mathbb{Z} \to \pi_1(UT\overline{S}) \xrightarrow{p_*} \pi_1(\overline{S}) \to 1$$

and so it suffices to show that $\widetilde{\Pi}(S) \cap \mathcal{P}(\ker(p_*)) = \{e\}$. Under $\mathcal{P}$, the generator of $\ker p_*$ is sent to T_Δ. As $\phi(\Delta) = -1$ and Δ was constructed by cutting along the nonseparating curve $b \subset \Sigma_{g,n}$, the twist-linearity formula shows that $\langle T_\Delta\rangle \cap \mathrm{Mod}_{g,n}[\phi] = \{e\}$ and the result follows. ∎

3.2. *Generating framed push subgroups*

We want to show that our finitely-generated subgroup $\mathcal{T}_\mathcal{C}$ contains a framed subsurface push subgroup $\widetilde{\Pi}(S)$ for a subsurface S that is "as large as possible." In Lemma 3.3 below, we show that this can be accomplished inductively by successively showing containments $\widetilde{\Pi}(S_i) \le \mathcal{T}_\mathcal{C}$ for an increasing union of subsurfaces $\cdots \subset S_i \subset S_{i+1} \subset \cdots$.

Lemma 3.3. *Let $S \subset \Sigma_{g,n}$ be a subsurface and let Δ be a boundary component of S such that $\phi(\Delta) = -1$, giving rise to the associated framed subsurface push subgroup $\widetilde{\Pi}(S)$. Let $a \subset \Sigma_{g,n}$ be an admissible curve disjoint from Δ such that $a \cap S$ is a single essential arc (it does not matter if a enters and exits S by the same or by different boundary components). Let $a' \subset S$ be an admissible curve satisfying $i(a, a') = 1$. Let S^+ be the subsurface given by a regular neighborhood of $S \cup a$. Then $\widetilde{\Pi}(S^+) \le \langle T_a, T_{a'}, \widetilde{\Pi}(S)\rangle$.*

Proof. Let $a'' \subset S^+$ be a curve such that

- $a \cup a'' \cup \Delta$ forms a pair of pants,

- $a'' \cap S$ is a single arc based at the same boundary component of S as $a \cap S$,
- a'' meets a' exactly once.

Such a curve can be constructed by taking the connected sum of a and Δ along an arc disjoint from a', for example (see Figure 12 and the accompanying discussion). By homological coherence (Lemma 2.4 (2)), a'' is admissible. Observe that $P_a := T_a T_{a''}^{-1} \in \widetilde{\Pi}(S^+)$.

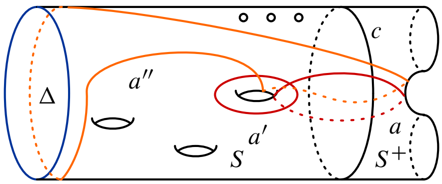

Fig. 1. The configuration discussed in Lemma 3.3, shown in the case where a enters and exits along the same boundary component c.

Moreover, we observe that since a' meets the pair of pants bounded by $a \cup a'' \cup \Delta$ in exactly one arc, either $P_a(a')$ or $P_a^{-1}(a')$ is contained in S; compare Figure 1. Without loss of generality, suppose it is $P_a(a')$. Now by the braid relation, we have $T_a T_{a'}(a) = a' \subset S$ and likewise

$$T_a T_{a'}(a'') = (T_a T_{a''}^{-1}) T_{a''} T_{a'}(a'') = P_a(a') \subset S.$$

Thus, we see that applying $T_a T_{a'}$ takes both a and a'' to admissible curves on S. In the case that $P_a^{-1}(a') \subset S$, a similar calculation can be performed to show that $T_a^{-1} T_{a'}^{-1}$ takes both a and a'' to admissible curves on S. Therefore,

$$(T_a T_{a'}) P_a (T_a T_{a'})^{-1} \in \widetilde{\Pi}(S)$$

and consequently $P_a \in \langle T_a, T_{a'}, \widetilde{\Pi}(S)\rangle$. Let $x_1, \dots, x_k \in \widetilde{\Pi}(S)$ be a generating set. The inclusion $S \hookrightarrow S^+$ induces an inclusion $\widetilde{\Pi}(S) \hookrightarrow \widetilde{\Pi}(S^+)$, and $\widetilde{\Pi}(S^+)$ is generated by $x_1, \dots, x_k$ and P_a (or P_a^{-1}). These elements are all contained in the group $\langle T_a, T_{a'}, \widetilde{\Pi}(S)\rangle$. ∎

Generation via networks. The inductive criterion of Lemma 3.3 leads to the notion of a *network*, which is a configuration of curves designed so that Lemma 3.3 can be repeatedly applied. Here we discuss the basic theory.

Definition 3.4 (Networks). Let S be a surface of finite type. For the purposes of the definition, punctures and boundary components are interchangeable; we convert both into boundary components. A *network* on S is any collection $\mathcal{N} = \{a_1, \dots, a_n\}$ of simple closed curves (*not* merely isotopy classes) on S such that $\#(a_i \cap a_j) \le 1$ for all pairs of curves $a_i, a_j \in \mathcal{N}$, and such that there are no triple intersections.

A network $\mathcal{N}$ has an associated *intersection graph* $\Lambda_{\mathcal{N}}$, whose vertices correspond to curves $x \in \mathcal{N}$, with vertices x, y adjacent if and only if $\#(x \cap y) = 1$. A network is said to be *connected* if $\Lambda_{\mathcal{N}}$ is connected, and *arboreal* if $\Lambda_{\mathcal{N}}$ is a tree. A network is *filling* if

$$S \setminus \bigcup_{a \in \mathcal{N}} a$$

is a disjoint union of disks and boundary-parallel annuli.

A network $\mathcal{N}$ determines a subgroup $\mathcal{T}_{\mathcal{N}} \le \mathrm{Mod}(\Sigma_{g,n})$ by taking the group generated by the Dehn twists about curves in $\mathcal{N}$:

$$\mathcal{T}_{\mathcal{N}} := \langle T_a \mid a \in \mathcal{N} \rangle.$$

The following appears in slightly modified form as [27, Lemma 9.4].

Lemma 3.5. *Let $S \subset \Sigma_{g,n}$ be a subsurface with a boundary component Δ satisfying $\phi(\Delta) = -1$. Let $\mathcal{N}$ be a network of admissible curves on S that is connected, arboreal, and filling, and suppose that there exist $a, a' \in \mathcal{N}$ such that $a \cup a' \cup \Delta$ forms a pair of pants. Then $\widetilde{\Pi}(S) \le \mathcal{T}_{\mathcal{N}}$.*

3.3. The key lemma

Proposition 3.10, to be stated below, gives a criterion for a group H to contain the admissible subgroup $\mathcal{T}_{\phi}$. It asserts that containing a framed subsurface push subgroup of the form $\widetilde{\Pi}(b)$ is "nearly sufficient." In preparation for this, we show here that $\mathcal{T}_{\mathcal{C}}$ contains such a subgroup. Ideally, we would like to use the network generation criterion (Lemma 3.5), but the configuration $\mathcal{C}$ does not satisfy the hypotheses and so more effort is required.

Lemma 3.6. *Let $\mathcal{C}$ be an E-arboreal spanning configuration of admissible curves, and let b be the curve indicated in Figure 2. Then $\widetilde{\Pi}(b) \le \mathcal{T}_{\mathcal{C}}$.*

We note first that by homological coherence, $\phi(b) = -1$ when oriented as in Figure 2. Our overall strategy is to iteratively apply Lemma 3.3, but the fact that b intersects some of the curves of $\mathcal{C}$ presents some difficulties.

The lemma will be proved in four steps. In each stage, we will consider a subconfiguration $\mathcal{C}_k \subset \mathcal{C}$ and the associated subsurface S_k spanned by these curves. We define $S'_k \subset S_k$ by removing a regular neighborhood of b, and we show that $\widetilde{\Pi}(S'_k) \le \mathcal{T}_{\mathcal{C}}$.

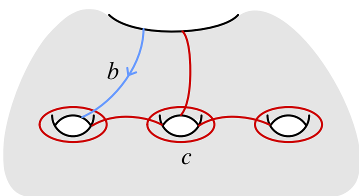

Fig. 2. The curve b of Lemma 3.6, shown in relation to the E_6 subgraph; b is constructed so as to be disjoint from all curves intersecting the central vertex c of the E_6 subgraph, but it may intersect other elements of $\mathcal{C}$ not pictured in the figure.

In the first step, recorded as Lemma 3.7 below, we take $\mathcal{C}_0$ to be the D_5 subconfiguration of the E_6 configuration. In the second step (Lemma 3.8), we take $\mathcal{C}_1 = E_6$ and in the third (Lemma 3.9) we take $\mathcal{C}_2$ to be the union of E_6 with all curves $c \in \mathcal{C}$ intersecting b; these are the most crucial steps of our argument because the ad-hoc methods employed here allow us to deal with those curves of $\mathcal{C}$ meeting b. Step 4 then deals with the rest of the surface ($\mathcal{C}_3 = \mathcal{C}$) by iteratively applying Lemma 3.3.

Step 1: D_5.

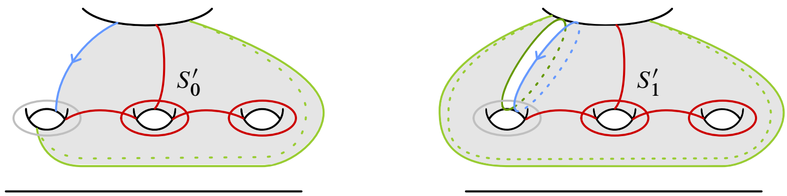

Fig. 3. Left: the surface $S_0 = S_0'$. Right: the surface S_1'.

Lemma 3.7. *Let $S_0' \subset \Sigma_{g,n}$ be the subsurface shown in Figure 3. Then $\widetilde{\Pi}(S_0') \leq \mathcal{T}_{\mathcal{C}}$.*

Proof. Let $\mathcal{C}_0$ be the network shown in Figure 3 consisting of the five red (dark) curves. This satisfies the hypotheses of Lemma 3.5, so that $\widetilde{\Pi}(S_0') \leq \mathcal{T}_{\mathcal{C}_0}$. Each element of $\mathcal{C}_0$ is an element of $\mathcal{C}$, so that the claim follows. ∎

Step 2: E_6.

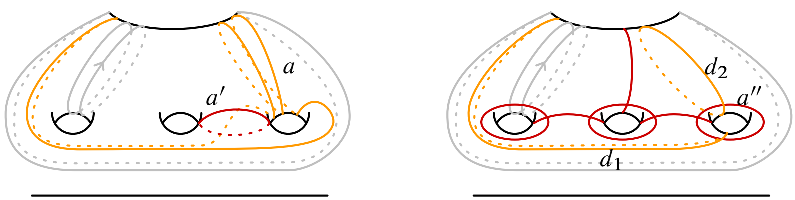

Fig. 4. Left: the curves a and a'; the latter is an element of the E_6 configuration inside $\mathcal{C}$. Right: the boundary components d_1, d_2 for the configuration of D_5 type, and the curve a'', also part of $E_6 \subset \mathcal{C}$.

Lemma 3.8. *Let $S_1' \subset \Sigma_{g,n}$ be the subsurface shown in Figure 3. Then $\widetilde{\Pi}(S_1') \leq \mathcal{T}_{\mathcal{C}}$.*

Proof. We appeal to Lemma 3.3. It suffices to find a curve $a \subset \Sigma_{g,n}$ such that (1) $a \cap S_0'$ is a single arc, (2) S_1 deformation-retracts onto $S_0' \cup a$, (3) there is a curve $a' \subset S_0'$ such that $i(a, a') = 1$ and $T_{a'} \in \mathcal{T}_{\mathcal{C}}$, and (4) $T_a \in \mathcal{T}_{\mathcal{C}}$. A curve a satisfying (1)–(3) is shown in Figure 4.

We claim that $T_a \in \mathcal{T}_{\mathcal{C}}$. To see this, we consider the right-hand portion of Figure 4. We see that five of the curves in the configuration of E_6 type determine a configuration of type D_5; the boundary components of the subsurface spanned by these curves are denoted d_1 and d_2. Applying the D_5 *relation* (see [5, Lemma 5.9]) to this configuration shows that $T_{d_1} T_{d_2}^3 \in \mathcal{T}_{\mathscr{C}}$. We then set

$$a = T_{d_1} T_{d_2}^3 (a''),$$

and as $a'' \in \mathscr{C}$, it follows that $T_a \in \mathcal{T}_{\mathscr{C}}$ as claimed. ∎

Step 3: Curves intersecting *b*.

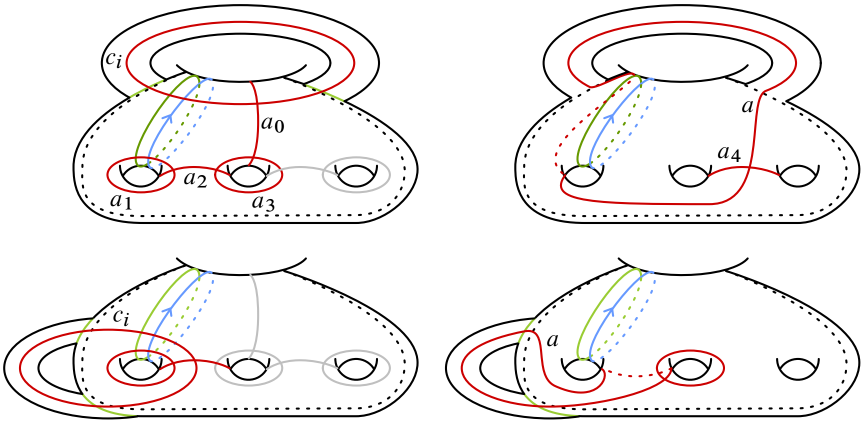

Fig. 5. Left: a curve $c_i \in \mathscr{C}_2$. Right: a twist a satisfying the hypotheses of Lemma 3.3. The top and bottom of the figure depicts the two possible intersection patterns for c_i with S_1'.

Lemma 3.9. *Let $\mathscr{C}_2$ be the configuration of curves given as the union of $E_6 \subset \mathscr{C}$ with all curves $c_i \in \mathscr{C}$ such that $i(c_i, b) \neq 0$. Let S_2 be the surface spanned by these curves, and let S_2' be obtained from S_2 by removing a neighborhood of b. Then $\widetilde{\Pi}(S_2) \leq \mathcal{T}_{\mathscr{C}}$.*

Proof. Recall that b is constructed so as to bound a pair of pants with a_0 and a_2. Since the intersection graph of $\mathscr{C}$ is a tree, a curve $c_i \in \mathscr{C}_2 \setminus E_6$ must be in one of the two configurations shown in Figure 5: it must intersect exactly one of the curves a_0 or a_2. Moreover, distinct $c_i, c_j \in \mathscr{C}_2 \setminus E_6$ must be pairwise disjoint. Thus we can attach the curves c_i in an arbitrary order to assemble S_2' from S_1', appealing to Lemma 3.3 at each step.

The right-hand portion of Figure 5 shows a curve a that satisfies the hypotheses of Lemma 3.3 for this pair of subsurfaces. The curve a is obtained from c_i via a sequence of twists about curves in $\mathscr{C}$:

$$a = T_{a_1} T_{a_2} T_{a_3} T_{a_0} (c_i)$$

in the top scenario, and

$$a = T_{a_1}^{-1} T_{a_2}^{-1} (c_i)$$

in the bottom. Thus, for each such curve $c_i \in \mathscr{C}_2 \setminus E_6$, the associated curve a satisfies $T_a \in \mathcal{T}_{\mathscr{C}}$. The claim now follows from repeated applications of Lemma 3.3. ∎

Step 4: Attaching the remaining curves. The proof of Lemma 3.6 now follows with no further special arguments.

Proof of Lemma 3.6. Step 3 (Lemma 3.9) shows that $\widetilde{\Pi}(S_2') \le \mathcal{T}_{\mathcal{C}}$, where S_2 is the span of E_6 and all curves intersecting a_0 and a_2, and S_2' is obtained from S_2 by removing a neighborhood of b. Let $c_i \in \mathcal{C} \setminus \mathcal{C}_2$ be adjacent to some element $c_j \in \mathcal{C}_2$. Then Lemma 3.3 applies directly to the pair $a = c_i$ and $a' = c_j$. We repeat this process next with curves c_i of graph distance 2 to $\mathcal{C}_2$, then graph distance 3, etc., until the vertices of $\mathcal{C} \setminus \mathcal{C}_2$ are exhausted. At the final stage, we have shown that $\widetilde{\Pi}(\Sigma_{g,n} \setminus \{b\}) = \widetilde{\Pi}(b) \le \mathcal{T}_{\mathcal{C}}$ as claimed. ■

3.4. *Finite generation of the admissible subgroup*

Proposition 3.10 below is taken from [27, Proposition 8.2]. There, it is formulated for r-spin structures on closed surfaces of genus $g \ge 5$, but the result and its proof hold *mutatis mutandis* for framings of $\Sigma_{g,n}$ with $g \ge 5$.

Proposition 3.10 (cf. [27, Proposition 8.2]). *Let ϕ be a framing of $\Sigma_{g,n}$ for $g \ge 5$. Let (a_0, a_1, b) be an ordered 3-chain of curves with $\phi(a_0) = \phi(a_1) = 0$ and $\phi(b) = -1$. Let $H \le \mathrm{Mod}(\Sigma_{g,n})$ be a subgroup containing T_{a_0}, T_{a_1}, and the framed subsurface push subgroup $\widetilde{\Pi}(b)$. Then H contains $\mathcal{T}_\phi$.*

Proof of Proposition 3.1. Since $\mathcal{T}_{\mathcal{C}} \le \mathcal{T}_\phi$ by construction, it suffices to apply Proposition 3.10 with the subgroup $H = \mathcal{T}_{\mathcal{C}}$. Lemma 3.6 asserts that $\widetilde{\Pi}(b) \le \mathcal{T}_{\mathcal{C}}$, so with reference to the labels in Figure 5, the chain (a_2, a_1, b) satisfies the hypotheses of Proposition 3.10. The result follows. ■

We observe that this argument can also be combined with results of our earlier paper [5] to give a vast generalization of the types of configurations which generate r-spin mapping class groups (see §2.1). In particular, the following result gives many new generating sets for the closed mapping class group Mod_g.

Corollary 3.11. *Let $\mathcal{C}$ denote a filling network of curves on a* closed *surface Σ_g with $g \ge 5$. Suppose that the intersection graph $\Lambda_{\mathcal{C}}$ is a tree which contains the E_6 Dynkin diagram as a subgraph and that $\mathcal{C}$ cuts the surface into n polygons with $4(k_1 + 1), \dots, 4(k_n + 1)$ sides. Set $r = \gcd(k_1, \dots, k_n)$. Then there exists an r-spin structure $\widehat{\phi}$ on Σ_g such that*

$$\mathrm{Mod}_g[\widehat{\phi}] = \langle T_c \mid c \in \mathcal{C} \rangle.$$

Proof. The r-spin structure $\widehat{\phi}$ is uniquely determined by stipulating that each curve $c \in \mathcal{C}$ is admissible. To see that the twists in $\mathcal{C}$ generate the stabilizer of this spin structure, we first observe that by [5, Proposition 6.1], the $\widehat{\phi}$-admissible subgroup $\mathcal{T}_{\widehat{\phi}}$ is equal to $\mathrm{Mod}_g[\widehat{\phi}]$. Therefore, it suffices to prove that $\langle T_c \mid c \in \mathcal{C} \rangle = \mathcal{T}_{\widehat{\phi}}$.

Let S denote a neighborhood of the curve system $\mathcal{C}$; by insisting that each curve of $\mathcal{C}$ is admissible in S, we see that S is naturally a surface of genus g equipped with a framing ϕ of signature $(-1-k_1, \dots, -1-k_n)$. Now by Proposition 3.1,

$$\mathcal{T}_\phi = \langle T_c \mid c \in \mathcal{C} \rangle$$

and as Σ_g is obtained from S by capping off each boundary component with a disk, we need only show that $\mathcal{T}_\phi$ surjects onto $\mathcal{T}_{\widehat{\phi}}$ under the capping homomorphism.

To show the desired surjection, consider any $\widehat{\phi}$-admissible curve c on Σ_g. Pick any $\widetilde{c}$ on S which maps to c under capping; then since $\widehat{\phi}$ is just the reduction of ϕ mod r, necessarily

$$\phi(\widetilde{c}) = rN \quad \text{for some } N \in \mathbb{Z}.$$

For each boundary component Δ_i of S, pick some loop γ_i based on Δ_i which intersects $\widetilde{c}$ exactly once. Now since $r = \gcd(k_1, \dots, k_n)$ there is some linear combination

$$m_1 k_1 + \cdots + m_n k_n = r$$

and so by the twist-linearity formula (Lemma 2.4 (1)), the curve

$$\widetilde{c}' = (\mathcal{P}(\gamma_1)^{m_1} \dots \mathcal{P}(\gamma_n)^{m_n})^{-N}(\widetilde{c})$$

must be ϕ-admissible, where $\mathcal{P}(\gamma_i)$ denotes the push of the boundary component Δ_i about γ_i.

But now $T_{\widetilde{c}'}$ is in $\mathcal{T}_\phi$ and $\mathcal{P}(\gamma_1)^{m_1} \dots \mathcal{P}(\gamma_n)^{m_n}$ is in the kernel of the boundary-capping map, and so the image of $T_{\widetilde{c}'}$ in Mod_g is the same as that of $T_{\widetilde{c}}$, which by construction is T_c. Hence $\mathcal{T}_\phi$ surjects onto $\mathcal{T}_{\widehat{\phi}}$, finishing the proof. ■

4. Separating twists and the single boundary case

4.1. *Separating twists*

We come now to the first of two sections dedicated to showing the equality $\mathcal{T}_\phi = \mathrm{Mod}_{g,n}[\phi]$. This will be accomplished by induction on n. In this section, we establish the base case $n = 1$, while in the next section we carry out the inductive step.

The base case $n = 1$ is in turn built around a close connection with the theory of r-spin structures on closed surfaces (see Definition 2.6) as studied in the prior papers [5,27]. We combine this work with a version of the Birman exact sequence (see (6)) to reduce the problem of showing $\mathcal{T}_\phi = \mathrm{Mod}_{g,1}[\phi]$ to the problem of showing that $\mathcal{T}_\phi$ contains a sufficient supply of Dehn twists about *separating* curves.

Below and throughout, the group $\mathcal{K}_{g,1}$ is defined to be the group generated by separating Dehn twists:

$$\mathcal{K}_{g,1} := \langle T_c \mid c \subset \Sigma_{g,1} \text{ separating} \rangle.$$

$\mathcal{K}_{g,1}$ is known as the *Johnson kernel*. It is a deep theorem of Johnson that $\mathcal{K}_{g,1}$ can be identified with the kernel of a certain “Johnson homomorphism” [17], but we will not need to pursue this any further here.

Proposition 4.1. *Fix $g \geq 5$, and let ϕ be a framing of $\Sigma_{g,1}$. Then $\mathcal{K}_{g,1} \leq \mathcal{T}_\phi$.*

A separating curve $c \subset (\Sigma_{g,1}, \phi)$ has two basic invariants. To define these, let $\mathrm{Int}(c)$ be the component of $\Sigma_{g,1} \setminus c$ that does not contain the boundary component of $\Sigma_{g,1}$. The first invariant of c is the *genus* $g(c)$, defined as the genus of $\mathrm{Int}(c)$. The second invariant of c is the *Arf invariant* $\mathrm{Arf}(c)$. When $g(c) \geq 2$, define $\mathrm{Arf}(c)$ to be the Arf invariant of $\phi|_{\mathrm{Int}(c)}$. As discussed in §2.2, the Arf invariant in genus 1 is special. For uniformity of notation later, if $g(c) = 1$ and $\mathrm{Arf}_1(\mathrm{Int}(c)) \neq 0$, define

$$\mathrm{Arf}(c) := \mathrm{Arf}_1(\phi|_{\mathrm{Int}(c)}) + 1 \pmod 2.$$

In the special case where $\mathrm{Arf}_1(\phi|_{\mathrm{Int}(c)}) = 0$, we declare $\mathrm{Arf}(c)$ to be the symbol 1^+ (for the purposes of arithmetic, we treat this as $1 \in \mathbb{Z}/2\mathbb{Z}$). Altogether, we define the *type* of a separating curve c to be the pair $(g(c), \mathrm{Arf}(c))$.

Lemma 4.2. *Let ϕ be a framing of a surface $\Sigma_{g,1}$. Let c be a separating curve of type (g, ε). For the pairs of (g, ε) listed below, the separating twist T_c is contained in $\mathcal{T}_\phi$:*

(1) $(1 + 4k, 1)$ *for* $k \geq 1$,

(2) $(2 + 4k, 1)$ *for* $k \geq 0$,

(3) $(3 + 4k, 0)$ *for* $k \geq 0$,

(4) $(4k, 0)$ *for* $k \geq 1$,

(5) $(1, 1^+)$,

(6) $(3, 1)$.

Proof. In cases (1)–(4), the Arf invariant of the surface $\mathrm{Int}(c)$ agrees with the Arf invariant of a surface of the same genus which supports a maximal chain of admissible curves. By the change-of-coordinates principle for framed surfaces (Proposition 2.15 (2)), every such surface supports a maximal chain of admissible curves. In case (5), by definition we have $\mathrm{Arf}_1(\phi|_{\mathrm{Int}(c)}) = 0$ and therefore every curve on $\mathrm{Int}(c)$ is admissible. Any 2-chain on $\mathrm{Int}(c)$ is therefore admissible.

For any of cases (1)–(5), applying the chain relation [11, Proposition 4.12] to the maximal admissible chain shows that the separating twist about the boundary component is an element of $\mathcal{T}_\phi$.

Consider now case (6), where $\mathrm{Int}(c)$ has genus 3 and Arf invariant 1. In this case, the framed change-of-coordinates principle implies that $\mathrm{Int}(c)$ supports a configuration $a_0, \ldots, a_5$ of admissible curves with intersection pattern given by the E_6 Dynkin diagram. By the "E_6 relation" [23, Theorem 1.4], T_c can be expressed as a product of the admissible twists $T_{a_0}, \ldots, T_{a_5}$. ∎

In the proof of Proposition 4.1, it will be important to understand the additivity properties of the Arf invariant when a surface is decomposed into subsurfaces. The following lemma is stated for only two subsurfaces in [26], but repeated application yields the statement recorded below.

Lemma 4.3 (cf. [26, Lemma 2.11]). *Suppose that (Σ, ϕ) is a framed surface and $\mathbf{c} = \bigcup c_i$ is a multicurve. Let $(S_j, \psi_j)_j$ denote the components of $\Sigma \setminus \mathbf{c}$, equipped with their induced framings. Then*

$$\mathrm{Arf}(\phi) = \sum_j \mathrm{Arf}(\psi_j) + \sum_{c_i \in \mathbf{c}} (\phi(c_i) + 1) \pmod 2.$$

In particular, if each c_i is separating and bounds a closed *subsurface on one side, then each $\phi(c_i)$ is odd and hence the Arf invariant is additive.*

The proof of Proposition 4.1 is built around the well-known *lantern relation*.

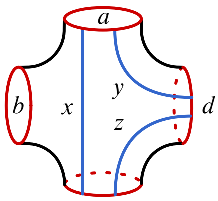

Fig. 6. The lantern relation. The arcs labeled x, y, z determine curves by taking a regular neighborhood of the arc and the incident boundary components.

Lemma 4.4 (Lantern relation). *For the curves a, b, c, d, x, y, z of Figure 6, there is a relation*

$$T_a T_b T_c T_d = T_x T_y T_z.$$

We will use the lantern relation to "manufacture" new separating twists using an initially limited set of twists.

Proof of Proposition 4.1. According to [18, Theorem 1], it suffices to show that $T_c \in \mathcal{T}_\phi$ for c a separating curve of genus 1 or 2. If c has type $(2, 1)$, then $T_c \in \mathcal{T}_\phi$ by Lemma 4.2. For the remaining types $(1, 1), (1, 0), (2, 0)$, we will appeal to a sequence of lantern relations (Configurations (A), (B), (C)) as shown in Figure 7. Each of the configurations below occupy a surface of genus ≤ 4, and by hypothesis, $g \ge 5$. Thus, in each configuration, the specified winding numbers do not constrain the Arf invariant. Therefore by the framed change-of-coordinates principle (Proposition 2.15), there is no obstruction to constructing such configurations. We also remark that we will use the additivity of the Arf invariant (Lemma 4.3) without comment throughout.

We say that separating curves $x, y \subset \Sigma_{g,1}$ are *nested* if either $\mathrm{Int}(x) \subset \mathrm{Int}(y)$ or $\mathrm{Int}(y) \subset \mathrm{Int}(x)$. Configuration (A) shows that $T_x T_y^{-1} \in \mathcal{T}_\phi$ for any x of type $(1, 0)$ and y of type $(2, 0)$ such that x and y are not nested.

Turning to Configuration (B), we apply the lantern relation to the curves on the subsurface bounded by a, b, z, w to find that

$$T_a T_b T_c^{-1} \in \mathcal{T}_\phi; \tag{4}$$

here c is any curve of type $(2, 0)$ and a, b have type $(1, 0)$ and are nested inside c. Let d be of type (2,0) and disjoint from c; hence d is not nested with a or b. Applying

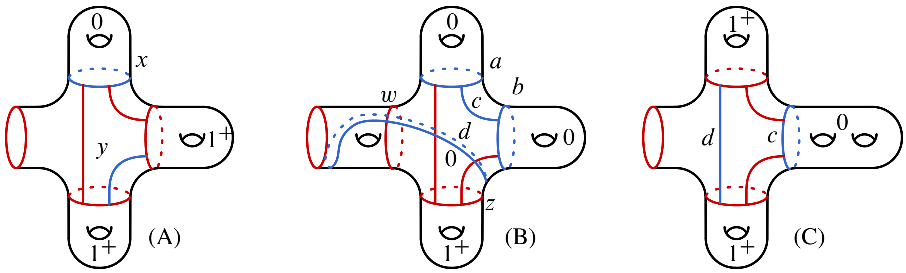

Fig. 7. The lantern relations used in the proof of Proposition 4.1. Curves and arcs colored red correspond to twists known to be in $\mathcal{T}_\phi$ from Lemma 4.2, while those in blue correspond to twists not yet known to be in $\mathcal{T}_\phi$. Numbers inside subsurfaces indicate Arf invariants.

Configuration (A) with $x = a$, $y = d$ and again with $x = b$, $y = d$ implies that

$$T_a T_d^{-1} \in \mathcal{T}_\phi \text{ and } T_b T_d^{-1} \in \mathcal{T}_\phi. \tag{5}$$

Combining (4) and (5) thus implies that $T_c^{-1} T_d^2 \in \mathcal{T}_\phi$ for an arbitrary pair c, d of curves of type $(2, 0)$.

Consider now Configuration (C). The associated lantern relation shows that $T_c T_d^{-1} \in \mathcal{T}_\phi$ for c, d again both of type $(2, 0)$. As also $T_c^{-1} T_d^2 \in \mathcal{T}_\phi$ by the above paragraph, it follows that $T_c \in \mathcal{T}_\phi$ for c an arbitrary curve of type $(2, 0)$.

Returning to Configuration (A), it now follows that $T_x \in \mathcal{T}_\phi$ for x an arbitrary curve of type $(1, 0)$. It remains only to show $T_a \in \mathcal{T}_\phi$ for a a curve of type $(1, 1)$. To obtain this, we return to Configuration (B), but replace the curve of type $(1, 1^+)$ with a general curve of type $(1, 1)$. The remaining twists in the lantern relation are now all known to be elements of $\mathcal{T}_\phi$, and hence curves of type $(1, 1)$ are elements of $\mathcal{T}_\phi$ as well. ∎

4.2. *The minimal stratum*

We are now prepared to prove the main result of the section.

Proposition 4.5. *Let $g \geq 5$ be given, and let ϕ be a framing of $\Sigma_{g,1}$. Then*

$$\mathcal{T}_\phi = \mathrm{Mod}_{g,1}[\phi].$$

As discussed above, this will be proved by relating the framing ϕ on $\Sigma_{g,1}$ to a $(2g-2)$-spin structure on Σ_g by way of a version of the Birman exact sequence. In the standard Birman exact sequence for the capping map $p : \Sigma_{g,1} \to \Sigma_g$, the kernel is given by the subgroup $\pi_1(UT\Sigma_g)$. In Lemma 4.6 below, the subgroup $H_g \leq \pi_1(UT\Sigma_g)$ is defined to be the preimage in $\pi_1(UT\Sigma_g)$ of $[\pi_1(\Sigma_g), \pi_1(\Sigma_g)] \leq \pi_1(\Sigma_g)$.

Lemma 4.6. *Let ϕ be a framing of $\Sigma_{g,1}$. Then there is a $(2g-2)$-spin structure $\widehat{\phi}$ on Σ_g such that the boundary-capping map $p : \mathrm{Mod}_{g,1} \to \mathrm{Mod}_g$ induces the following exact sequence:*

$$1 \to H_g \to \mathrm{Mod}_{g,1}[\phi] \to \mathrm{Mod}_g[\widehat{\phi}]. \tag{6}$$

Proof. The framing ϕ determines a nonvanishing vector field on $\Sigma_{g,1}$. Capping the boundary, this can be extended to a vector field on Σ_g with a single zero. This vector field gives rise to the $(2g-2)$-spin structure $\widehat{\phi}$ (see §2.1), and if $f \in \mathrm{Mod}_{g,1}$ preserves ϕ, then necessarily $p(f)$ preserves $\widehat{\phi}$.

It remains to show that $\ker(p) \cap \mathrm{Mod}_{g,1}[\phi] = H_g$. Since H_g consists of disk pushes about homologically trivial, hence separating, curves, we have $H_g \le \mathcal{K}_{g,1} \le \mathrm{Mod}_{g,1}[\phi]$.

For the converse, we first express the action of a simple based loop $\gamma \in \pi_1(UT\Sigma_g)$ on the winding number of an arbitrary simple closed curve a. Let γ_L (resp. γ_R) denote the curves on $\Sigma_{g,1}$ lying to the left (resp. right) of γ. Then γ acts via the mapping class $\mathcal{P}(\gamma) = T_{\gamma_L} T_{\gamma_R}^{-1}$. The twist-linearity formula (Lemma 2.4 (1)) as applied to $\mathcal{P}(\gamma)$ shows that

$$\begin{aligned} \phi(\mathcal{P}(\gamma)(a)) &= \phi(a) + \langle a, \gamma_L\rangle \phi(\gamma_L) - \langle a, \gamma_R\rangle \phi(\gamma_R) \\ &= \phi(a) + \langle a, \gamma\rangle (\phi(\gamma_L) - \phi(\gamma_R)) \\ &= \phi(a) + \langle a, \gamma\rangle (2-2g). \end{aligned} \tag{7}$$

Here, the second equality holds since γ_L, γ_R, and γ all determine the same homology class, and the third equality holds by homological coherence (Lemma 2.4 (2)), since $\gamma_L \cup \gamma_R \cup \partial\Sigma_{g,1}$ cobounds a pair of pants and necessarily $\phi(\Delta_1) = 1-2g$.

We now use this formula to deduce that if $\mathcal{P}(\gamma) \in \mathrm{Mod}_{g,1}[\phi]$, then $\gamma \in H_g$. To see this, let γ be an arbitrary curve, not necessarily simple, and factor $\gamma = \gamma_1 \ldots \gamma_n$ with each γ_i simple. Since $\mathcal{P}(\gamma_i)$ acts trivially on homology, there is an equality $[\mathcal{P}(\gamma_1 \ldots \gamma_i)(a)] = [a]$ of elements of $H_1(\Sigma_{g,1}; \mathbb{Z})$ for each $i = 1, \ldots, n$, and hence

$$\langle \mathcal{P}(\gamma_1 \ldots \gamma_i)(a), \gamma_{i+1}\rangle = \langle a, \gamma_{i+1}\rangle.$$

Thus applying (7) successively with $\mathcal{P}(\gamma_i)$ acting on $\mathcal{P}(\gamma_1 \ldots \gamma_{i-1})(a)$ for $i = 1, \ldots, n$ shows that

$$\phi(\mathcal{P}(\gamma)(a)) = \phi(a) + \langle a, \gamma\rangle (2-2g). \tag{8}$$

If γ is not contained in H_g, then there exists some simple curve a such that $\langle a, \gamma\rangle \neq 0$. Therefore, (7) shows that $\phi(\mathcal{P}(\gamma)(a)) \neq \phi(a)$ and hence $\mathcal{P}(\gamma) \notin \mathrm{Mod}_{g,1}[\phi]$. ∎

Proof of Proposition 4.5. Consider the forgetful map $p : \mathrm{Mod}_{g,1} \to \mathrm{Mod}_g$, with kernel $\ker(p) = \pi_1(UT\Sigma_g)$. By Lemma 4.6, it suffices to see that

$$p(\mathcal{T}_\phi) = \mathrm{Mod}_g[\widehat{\phi}] \quad \text{and} \quad \mathcal{T}_\phi \cap \pi_1(UT\Sigma_g) = H_g.$$

To see that $p(\mathcal{T}_\phi) = \mathrm{Mod}_g[\widehat{\phi}]$, we appeal to [5, Theorem B]. By the framed change-of-coordinates principle (Proposition 2.15), there exists a configuration of admissible curves on $\Sigma_{g,1}$ (in the notation of [5, Definition 3.11]) of type $\mathsf{C}(2g-2, \mathrm{Arf}(\phi))$. Any such configuration satisfies the hypotheses of [5, Theorem B], showing that $p(\mathcal{T}_\phi) = \mathrm{Mod}_g[\widehat{\phi}]$.

The containment $H_g \le \mathcal{T}_\phi$ will follow from the work of §4.1. By Proposition 4.1, we have $\mathcal{K}_{g,1} \le \mathcal{T}_\phi$. According to [25, Theorem 4.1],

$$\mathcal{K}_{g,1} \cap \pi_1(UT\Sigma_g) = H_g,$$

showing the claim. ∎

We observe that this proof also shows that (6) can be upgraded to the exact sequence

$$1 \to H_g \to \mathrm{Mod}_{g,1}[\phi] \to \mathrm{Mod}_g[\widehat{\phi}] \to 1.$$

5. The restricted arc graph and general framings

The goal of this section is to prove that for $g \geq 5$, given an arbitrary surface $\Sigma_{g,n}$ equipped with a relative framing ϕ, we have $\mathrm{Mod}_{g,n}[\phi] = \mathcal{T}_\phi$. We will argue by induction on n. The base case $n = 1$ was established in Proposition 4.5. To induct, we study the action of $\mathrm{Mod}_{g,n}[\phi]$ on a certain subgraph of the arc graph, and identify the stabilizer of a vertex with a certain $\mathrm{Mod}_{g,n-1}[\phi']$. In §5.1, we introduce the *s-restricted arc graph* $\mathcal{A}^s(\phi; p, q)$ and show that $\mathrm{Mod}_{g,n}[\phi]$ acts transitively on vertices and edges. In §5.3, we prove that $\mathcal{A}^s(\phi; p, q)$ is connected, modulo a surgery argument. In §5.4, we prove this "admissible surgery lemma." Finally, in §5.5 we use these results to prove that $\mathrm{Mod}_{g,n}[\phi] = \mathcal{T}_\phi$ for $g \geq 5$ and $n \geq 0$ arbitrary.

5.1. The restricted arc complex

We must first clarify some conventions and terminology about configurations of arcs. If $a \subset \Sigma_{g,n}$ is an arc and $c \subset \Sigma_{g,n}$ is a curve, then the geometric intersection number $i(a, c)$ is defined as for pairs of curves: $i(a, c)$ denotes the minimum number of intersections of transverse representatives of the isotopy classes of a and c. If a, b are both arcs (possibly based at one or more common point), we define $i(a, b)$ to be the minimum number of intersections of transverse representatives of the isotopy classes of a, b *on the interiors of a, b*. In other words, intersections at common endpoints are not counted. We say that arcs a, b are *disjoint* if $i(a, b) = 0$, so that "disjoint" properly means "disjoint except at common endpoints." As usual, we say that an arc a is *nonseparating* if the complement $\Sigma_{g,n} \setminus a$ is connected, and we say that a pair of arcs a, b is *mutually nonseparating* if the complement $\Sigma_{g,n} \setminus \{a, b\}$ is connected (possibly after passing to well-chosen representatives of the isotopy classes in order to eliminate inessential components).

Our objective is to identify a suitable subgraph of the arc graph on which $\mathrm{Mod}_{g,n}[\phi]$ acts transitively. Before presenting the full definition (see Definition 5.2 below), we first provide a motivating discussion. By definition, an element of $\mathrm{Mod}_{g,n}[\phi]$ must preserve the winding number of every arc, and so to ensure transitivity we must restrict the vertices of this subcomplex to be arcs of a fixed winding number $s \in \mathbb{Z} + \frac{1}{2}$. However, this alone is insufficient, as Lemma 5.1 below makes precise.

If α and β are two disjoint (legal) arcs which connect the same points p and q on boundary components Δ_p and Δ_q, then the action of $\mathrm{Mod}_{g,n}[\phi]$ must preserve the winding number of each boundary component of a neighborhood of $\Delta_p \cup \alpha \cup \beta \cup \Delta_q$. The winding numbers of these curves depends on the ϕ values of α, β, Δ_p, and Δ_q, but also on the configuration of α and β.

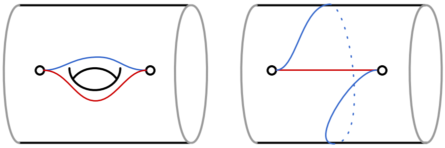

Fig. 8. Sidedness. Left: a one-sided pair. Right: a two-sided pair.

To that end, we say that a pair of arcs $\{\alpha, \beta\}$ as above is *one-sided* (respectively, *two-sided*) if α leaves p and enters q on the same side (respectively, opposite sides) of β for every disjoint realization of α and β on Σ. See Figure 8.

A quick computation yields an equivalent formulation in terms of winding numbers, which for clarity of exposition we state only in the case when $\phi(\alpha) = \phi(\beta)$.

Lemma 5.1. *Let $\{\alpha, \beta\}$ be a pair of arcs as above with $\phi(\alpha) = \phi(\beta)$. Let $c^\pm$ denote the two curves forming the boundary of a small regular neighborhood of $\Delta_p \cup \alpha \cup \beta \cup \Delta_q$, oriented so that the subsurface containing Δ_p and Δ_q lies on their right. Then $\{\alpha, \beta\}$ is*

- *one-sided if and only if $\{\phi(c^+), \phi(c^-)\} = \{1, \phi(\Delta_p) + \phi(\Delta_q) + 1\}$,*
- *two-sided if and only if $\{\phi(c^+), \phi(c^-)\} = \{\phi(\Delta_p) + 1, \phi(\Delta_q) + 1\}$.*

In particular, if c is an admissible curve with $i(\alpha, c) = 1$, then $\{\alpha, T_c(\alpha)\}$ is two-sided.

Proof. This essentially follows by inspection of Figure 8. A given curve $c^\pm$ can be decomposed into a sequence of subintervals – we amalgamate the contributions each makes to the winding number integral. First, $c^\pm$ follows (without loss of generality) α forwards from $\alpha(\varepsilon)$ to $\alpha(1-\varepsilon)$ for suitably small ε, adding $\phi(\alpha)$ to the integral computing $\phi(c^\pm)$. The curve then either turns the "short way" around to follow β back, which adds $1/2$ to the winding number, or else takes the "long way," turning $90°$ positively (adding $1/4$), then following Δ_q around (adding $\phi(\Delta_q)$), and then turning an additional $90°$, adding an additional $1/4$; the "long way" thus adds a total of $1/2 + \phi(\Delta_q)$. The curve now follows β backwards from $\beta(1-\varepsilon)$ to $\beta(\varepsilon)$, subtracting $\phi(\beta) = \phi(\alpha)$ from $\phi(c^\pm)$, and then either takes the short way or the long way around Δ_p. In the one-sided case, the curves $c^\pm$ take the same way (long or short) at both ends, and in the two-sided case, they each take the long way once and the short way once. ■

Having identified sidedness as a further obstruction to transitivity, we come to the definition of the complex under discussion. For any $s \in \mathbb{Z} + \frac{1}{2}$, we say that an arc α is an *s-arc* if $\phi(\alpha) = s$.

Definition 5.2. Let $(\Sigma_{g,n}, \phi)$ be a framed surface with $n \geq 2$. Suppose that p and q are legal basepoints on distinct boundary components Δ_p and Δ_q and fix some $s \in \mathbb{Z} + \frac{1}{2}$. Then the *restricted s-arc graph* $\mathcal{A}^s_\pm(\phi; p, q)$ is defined as follows:

- A vertex of $\mathcal{A}^s_\pm(\phi; p, q)$ is an isotopy class α of s-arcs connecting p and q.
- Two vertices α and β are connected by an edge if they are disjoint and mutually non-separating.

The *two-sided restricted s-arc graph* $\mathcal{A}^s(\phi; p, q)$ is the subgraph of $\mathcal{A}^s_{\pm}(\phi; p, q)$ such that:

- The vertex set of $\mathcal{A}^s(\phi; p, q)$ is the same as that of $\mathcal{A}^s_{\pm}(\phi; p, q)$.
- Two arcs α and β are connected in $\mathcal{A}^s(\phi; p, q)$ if and only if they are connected in $\mathcal{A}^s_{\pm}(\phi; p, q)$ and the pair $\{\alpha, \beta\}$ is two-sided.

5.2. *Transitivity*

In this subsection we prove that the action of $\mathrm{Mod}_{g,n}[\phi]$ on $\mathcal{A}^s(\phi; p, q)$ is indeed transitive on both edges and vertices. The definition of $\mathcal{A}^s(\phi; p, q)$ above was rigged so that the proof of Lemma 5.3 follows as an extended consequence of the framed change-of-coordinates principle (Proposition 2.15). The length of the proof is thus a consequence more of careful bookkeeping than genuine depth.

Lemma 5.3. *The action of* $\mathrm{Mod}_{g,n}[\phi]$ *on* $\mathcal{A}^s(\phi; p, q)$ *is transitive on vertices and on edges.*

We caution the reader that the action of $\mathrm{Mod}_{g,n}[\phi]$ is *not* transitive on *oriented* edges of $\mathcal{A}^s(\phi; p, q)$.

Proof of Lemma 5.3. We begin by showing that edge transitivity implies vertex transitivity. Suppose that α and β are two vertices of $\mathcal{A}^s(\phi; p, q)$; we will exhibit an element of $\mathrm{Mod}_{g,n}[\phi]$ taking α to β.

By the framed change-of-coordinates principle (Proposition 2.15), there is some admissible curve c which meets α exactly once, and so by Lemma 5.1 the arc $T_c(\alpha)$ is adjacent to α in $\mathcal{A}^s(\phi; p, q)$. Now choose some $\gamma \in \mathcal{A}^s(\phi; p, q)$ adjacent to β. By edge transitivity, there exists a $g \in \mathrm{Mod}_{g,n}[\phi]$ which takes the $\{\alpha, T_c(\alpha)\}$ edge to the $\{\beta, \gamma\}$ edge. If $g(\alpha) = \beta$, then we are done. Otherwise, $g(T_c(\alpha)) = \beta$, and since c is admissible, $gT_c \in \mathrm{Mod}_{g,n}[\phi]$.

It remains to establish edge transitivity. Up to relabeling, we may assume that $p \in \Delta_1$ and $q \in \Delta_2$. Suppose that $\alpha = \{\alpha_1, \alpha_2\}$ and $\beta = \{\beta_1, \beta_2\}$ are two edges of $\mathcal{A}^s(\phi; p, q)$. We also assume that α_1 leaves p from the right-hand side of α_2 and enters q to the left, and the same for β_1 and β_2.

For each $\bullet \in \{\alpha, \beta\}$, let $c_\bullet^\pm$ denote the two boundary components of a neighborhood of $\bullet \cup \Delta_1 \cup \Delta_2$. Let $X_\bullet$ (respectively $Y_\bullet$) denote the component of $\Sigma \setminus c_\bullet^\pm$ containing (respectively, not containing) $\bullet$, equipped with the induced framing $\xi_\bullet$ (respectively, $\eta_\bullet$). Finally, orient each curve of $c_\bullet^\pm$ so that $Y_\bullet$ lies on its left-hand side. See Figure 9.

Since both α and β are two-sided, Lemma 5.1 implies that

$$\{\phi(c_\bullet^+), \phi(c_\bullet^-)\} = \{\phi(\Delta_1) + 1, \phi(\Delta_2) + 1\} \tag{9}$$

for $\bullet \in \{\alpha, \beta\}$ and we fix the convention that $\phi(c_\bullet^+) = \phi(\Delta_1) + 1$.

The proof now follows by building homeomorphisms $X_\alpha \to X_\beta$ and $Y_\alpha \to Y_\beta$ and gluing them together. To that end, we must first describe these subsurfaces in more detail.

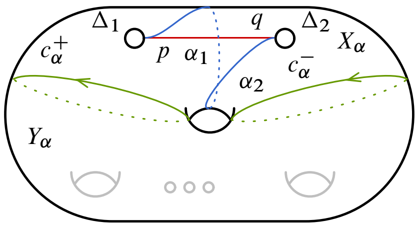

Fig. 9. The curves and subsurfaces determined by a two-sided pair of disjoint s-arcs.

Distinguished arcs in $X_\bullet$. Recall that $c_\bullet^\pm$ are defined as the boundary components of a neighborhood of $\bullet \cup \Delta_1 \cup \Delta_2$. If this neighborhood is taken to be very small (with respect to some auxiliary metric on Σ), then away from p and q the framing restricted to $c_\bullet^\pm$ looks like the framing on segments of $\bullet \cup \Delta_1 \cup \Delta_2$. In particular, for each point $p' \neq p$ of Δ_1 with an orthogonally inward- or outward-pointing framing vector there is a corresponding point of $c_\bullet^+$ with an orthogonally inward- or outward-pointing framing vector. The analogous statement of course also holds for $q \neq q' \in \Delta_2$ and $c_\bullet^-$.

Pick points p' and $p'' \neq p$ on Δ_1 such that the framing vector at p' points orthogonally outwards and the framing vector at p'' points orthogonally inwards. For the sake of concreteness, we will assume that $\phi(\Delta_1)$ is negative and take p' (respectively p'') so that the arc γ' (respectively γ'') of Δ_1 which runs clockwise connecting p to p' (p'') has winding number $-1/2$ (respectively -1). When $\phi(\Delta_1)$ is positive, the proof is identical except the arcs γ' and γ'' will have winding numbers $1/2$ and 1, respectively.

Now let $x_\bullet^+$ and $y_\bullet^+$ denote the corresponding points of $c_\bullet^+$; by construction, the framing vectors at these points point orthogonally into $X_\bullet$ and $Y_\bullet$, respectively. Using q' and $q'' \neq q$ on Δ_2, one may similarly construct $x_\bullet^-$ and $y_\bullet^-$. See Figure 10.

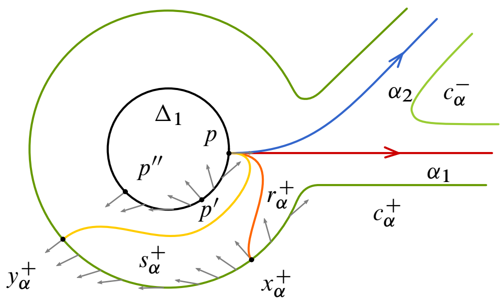

Fig. 10. Distinguished points and arcs in a neighborhood of $\alpha \cup \Delta_1$.

By our choice of p' and p'', one may observe that there exist arcs $r_\bullet^\pm$ from p to $x_\bullet^\pm$ with

$$\phi(r_\bullet^+) = -1/2 \quad \text{and} \quad \phi(r_\bullet^-) = s.$$

Similarly, there exist $s_\bullet^\pm$ from p to $y_\bullet^\pm$ with

$$\phi(s_\bullet^+) = -1 \quad \text{and} \quad \phi(s_\bullet^-) = s - 1/2.$$

Now $\{\bullet_1, r_\bullet^+, r_\bullet^-\}$ forms a distinguished geometric basis for $X_\bullet$, and hence one can compute that

$$\mathrm{Arf}(\xi_\bullet) = (s+1/2)(\phi(\Delta_2)+1) + (0)(\phi(\Delta_1)+2) + (s+1/2)(\phi(\Delta_2)+2), \tag{10}$$

which in particular does not depend on $\bullet \in \{\alpha, \beta\}$.

Building homeomorphisms on subsurfaces. By construction both Y_α and Y_β are homeomorphic to $\Sigma_{g-1,n}$. By (9) their boundary signatures agree:

$$\mathrm{sig}(\eta_\alpha) = (\phi(\Delta_1)+1, \phi(\Delta_2)+1, \phi(\Delta_3), \ldots, \phi(\Delta_n)) = \mathrm{sig}(\eta_\beta).$$

Moreover, by the additivity of the Arf invariant (Lemma 4.3) together with (9) and (10), we have that

$$\begin{aligned}\mathrm{Arf}(\eta_\alpha) &= \mathrm{Arf}(\phi) + \mathrm{Arf}(\xi_\alpha) + \phi(\Delta_1) + \phi(\Delta_2)\\ &= \mathrm{Arf}(\phi) + \mathrm{Arf}(\xi_\beta) + \phi(\Delta_1) + \phi(\Delta_2) = \mathrm{Arf}(\eta_\beta) \pmod 2\end{aligned}$$

and so by the classification of $\mathrm{Mod}_{g,n}$ orbits of framed surfaces (Proposition 2.10) there is a homeomorphism $f_Y : Y_\alpha \to Y_\beta$ such that $f_Y^*(\eta_\beta) = \eta_\alpha$. Moreover, $f_Y(c_\alpha^\pm) = c_\beta^\pm$ and in fact $f_Y(y_\alpha^\pm) = y_\beta^\pm$.

Now in order to extend f_Y to a self-homeomorphism of Σ which takes α to β, we need only specify a homeomorphism f_X of X_α with X_β. This can be done easily by observing that $\alpha \cup r_\alpha^\pm$ cuts X_α into disks with the same combinatorics as $\beta \cup r_\beta^\pm$ cuts X_β, and hence there is a unique homeomorphism $f_X : X_\alpha \to X_\beta$ which takes α to β and $r_\alpha^\pm$ to $r_\beta^\pm$.

Pasting f_X and f_Y together without twisting around $c_\alpha^\pm$, we therefore get a homeomorphism $\widetilde{f} : \Sigma \to \Sigma$ which takes α to β.

Preserving the framing. It remains to show that $\widetilde{f}$ preserves the framing ϕ. Choose a distinguished geometric basis

$$\mathcal{B}_\beta = \{x_1, y_1, \ldots, x_{g-1}, y_{g-1}\} \cup \{a_2, \ldots, a_n\}$$

for Y_β such that all the arcs a_i of $\mathcal{B}_\beta$ emanate from $y_\beta^+ \in c_\beta^+$. By convention, suppose that a_2 runs from y_β^+ to y_β^-, and by twisting around c_β^+ if necessary, suppose that a_2 emerges to the left of all other a_i. Then $\mathcal{B}_\beta$ extends to a distinguished geometric basis of Σ in the following way:

$$\widetilde{\mathcal{B}}_\beta = \{x_1, y_1, \ldots, x_{g-1}, y_{g-1}, c_\beta^-, a_2 \cdot \overline{(s_\beta^-)} \cdot s_\beta^+\} \cup \{\beta_2, s_\beta^+ \cdot a_3, \ldots, s_\beta^+ \cdot a_n\}$$

where $a \cdot b$ represents the concatenation of the arcs a and b and $\overline{a}$ represents the arc a traveled backwards. See Figure 11. By concatenating with $s_\alpha^\pm$ arcs, the basis $f_Y^{-1}(\mathcal{B})$ on Y_α also extends to a basis of Σ in a similar fashion:

$$\begin{aligned}\widetilde{\mathcal{B}}_\alpha = \{&f_Y^{-1}(x_1), f_Y^{-1}(y_1), \ldots, f_Y^{-1}(x_{g-1}), f_Y^{-1}(y_{g-1}), c_\alpha^-, f_Y^{-1}(a_2) \cdot \overline{(s_\alpha^-)} \cdot s_\alpha^+\}\\ &\cup \{\alpha_2, s_\alpha^+ \cdot f_Y^{-1}(a_3), \ldots, s_\alpha^+ \cdot f_Y^{-1}(a_n)\}\end{aligned} \tag{11}$$

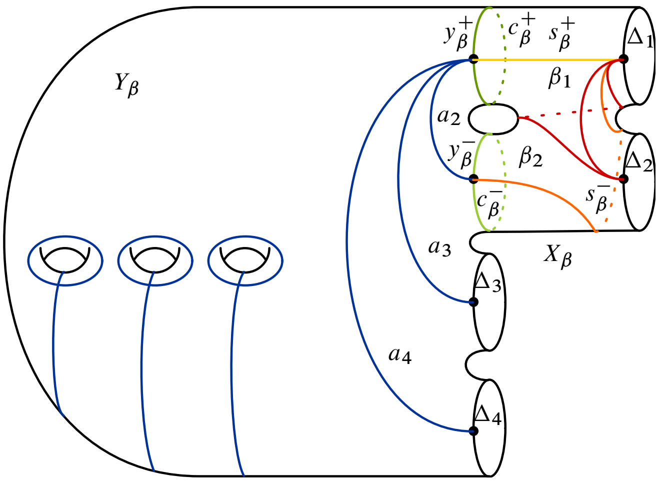

Fig. 11. Extending a distinguished geometric basis from Y_β to Σ by using the arcs $s_\beta^\pm$ of X_β.

Now by construction we have $f_X(s_\alpha^\pm) = s_\beta^\pm$ and $\phi(s_\alpha^\pm) = s - 1/2 = \phi(s_\beta^\pm)$, so for each element $x \in \widetilde{\mathcal{B}}_\alpha$,

$$\widetilde{f}(x) \in \widetilde{\mathcal{B}}_\beta \quad \text{and} \quad \phi(\widetilde{f}(x)) = \phi(x).$$

Therefore $\widetilde{f}$ preserves the winding numbers of a distinguished geometric basis, and so by Remark 2.7, $\widetilde{f} \in \mathrm{Mod}_{g,n}[\phi]$. ∎

5.3. Connectedness

Lemma 5.4. *Let $(\Sigma_{g,n}, \phi)$ be a framed surface with $g \geq 5$ and $n \geq 2$. Let p, q be distinct boundary components, and let $s \in \mathbb{Z}$ be arbitrary. Then $\mathcal{A}^s(\phi; p, q)$ is connected.*

This will require the preliminary Lemmas 5.5, 5.6, and 5.7. The first of these was proved in [27]. There it was formulated only for closed surfaces, but the same proof applies for surfaces with an arbitrary number of punctures and boundary components.

Lemma 5.5 (cf. [27, Lemma 7.3]). *Let $g \geq 5$ and $n \geq 0$ be given. Let S and S' be subsurfaces of $\Sigma_{g,n}$, each homeomorphic to $\Sigma_{2,1}$. Then there is a sequence $S = S_0, \ldots, S_n = S'$ of subsurfaces of $\Sigma_{g,n}$ such that S_{i-1} and S_i are disjoint and $S_i \cong \Sigma_{2,1}$ for all $i = 1, \ldots, n$.*

Lemma 5.6 (Admissible surgery). *Fix $g \geq 3, n \geq 2$, and let $(\Sigma_{g,n}, \phi)$ be a framed surface with distinguished legal basepoints p and q on boundary components Δ_p and Δ_q. Let $S \subset \Sigma_{g,n}$ be a subsurface homeomorphic to $\Sigma_{2,1}$ (necessarily not containing p or q). Let η be an s-arc connecting p, q that is disjoint from S. Let $x \subset \Sigma_{g,n}$ be either a separating curve or an arc connecting p to q, in either case disjoint from S. Then there is a path $\eta = \eta_0, \ldots, \eta_k$ in $\mathcal{A}^s_\pm(\phi; p, q)$ such that $i(\eta_k, x) = 0$.*

Lemma 5.7. *With hypotheses as above, if $\mathcal{A}^s_\pm(\phi; p, q)$ is connected, then also $\mathcal{A}^s(\phi; p, q)$ is connected.*

The proofs of Lemmas 5.6 and 5.7 are deferred to follow the proof of Lemma 5.4. To prove Lemma 5.4, we first introduce the notion of a "connected sum" of curves.

Connected sums. We recall the notion of a "connected sum" as discussed in [27, Section 3.2] and [5, Definition 6.18] (in the former this is called a "curve-arc sum"). Let a be an oriented curve or arc, b be an oriented curve, and ε be an embedded arc connecting the left sides of a and b and otherwise disjoint from $a \cup b$ (if a is an arc we require $\varepsilon \cap a$ to be a point on the interior of a). If a is a curve (resp. arc), the *connected sum* $a +_\varepsilon b$ is the curve (resp. arc) obtained by dragging a across b along the path ε; see Figure 12.

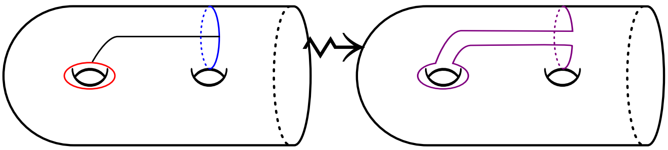

Fig. 12. The connected sum operation.

Lemma 5.8 (cf. [27, Lemma 3.13]). *Let a, b, ε be as above and let ϕ be a relative winding number function. Then*

$$\phi(a +_\varepsilon b) = \phi(a) + \phi(b) + 1.$$

Proof of Lemma 5.4. Following Lemma 5.7, it suffices to show that $\mathcal{A}^s_\pm(\phi; p, q)$ is connected. Let α and ω be s-arcs. Let $S_\alpha \cong \Sigma_{2,1}$ be disjoint from α, and likewise choose $S_\omega \cong \Sigma_{2,1}$ disjoint from ω. By Lemma 5.5, there is a sequence $S_\alpha = S_0, \dots, S_n = S_\omega$ of subsurfaces such that $S_i \cong \Sigma_{2,1}$ and such that S_{i-1} and S_i are disjoint for all $i = 1, \dots, n$. We apply the Admissible Surgery Lemma (Lemma 5.6) taking $(S, \eta, x) = (S_\alpha, \alpha, \partial S_1)$. This gives a path $\alpha = \alpha_0, \dots, \alpha_m$ in $\mathcal{A}^s_\pm(\phi; p, q)$ such that α_m is disjoint from S_1. We now repeat this process for each S_i $(i \geq 1)$, finding intermediate paths of s-arcs, beginning with one disjoint from S_i and ending with one disjoint from S_{i+1}.

At the end of this process we have produced a path of s-arcs $\alpha, \dots, \psi$ with the final arc ψ disjoint from S_ω. To complete the argument we apply the Admissible Surgery Lemma one final time with $(S, \eta, x) = (S_\omega, \psi, \omega)$. This produces a path $\psi = \psi_0, \dots, \psi_k$ in $\mathcal{A}^s_\pm(\phi; p, q)$ with $i(\psi_k, \omega) = 0$. If $\psi_k \cup \omega$ is nonseparating, then ψ_k and ω are adjacent in $\mathcal{A}^s_\pm(\phi; p, q)$, completing the path from α to ω. If $\psi_k \cup \omega$ is separating, then at least one side of the complement has genus $h \geq 2$, and thus there exists a nonseparating oriented curve d disjoint from $\psi_k \cup \omega$ that satisfies $\phi(d) = -1$. Define $\psi_{k+1} = \psi_k +_\varepsilon d$ for a suitable arc ε. Then $\psi_k, \psi_{k+1}, \omega$ is a path in $\mathcal{A}^s_\pm(\phi; p, q)$, completing the argument in this case. ∎

5.4. *Proof of the Admissible Surgery Lemma*

The proof will require the preliminary result of Lemma 5.9 below.

A change-of-coordinates lemma. We study the existence of suitable curves on genus 2 subsurfaces. The proof of Lemma 5.9 is a standard appeal to the framed change-of-coordinates principle (Proposition 2.15).

Lemma 5.9. *Let* $(\Sigma_{2,1}, \phi)$ *be a framed surface, and let* α *be a nonseparating, properly embedded arc on* $\Sigma_{2,1}$*. For* $t \in \mathbb{Z}$ *arbitrary, there is an oriented nonseparating curve* $c_t \subset \Sigma_{2,1}$ *such that* $\phi(c_t) = t$ *and such that* $\langle \alpha, c_t \rangle = 1$.

Proof of Lemma 5.6. The idea is to perform a sequence of surgeries on η in order to successively reduce $i(\eta, x)$. Such surgeries will alter the winding number, but this will be repaired by using the "unoccupied" subsurface S to fix the winding number while preserving the intersection pattern with x. The care we take below in selecting a suitable location for surgery ensures that the intersection pattern with S remains unaltered. Throughout the proof we will refer to Figure 13.

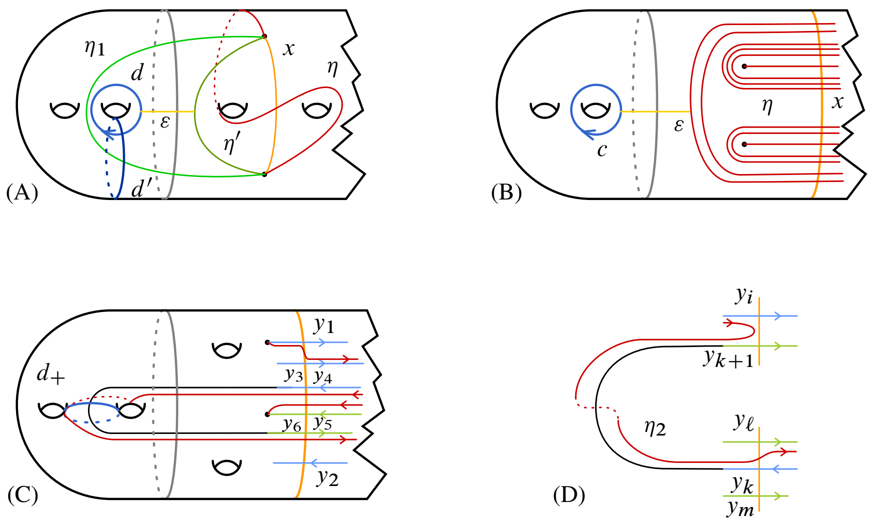

Fig. 13. (A): The case $i(\eta, x) = 1$, illustrated for x an arc. (B): The construction of η_1. (C): The surgery procedure on adjacent initial points (η_1 and η_2 are shown, but η_2' is not). (D): The surgery procedure when the crossings alternate between initial and terminal. In (C, D), we have used blue to indicate initial points and green to indicate terminal points.

Low intersection number. If $i(\eta, x) = 0$ there is nothing more to be done. If $i(\eta, x) = 1$, then there exists an arc η' connecting p to q that is disjoint from $\eta \cup x \cup S$, and such that $\Sigma_{g,n} \setminus \{\eta \cup x \cup \eta'\}$ is connected. See Figure 13 (A). Let $d \subset S$ be an oriented nonseparating curve satisfying $\phi(d) = s - \phi(\eta') - 1$. Let ε be an arc disjoint from $\eta \cup x$ and such that $i(\varepsilon, \partial S) = 1$ that connects η' to the left side of d, and let $\eta_1 = \eta' +_\varepsilon d$. By Lemma 5.8, $\phi(\eta_1) = \phi(\eta') + \phi(d) + 1 = s$. Since there exists a curve $d' \subset S$ such that $i(d', \eta_1) = 1$ and $i(d', \eta) = 0$, it follows that $\eta \cup \eta_1$ is nonseparating, completing the argument in the case $i(\eta, x) = 1$.

The general case: outline. We now consider the case $i(\eta, x) = N \geq 2$. We will first pass to an adjacent s-arc η_1 that enters and exits S exactly once. We will use this in combination with a surgery argument to produce an s-arc η_2 that is adjacent to η_1, satisfies $i(\eta_2, x) < i(\eta, x)$, and also enters and exits S once. As the above arguments (treating the cases $N \leq 1$) can easily be adapted to the situation where η passes once through S, this will complete the proof.

First steps; initial and terminal points. Let $c \subset S$ be an oriented nonseparating curve satisfying $\phi(c) = -1$. Let ε be an arc disjoint from x and such that $i(\varepsilon, \partial S) = 1$ connecting η to the left side of c, and let $\eta_1 = \eta +_\varepsilon c$. See Figure 13 (B). By Lemma 5.8, η_1 is an s-arc, and by construction, $\eta \cup \eta_1$ is isotopic to c and is therefore nonseparating.

Enumerate the intersection points of $\eta_1 \cap x$ as $y_1, \dots, y_N$, numbered consecutively as η_1 runs from p to q; further set $y_0 = p$ and $y_{N+1} = q$. For some $0 \leq k \leq N$, the arc η_1 leaves y_k, enters S, and crosses back through y_{k+1}. The points $y_0, \dots, y_k$ are called *initial*, and the points $y_{k+1}, \dots, y_{N+1}$ are called *terminal*. We say that y_i and y_j are x-*adjacent* if y_i and y_j appear consecutively when running along x (in either direction).

Case 1: adjacent initial/terminal points. Suppose first that there is a pair of x-adjacent points y_i, y_j that are either both initial or both terminal (if x is a curve we consider $1 \leq i < j \leq N$, but if x is an arc, the surgeries we describe below will work for all $0 \leq i < j \leq N + 1$). In this case, let η_2' be obtained from η_1 by following η_1 from p to y_i, then along x to y_j, then finally along η_1 from y_j to q. See Figure 13 (C). Note first that $i(\eta_1, \eta_2') = 0$ and that $i(\eta_2', x) < i(\eta_1, x)$. It remains to alter η_2' to an arc η_2 that is also disjoint from η_1 but with $\phi(\eta_2) = s$ and $\eta_1 \cup \eta_2$ nonseparating, i.e., such that η_1, η_2 is an edge in $\mathcal{A}^s_\pm(\phi; p, q)$.

The method will be to find a curve on S to twist along to correct the winding number of η_2', but care must be taken to ensure that the twisted arc remains disjoint from η_1. Push η_2' off of η_1 so that it runs parallel to η_1 except at the location of the surgery. As η_2' and η_1 run along the segment between y_k and y_{k+1} through S, the push-off of η_2' lies to the left or to the right of η_1 in the direction of travel. We call the former case *positive position* and the latter *negative position*. If c is a curve with $i(c, \eta_2') = 1$, observe that $T_c^\pm(\eta_2')$ is disjoint from η_1 so long as the sign of the twist coincides with the sign of the position.

Define $t := \phi(\eta_2')$. By Lemma 5.9, there are nonseparating curves $d_\pm \subset S$ such that $\phi(d_\pm) = \pm(s - t)$ and such that $\langle \eta_2', d_\pm \rangle = 1$. Set $\eta_2 = T^\pm_{d_\pm}(\eta_2')$, where the sign depends on the sign of the position η_2'. Then η_2 is an s-curve adjacent to η_1 in $\mathcal{A}^s(\phi; p, q)$ and $i(\eta_2, x) < i(\eta_1, x)$.

Case 2: alternating initial/terminal points. It remains to consider the case where every pair of x-adjacent points y_i, y_j has one initial and one terminal element. Here there are two possibilities to consider: either there is exactly one terminal point (and hence $N = 2$), or else at least two. If there is exactly one terminal point and x is an arc, then necessarily this terminal point is q. Then the unique initial point is p, and hence η_1 is disjoint from x except at endpoints and there is nothing left to be done. If x is a separating curve, then necessarily there are at least two terminal points, since if η_1 crosses into the subsurface bounded by x at a terminal point, it must necessarily exit through another terminal point.

We therefore assume that every pair of x-adjacent points y_i, y_j have one initial and one terminal element and that there are at least two terminal points. The first terminal point y_{k+1} is x-adjacent to two distinct initial points y_i, y_j $(i < j \leq k)$, and likewise the last initial point y_k is adjacent to two distinct terminal points y_ℓ, y_m $(k + 1 \leq \ell < m)$.

A suitable surgery is illustrated in Figure 13 (D). The surgered arc η_2' begins by following η_1 forwards from p to y_i, then along x from y_i to y_{k+1}, continuing *backwards* along η_1 from y_{k+1} to y_k. At this point there is a choice: do we follow x to y_ℓ or y_m (in both cases continuing from here forwards along η_1 to q)? The orientation of η_1 endows each y_i with a left and right side. If y_i is adjacent to the left side of y_{k+1}, we continue η_2' to whichever of y_ℓ, y_m lies to the right of y_k, and if y_i lies to the right of y_{k+1}, we continue η_2' to the point y_ℓ, y_m to the left. Observe that $i(\eta_2', x) < i(\eta_1, x)$, even in the exceptional case where the chosen terminal point y_ℓ happens to be y_{k+1}.

The construction of η_2' above facilitates the next step of the argument, which is to adjust η_2' to an s-arc η_2 adjacent to η_1 in $\mathcal{A}^s_\pm(\phi; p, q)$. As in the prior case, let $\phi(\eta_2') = t$, and select (by Lemma 5.9) nonseparating curves $d_\pm \subset S$ such that $\phi(d_\pm) = \pm(s - t)$ and $\langle \eta, d_\pm \rangle = 1$. If y_i is adjacent to the left side of y_{k+1}, define

$$\eta_2 = T_{d_-}^{-1}(\eta_2'),$$

and otherwise define

$$\eta_2 = T_{d_+}(\eta_2').$$

In both cases, η_2 is an s-arc adjacent to η_1 in $\mathcal{A}^s_\pm(\phi; p, q)$ and $i(\eta_2, x) < i(\eta_1, x)$. ■

Having established the admissible surgery lemma (Lemma 5.6), it remains only to give the proof of Lemma 5.7, showing that connectedness of the restricted s-arc graph $\mathcal{A}^s_\pm(\phi; p, q)$ implies the connectedness of the two-sided restricted s-arc graph $\mathcal{A}^s(\phi; p, q)$.

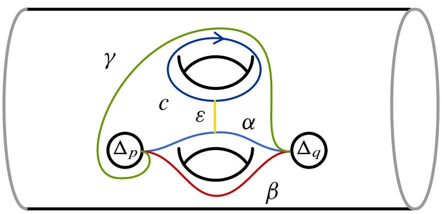

Fig. 14. Connecting a one-sided pair $\{\alpha, \beta\}$ via two-sided pairs $\{\alpha, \gamma\}$ and $\{\gamma, \beta\}$.

Proof of Lemma 5.7. We will refer to Figure 14 throughout. It suffices to show that if $\{\alpha, \beta\}$ is a one-sided edge in $\mathcal{A}^s_\pm(\phi; p, q)$, there is a path in $\mathcal{A}^s(\phi; p, q)$ connecting α to β. Without loss of generality, suppose that α exits p and enters q on the left-hand side of β.

As $\Sigma \setminus (\alpha \cup \beta)$ is a framed surface with genus $g - 1 \geq 2$ and n boundary components, we can apply the framed change-of-coordinates principle (Proposition 2.15) to deduce that there exists a nonseparating curve c on Σ, disjoint from $\alpha \cup \beta$, such that

$$\phi(c) = s - \phi(\alpha) - \phi(\Delta_p) - 1.$$

As $\alpha \cup \beta$ and c are both nonseparating, there exists an arc ε from the left-hand side of α to the left-hand side of c. Therefore, by Lemmas 5.8 and 2.4 (1), we see that for $\gamma := T_{\Delta_p}^{-1}(\alpha +_\varepsilon c)$,

$$\phi(\gamma) = \phi(\alpha +_\varepsilon c) + \phi(\Delta_p) = \phi(c) + \phi(\alpha) + 1 + \phi(\Delta_p) = s.$$

Now since $\alpha +_\varepsilon c$ leaves p and enters q on the left-hand side of α (by construction), we see that γ leaves p to the right of both α and β, but enters q on the left of both α and β. Therefore $\{\alpha, \gamma\}$ and $\{\gamma, \beta\}$ are edges in $\mathcal{A}^s(\phi; p, q)$. ■

5.5. *The inductive step*

Having completed the proof of the Admissible Surgery Lemma, we can proceed with the proof of Theorem B. Our objective in this subsection is Proposition 5.11, which shows that $\mathrm{Mod}_{g,n}[\phi]$ coincides with the admissible subgroup $\mathcal{T}_\phi$. We follow a standard technique of geometric group theory.

Lemma 5.10. *Let G be a group acting on a connected graph X. Suppose that G acts transitively on vertices and edges of X. For a vertex v, let G_v denote the stabilizer of v. Let e be an edge connecting vertices v, w, and let $h \in G$ satisfy $h(w) = v$. Then $G = \langle G_v, h\rangle$.*

Proof. This is very similar to [11, Lemma 4.10]. The argument there can easily be adapted to prove this slightly stronger statement. ■

Proposition 5.11. *Let $g \geq 5$ and $n \geq 1$ be given, and consider a surface $\Sigma_{g,n}$ equipped with a framing ϕ of holomorphic or meromorphic type. Then*

$$\mathcal{T}_\phi = \mathrm{Mod}_{g,n}[\phi].$$

Proof. We argue by induction on the number n of boundary components. The base case $n = 1$ was established above as Proposition 4.5. To proceed, we appeal to Lemma 5.10, taking $G = \mathrm{Mod}_{g,n}[\phi]$ and $X = \mathcal{A}^s(\phi; p, q)$ for $s \in \mathbb{Z} + \frac{1}{2}$ arbitrary.

Lemmas 5.3 and 5.4 combine to show that the hypotheses of Lemma 5.10 are satisfied for $G = \mathrm{Mod}_{g,n}[\phi]$ and $X = \mathcal{A}^s(\phi; p, q)$. Let α be any s-arc connecting p to q. By the framed change-of-coordinates principle (Proposition 2.15), there exists an admissible curve a such that $i(\alpha, a) = 1$. The arc $T_a(\alpha)$ is disjoint from α, the adjacency is two-sided, and the union $T_a(\alpha) \cup \alpha$ is nonseparating. As a is admissible, it follows that $\phi(T_a(\alpha)) = \phi(\alpha) = s$, so that $T_a(\alpha)$ is adjacent to α in $\mathcal{A}^s(\phi; p, q)$.

By Lemma 5.10, it now follows that $\mathrm{Mod}_{g,n}[\phi]$ is generated by T_a and the stabilizer $(\mathrm{Mod}_{g,n}[\phi])_\alpha$. By hypothesis, $T_a \in \mathcal{T}_\phi$. To complete the inductive step, it remains to see that $(\mathrm{Mod}_{g,n}[\phi])_\alpha \le \mathcal{T}_\phi$. Let Δ be the boundary of a neighborhood of $\alpha \cup \Delta_p \cup \Delta_q$, and consider the subsurface $\Sigma_{g,n-1} \le \Sigma_{g,n}$ obtained by ignoring this neighborhood; it inherits a canonical framing ϕ' from the framing ϕ of $\Sigma_{g,n}$. The inclusion of framed surfaces $(\Sigma_{g,n-1}, \phi') \to (\Sigma_{g,n}, \phi)$ induces inclusions

$$\mathrm{Mod}_{g,n-1}[\phi'] \hookrightarrow (\mathrm{Mod}_{g,n}[\phi])_\alpha, \quad \mathcal{T}_{\phi'} \hookrightarrow \mathcal{T}_\phi.$$

By the inductive hypothesis, $\mathrm{Mod}_{g,n-1}[\phi'] \cong \mathcal{T}_{\phi'}$. On the other hand, it is easy to see that

$$(\mathrm{Mod}_{g,n})_\alpha \cap \mathrm{Mod}_{g,n}[\phi] \cong \mathrm{Mod}_{g,n-1}[\phi'],$$

and hence the inclusion $\mathrm{Mod}_{g,n-1}[\phi'] \hookrightarrow (\mathrm{Mod}_{g,n}[\phi])_\alpha$ is an isomorphism. The result follows. ∎

5.6. *Completing the proof of Theorem* B

Theorem B has two assertions. Part (I) asserts that if ϕ is a framing of holomorphic type, then $\mathrm{Mod}_{g,n}[\phi]$ is generated by the Dehn twists about an E-arboreal spanning configuration of admissible curves. This claim follows immediately from the work we have done: by Proposition 3.1, the admissible subgroup $\mathcal{T}_\phi$ is generated by this collection of twists, and the claim now follows from Proposition 5.11.

It therefore remains to establish claim (II) of Theorem B. We recall the statement. Recall (from the paragraph preceding the statement of Theorem B) the notion of an *h-assemblage of type E*: begin with a collection $\mathcal{C}_1 = \{c_1, \dots, c_k\}$ forming an E-arboreal spanning configuration on a subsurface $S \subset \Sigma_{g,n}$ of genus h, and then successively add in curves $c_{k+1}, \dots, c_\ell$ to S, at each stage attaching a new 1-handle to the subsurface. Theorem B (II) asserts that if $\mathcal{C}$ is an h-assemblage of type E for $h \ge 5$, and if each $c \in \mathcal{C}$ is admissible for some framing ϕ (of holomorphic or meromorphic type), then $\mathrm{Mod}(\Sigma_{g,n})[\phi]$ is generated by the finite collection $\{T_c \mid c \in \mathcal{C}\}$ of Dehn twists.

Before proceeding to the (short) proof below, we offer a comment on why we work in such generality. There are two reasons: one of necessity, the other of convenience. On the one hand, the homological coherence criterion (Lemma 2.4 (2)) implies that if $\mathcal{C}$ is an arboreal spanning configuration of *admissble* curves on the framed surface (Σ, ϕ), then necessarily ϕ is of holomorphic type (see Lemma 5.12 below). Thus for meromorphic type we must consider generating sets built on something more general than arboreal spanning configurations. Secondly, while the results of this paper (specifically Theorem A) do not require assemblages, for other applications (especially [24]), this more general framework is essential.

Above we asserted that if the framed surface (Σ, ϕ) admits an E-arboreal spanning configuration, then ϕ is necessarily of holomorphic type. In the course of proving Theorem B (II), we will make use of this fact.

Lemma 5.12. *Let (Σ, ϕ) be a framed surface. Suppose that $\mathcal{C} = \{c_1, \dots, c_k\}$ is an arboreal spanning configuration of admissible curves. Then ϕ is of holomorphic type.*

Proof. The simplest proof uses the perspective of translation surfaces as discussed below in §7. We employ the Thurston–Veech construction (see, e.g. [5, Section 3.3]). Given a configuration of admissible curves whose intersection graph is a tree, this produces a translation surface on which each $c_i \in \mathcal{C}$ is represented as the core of a cylinder. The framing ϕ is incarnated as the framing associated to the horizontal vector field. The framing on a translation surface is necessarily of holomorphic type; the result follows. ∎

We have rigged the definition of an h-assemblage of type E so as to make Theorem B (II) an immediate corollary of the following "stabilization lemma." Below, we use ϕ to denote both a framing of $\Sigma_{g,n}$ and the induced framing on subsurfaces (the latter was denoted ϕ' above).

Lemma 5.13. *Let $(\Sigma_{g,n}, \phi)$ be a framed surface of holomorphic or meromorphic type. Let $S \subset \Sigma_{g,n}$ be a subsurface of genus at least 5 and let $a \subset \Sigma_{g,n}$ be an admissible curve such that $a \cap S$ is a single arc; let S^+ denote a regular neighborhood of $S \cup a$. Then*

$$\mathrm{Mod}(S^+)[\phi] = \langle \mathrm{Mod}(S)[\phi], T_a \rangle.$$

Proof. Let $\mathcal{T}_\phi(S^+)$ denote the admissible subgroup of $\mathrm{Mod}(S^+)[\phi]$. Following Proposition 5.11, it suffices to show that

$$\mathcal{T}_\phi(S^+) \le \langle \mathrm{Mod}(S)[\phi], T_a \rangle.$$

To do so, we appeal to the methods of §3, specifically Proposition 3.10 and Lemma 3.3. Let $b \subset S$ be an arbitrary oriented nonseparating curve satisfying $\phi(b) = -1$, and consider the subsurface push subgroup

$$\widetilde{\Pi}(b) = \widetilde{\Pi}(S^+ \setminus b).$$

By the framed change-of-coordinates principle (Proposition 2.15), b can be extended to a 3-chain (a_0, a_1, b) with a_0, a_1 admissible curves contained in S. Proposition 3.10 asserts that

$$\mathcal{T}_\phi(S^+) \le \langle T_{a_0}, T_{a_1}, \widetilde{\Pi}(b) \rangle.$$

As a_0, a_1 are admissible curves on S, by hypothesis, $T_{a_0}, T_{a_1} \in \mathrm{Mod}(S)$, and so it remains to show that $\widetilde{\Pi}(b) \le \langle \mathrm{Mod}(S)[\phi], T_a \rangle$.

To see this, we observe that $\widetilde{\Pi}(S \setminus b) \le \mathrm{Mod}(S)[\phi]$. Again by the framed change-of-coordinates principle, there is an admissible curve $a' \subset S \setminus b$ such that $i(a', a) = 1$. Appealing to Lemma 3.3, it follows that $\widetilde{\Pi}(b)$ is contained in the group $\langle \widetilde{\Pi}(S \setminus b), T_a, T_{a'} \rangle$, and this latter group is contained in $\langle \mathrm{Mod}(S)[\phi], T_a \rangle$. ∎

Proof of Theorem B (II). Let $\mathcal{C} = \{c_1, \dots, c_k, c_{k+1}, \dots, c_\ell\}$ be an h-assemblage of type E with $h \ge 5$. Assume further that each c_i is admissible for the framing ϕ. Recall that for $1 \le j \le \ell$, the regular neighborhood of the curves $c_1, \dots, c_j$ is denoted by S_j. By hypothesis, $\{c_1, \dots, c_k\}$ forms an E-arboreal spanning configuration on S_k, and the genus of S_k is $h \geq 5$. In particular, the restriction of ϕ to S_k is of holomorphic type (Lemma 5.12). By Theorem B (I), it follows that $\mathrm{Mod}(S_k)[\phi]$ is contained in the group $\mathcal{T}(\mathcal{C}) = \langle T_{c_i} \mid c_i \in \mathcal{C} \rangle$.

We now argue by induction. Supposing $\mathrm{Mod}(S_j)[\phi] \leq \mathcal{T}(\mathcal{C})$ for some $j \geq k$, it follows from Lemma 5.13 that since S_{j+1} is the stabilization of S_j along c_{j+1}, also $\mathrm{Mod}(S_{j+1})[\phi] \leq \mathcal{T}(\mathcal{C})$. As $S_\ell = \Sigma_{g,n}$ by assumption, the result follows. ■

6. Other framed mapping class groups

In this section we leverage our work on the framed mapping class group $\mathrm{Mod}_{g,n}[\phi]$ in order to study two variants we will encounter in our investigation of the monodromy groups of strata of abelian differentials. The most straightforward variant we consider is the stabilizer of the framing up to *absolute* isotopy, i.e., where the isotopy is not necessarily trivial on the boundary. In this case, we will see that there is a sensible theory even when boundary components are replaced by marked points. We carry this out in §6.2, culminating in Proposition 6.10.

Our analysis of this case is built on a study of an intermediate refinement we call the "pronged mapping class group." This group was introduced in a slightly different form in [1], where it was called the "mapping class group rel boundary" of the surface. In §6.1, we lay out the basic theory of prong structures, pronged mapping class groups, and framings on pronged surfaces, leading to the structural result of Proposition 6.7. The material in §6.2 then follows as an easy corollary.

§6.3 contains an analysis of the relationship between the (relative) framed mapping class group $\mathrm{Mod}_{g,n}[\phi]$ and its absolute counterpart $\mathrm{PMod}_g^n[\bar{\phi}]$. The main result here is Proposition 6.14, which identifies an obstruction for $\mathrm{Mod}_{g,n}[\phi]$ to surject onto $\mathrm{PMod}_g^n[\bar{\phi}]$.

Remark 6.1. For clarity of exposition, we restrict our attention throughout this section to framings ϕ of holomorphic type. Similar statements hold for arbitrary framings but the corresponding statements become somewhat messier; see Remarks 6.9 and 6.13.

6.1. Pronged surfaces and pronged mapping class groups

In our study of the monodromy of strata of abelian differentials, we will encounter a variant of a puncture/boundary component known as a *prong structure*. Here we outline the basic theory of surfaces with prong structure and their mapping class groups.

Definition 6.2 (Prong structure, pronged mapping class group). Let Σ be a surface of genus g equipped with a Riemannian metric and $p_i \in \Sigma$ a marked point. A *prong point of order k_i* at p_i is a choice of k_i distinct unit vectors (*prongs*) $v_1, \dots, v_{k_i} \in T_{p_i}\Sigma$ spaced at equal angles. With this data specified, we will write $\vec{p}_i$ to refer to the set of prongs based at p_i, and $\vec{P}$ to indicate a set of prong points $\{\vec{p}_1, \dots, \vec{p}_n\}$ with underlying points $P = \{p_1, \dots, p_n\}$.

Let $\vec{P}$ be a collection of prong points. Define $\mathrm{Diff}^+(\Sigma; \vec{P})$ to be the group of orientation-preserving diffeomorphisms of Σ that preserve each prong point (elements must fix the *points* $p_1, \ldots, p_n$ pointwise, but can induce a (necessarily cyclic) permutation of the overlying tangent vectors). The *pronged mapping class group* is then defined as

$$\mathrm{Mod}(\Sigma, \vec{P}) := \pi_0(\mathrm{Diff}^+(\Sigma; \vec{P})).$$

In the interest of compressing notation, this will often be written simply $\mathrm{Mod}^*_{g,n}$, with the underlying prong structure understood from context.

Prongs vs. boundary components. Here we outline the relationship between prongs and boundary components. First note that a prong point of order 1 is simply a choice of fixed tangent vector. Also recall that in the case where all prong points have order 1, there is a natural isomorphism

$$\mathrm{Mod}(\Sigma, \vec{P}) \cong \mathrm{Mod}(\Sigma^*),$$

where $\Sigma^* \cong \Sigma_{g,n}$ is the surface with n boundary components obtained by performing a real oriented blow-up at each p_i (see Construction 7.4 below). More generally, let $(\Sigma, \vec{P})$ be an arbitrary pronged surface with $\vec{p}_i$ a prong point of order k_i. Let $\mu_k \le \mathbb{C}^\times$ denote the group of kth roots of unity, and define the "prong rotation group"

$$PR := \prod_{i=1}^{n} \mu_{k_i}.$$

For each $1 \le i \le n$, there is a map $D_i : \mathrm{Mod}^*_{g,n} \to \mu_{k_i}$ given by taking the rotational part of the derivative at $T_{p_i}\Sigma \cong \mathbb{C}$. We define

$$D := \prod_{i=1}^{n} D_i : \mathrm{Mod}^*_{g,n} \to PR$$

to be the product. Then D induces the following short exact sequence (cf. [1, (6)]):

$$1 \to \mathrm{Mod}_{g,n} \to \mathrm{Mod}^*_{g,n} \to PR \to 1. \tag{12}$$

Fractional twists. There is an explicit set-theoretic splitting of (12). We define a *fractional twist* $T_{\vec{p}_i}$ at the prong point $\vec{p}_i$ of order k_i to be the mapping class specified in a local complex coordinate $z \le 1$ near p_i by

$$T_{\vec{p}_i}(z) = z e^{(2\pi i(1-|z|))/k_i}.$$

Intuitively, $T_{\vec{p}_i}$ acts by applying a "screwing motion" of angle $2\pi/k_i$ at p_i, viewing a small neighborhood of p_i as being constructed from an elastic material connected to a rigid immobile boundary component. It is then clear that PR embeds into $\mathrm{Mod}^*_{g,n}$ as the set of fractional twists $\{T_{\vec{p}_1}^{j_1} \ldots T_{\vec{p}_n}^{j_n} \mid 0 \le j_i < k_i\}$. Define

$$FT = \langle T_{\vec{p}_i} \mid 1 \le i \le n \rangle \le \mathrm{Mod}^*_{g,n}$$

to be the group generated by the fractional twists.

On the blow-up Σ^*, the fractional twist $T_{\vec{p}_i}$ acts as a fractional rotation of the corresponding boundary component. It will be convenient to introduce the following notation: if Δ_i is the boundary component corresponding to $\vec{p}_i$, write $T_{\Delta_i}^{1/k_i}$ in place of $T_{\vec{p}_i}$. Note that $T_{\vec{p}_i}^{k_i}$ can be identified with the full Dehn twist about Δ_i, so that the equation

$$(T_{\Delta_i}^{1/k_i})^{k_i} = T_{\Delta_i}$$

holds as it should. We will therefore allow T_{Δ_i} to assume fractional exponents with denominator k_i.

Relative framings of pronged surfaces. Let $(\Sigma, \vec{P})$ be a pronged surface. For simplicity, we will formulate the discussion in this paragraph in terms of the blow-up Σ^*. Consider now a nonvanishing vector field ξ on Σ^*. As in §2, when a Riemannian metric is fixed, ξ gives rise to a framing ϕ of Σ^*. In the presence of a prong structure, we will impose a further "compatibility" requirement on ϕ at the boundary.

Definition 6.3 (Compatible framing). Let $\vec{p}_i$ be a prong point of order k_i on Σ and let Δ_i be the associated boundary component of Σ^*. There is a canonical identification $\Delta_i \cong UT_{p_i}\Sigma$ between Δ_i and the space of unit tangent directions at p_i. We say that a framing ϕ is *compatible with* $\vec{p}_i$ if the following conditions hold:

(1) For $v \in UT_{p_i}\Sigma$, the framing vector $\phi(v)$ is orthogonally inward-pointing on Σ^* if and only if v is a prong.

(2) The restriction of ϕ to $UT_{p_i}\Sigma$ is invariant under the action of μ_{k_i} on $UT_{p_i}\Sigma$.

If $\vec{P}$ is a prong structure, we say that ϕ is compatible with $\vec{P}$ if ϕ is compatible with each $\vec{p} \in \vec{P}$ in the above sense.

Remark 6.4. Observe that if ϕ is compatible with a prong point $\vec{p}_i$ of order k_i, then the winding number of the associated boundary component Δ_i of Σ^* is only determined up to sign: $\phi(\Delta_i) = \pm k_i$, depending on which way ϕ turns between prong points (recall the standing convention that boundary components are oriented with the interior of the surface lying to the left). Throughout, we will assume that ϕ is of holomorphic type so that $\phi(\Delta_i) < 0$ for all i.

Observe that the notion of relative isotopy of framings still makes sense on a pronged surface. If ϕ is compatible with $\vec{P}$, then there is a well-defined action of $\mathrm{Mod}_{g,n}^*$ on the set of relative isotopy classes of framings. Exactly as in §2, we define the *framed mapping class group* of the framed pronged surface (Σ^*, ϕ) to be the stabilizer of ϕ:

$$\mathrm{Mod}_{g,n}^*[\phi] = \{f \in \mathrm{Mod}_{g,n}^* \mid f \cdot \phi = \phi\}.$$

Winding number functions. Exactly as in the setting of §2, relative isotopy classes of framings on pronged surfaces are in bijection with suitably defined winding number functions. The definition of winding number of a closed curve needs no modification. To set up a theory of winding numbers for arcs on pronged surfaces, we adopt the natural counterparts of the definitions of legal basepoint and legal arc from §2.

Suppose $(\Sigma, \vec{P})$ is a pronged surface equipped with a compatible framing ϕ. Let $\vec{p}_i$ be a prong point of order k_i. The prongs $v_1, \dots, v_{k_i} \in T_{p_i}\Sigma$ correspond to k_i distinct points on the corresponding boundary component Δ_i of Σ^*, and by the compatibility assumption, each v_i is a legal basepoint in the sense of §2. We say that an arc α on the pronged surface $(\Sigma, \vec{P})$ (or equivalently on the blow-up Σ^*) is *legal* if α is properly embedded, each endpoint is some legal basepoint on Σ^*, and if $\alpha'(0)$ is orthogonally inward-pointing and $\alpha'(1)$ is orthogonally outward-pointing. In short, the theory of legal arcs on pronged surfaces differs from the theory on surfaces with boundary only in that we allow arcs to be based at *any* legal basepoint, not at one fixed point per boundary component. Fractional twists about the boundary may change the basepoint, and so we must consider all legal basepoints at once, instead of a single one.

Under this definition, legal arcs have half-integral winding number as before, and moreover $\mathrm{Mod}^*_{g,n}$ acts on the set of isotopy classes of relative arcs. The twist-linearity formula (Lemma 2.4 (1)) generalizes to fractional twists as follows (the proof is straightforward and is omitted).

Lemma 6.5. *Let a be a legal arc on a pronged surface $(\Sigma, \vec{P})$ equipped with a compatible framing ϕ. If a has an endpoint at a legal basepoint on Δ_i, then*

$$\phi(T_{\Delta_i}^{1/k_i}(a)) = \phi(a) \pm 1,$$

with the sign positive if and only if a is oriented so as to be incoming at Δ_i.

We also have the following straightforward extension of Lemma 2.2.

Lemma 6.6. *Let ϕ and ψ be two framings of the pronged surface $(\Sigma, \vec{P})$, both compatible with the prong structure. Then ϕ and ψ are relatively isotopic if and only if the associated relative winding number functions are equal. Moreover, $\phi = \psi$ as relative winding number functions if and only if $\phi(b) = \psi(b)$ for all elements b of a distinguished geometric basis $\mathcal{B}$.*

We come now to the main result of the section. This describes the relationship between the stabilizer $\mathrm{Mod}^*_{g,n}[\phi]$ in the pronged mapping class group, and its subgroup $\mathrm{Mod}_{g,n}[\phi]$ where each prong is required to be individually fixed. In order to do so, we define a certain subgroup of PR. For convenience, we will switch to additive notation and identify $\mu_k \cong \mathbb{Z}/k\mathbb{Z}$, writing

$$PR = \left\{ \sum c_i e_i \;\middle|\; c_i \in \mathbb{Z}/k_i\mathbb{Z} \right\}.$$

We furthermore write $\sum'$ to indicate a sum over all indices i such that k_i is even. Then define

$$PR' := \left\{ \sum c_i e_i \in PR \;\middle|\; {\sum}' c_i \equiv 0 \pmod 2 \right\}. \tag{13}$$

Observe that if all k_i are odd then $PR' = PR$.

Proposition 6.7. *Let $(\Sigma, \vec{P})$ be a pronged surface and let ϕ be a compatible framing. Then the map $D : \mathrm{Mod}^*_{g,n} \to PR$ induces the short exact sequence*

$$1 \to \mathrm{Mod}_{g,n}[\phi] \to \mathrm{Mod}^*_{g,n}[\phi] \to PR' \to 1. \tag{14}$$

Before we begin the proof, we introduce the notion of an *auxiliary curve*.

Definition 6.8 (Auxiliary curves). Let $(\Sigma_{g,n}, \phi)$ be a framed surface with boundary components $\Delta_1, \dots, \Delta_n$. An *auxiliary curve for* Δ_k is a separating curve d_k such that d_k separates Δ_k from the remaining boundary components, and such that $\phi(d_k) = \pm 1$ or ± 2 according to whether $\phi(\Delta_k)$ is odd or even. An *auxiliary curve for* Δ_i *and* Δ_j is any separating curve $c_{i,j}$ that separates the boundary components Δ_i, Δ_j from the remaining components.

Proof of Proposition 6.7. By definition,

$$\mathrm{Mod}_{g,n}[\phi] := \mathrm{Mod}_{g,n} \cap \mathrm{Mod}^*_{g,n}[\phi],$$

and $\mathrm{Mod}_{g,n}$ is the kernel of D. This establishes exactness at $\mathrm{Mod}_{g,n}[\phi]$ and at $\mathrm{Mod}^*_{g,n}[\phi]$.

It remains to be seen that $D(\mathrm{Mod}^*_{g,n}[\phi]) = PR'$. We first show that $PR' \le D(\mathrm{Mod}^*_{g,n}[\phi])$ by explicit construction. Observe that PR' is generated by elements of three kinds:

(G1) e_i for i such that k_i is odd,

(G2) $2e_i$ for i such that k_i is even,

(G3) $m(e_i + e_j)$ for m odd and i, j such that k_i and k_j are even.

Let $\Delta_i \subset \partial\Sigma^*$ be given. Choose a curve d_i in the following way: if k_i is odd, pick $d_i \subset \Sigma^*$ such that d_i separates Δ_i from all remaining boundary components and the subsurface bounded by Δ_i and d_i has genus $(k_i + 1)/2$. If k_i is even, d_i may be defined identically except that $\Delta_i \cup d_i$ must cobound a surface of genus $(k_i + 2)/2$. By Remark 6.4, if ϕ is compatible with $\vec{P}$, then ϕ is of holomorphic type, and hence $k_i \le 2g - 1$ for all i. Thus the genus of the surface cobounded by $\Delta_i \cup d_i$ is at most g, and hence such d_i exist for all boundary components Δ_i. In both cases, orient d_i so that Δ_i is on its left.

By Remark 6.4 and the homological coherence property (Lemma 2.4 (2)), if k_i is odd, then $\phi(d_i) = -1$, and if k_i is even, then $\phi(d_i) = -2$. Therefore, d_i is an auxiliary curve for Δ_i.

If k_i is odd, we define an *auxiliary twist of type* 1 to be the mapping class

$$A_i := T_{\Delta_i}^{1/k_i} T_{d_i}^{-1}$$

(where d_i is as above), and if k_i is even, we define it to be

$$A_i := T_{\Delta_i}^{2/k_i} T_{d_i}^{-1}.$$

Observe that $D(A_i) = e_i$ if k_i is odd, and $D(A_i) = 2e_i$ if k_i is even. By the twist-linearity formula (Lemma 2.4 (1)) and its extension to fractional twists (Lemma 6.5), one verifies that $A_i \in \mathrm{Mod}^*_{g,n}[\phi]$.

Thus we have exhibited generators of the form (G1), (G2) for PR'. It remains to construct elements mapping to generators of type (G3). Let k_i, k_j be even and choose an auxiliary curve $c_{i,j}$ for Δ_i and Δ_j. By homological coherence (Lemma 2.4 (2)), $\phi(c_{i,j})$

is odd. Define the *auxiliary twist of type* 2 to be the mapping class

$$B_{i,j} := T_{\Delta_i}^{\phi(c_{i,j})/k_i} T_{\Delta_j}^{\phi(c_{i,j})/k_j} T_{c_{i,j}}^{-1}.$$

Then $D(B_{i,j}) = \phi(c_{i,j})(e_i + e_j)$ represents a generator of type (G3), and as before, the twist-linearity formula shows that $B_{i,j} \in \mathrm{Mod}^*_{g,n}[\phi]$.

We now establish the converse assertion $D(\mathrm{Mod}^*_{g,n}[\phi]) \le PR'$. For this, we recall that there is a set-theoretic splitting $s : PR \to \mathrm{Mod}^*_{g,n}$ given by fractional twists, and define the set-theoretic retraction $r : \mathrm{Mod}^*_{g,n} \to \mathrm{Mod}_{g,n}$ by $r(f) = sD(f^{-1})f$.

Let $f \in \mathrm{Mod}^*_{g,n}[\phi]$ be given. Then $r(f) \in \mathrm{Mod}_{g,n}$ by construction. The Arf invariant classifies orbits of relative framings under the action of $\mathrm{Mod}_{g,n}$, and hence we must have

$$\mathrm{Arf}(r(f)\cdot\phi) = \mathrm{Arf}(\phi).$$

Let $\mathcal{B} = \{x_1, \dots, y_g\} \cup \{a_2, \dots, a_n\}$ be a distinguished geometric basis. By hypothesis, $\phi(f \cdot b) = \phi(b)$ for all $b \in \mathcal{B}$. Also note that

$$\phi(T_{\Delta_i}^{1/k_i}(a_i)) = \phi(a_i) + 1$$

while fixing the ϕ values of all other elements of $\mathcal{B}$. Likewise,

$$\phi(T_{\Delta_1}^{1/k_1}(a_i)) = \phi(a_i) - 1$$

for $i = 2, \dots, n$, while $T_{\Delta_1}^{1/k_1}$ fixes the ϕ value of each of the curves $x_1, \dots, y_g$. Since ϕ is a framing, we have $\sum(\varepsilon_i k_i - 1) = 2g - 2$ for some $\varepsilon_i \in \{\pm 1\}$ (compare Remark 6.4). In particular, we have that $\sum(k_i + 1) \equiv 0 \pmod 2$. Considering the Arf invariant formula (1), it follows that

$$\begin{aligned}
\mathrm{Arf}(f\cdot\phi) - \mathrm{Arf}(r(f)\cdot\phi) &= \mathrm{Arf}(\phi) - \mathrm{Arf}(sD(f^{-1})\cdot\phi)\\
&= \sum_{i=2}^{n} (\phi(a_i) - \phi(sD(f)a_i))(k_i + 1)\\
&= \sum_{i=2}^{n} (D_1(f) - D_i(f))(k_i + 1)\\
&\equiv \sum_{i=1}^{n} D_i(f)(k_i + 1) \equiv {\sum}' D_i(f) \pmod 2
\end{aligned}$$

where the penultimate equality follows because $\sum_{i=2}^{n}(k_i + 1)D_1 \equiv (k_1 + 1)D_1 \pmod 2$.

But now f preserves ϕ by construction, hence $\mathrm{Arf}(f\cdot\phi) = \mathrm{Arf}(\phi) = \mathrm{Arf}(r(f)\cdot\phi)$. Therefore we see that ${\sum}' D_i(f) = 0$, i.e., $D(f) \in PR'$. ∎

Remark 6.9. As noted above, this theory generalizes to arbitrary framings compatible with a prong structure. When boundary components have positive winding number the signs of the formula in Lemma 6.5 reverse. More substantially, the proof of Proposition 6.7 must be altered, for there do not exist auxiliary curves for a boundary component Δ of arbitrary winding number. Instead, one must take a combination of twists on curves separating Δ from the other boundary components (together with a fractional twist about Δ) to produce the generators (G1) and (G2).

If one wishes to include boundary components of winding number 0 in this theory, the easiest method is to introduce them separately and consider framings on surfaces with both boundary (of winding number 0) and prongs. In this case, the corresponding factor in the prong rotation group is trivial (see below), and boundary twists about winding number 0 curves are all in the stabilizer of the framing.

6.2. *Pointed surfaces and absolute framed mapping class groups*

The second variant of a framed mapping class group we consider is the most coarse. As in the pronged setting, we consider a closed genus g surface Σ with a collection $P = \{p_1, \dots, p_n\}$ of marked points, but we do not equip each p_i with the structure of a prong point. Thus the mapping class group acting up to isotopy on (Σ, P) is the familiar punctured mapping class group Mod_g^n.

Forgetting the prong structure induces the following short exact sequence of mapping class groups (see [1, Lemma 2.4]):

$$1 \to FT \to \mathrm{Mod}_{g,n}^* \to \mathrm{PMod}_g^n \to 1, \tag{15}$$

where we recall that FT is the group generated by the fractional twists at each p_i.

Suppose that ξ is a vector field on Σ vanishing only at P. Then ξ determines a framing of the punctured surface $\Sigma \setminus P$. Since we do not fix any boundary data, the notion of relative isotopy is ill-defined. To emphasize this, we use the term *absolute* in this setting, and so we speak of *absolute isotopy classes* of framings, *absolute winding number functions* (which measure winding numbers only of oriented simple closed curves, not arcs), and the *absolute framed mapping class group*. In this language, the pointed mapping class group Mod_g^n acts on the set of isotopy classes of absolute framings, or equivalently on the set of absolute winding number functions. If $\bar{\phi}$ is an absolute framing/winding number function, we write $\mathrm{Mod}_g^n[\bar{\phi}]$ to denote the stabilizer, and $\mathrm{PMod}_g^n[\phi]$ for the finite-index subgroup where each marked point is individually fixed. The main result concerning $\mathrm{Mod}_g^n[\bar{\phi}]$ that we will need is the following.

Theorem 6.10. *Let $(\Sigma, \vec{P})$ be a pronged surface equipped with a compatible framing ϕ. The forgetful map $p : (\Sigma, \vec{P}) \to (\Sigma, P)$ induces a surjection*

$$p_* : \mathrm{Mod}_{g,n}^*[\phi] \to \mathrm{PMod}_g^n[\bar{\phi}].$$

Proof. Let $\bar{f} \in \mathrm{PMod}_g^n[\bar{\phi}]$ be given, and choose a lift $f \in \mathrm{Mod}_{g,n}^*$. The set of lifts is a torsor on the kernel FT of the forgetful map $p_* : \mathrm{Mod}_{g,n}^* \to \mathrm{PMod}_g^n$. The group of fractional twists FT preserves all absolute winding numbers. If $\mathcal{B} = \{x_1, \dots, y_g\} \cup \{a_2, \dots, a_n\}$ is a distinguished geometric basis, then by Lemma 6.5, FT acts transitively on the set of values $(\phi(a_2), \dots, \phi(a_n))$. Thus there is $g \in FT$ such that $\phi(gf(b)) = \phi(b)$ for all $b \in \mathcal{B}$. By Lemma 6.6, gf is an element of $\mathrm{Mod}_{g,n}^*[\phi]$, and by construction $p_*(gf) = \bar{f}$. ∎

Combining this result with Proposition 6.7 yields an explicit generating set for $\mathrm{PMod}_g^n[\bar{\phi}]$.

Definition 6.11. Let $\Sigma_{g,n}$ have boundary components $\Delta_1, \dots, \Delta_n$. An *auxiliary curve system* is a collection of the following auxiliary curves:

- auxiliary curves $c_{i,j}$ for all pairs i, j such that both $\phi(\Delta_i), \phi(\Delta_j)$ are even,
- auxiliary curves d_k for all indices k.

No requirements are imposed on the intersection pattern of curves in an auxiliary curve system.

Corollary 6.12. *Let $\underline{\kappa}$ be a partition of $2g-2$ by positive integers. Let ϕ be a relative framing with signature $-1-\underline{\kappa}$ and let $\bar{\phi}$ denote the absolute framing induced on Σ_g^n by capping off the boundary components of $\Sigma_{g,n}$ with punctured disks. Then $\mathrm{PMod}_g^n[\bar{\phi}]$ is generated by $p_*(\mathrm{Mod}_{g,n}[\phi])$ together with the twists about an auxiliary curve system $\mathcal{A}$.*

In particular, if $g \geq 5$, then $\mathrm{PMod}_g^n[\bar{\phi}]$ is generated by the Dehn twists in the curves of $\mathcal{C} \cup \mathcal{A}$, where $\mathcal{C}$ is as in Figure 15.

Proof. As in the proof of Proposition 6.7, the group of auxiliary twists $\langle A_k, B_{i,j}\rangle$ surjects onto PR' and hence by Proposition 6.7,

$$\langle \mathrm{Mod}_{g,n}[\phi], A_k, B_{i,j}\rangle = \mathrm{Mod}_{g,n}^*[\phi].$$

Now it remains to observe that by Theorem 6.10 this group surjects onto $\mathrm{PMod}_g^n[\bar{\phi}]$ and that by construction $p_*(A_k) = T_{d_k}^{-1}$ and $p_*(B_{i,j}) = c_{i,j}^{-1}$. ■

Remark 6.13. As observed in Remark 6.9, auxiliary curve systems do not always exist for arbitrary framings. The substitutions outlined there can similarly be used to give a generating set for arbitrary $\mathrm{PMod}_g^n[\bar{\phi}]$ in terms of $p_*(\mathrm{Mod}_{g,n}[\phi])$ and combinations of separating twists.

6.3. *The image of the relatively framed mapping class group*

The final result we consider here determines the image of $\mathrm{Mod}_{g,n}[\phi]$ in $\mathrm{PMod}_g^n[\bar{\phi}]$ induced by the boundary-capping map $\Sigma_{g,n} \to \Sigma_g^n$. Again, we restrict to the case when ϕ is of holomorphic type; the corresponding statements and proofs for framings of arbitrary type are left to the interested reader (one needs only change the signs of some generators).

Proposition 6.14. *The image of $\mathrm{Mod}_{g,n}[\phi]$ in $\mathrm{PMod}_g^n[\bar{\phi}]$ is a normal subgroup with quotient isomorphic to $PR'/\langle(1, \dots, 1)\rangle$.*

Proof. By Proposition 6.7, $\mathrm{Mod}_{g,n}[\phi] \lhd \mathrm{Mod}_{g,n}^*[\phi]$ is a normal subgroup with quotient PR' induced by the "prong rotation map"

$$D : \mathrm{Mod}_{g,n}^* \to PR.$$

By Theorem 6.10, the boundary-capping map $p_* : \mathrm{Mod}^*_{g,n} \to \mathrm{PMod}^n_g$ restricts to a surjection $p_* : \mathrm{Mod}^*_{g,n}[\phi] \to \mathrm{PMod}^n_g[\bar{\phi}]$. Thus the image of $\mathrm{Mod}_{g,n}[\phi]$ is a normal subgroup of $\mathrm{PMod}^n_g[\bar{\phi}]$.

The quotient $\mathrm{PMod}^n_g[\bar{\phi}]/p_*(\mathrm{Mod}_{g,n}[\phi])$ can be identified by the isomorphism theorems. Suppose that G is a group and N_1, N_2 are normal subgroups. Then N_1N_2 is normal in G, and by the third isomorphism theorem,

$$(G/N_1)/(N_1N_2/N_1) \cong G/(N_1N_2) \cong (G/N_2)/(N_1N_2/N_2).$$

We apply this here with

$$G = \mathrm{Mod}^*_{g,n}[\phi], \quad N_1 = \ker p_*, \quad N_2 = \mathrm{Mod}_{g,n}[\phi].$$

Then

$$G/N_1 \cong \mathrm{PMod}^n_g[\bar{\phi}], \quad G/N_2 \cong PR',$$
$$N_1N_2/N_1 \cong p_*(N_2) \cong p_*(\mathrm{Mod}_{g,n}[\phi]), \quad N_1N_2/N_2 \cong D(N_1) \cong D(\ker p_*).$$

Altogether,

$$\mathrm{PMod}^n_g[\bar{\phi}]/p_*(\mathrm{Mod}_{g,n}[\phi]) \cong PR'/D(\ker p_*).$$

To complete the argument, it therefore suffices to show that $D(\ker p_*) \cong \langle(1,\dots,1)\rangle$. According to (15), the kernel of p_* on $\mathrm{Mod}^*_{g,n}$ is the group FT of fractional twists. Thus we must identify $FT \cap \mathrm{Mod}^*_{g,n}[\phi]$. We claim that

$$FT \cap \mathrm{Mod}^*_{g,n}[\phi] \cong \mathbb{Z}$$

generated by the fractional twist $\prod_{i=1}^n T_{\Delta_i}^{1/k_i}$. Note that

$$D\Big(\prod_{i=1}^n T_{\Delta_i}^{1/k_i}\Big) = (1,\dots,1),$$

so that showing this isomorphism will complete the argument.

To show this claim, consider an arbitrary element

$$f := \prod_{i=1}^n T_{\Delta_i}^{a_i/k_i}$$

of $FT \cap \mathrm{Mod}^*_{g,n}[\phi]$. Let $\alpha_{i,j}$ be a legal arc connecting Δ_i to Δ_j. By Lemma 6.5,

$$\phi(f(\alpha_{i,j})) - \phi(\alpha_{i,j}) = a_j - a_i,$$

so that if $f \in \mathrm{Mod}^*_{g,n}[\phi]$, necessarily $a_i = a_j$ for all pairs i, j as claimed. ■

We now unravel the condition that $PR'/\langle(1,\dots,1)\rangle$ is trivial using some elementary group theory.

Corollary 6.15. *Let $\underline{\kappa}$ be a partition of $2g-2$ and write $\underline{\kappa} = (\eta_1, \dots, \eta_p, \upsilon_1, \dots, \upsilon_q)$ where η_i are even, υ_j are odd, and $p+q=n$. Then $\mathrm{Mod}_{g,n}[\phi]$ surjects onto $\mathrm{PMod}_g^n[\bar{\phi}]$ if and only if $q \leq 2$ and*

$$\left\{\eta_1+1, \dots, \eta_p+1, \frac{\upsilon_1+1}{2}, \dots, \frac{\upsilon_q+1}{2}\right\}$$

are pairwise coprime.

Proof. By Proposition 6.14, it suffices to determine when PR' is cyclic (with generator $(1, \dots, 1)$).

Using the additive notation introduced above, write

$$PR_e := \left\{\sum_{j=1}^{q} c_j e_j \;\middle|\; c_j \in \mathbb{Z}/(\upsilon_j+1)\mathbb{Z}\right\}, \quad PR_o := \left\{\sum_{i=1}^{p} c_i e_i \;\middle|\; c_i \in \mathbb{Z}/(\eta_i+1)\mathbb{Z}\right\}.$$

Even though PR_e is a product over the odd υ_j's, our notation reflects the fact that each of its factors has even order.

Now by definition $PR = PR_e \times PR_o$, and likewise we can write $PR' = PR'_e \times PR_o$ where

$$PR'_e := \left\{\sum_j c_j e_j \in PR_e \;\middle|\; \sum c_j \equiv 0 \pmod 2\right\}$$

as in (13). We observe that PR' is cyclic if and only if PR_o and PR'_e are cyclic of coprime order, and that PR_o is cyclic if and only if the set of η_i+1 are all pairwise coprime.

Suppose that PR'_e is cyclic. If PR_e (and hence also PR'_e) is trivial, then the claim holds. Otherwise, there is a short exact sequence

$$1 \to PR'_e \to PR_e \to \mathbb{Z}/2 \to 1.$$

It follows that PR_e is either cyclic of order $2|PR'_e|$ or else is isomorphic to $\mathbb{Z}/2 \times PR'_e$. Now as each factor of PR_e is a cyclic group of even order, this implies that PR_e has at most two factors, i.e., $q \leq 2$. Since PR_e is a product over the odd υ_j's and we have $\sum \eta_i + \sum \upsilon_j = 2g-2$, this implies that q is even and is therefore either 0 or 2. So if PR_e is nontrivial it must be isomorphic to $(PR'_e) \times \mathbb{Z}/2$. Necessarily then $(\upsilon_1+1)/2$ and $(\upsilon_2+1)/2$ are coprime and

$$PR'_e \cong \mathbb{Z}/\left(\frac{\upsilon_1+1}{2}\right)\mathbb{Z} \times \mathbb{Z}/\left(\frac{\upsilon_2+1}{2}\right)\mathbb{Z}.$$

The remaining hypothesis that PR_o and PR'_e be cyclic of coprime order readily implies the claim that the elements of

$$\left\{\eta_1+1, \dots, \eta_p+1, \frac{\upsilon_1+1}{2}, \dots, \frac{\upsilon_q+1}{2}\right\}$$

are pairwise coprime as required. ∎

7. Monodromy of strata

In this section, we fuse our discussion of framings and mapping class groups with the theory of abelian differentials to deduce Theorem A from Theorem B.

We begin in §7.1 by collecting basic results about abelian differentials and their strata. We also record Kontsevich and Zorich's seminal classification (Theorem 7.2) of the connected components of strata, together with a slight variation in which one labels the zeros of the differential (Lemma 7.3).

With these foundations laid, we proceed to discuss the relationship between abelian differentials and framings on punctured, bordered, and pronged surfaces in §7.2. This section also contains a detailed description of the real oriented blow-up of an abelian differential along its zero locus (Construction 7.4). The main result of this subsection is Boissy's classification of the components of strata of *prong-marked* differentials (Theorem 7.5) and its implications for monodromy (Corollary 7.6).

Mapping class groups enter the picture in §7.3, in which we define a family of coverings of strata (first introduced in [1]) whose deck groups are mapping class groups of punctured, bordered, and pronged surfaces (see Diagram (16)). Using the relations between these spaces, we then prove that the monodromy of each covering must preserve the appropriate framing datum (Lemma 7.10 and Corollaries 7.11 and 7.12).

We establish the reverse inclusions (that the monodromy is the entire stabilizer of the framing) in §7.4 as Theorems 7.13, 7.14, and A.

7.1. Abelian differentials

An *abelian differential* ω is a holomorphic 1-form on a Riemann surface X. The collection of all abelian differentials forms a vector bundle (in the orbifold sense) $\Omega\mathcal{M}_g$ over the moduli space of curves. The complement of the zero section of this bundle is naturally partitioned into disjoint subvarieties called *strata* which have a fixed number and order of zeros. For $\underline{\kappa} = (\kappa_1, \dots, \kappa_n)$, we will let $\Omega\mathcal{M}_g(\underline{\kappa})$ denote the space of all pairs (X, ω) where ω is an abelian differential on X which has zeros of order $\kappa_1, \dots, \kappa_n$. Throughout this section, we will use Z to denote the set of zeros of ω.

Away from its zeros, an abelian differential has canonical local coordinates in which it can be written as dz; the transition maps between these coordinate charts are translations, so the data of an abelian differential ω on a Riemann surface X is also sometimes called a *translation structure*. Pulling back the Euclidean metric of $\mathbb{C}$ along the canonical coordinates defines a flat metric on X with a cone point of angle $2\pi(k+1)$ at each zero of order k.

Every abelian differential ω also defines a horizontal vector field $H_\omega := 1/\omega$ on X with singularities of order $-\kappa_1, \dots, -\kappa_n$, and hence gives rise to a prong structure $\vec{Z}$ with a prong point of order $\kappa_i + 1$ at the ith zero of ω. Forgetting the prong structure and the marked points, the differential induces a $\gcd(\underline{\kappa})$-spin structure ϕ on X (see Definition 2.6 and the discussion which follows it). In particular, if $\gcd(\underline{\kappa})$ is even, then there is a well-defined mod-2 reduction of the spin structure.

Remark 7.1. For a more thorough treatment of the relationship between (higher) spin structures and abelian differentials, the reader is directed to [4] and [5].

While it does not make sense to compare spin structures on different (unmarked) surfaces, when $\gcd(\underline{\kappa})$ is even the Arf invariant of ϕ is well-defined even without choice of marking. Moreover, $\mathrm{Arf}(\phi)$ is invariant under deformation and classifies the nonhyperelliptic components of $\Omega\mathcal{M}_g(\underline{\kappa})$.

Theorem 7.2 ([21, Theorem 1]). *Let $g \geq 4$ and $\underline{\kappa} = (\kappa_1, \ldots, \kappa_n)$ be a partition of $2g-2$. Then $\Omega\mathcal{M}_g(\underline{\kappa})$ has at most three components:*

- *If $\underline{\kappa} = (2g-2)$ or $(g-1, g-1)$ then there is a unique component of $\Omega\mathcal{M}_g(\underline{\kappa})$ that consists entirely of hyperelliptic differentials.*[4]
- *If $\gcd(\underline{\kappa})$ is even then there are two components containing nonhyperelliptic differentials, classified by the Arf invariant of their induced 2-spin structure.*
- *If $\gcd(\underline{\kappa})$ is odd then there is a unique component containing nonhyperelliptic differentials.*

We will focus our attention on the nonhyperelliptic components of $\Omega\mathcal{M}_g(\underline{\kappa})$; the hyperelliptic components are $K(G,1)$'s for (finite extensions of) spherical braid groups [22, §1.4] and therefore their monodromy can be understood entirely through Birman–Hilden theory. See [4, §2] for a more thorough discussion.

We will often find it convenient to label the zeros of an abelian differential so as to distinguish them. The corresponding stratum $\Omega\mathcal{M}_g^{\mathrm{lab}}(\underline{\kappa})$ of abelian differentials with labeled singularities is clearly a finite cover of $\Omega\mathcal{M}_g(\underline{\kappa})$ with deck group

$$\mathrm{Sym}(\underline{\kappa}) = \prod_{j=1}^{m} \mathrm{Sym}(r_j)$$

where $\underline{\kappa} = (k_1^{r_1}, \ldots, k_m^{r_m})$ and $\mathrm{Sym}(n)$ is the symmetric group on n letters. Moreover, it is not hard to show that each preimage of a connected component of $\Omega\mathcal{M}_g(\underline{\kappa})$ is itself connected (see, e.g., [2, Proposition 4.1]), and hence the monodromy of this covering map is the entirety of the deck group.

Lemma 7.3. *Let $\mathcal{H}$ be a component of $\Omega\mathcal{M}(\underline{\kappa})$. Then the monodromy homomorphism*

$$\pi_1^{\mathrm{orb}}(\mathcal{H}) \to \mathrm{Sym}(\underline{\kappa})$$

associated to the covering $\Omega\mathcal{M}_g^{\mathrm{lab}}(\underline{\kappa}) \to \Omega\mathcal{M}_g(\underline{\kappa})$ is surjective.

7.2. *From differentials to framings*

To relate our results on framed mapping class groups to strata, we must first understand the different types of framings which an abelian differential induces.

[4]Recall that an abelian differential is *hyperelliptic* if it arises as the global square root of a quadratic differential on $\widehat{\mathbb{C}}$ with at worst simple poles.

As observed above, the horizontal vector field of an abelian differential $(X, \omega) \in \Omega\mathcal{M}_g(\underline{\kappa})$ induces a $\gcd(\underline{\kappa})$-spin structure on the underlying closed surface X. Moreover, since ω has canonical coordinates (in which it looks like dz) away from its zeros, we see that ω induces a trivialization of $T(X \setminus Z)$ and hence an absolute framing of $X \setminus Z$ (in the sense of §6.2).

To obtain a relative framing from ω, we must first identify the surface X^* which is to be framed. Informally, this "real oriented blow-up" X^* is obtained by replacing each zero of ω by the circle of directions at that point; the horizontal vector field then extends by continuity along rays to a vector field (and eventually, a framing) on X^* whose boundary data can be read off from the order of the singularities. For more on the real oriented blow-up construction in the context of translation surfaces, see [1, §2.5].

Construction 7.4 (Real oriented blow-ups). We begin by first describing the real oriented blow-up of $0 \in \mathbb{C}$; this toy example will provide a local model for the blow-up of a translation surface. Equipping $\mathbb{C}$ with polar coordinates $z = re^{i\theta}$ gives a parametrization of $\mathbb{C} \setminus \{0\}$ by the infinite open half-cylinder $\mathbb{R}_{>0} \times [0, 2\pi]/(0 = 2\pi)$. The *real oriented blow-up* of $0 \in \mathbb{C}$ is the closed half-cylinder $\mathbb{R}_{\geq 0} \times [0, 2\pi]/(0 = 2\pi)$, which has a natural surjective map onto $\mathbb{C}$ extending polar coordinates. The fiber of this map above 0 is therefore identified with the circle of directions at 0.

To blow up a cone point, let $k \geq 1$ and consider the branched cover of $\mathbb{C}$ given by $z \mapsto z^k$. The Euclidean metric of $\mathbb{C}$ pulls back to a cone metric with cone angle $2k\pi$ at 0, and similarly the polar parametrization of $\mathbb{C} \setminus \{0\}$ pulls back to a parametrization by $\mathbb{R}_{>0} \times [0, 2k\pi]/(0 = 2k\pi)$. Therefore, the blow-up of a cone point of angle $2k\pi$ is the corresponding closed cylinder $\mathbb{R}_{\geq 0} \times [0, 2k\pi]/(0 = 2k\pi)$, and the fiber above 0 corresponds to the $2k\pi$'s worth of directions at 0.

Now suppose that $(X, \omega) \in \Omega\mathcal{M}_g^{\mathrm{lab}}(\underline{\kappa})$ with zeros at points $p_1, \ldots, p_n$. The *real oriented blow-up* X^* of X is the space obtained after blowing up each cone point p_i via the above construction.

Observe that X^* is naturally a surface of the same genus as X with boundary components $\Delta_1, \ldots, \Delta_n$, the ith of which comes with an identification with the $(k_i + 1)$-fold cover of the circle (which is of course just a circle itself, equipped with a cyclic symmetry of order $k_i + 1$).

Moreover, the unit horizontal vector field H of ω induces a (nonvanishing) unit vector field H^* on X^* by extending H continuously along rays into each cone point. For each boundary component Δ_i, the vector field $H^*|_{\Delta_i}$ is invariant under the cyclic symmetry described above, and its winding number is $-1 - k_i$.

Hence H^* induces a framing ϕ of X^* with boundary signature

$$\mathrm{sig}(\phi) = (-1 - k_1, \ldots, -1 - k_n)$$

which is compatible (in the sense of Definition 6.3) with the prong structure $(X, \vec{Z})$ induced by ω.

Prong markings. While the blow-up X^* of (X,ω) is topologically a surface with boundary, it is more accurate to view X^* as (the blow-up of) a surface with prong structure. In particular, there exist loops[5] in $\Omega\mathcal{M}_g^{\text{lab}}(\underline{\kappa})$ which rotate the prongs of $(X,\vec{Z})$ and hence act by fractional twists on ∂X^*. This phenomenon can be interpreted as the monodromy of $\Omega\mathcal{M}_g^{\text{lab}}(\underline{\kappa})$ taking values in $\text{Mod}_{g,n}^*$ rather than $\text{Mod}_{g,n}$ (see §7.3 below).

By passing to a finite cover of $\Omega\mathcal{M}_g^{\text{lab}}(\underline{\kappa})$ which remembers more information, we may therefore constrain the monodromy to lie in $\text{Mod}_{g,n}$. To that end, define a *prong marking* of an abelian differential ω to be a labeling of its zeros together with a choice of positive horizontal separatrix at each zero (these are also sometimes called *framings of* ω, as in [3]). In terms of the prong structure $\vec{Z}$ on X induced by ω, a prong marking chooses a prong at each zero.

The (components of the) space $\Omega\mathcal{M}_g^{\text{pr}}(\underline{\kappa})$ of prong-marked abelian differentials are finite covers of (the components of) a stratum $\Omega\mathcal{M}_g^{\text{lab}}(\underline{\kappa})$. Moreover, the deck transformations of this covering rotate the choice of specified prong at each zero and hence the deck group is exactly *PR*.

Any loop in $\Omega\mathcal{M}_g^{\text{pr}}(\underline{\kappa})$ preserves the prong marking and so acts as the identity on ∂X^* by the correspondence outlined in §6.1. Therefore, the real oriented blow-up of a prong-marked abelian differential can be consistently interpreted as a surface with boundary (and hence the monodromy of $\Omega\mathcal{M}_g^{\text{pr}}(\underline{\kappa})$ is in $\text{Mod}_{g,n}$, as we will see in §7.3).

We now record Boissy's classification of the components of $\Omega\mathcal{M}^{\text{pr}}(\underline{\kappa})$. Our statement of the following theorem looks rather different than that which appears in [3]; we reconcile these differences in Remark 7.7 at the end of this section.

Theorem 7.5 (cf. [3, Theorem 1.3]). *Suppose that $\mathcal{H}$ is a nonhyperelliptic connected component of $\Omega\mathcal{M}_g^{\text{lab}}(\underline{\kappa})$. Then the preimage of $\mathcal{H}$ in $\Omega\mathcal{M}_g^{\text{pr}}(\underline{\kappa})$ has*

- *one connected component if* $\gcd(\underline{\kappa})$ *is even, and*
- *two connected components if* $\gcd(\underline{\kappa})$ *is odd, distinguished by the generalized Arf invariant* (1) *of the relative framing on the real oriented blow-up.*

Combining this classification with Theorem 7.2 immediately implies that for $g \geq 4$ there are exactly two nonhyperelliptic components of $\Omega\mathcal{M}_g^{\text{pr}}(\underline{\kappa})$, classified by generalized Arf invariant (compare with Proposition 2.10).

Translating Theorem 7.5 into the action of the deck group *PR* therefore identifies the monodromy of the covering $\Omega\mathcal{M}_g^{\text{pr}}(\underline{\kappa}) \to \Omega\mathcal{M}_g^{\text{lab}}(\underline{\kappa})$:

Corollary 7.6. *Suppose that $g \geq 3$ and let $\mathcal{H}$ be a nonhyperelliptic connected component of $\Omega\mathcal{M}_g^{\text{lab}}(\underline{\kappa})$. Then the image of the monodromy homomorphism $\pi_1^{\text{orb}}(\mathcal{H}) \to PR$ is exactly the subgroup PR' of* (13).

In particular, when $\gcd(\underline{\kappa})$ is even the monodromy is all of *PR*.

[5] For example, the SO(2) action on $\Omega\mathcal{M}_g^{\text{lab}}(\underline{\kappa})$ [1, §2.10].

Proof. This is an easy consequence of Lemma 6.5 together with the formula for the generalized Arf invariant (1); simply observe that a prong rotation changes the winding number of every arc incident to the prong point by ± 1.

More explicitly, pick $\omega \in \mathcal{H}$ and a preimage $\omega_1^{\mathrm{pr}} \in \Omega\mathcal{M}_g^{\mathrm{pr}}(\underline{\kappa})$. Choose a distinguished geometric basis $\mathcal{B}_1$ on X^* with its basepoints specified by ω_1^{pr}. Then given another preimage ω_2^{pr} of the same $\omega \in \mathcal{H}$ there is a corresponding basis $\mathcal{B}_2$ which differs from $\mathcal{B}_1$ by a fractional multitwist τ (this choice is not unique, but is determined up to full twists about the boundary of the blow-up X^* of ω).

Computing the Arf invariants of ω_1^{pr} and ω_2^{pr} with respect to these bases, we see that the two Arf invariants agree (and hence ω_1^{pr} and ω_2^{pr} live in the same component of $\Omega\mathcal{M}_g^{\mathrm{pr}}(\underline{\kappa})$) if and only if

$$\sum_{i=2}^{n} \phi(a_i)\kappa_i = \sum\nolimits' \phi(a_i) = \sum\nolimits' \phi(\tau(a_i)) = \sum_{i=2}^{n} \phi(\tau(a_i))\kappa_i \pmod 2.$$

In particular, this equality holds if and only if $D(\tau) \in PR'$.

We observe that the choice of τ does not matter as any two choices τ and τ' differ by full twists about the boundary of X^*: twists about boundary components with even winding number do not change the parity of $\phi(\tau(a_i))$, while twists about odd boundaries change parity of arcs which do not end up contributing to $\sum'$. ∎

Remark 7.7. In addition to the differences in terminology, the statement and proof of Theorem 7.5 in [3] distinguishes the two nonhyperelliptic components of $\Omega\mathcal{M}_g^{\mathrm{pr}}(\underline{\kappa})$ by a different invariant. There, Boissy differentiates the two by first choosing a set of arcs on $\omega \in \Omega\mathcal{M}_g^{\mathrm{pr}}(\underline{\kappa})$ which pair up the odd order zeros, are transverse to the horizontal foliation, and are tangent to the specified prongs. Applying the "parallelogram construction" of [10] to ω along these arcs results in a new differential ω' of higher genus with all even zeros; the Arf invariant of the 2-spin structure induced by ω' then distinguishes the components of $\Omega\mathcal{M}_g^{\mathrm{pr}}(\underline{\kappa})$.

The reader may verify that Boissy's invariant coincides with the generalized Arf invariant by computing the contribution to $\mathrm{Arf}(\omega')$ of each new handle and comparing it to the corresponding term in the expression for $\mathrm{Arf}(\phi)$ (where ϕ is interpreted as a framing of X^*).

7.3. Markings and monodromy

In order to compare the framings induced by differentials on different Riemann surfaces, we pull them back to framings of a reference topological surface. To that end, we need to understand markings of X, X^*, and $(X, \vec{Z})$, together with the corresponding spaces of marked differentials.

The coarsest type of marking data we consider in this section are homeomorphisms from a closed surface Σ_g to X which take a specified set of (labeled) marked points P to the (labeled) zeros Z of ω. With this data, we can define the corresponding space $\Omega\mathcal{T}_g^{\mathrm{lab}}(\underline{\kappa})$ of marked differentials with marked points as the space of triples (X, ω, f), where $(X,\omega) \in \Omega\mathcal{M}_g^{\mathrm{lab}}(\underline{\kappa})$ and $f : (\Sigma, P) \to (X, Z)$ is an isotopy class of homeomorphisms of pairs. The space of marked differentials with marked points is naturally a (disconnected, orbifold) covering space of $\Omega\mathcal{M}_g^{\mathrm{lab}}(\underline{\kappa})$ whose deck transformations correspond to changing the marking, so its deck group is PMod_g^n.

On the other end of the spectrum, we may also mark a differential $\omega \in \Omega\mathcal{M}_g^{\mathrm{lab}}(\underline{\kappa})$ by a pronged surface. Fix a topological pronged surface $(\Sigma, \vec{P})$ with prong points of order $\kappa_1 + 1, \ldots, \kappa_n + 1$, and recall that any differential ω naturally equips its underlying surface X with a prong structure $\vec{Z}$ of the same prong type. Then a *marking* of (X,ω) by $(\Sigma, \vec{P})$ is a diffeomorphism of pairs f from (Σ, P) to (X, Z) such that Df takes the prong structure of $\vec{P}$ to that of $\vec{Z}$. Equivalently, f is a diffeomorphism from the real oriented blow-up Σ^* of Σ to X^* which takes the distinguished points[6] of $\partial\Sigma^*$ to those of ∂X^*.

We may now record the definition of the corresponding space of marked differentials:

Definition 7.8 (cf. [1, §2.9]). $\Omega\mathcal{T}_g^{\mathrm{pr}}(\underline{\kappa})$ is the space of marked differentials with a prong-marking, that is, the space of triples (X,ω,f) where (X,ω) is an abelian differential in $\Omega\mathcal{M}_g^{\mathrm{lab}}(\underline{\kappa})$ and $f : (\Sigma, \vec{P}) \to (X, \vec{Z})$ is an isotopy class of marking of pronged surfaces.

Forgetting the prong structure induces a covering map from $\Omega\mathcal{T}_g^{\mathrm{pr}}(\underline{\kappa})$ to $\Omega\mathcal{T}_g^{\mathrm{lab}}(\underline{\kappa})$, the deck group of which is exactly FT. Similarly, forgetting the marking of the surface but *remembering the marking of its boundary* (i.e., remembering $\partial f : \partial\Sigma_{g,n} \to \partial X^*$) induces a map from $\Omega\mathcal{T}_g^{\mathrm{pr}}(\underline{\kappa})$ to $\Omega\mathcal{M}_g^{\mathrm{pr}}(\underline{\kappa})$. Since this map remembers the boundary marking, hence the specified prong, the deck group of this covering is thus a change-of-marking group which preserves the boundary pointwise.

Lemma 7.9 ([1, Corollary 2.7]). *The space $\Omega\mathcal{T}_g^{\mathrm{pr}}(\underline{\kappa})$ is a (disconnected, orbifold) covering of both $\Omega\mathcal{M}_g^{\mathrm{pr}}(\underline{\kappa})$ and $\Omega\mathcal{T}_g^{\mathrm{lab}}(\underline{\kappa})$. Moreover, the deck group of the former covering is* $\mathrm{Mod}_{g,n}$, *and of the latter FT.*

Putting this lemma together with either (12) or (15), we see that the deck group of the covering $\Omega\mathcal{T}_g^{\mathrm{pr}}(\underline{\kappa}) \to \Omega\mathcal{M}_g^{\mathrm{lab}}(\underline{\kappa})$ is exactly the pronged mapping class group $\mathrm{Mod}_{g,n}^*$.

We summarize the relationship between all of these spaces (and their deck groups) in the following diagram, in which arrow labels correspond to the deck group of the covering:

$$\begin{tikzcd}
\Omega\mathcal{T}_g^{\mathrm{pr}}(\underline{\kappa}) \arrow[r, "FT"] \arrow[d, "\mathrm{Mod}_{g,n}"] \arrow[dr, "\mathrm{Mod}_{g,n}^*"] & \Omega\mathcal{T}_g^{\mathrm{lab}}(\underline{\kappa}) \arrow[d, "\mathrm{PMod}_g^n"] \arrow[dr, "\mathrm{Mod}_g^n"] & \\
\Omega\mathcal{M}_g^{\mathrm{pr}}(\underline{\kappa}) \arrow[r, "PR"] & \Omega\mathcal{M}_g^{\mathrm{lab}}(\underline{\kappa}) \arrow[r, "\mathrm{Sym}(\underline{\kappa})"] & \Omega\mathcal{M}_g(\underline{\kappa})
\end{tikzcd} \tag{16}$$

Observe that the leftmost triangle demonstrates the exact sequence (12), while the center triangle corresponds to (15).

[6] Recall that $\partial\Sigma^*$ can be identified with the circle of directions above each blown-up point, hence a prong point of order k corresponds to a boundary component with k distinguished points.

Constraining monodromy. While the deck groups of the coverings in (16) are easy to describe, their elements generally permute the components of the corresponding covers. We now shift our focus to the *stabilizer* of a component of one of these covers; for concreteness, throughout the rest of this section we will focus on the covering $\Omega\mathcal{T}_g^{\mathrm{pr}}(\underline{\kappa}) \to \Omega\mathcal{M}_g^{\mathrm{lab}}(\underline{\kappa})$ and use this discussion to deduce the corresponding results for each of the intermediate coverings.

We observe that by the path-lifting property, understanding the stabilizer of a component of $\Omega\mathcal{T}_g^{\mathrm{pr}}(\underline{\kappa})$ is equivalent to understanding the monodromy of the component of $\Omega\mathcal{M}_g^{\mathrm{lab}}(\underline{\kappa})$ which it covers (cf. [5, Proposition 3.7]).

Monodromy groups are always only defined up to conjugacy, so we fix some reference marking $f : (\Sigma, \vec{P}) \to (X, \vec{Z})$ (equivalently, a lift of a basepoint in $\Omega\mathcal{M}_g^{\mathrm{lab}}(\underline{\kappa})$ to $\Omega\mathcal{T}^{\mathrm{pr}}(\underline{\kappa})$). Pulling back the framing of $(X, \vec{Z})$ induced by ω along f induces a framing ϕ of Σ compatible with $\vec{P}$. Path-lifting then allows us to place the following constraint on the monodromy, which is just a version of [4, Corollary 4.8] adapted to the setting of framings rather than r-spin structures.

Lemma 7.10. *Let $\mathcal{H}$ be a component of $\Omega\mathcal{M}_g^{\mathrm{lab}}(\underline{\kappa})$ and fix some basepoint $(X, \omega) \in \mathcal{H}$. Then the monodromy of the covering of $\mathcal{H}$ by (a component of) $\Omega\mathcal{T}_g^{\mathrm{pr}}(\underline{\kappa})$ preserves the induced framing ϕ on the pronged surface X. In other words, the image of*

$$\rho : \pi_1^{\mathrm{orb}}(\mathcal{H}, (X, \omega)) \to \mathrm{Mod}_{g,n}^*$$

is contained inside of $\mathrm{Mod}_{g,n}^*[\phi]$.

Proof. Let f be a marking of $(X, \vec{Z})$ by $(\Sigma, \vec{P})$ and let γ be a loop in $\mathcal{H}$ based at (X, ω).

By Lemma 7.9, the loop γ lifts to a path $\tilde{\gamma}$ in $\Omega\mathcal{T}_g^{\mathrm{pr}}(\underline{\kappa})$, and hence a path of marked prong-marked abelian differentials (X_t, ω_t, f_t) with horizontal vector fields H_t. Pulling back H_t by f_t yields a continuous family of vector fields on $(\Sigma, \vec{P})$, all compatible with the prong structure. Let ϕ_t denote the associated framing of $(\Sigma, \vec{P})$.

Then since the vector fields vary continuously (and do not vanish on $\Sigma \setminus P$), the winding number $\phi_t(a)$ of every simple closed curve or legal arc a is continuous in t. However, $\phi_t(a)$ takes values only in $\mathbb{Z}$ or $\mathbb{Z} + \frac{1}{2}$ and must therefore be constant over the entire path $\tilde{\gamma}$. Thus $\rho(\gamma)$ preserves the winding number of every simple closed curve and legal arc, hence the entire framing. ∎

In view of the sequences (14) and (15), together with (16), this implies the following two results, where ϕ and $\bar{\phi}$ denote the relative and absolute framings induced on X^* and $X \setminus Z$, respectively.

Corollary 7.11. *The monodromy of the covering $\Omega\mathcal{T}_g^{\mathrm{pr}}(\underline{\kappa}) \to \Omega\mathcal{M}_g^{\mathrm{pr}}(\underline{\kappa})$ lies inside of* $\mathrm{Mod}_{g,n}[\phi]$.

Corollary 7.12. *The monodromy of the covering $\Omega\mathcal{T}_g^{\mathrm{lab}}(\underline{\kappa}) \to \Omega\mathcal{M}_g^{\mathrm{lab}}(\underline{\kappa})$ lies inside of* $\mathrm{PMod}_g^n[\bar{\phi}]$.

We note that these corollaries can also be proven directly using the same argument as Lemma 7.10.

For brevity, in what follows we use $\mathcal{G}^*$, $\mathcal{G}^{\mathrm{pr}}$, and $\mathcal{G}^{\mathrm{lab}}$ to denote the monodromy groups of the coverings appearing in Lemma 7.10 and Corollaries 7.11 and 7.12, respectively.

7.4. *Generating framed mapping class groups*

Now that we have shown that the monodromy of each of the coverings under consideration stabilizes the relevant framing, we can use Theorem B to show that the group is the *entire* stabilizer of the framing:

Theorem 7.13. *Suppose that $g \geq 5$ and let $\mathcal{H}$ be a nonhyperelliptic component of $\Omega\mathcal{M}_g^{\mathrm{pr}}(\underline{\kappa})$. Then*

$$\mathcal{G}^{\mathrm{pr}}(\mathcal{H}) \cong \mathrm{Mod}_{g,n}[\phi]$$

where ϕ is the relative framing of the real oriented blow-up X^* induced by the horizontal vector field on any $(X, \omega) \in \Omega\mathcal{M}_g^{\mathrm{pr}}(\underline{\kappa})$.

Assuming this theorem, we can leverage our understanding of the relationship between framing stabilizers to characterize both $\mathcal{G}^*$ and $\mathcal{G}^{\mathrm{lab}}$.

Theorem 7.14. *Suppose that $g \geq 5$ and let $\mathcal{H}$ be a nonhyperelliptic component of $\Omega\mathcal{M}_g^{\mathrm{lab}}(\underline{\kappa})$. Then*

$$\mathcal{G}^*(\mathcal{H}) \cong \mathrm{Mod}_{g,n}^*[\phi]$$

where ϕ is the relative framing of a pronged surface induced by the horizontal vector field of any $(X, \omega) \in \Omega\mathcal{M}_g^{\mathrm{lab}}(\underline{\kappa})$.

Proof. By Lemma 7.10, $\mathcal{G}^* \leq \mathrm{Mod}_{g,n}^*[\phi]$. By Theorem 7.5, $\mathcal{G}^*$ surjects onto PR', and so by Theorem 7.13 and Proposition 6.7, it follows that $\mathrm{Mod}_{g,n}^*[\phi] \leq \mathcal{G}^*$. ■

Forgetting the prongs, we can push this result down to the mapping class group with marked points to complete the proof of Theorem A.

Proof of Theorem A. Observe that $\mathcal{G}^{\mathrm{lab}}$ is the image of $\mathcal{G}^*$ under the surjection of (15). Therefore, Theorems 7.14 and 6.10 together imply that

$$\mathcal{G}^{\mathrm{lab}} = \mathrm{PMod}_g^n[\bar{\phi}].$$

Combining this result with Lemma 7.3 and the short exact sequence

$$1 \to \mathrm{PMod}_g^n[\bar{\phi}] \to \mathrm{Mod}_g^n[\bar{\phi}] \to \mathrm{Sym}(\underline{\kappa}) \to 1$$

completes the proof of the theorem. ■

Cylinder shears and prototypes. It therefore remains to show that $\mathrm{Mod}_{g,n}[\phi] \leq \mathcal{G}^{\mathrm{pr}}$. In order to demonstrate this, we will need to build a collection of loops of abelian differentials with prescribed monodromy.

Recall that a *cylinder* on an abelian differential is an embedded Euclidean cylinder with no singularities in its interior. Shearing the cylinder while leaving the rest of the surface fixed results in a loop in a stratum whose monodromy is the Dehn twist of the core curve of the cylinder.

Lemma 7.15 (cf. [4, Lemma 6.2]). *Let $\mathcal{H}$ be a component of $\Omega\mathcal{M}_g^{\mathrm{pr}}(\underline{\kappa})$ and suppose that $\omega \in \mathcal{H}$ has a cylinder with core curve c. Then $T_c \in \mathcal{G}^{\mathrm{pr}}$.*

Now that we can use cylinder shears to exhibit Dehn twists, it remains to show that there exist differentials in a stratum with a configuration of cylinders to which we can apply Theorem B.

Lemma 7.16 ([5, Lemma 3.14]). *Suppose that $g \geq 5$ and $\underline{\kappa} = (\kappa_1, \dots, \kappa_n)$ is a partition of $2g-2$. Let $\mathcal{H}$ be a nonhyperelliptic component of $\Omega\mathcal{M}_g(\underline{\kappa})$. Set $\mathcal{C}$ to be the curve system specified in Figure 15, of type 1 if $\gcd(\kappa_1, \dots, \kappa_n)$ is even and*

$$\mathrm{Arf}(\mathcal{H}) = \begin{cases} 1, & g \equiv 0, 3 \pmod 4, \\ 0, & g \equiv 1, 2 \pmod 4, \end{cases}$$

and of type 2 otherwise. Then there exists some $\omega \in \mathcal{H}$ whose horizontal and vertical cylinders are exactly the curves of $\mathcal{C}$.

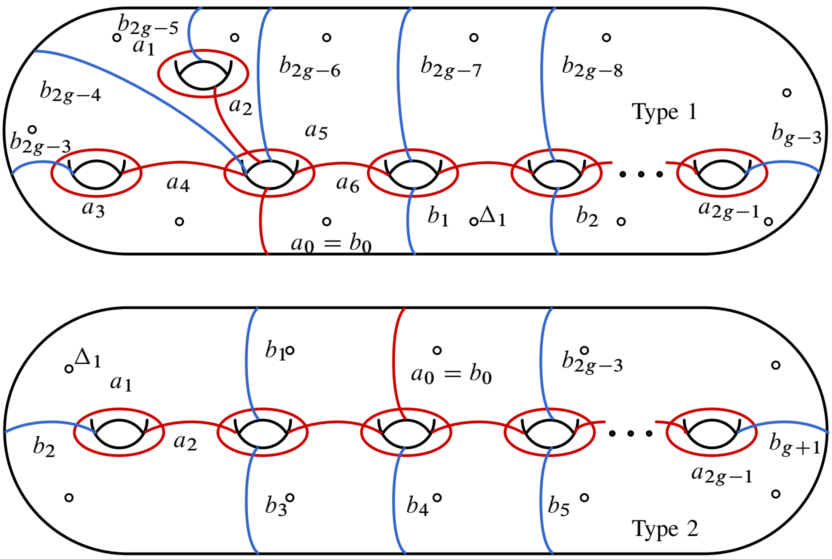

Fig. 15. Configurations of types 1 and 2 determining generating sets for $\mathrm{Mod}_{g,n}[\phi]$. We label the boundary components Δ_i for $i = 0, \dots, 2g-3$, with Δ_i positioned between b_i and b_{i+1} (for clarity most of the labels have been omitted). Given a partition $\underline{\kappa} = (\kappa_1, \dots, \kappa_n)$ of $2g-2$, only the n boundary components Δ_{i_ℓ} for i_ℓ of the form $\sum_{j=1}^{\ell} \kappa_j$ (interpreted mod $2g-2$) are included, and Δ_{i_ℓ} is assigned the signature $-1-\kappa_\ell$. Under this scheme, each complementary region determined by $\mathcal{C}$ contains exactly one boundary component, and the signatures are arranged so that each curve in $\mathcal{C}$ is admissible.

By labeling a separatrix at each zero, these prototype surfaces also give rise to prong-marked abelian differentials with specified Arf invariant.

Lemma 7.17. *Let g, $\underline{\kappa}$, $\mathcal{H}$, and $\mathcal{C}$ be as above. Let $\mathcal{H}_{\mathrm{pr}}$ be a* (*nonhyperelliptic*) *component of the stratum of prong-marked abelian differentials which covers $\mathcal{H}$. Then there exists some $\omega \in \mathcal{H}_{\mathrm{pr}}$ whose horizontal and vertical cylinders are exactly the curves of $\mathcal{C}$.*

Proof. Let $\omega \in \mathcal{H}$ be the differential constructed in Lemma 7.16. If $\gcd(\underline{\kappa})$ is even, then by Theorem 7.5 the entire preimage of $\mathcal{H}$ in $\Omega\mathcal{M}_g^{\mathrm{lab}}(\underline{\kappa})$ is $\mathcal{H}_{\mathrm{pr}}$ and so any prong marking of ω lives in $\mathcal{H}_{\mathrm{pr}}$. If $\gcd(\underline{\kappa})$ is odd, then choose an arbitrary prong marking ω_{pr} of ω and let p_i be some zero of odd order (corresponding to a prong point of even order).

Now by Corollary 7.6, since $T_{\vec{p}_i} \notin PR'$ the differentials ω_{pr} and $T_{\vec{p}_i}(\omega_{\mathrm{pr}})$ must lie in different components of $\Omega\mathcal{M}_g^{\mathrm{pr}}(\underline{\kappa})$. Since there are only two components (by Theorem 7.5), one of $\{T_{\vec{p}_i}(\omega_{\mathrm{pr}}), \omega_{\mathrm{pr}}\}$ must be in $\mathcal{H}_{\mathrm{pr}}$. ∎

It is now completely straightforward to deduce that the monodromy group of a component $\mathcal{H}_{\mathrm{pr}}$ of $\Omega\mathcal{M}_g^{\mathrm{pr}}(\underline{\kappa})$ is exactly the stabilizer of the relative framing induced on the real oriented blow-up.

Proof of Theorem 7.13. Let $(X, \omega) \in \mathcal{H}_{\mathrm{pr}}$ be the prototype surface built in Lemma 7.17 above. Observe that each configuration of curves specified in Figure 15 spans the indicated surface, has intersection graph a tree, and contains E_6 as a subgraph. Then by Theorem B, Lemma 7.15, and Corollary 7.11, we have

$$\mathrm{Mod}_{g,n}[\phi] = \langle T_c \mid c \in \mathcal{C} \rangle \le \mathcal{G}^{\mathrm{pr}} \le \mathrm{Mod}_{g,n}[\phi],$$

completing the proof of Theorem 7.13. ∎

8. Further corollaries

In this final section we collect some further corollaries of the work we have done.

8.1. Classification of components

The monodromy computations in Theorems A, 7.13, and 7.14 lead to the following classification of the nonhyperelliptic components of strata of marked differentials (cf. [5, Theorem A]):

Corollary 8.1. *There is a bijection between the nonhyperelliptic components of $\Omega\mathcal{T}_g^{\mathrm{lab}}(\underline{\kappa})$ and the isotopy classes of absolute framings of Σ_g^n with signature $-1-\underline{\kappa}$.*

If $\gcd(\underline{\kappa})$ is odd then the permutation action of Mod_g^n is transitive, while if $\gcd(\underline{\kappa})$ is even there are two orbits, classified by the Arf invariant.

Corollary 8.2. *There is a bijection between the nonhyperelliptic components of $\Omega\mathcal{T}_g^{\mathrm{pr}}(\underline{\kappa})$ and the* (*relative*) *isotopy classes of relative framings of $\Sigma_{g,n}$ with signature $-1-\underline{\kappa}$.*

The action of $\mathrm{Mod}^*_{g,n}$ *is transitive if* $\gcd(\underline{\kappa})$ *is odd and has two orbits if* $\gcd(\underline{\kappa})$ *is even, classified by the Arf invariant of the absolute framing. The action of* $\mathrm{Mod}_{g,n}$ *has two orbits no matter the parity of* $\gcd(\underline{\kappa})$*, classified by the generalized Arf invariant.*

The proofs of both corollaries are simply a consequence of the orbit-stabilizer theorem applied to the PMod_g^n action on the set of components of $\Omega\mathcal{T}_g^{\mathrm{lab}}(\underline{\kappa})$, respectively the $\mathrm{Mod}_{g,n}$ action on the components of $\Omega\mathcal{T}_g^{\mathrm{pr}}(\underline{\kappa})$, together with the classification of orbits (Theorems 7.2 and 7.5, respectively).

8.2. *Cylinder shears and fundamental groups of strata*

For a component $\mathcal{H}$ of a general stratum $\Omega\mathcal{M}_g(\underline{\kappa})$, no explicit set of generators for $\pi_1^{\mathrm{orb}}(\mathcal{H})$ is known. Cylinder shears play a role analogous to Dehn twists in the theory of the mapping class group, and it is natural to wonder about the extent to which shears generate $\pi_1^{\mathrm{orb}}(\mathcal{H})$. If the partition $\underline{\kappa}$ contains any repeated elements, the corresponding zeros can be exchanged, but this certainly cannot be accomplished using shears. Thus one must first pass to a "labeled stratum," i.e., a component $\mathcal{H}_{\mathrm{lab}}$ of the cover $\Omega\mathcal{M}_g^{\mathrm{lab}}(\underline{\kappa})$. Even here, the work of Boissy (in the guise of Corollary 7.6) implies that $\pi_1^{\mathrm{orb}}(\mathcal{H}_{\mathrm{lab}})$ is *never* generated by shears alone, since shears map trivially onto the prong rotation group PR. However, prong rotation is not detected by the monodromy representation $\rho : \pi_1^{\mathrm{orb}}(\mathcal{H}_{\mathrm{lab}}) \to \mathrm{PMod}_g^n[\bar{\phi}]$, only by the refinement whose target is the pronged mapping class group.

As a corollary of our monodromy computations and Corollary 6.15, we find that the arithmetic of the partition $\underline{\kappa}$ provides an obstruction for the subgroup of $\pi_1^{\mathrm{orb}}(\mathcal{H}_{\mathrm{lab}})$ generated by cylinder shears to generate the monodromy group in PMod_g^n. Thus, the prong rotation group "leaves a trace" in the group $\mathrm{PMod}_g^n[\bar{\phi}]$, even though there is no way of measuring prong rotation in $\mathrm{PMod}_g^n[\bar{\phi}]$ directly.

Corollary 8.3. *Let* $\underline{\kappa}$ *be a partition of* $2g-2$ *for* $g \geq 5$ *and let* $\mathcal{H}_{\mathrm{lab}}$ *be a nonhyperelliptic component of the stratum* $\Omega\mathcal{M}_g^{\mathrm{lab}}(\underline{\kappa})$*. Write* $\underline{\kappa} = (\eta_1, \ldots, \eta_p, \upsilon_1, \ldots, \upsilon_q)$ *where* η_i *are even,* υ_j *are odd, and* $p + q = n$*. If* $q > 2$*, or if the elements of*

$$\left\{\eta_1 + 1, \ldots, \eta_p + 1, \frac{\upsilon_1 + 1}{2}, \ldots, \frac{\upsilon_q + 1}{2}\right\}$$

are not pairwise coprime, then the subgroup of $\pi_1^{\mathrm{orb}}(\mathcal{H}_{\mathrm{lab}})$ *generated by cylinder shears does not surject onto* $\mathcal{G}^{\mathrm{lab}}$.

Proof. Choose an arbitrary component $\mathcal{H}_{\mathrm{pr}}$ of the preimage of $\mathcal{H}_{\mathrm{lab}}$ in $\Omega\mathcal{M}_g^{\mathrm{pr}}(\underline{\kappa})$ and let ϕ denote the induced relative framing. Let $\mathcal{C}$ denote the subgroup of $\pi_1^{\mathrm{orb}}(\mathcal{H}_{\mathrm{lab}})$ generated by cylinder shears (or rather, by elements which are conjugate to cylinder shears by some path connecting a basepoint (X, ω) to the surface (Y, ω) realizing the relevant cylinder). As cylinder shears preserve prong-markings, we see that $\mathcal{C} \leq \pi_1^{\mathrm{orb}}(\mathcal{H}_{\mathrm{pr}})$.

Recall that $\rho_{\mathrm{pr}} : \pi_1^{\mathrm{orb}}(\mathcal{H}_{\mathrm{pr}}) \to \mathrm{Mod}_{g,n}$ denotes the monodromy representation of $\mathcal{H}_{\mathrm{pr}}$ with image $\mathcal{G}^{\mathrm{pr}}$, and that $\mathcal{G}^{\mathrm{lab}}$ denotes the image of the monodromy representation of $\mathcal{H}_{\mathrm{lab}}$

into PMod_g^n. Now Theorem A finds that $\mathcal{G}^{\mathrm{lab}} = \mathrm{PMod}_g^n[\bar{\phi}]$, and Theorem 7.13 finds that $\rho_{\mathrm{pr}}(\mathcal{C}) \leq \mathcal{G}^{\mathrm{pr}} = \mathrm{Mod}_{g,n}[\phi]$.[7] By Corollary 6.15, under our current hypotheses, the map $\mathrm{Mod}_{g,n}[\phi] \to \mathrm{PMod}_g^n[\bar{\phi}]$ is never surjective. Therefore, the image of $\mathcal{C}$ in $\mathcal{G}^{\mathrm{lab}}$ is a strict subgroup. ■

8.3. Change of coordinates for saddle connections

In this section, we use the machinery of prong-markings and the framed change-of-coordinates principle to prove Corollary 1.3 (realization of arcs as saddles). As in Corollary 1.2, the idea is to use the framed change-of-coordinates principle to take a given arc to some target arc which is realized as a saddle connection. However, unlike cylinders, saddle connections do not share a common winding number. The main difficulty in the proof of Corollary 1.3 is therefore to construct a sufficiently large set of saddle connections to play the role of target arcs (Lemma 8.5).

We begin by clarifying some conventions with regards to arcs on surfaces with boundary versus surfaces with marked points. Recall that if $(\Sigma_{g,n}, \phi)$ is a relatively framed surface, then we have fixed once and for all a legal basepoint on each boundary component and we only consider arcs ending on this prescribed basepoint. When $(\Sigma_{g,n}, \phi)$ arises from the blow-up of a prong-marked abelian differential, this is equivalent to stipulating that arcs must leave and enter the zeros with prescribed tangent directions.

Upon capping each boundary component of $\Sigma_{g,n}$ with a punctured disk (equivalently, forgetting the prong structure coming from the differential), each (relative) isotopy class of arc a on $\Sigma_{g,n}$ projects to an (absolute) isotopy class of arc on Σ_g^n, which we will denote by $\pi(a)$. Observe that the map π is not injective; elements of its fibers are related by Dehn twists about the endpoint(s) of the arc.

Saddle connections on one-cylinder differentials. In order to exhibit the desired collection of saddles, it will be convenient to use different model surfaces than the ones introduced in §7.4.

To that end, recall that an abelian differential is called *Jenkins–Strebel* if it is completely horizontally periodic, i.e., if it can be written as the union of closed horizontal cylinders. We will in particular be interested in those which can be obtained by identifying boundary edges of a single cylinder, called *one-cylinder* Jenkins–Strebel differentials.

The existence result we will use is the following; see [31, Section 2] for an explicit construction.

Proposition 8.4. *There exists a one-cylinder Jenkins–Strebel differential* (Y, η) *in every* (*nonempty*) *component of every stratum* $\Omega\mathcal{M}_g(\underline{\kappa})$.

As in the proof of Lemma 7.17, we can also upgrade these η to yield one-cylinder Jenkins–Strebel differentials in each component of $\Omega\mathcal{M}_g^{\mathrm{lab}}(\underline{\kappa})$ and $\Omega\mathcal{M}_g^{\mathrm{pr}}(\underline{\kappa})$ by labeling zeros and prongs, respectively.

[7] In fact, the proof of Theorem 7.13 shows that the image of $\mathcal{C}$ under the monodromy map is all of $\mathrm{Mod}_{g,n}[\phi]$.

Using these new model differentials, we may now exhibit saddle connections which have (a preimage under π which has) arbitrary winding number.

Lemma 8.5. *Let $\mathcal{H}_{\mathrm{pr}}$ be a component of $\Omega\mathcal{M}_g^{\mathrm{pr}}(\underline{\kappa})$ and let p, q be two distinct zeros. Then for every $s \in \mathbb{Z} + \frac{1}{2}$, there is a differential $(Y, \eta) \in \mathcal{H}_{\mathrm{pr}}$ and nonseparating arc a on Y^* from Δ_p to Δ_q (based at the legal basepoints prescribed by the prong-marking) such that*

(1) *$\phi(a) = s$ for the relative framing ϕ induced by $1/\eta$,*

(2) *the geodesic representative of $\pi(a)$ on $(Y, \mathrm{Zeros}(\eta))$ is a saddle connection.*

Proof. As twisting around Δ_q changes $\phi(a)$ by $\pm\phi(\Delta_q)$, it suffices to prove that there exist such arcs for each residue class modulo $\phi(\Delta_q)$.

Let (Y, η) be a one-cylinder Jenkins–Strebel differential in $\mathcal{H}_{\mathrm{pr}}$ (which exists by Proposition 8.4). Now by definition, we can write $Y = C/\sim$, where $C = S^1 \times [0, 1]$ is a closed cylinder and $\sim$ identifies segments of the "top boundary" $S^1 \times \{1\}$ with segments of the "bottom boundary" $S^1 \times \{0\}$. Let $Q : C \to Y$ denote the quotient map.

We observe that our choice of prong at p determines a unique half-separatrix on (Y, η) and hence a pair of half-segments in ∂C. In particular, the prong-marking determines a unique point $\tilde{p} \in Q^{-1}(p) \cap (S^1 \times \{0\})$.

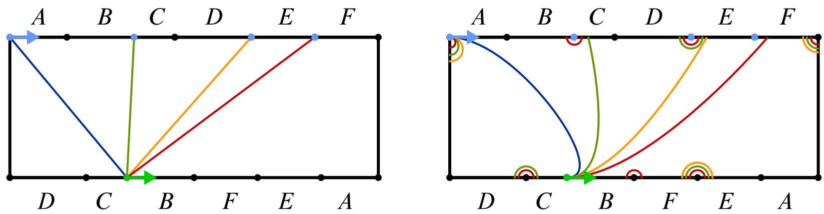

Fig. 16. A one-cylinder Jenkins–Strebel differential in $\Omega\mathcal{M}_3(3, 1)$ and saddle connections on it. On the left the arcs have been realized geodesically; on the right they have been realized with prescribed tangential data.

Now consider the set $\mathcal{A}$ of all arcs in C which start at $\tilde{p}$ and end at a point of $Q^{-1}(q) \cap S^1 \times \{1\}$; up to Dehn twisting along the core curve of C, there are exactly $\phi(\Delta_q)$ such arcs (see Figure 16). Moreover, since the arcs of $\mathcal{A}$ are each realized as straight line segments on C, the arcs of $\pi(Q(\mathcal{A}))$ are all realized as saddle connections on (Y, η). Isotoping these arcs to leave p and enter q with the prescribed tangential data, we may measure the winding numbers of $a \in Q(\mathcal{A})$. Careful inspection of Figure 16 shows that

$$\{\phi(a) \bmod \phi(\Delta_q) : a \in Q(\mathcal{A})\} = \left\{\frac{1}{2}, \frac{3}{2}, \ldots, \frac{2\phi(\Delta_q) - 1}{2}\right\},$$

finishing the proof of the lemma. ∎

Now that we have a sufficiently large collection of target saddle connections, we may apply the framed change-of-coordinates principle to deduce Corollary 1.3.

Proof of Corollary 1.3. Let $\bar{a}$ be an arc on (X, ω) with endpoints p and q. Choose an arbitrary prong-marking of (X, ω) and an arc a on X^* such that $\pi(a) = \bar{a}$. Let (Y, η) be the one-cylinder Jenkins–Strebel differential from Lemma 8.5 in the same component of $\Omega\mathcal{M}_g^{\mathrm{pr}}(\underline{\kappa})$ as (X,ω) and let ϕ denote the induced relative framing on Y^*.

Choose a path α connecting (X,ω) to (Y,η) and let $\alpha_*(a)$ denote the parallel transport of a along α (equivalently, lift α to a path in $\Omega\mathcal{T}_g^{\mathrm{pr}}(\underline{\kappa})$ and use the marking maps). Then by Lemma 8.5 there is an arc b on Y^* with $\phi(b) = \phi(\alpha_*(a))$ and so that $\pi(b)$ is a saddle connection on (Y,η).

Now by the framed change-of-coordinates principle (Proposition 2.15), there is an element of $\mathrm{Mod}_{g,n}[\phi]$ taking $\alpha_*(a)$ to b. By Theorem 7.13, this element can be represented by a loop β in $\mathcal{H}_{\mathrm{pr}}$. The concatenated path $\alpha \cdot \beta$ therefore takes a on X^* to b on Y^*, and so its projection to $\Omega\mathcal{M}_g(\underline{\kappa})$ is the desired path. ■

Appendix A. Summary of generating set results

In §§A.1–A.4 of this appendix we collect the results concerning generating sets for the various framed mapping class groups studied throughout the paper. In §A.5, we illustrate these results with an example. Finally, in §A.6, we examine the extent to which these generating sets are minimal. For the sake of readability and relative self-containment, we have repeated and reproduced what is necessary to understand the statements, although the reader may need to consult §§1, 2 and/or 6 of the paper for further background.

A.1. The framed mapping class group $\mathrm{Mod}_{g,n}[\phi]$

Here we consider a surface $\Sigma_{g,n}$ of genus $g \geq 5$ with $n \geq 1$ boundary components, and a relative isotopy class of framing ϕ of $\Sigma_{g,n}$ of either holomorphic or meromorphic type. Recall that a framing ϕ is of *holomorphic type* if each boundary component has negative winding number (when oriented with the surface to the left), and is of *meromorphic type* otherwise.

Theorem B asserts that $\mathrm{Mod}_{g,n}[\phi]$ admits generating sets consisting of $2g+n-1$ Dehn twists; in fact, many combinatorially distinct such sets. We recall here the notions of *E-arboreal spanning configurations* and *h-assemblage* that underpin the theory of generating sets. Let $\mathcal{C} = \{c_1, \ldots, c_k\}$ be a collection of curves on a surface $\Sigma_{g,n}$, pairwise in minimal position, with the property that the geometric intersection number $i(c_i, c_j)$ is at most 1 for all pairs $c_i, c_j \in \mathcal{C}$. Associated to such a configuration is its *intersection graph* $\Lambda_{\mathcal{C}}$, whose vertices correspond to the elements of $\mathcal{C}$, with c_i and c_j joined by an edge whenever $i(c_i, c_j) = 1$. Such a configuration $\mathcal{C}$ *spans* $\Sigma_{g,n}$ if there is a deformation retraction of $\Sigma_{g,n}$ onto the union of the curves in $\mathcal{C}$. We say that $\mathcal{C}$ is *arboreal* if the intersection graph $\Lambda_{\mathcal{C}}$ is a tree, and *E-arboreal* if $\Lambda_{\mathcal{C}}$ moreover contains the E_6 Dynkin diagram as a subgraph.

An *h-assemblage of type E* on $\Sigma_{g,n}$ is a set of curves

$$\mathcal{C} = \{c_1, \ldots, c_{2h+m-1}, c_{2h+m}, \ldots, c_{2g+n-1}\}$$

such that (1) $\mathcal{C}_1 = \{c_1, \ldots, c_{2h+m-1}\}$ is an E-arboreal spanning configuration on a subsurface $S \subset \Sigma_{g,n}$ of genus $g(S) = h$ with m boundary components, (2) for $j \geq k$, let

S_j denote a regular neighborhood of the curves $\{c_1, \ldots, c_j\}$; then for $j \geq 2h + m$, we require that $c_j \cap S_{j-1}$ be a single arc (possibly, but not necessarily, entering and exiting along the same boundary component of S_j), and (3) $S_{2g+n-1} = \Sigma_{g,n}$. In other words, an assemblage of type E is built from an E-arboreal spanning configuration on a subsurface by sequentially attaching (neighborhoods of) further curves, decreasing the Euler characteristic by exactly 1 at each stage but otherwise allowing the new curves to intersect individual old curves arbitrarily.

Theorem B shows that E-arboreal spanning configurations and h-assemblages consisting of admissible curves generate the relatively framed mapping class group.

Theorem B. *Let $\Sigma_{g,n}$ be a surface of genus $g \geq 5$ with $n \geq 1$ boundary components.*

(I) *Suppose ϕ is a framing of $\Sigma_{g,n}$ of holomorphic type. Let $\mathcal{C} = \{c_1, \ldots, c_{2g+n-1}\}$ be an E-arboreal spanning configuration of curves on $\Sigma_{g,n}$ such that $\phi(c) = 0$ for all $c \in \mathcal{C}$. Then*

$$\mathrm{Mod}_{g,n}[\phi] = \langle T_c \mid c \in \mathcal{C} \rangle.$$

(II) *If ϕ is an arbitrary framing (of holomorphic or meromorphic type) and $\mathcal{C} = \{c_1, \ldots, c_{2g+n-1}\}$ is an h-assemblage of type E for $h \geq 5$ of curves such that $\phi(c) = 0$ for all $c \in \mathcal{C}$, then*

$$\mathrm{Mod}_{g,n}[\phi] = \langle T_c \mid c \in \mathcal{C} \rangle.$$

A.2. Framed mapping class groups of pronged surfaces $\mathrm{Mod}^*_{g,n}[\phi]$

Here we consider a surface $(\Sigma, \vec{P})$ equipped with a finite set of n prong points (Definition 6.2). Suppose ϕ is a framing of holomorphic type compatible with $\vec{P}$ (Definition 6.3). We obtain the framed pronged mapping class group $\mathrm{Mod}^*_{g,n}[\phi]$. To state a generating set for $\mathrm{Mod}^*_{g,n}[\phi]$, we recall the definition of an *auxiliary curve* from Definition 6.8 as well as the definition of the *auxiliary twists* of types 1 and 2 from the proof of Proposition 6.7 and an *auxiliary curve system* from Definition 6.11.

Definition (Auxiliary curve, auxiliary twist, auxiliary curve system)**.** Let $(\Sigma_{g,n}, \phi)$ be a framed surface with boundary components $\Delta_1, \ldots, \Delta_n$ arising as the blow-up of a set of prong points $\vec{P} = \{\vec{p}_1, \ldots, \vec{p}_n\}$, with $\vec{p}_i$ a prong point of order k_i.

An *auxiliary curve for* Δ_i is a separating curve d_i such that d_i separates Δ_i from the remaining boundary components, and with $\phi(d_i) = \pm 1$ or ± 2 according to whether $k_i = \phi(\Delta_i)$ is odd or even. If k_i is odd, we define an *auxiliary twist of type* 1 to be the mapping class

$$A_i := T_{\Delta_i}^{1/k_i} T_{d_i}^{-1}$$

(where d_i is as above), and if k_i is even, we define it to be

$$A_i := T_{\Delta_i}^{2/k_i} T_{d_i}^{-1}.$$

An *auxiliary curve for* Δ_i *and* Δ_j is any separating curve $c_{i,j}$ that separates the boundary components Δ_i, Δ_j from the remaining components. Define the *auxiliary twist of type* 2 to be the mapping class

$$B_{i,j} := T_{\Delta_i}^{\phi(c_{i,j})/k_i} T_{\Delta_j}^{\phi(c_{i,j})/k_j} T_{c_{i,j}}^{-1}.$$

An *auxiliary curve system* is a collection of the following auxiliary curves:

- auxiliary curves $c_{i,j}$ for all pairs i, j such that both $\phi(\Delta_i), \phi(\Delta_j)$ are even,
- auxiliary curves d_k for all indices k.

No requirements are imposed on the intersection pattern of curves in an auxiliary curve system.

The following generating set for $\mathrm{Mod}_{g,n}^*[\phi]$ was obtained in the course of proving Corollary 6.12.

Proposition A.1. *Let $\vec{P} = \{\vec{p}_1, \ldots, \vec{p}_n\}$ be a set of $n \geq 1$ prong points on Σ_g for $g \geq 5$, and let ϕ be a framing of $(\Sigma_g, \vec{P})$ of holomorphic type compatible with $\vec{P}$. Let $\mathcal{A} = \{c_{i,j}, d_k\}$ be an auxiliary curve system, and let $\mathcal{C} = \{c_1, \ldots, c_{2g+n-1}\}$ be any set of curves forming a generating system for the relative framed mapping class group $\mathrm{Mod}_{g,n}[\phi]$ as in Theorem B. Then $\mathrm{Mod}_{g,n}^*[\phi]$ is generated by the Dehn twists about the curves in $\mathcal{C}$ along with the auxiliary twists in the set $\mathcal{A}$.*

Remark A.2. On the pronged surface, the auxiliary twists A_i and $B_{i,j}$ are actually “fractional multitwists,” not genuine single Dehn twists.

A.3. The absolute framed mapping class group $\mathrm{PMod}_g^n[\bar{\phi}]$

Here we consider a closed surface Σ_g^n of genus $g \geq 5$ equipped with a set of $n \geq 1$ marked points. We equip Σ_g^n with a framing $\bar{\phi}$ of holomorphic type. In an obvious way, we can consider $\bar{\phi}$ as a framing ϕ on an associated surface $\Sigma_{g,n}$ with boundary. The following is a slight reformulation of Corollary 6.12.

Corollary 6.12. *Let $\mathcal{C}$ be any set of curves forming a generating system for the relative framed mapping class group $\mathrm{Mod}_{g,n}[\phi]$, and let $\mathcal{A}$ be an auxiliary curve system for ϕ. Then $\mathrm{PMod}_g^n[\bar{\phi}]$ is generated by the Dehn twists along the curves of $\mathcal{C} \cup \mathcal{A}$.*

A.4. r-spin mapping class groups on closed surfaces

Here we consider a closed surface Σ_g of genus $g \geq 5$ equipped with an r-spin structure $\widehat{\phi}$ for some r dividing $2g-2$. The following is a slight reformulation of Corollary 3.11.

Corollary 3.11. *Let $\mathcal{C}$ denote a filling network of curves on a* closed *surface Σ_g with $g \geq 5$. Suppose that the intersection graph $\Lambda_{\mathcal{C}}$ is a tree which contains the E_6 Dynkin diagram as a subgraph and that $\mathcal{C}$ cuts the surface into n polygons with*

$4(k_1+1), \ldots, 4(k_n+1)$ sides. Set $r = \gcd(k_1, \ldots, k_n)$ and let $\widehat{\phi}$ be the r-spin structure determined by the condition that each curve $c \in \mathcal{C}$ be admissible. Then

$$\mathrm{Mod}_g[\widehat{\phi}] = \langle T_c \mid c \in \mathcal{C} \rangle.$$

In particular, if $r = 1$ then $\{T_c\}$ generate the entire mapping class group Mod_g.

A.5. An example

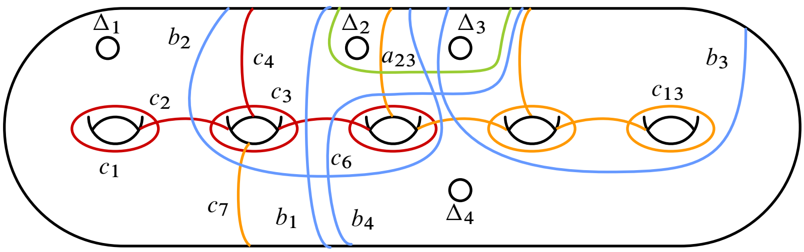

Fig. 17

Consider the surface $\Sigma_{5,4}$ shown in Figure 17. We see that the red curves $(c_1, \ldots, c_6)$ form an E_6 configuration, and the configuration $\mathcal{C} = \{c_1, \ldots, c_{13}\}$ of red and orange curves is E-arboreal and spanning. By Theorem B, $\mathrm{Mod}(\Sigma_{5,4})[\phi]$ is generated by the Dehn twists about curves in $\mathcal{C}$, where ϕ is the framing determined by the condition that each curve in $\mathcal{C}$ is admissible.

We compute that

$$\phi(\Delta_1) = -3, \quad \phi(\Delta_2) = \phi(\Delta_3) = -2, \quad \phi(\Delta_4) = -5.$$

We next suppose that each Δ_i is replaced with a prong $\vec{p}_i$ of order $k_i = -\phi(\Delta_i)$. Then the green curve a_{23} forms an auxiliary curve for the prongs $\vec{p}_2, \vec{p}_3$, both of even order, and the blue curves $b_1, \ldots, b_4$ form a set of auxiliary curves of type 1 for the associated prongs. By Proposition A.1, the pronged framed mapping class group $\mathrm{Mod}_{5,4}^*[\phi]$ is generated by the Dehn twists about the curves in $\mathcal{C}$ along with the auxiliary twists associated to $a_{23}, b_1, \ldots, b_4$.

We suppose next that the prongs $\vec{p}_1, \ldots, \vec{p}_4$ are replaced with the corresponding set $p_1, \ldots, p_4$ of ordinary marked points. Corollary 6.12 shows that the absolute framed mapping class group $\mathrm{PMod}(\Sigma_5^4)[\phi]$ is generated by the Dehn twists about the curves $c_1, \ldots, c_{13}, a_{23}, b_1, \ldots, b_4$.

If one forgets the marked points $p_1, \ldots, p_4$ entirely, the corresponding decomposition of the closed surface Σ_5 consists of four polygons with $12, 8, 8, 20$ sides, respectively. The gcd r as computed in Corollary 3.11 is then 1; it is not hard to see that the twists about the curves $c_1, \ldots, c_{13}$ generate the mapping class group $\mathrm{Mod}(\Sigma_5)$.

A.6. Minimality

Here we examine the extent to which the above generating sets are *minimal*. In general, the minimal number of generators depends strongly on the conjugacy class(es) involved – e.g., the full mapping class group $\mathrm{Mod}(\Sigma_g)$ requires $2g+1$ Dehn twists to generate [11, Proposition 6.5] but can be generated by as few as two finite-order elements [30]. Here we consider only generation by Dehn twists. The cases of pronged surfaces and punctured surfaces with absolute framings are somewhat more intricate, owing to the necessity of including additional auxiliary twists, and we do not consider them here. On the other hand, the results we can obtain for relatively framed mapping class groups and for r-spin mapping class groups are nearly optimal. The arguments below are generally modeled on [11, Proposition 6.5], although Proposition A.4 makes the connection with spin structures more explicit.

Proposition A.3. *The framed mapping class group* $\mathrm{Mod}(\Sigma_{g,n})[\phi]$ *is not generated by fewer than* $2g+n-1$ *Dehn twists. Thus the generating sets of Theorem* B *are minimal.*

Proof. We consider the action of the group $\mathrm{Mod}(\Sigma_{g,n})[\phi]$ on the relative homology $H_1(\Sigma_{g,n}, \partial\Sigma_{g,n}; \mathbb{Q})$. The action of the Dehn twist T_c on $H_1(\Sigma_{g,n}, \partial\Sigma_{g,n}; \mathbb{Q})$ is given by the formula

$$T_c(x) = x + \langle x, [c]\rangle[\bar{c}], \tag{17}$$

where $[c]$ (resp. $[\bar{c}]$) denotes the class of the curve c in $H_1(\Sigma_{g,n}; \mathbb{Q})$ (resp. its image in $H_1(\Sigma_{g,n}, \partial\Sigma_{g,n}; \mathbb{Q})$), and

$$\langle \cdot, \cdot \rangle : H_1(\Sigma_{g,n}, \partial\Sigma_{g,n}; \mathbb{Q}) \otimes H_1(\Sigma_{g,n}; \mathbb{Q}) \to \mathbb{Q}$$

is induced by the intersection pairing between absolute and relative classes.

For a given $[c] \in H_1(\Sigma_{g,n}; \mathbb{Q})$, the function $\langle \cdot, [c]\rangle$ is a linear functional on $H_1(\Sigma_{g,n}, \partial\Sigma_{g,n}; \mathbb{Q})$ and thus acts trivially on a subspace of codimension at most 1. The space $H_1(\Sigma_{g,n}, \partial\Sigma_{g,n}; \mathbb{Z})$ has dimension $2g+n-1$, and so any collection of fewer than $2g+n-1$ Dehn twists necessarily acts trivially on some nontrivial subspace.

On the other hand, the action of $\mathrm{Mod}(\Sigma_{g,n})[\phi]$ fixes no such subspace. Any element $x \in H_1(\Sigma_{g,n}, \partial\Sigma_{g,n}; \mathbb{Q})$ can be represented by a single simple closed curve or arc α, in either case nonseparating, equipped with a rational weight. By the framed change-of-coordinates principle, one can construct an admissible curve β crossing transversely through α exactly once. Then (17) shows that the Dehn twist about β does not fix x. ■

We next consider minimality in the r-spin case. An Euler characteristic argument shows that one needs at least $2g$ curves in order to cut Σ_g into a union of one or more polygons. In the case $r = 2g-2$, for either value of the Arf invariant, there are configurations of $2g$ curves whose complement is a single polygon (see, e.g., [5, Theorem C]), but at least $2g+1$ curves are necessary to divide Σ_g into at least two polygons and hence make r a proper divisor of $2g-2$ in Corollary 3.11, at least if they are arranged in a network whose intersection graph is a tree. Considering the partition $\{r, 2g-2-r\}$

of $2g-2$, one can again use the construction detailed in [5, Theorem C] to create configurations of exactly $2g+1$ curves that generate $\mathrm{Mod}(\Sigma_g)[\hat{\phi}]$. Thus, for any r-spin structure $\hat{\phi}$, we find generating sets for $\mathrm{Mod}(\Sigma_g)[\hat{\phi}]$ consisting of $2g$ generators (in the case $r = 2g-2$) or else of $2g+1$ generators (for $r < 2g-2$). Proposition A.4 below shows that these generating sets are minimal in the cases r odd or $r = 2g-2$, and are nearly minimal otherwise.

Proposition A.4. *Let $\hat{\phi}$ be an r-spin structure. For r even,* $\mathrm{Mod}(\Sigma_g)[\hat{\phi}]$ *is generated by no fewer than* $2g$ *Dehn twists, and for r odd,* $\mathrm{Mod}(\Sigma_g)[\hat{\phi}]$ *is generated by no fewer than* $2g+1$ *Dehn twists.*

Proof. We follow the same outline as in Proposition A.3; for the sake of subsequent arguments we switch coefficients from $\mathbb{Q}$ to $\mathbb{F}_2$. By considering the action of a Dehn twist on $H_1(\Sigma_g;\mathbb{F}_2)$, one concludes as above that at least $2g$ Dehn twists are necessary, and that moreover the corresponding set of homology classes must span $H_1(\Sigma_g;\mathbb{F}_2)$. Now suppose that r is odd. Let $c_1,\dots,c_{2g}$ be any set of admissible curves. We assume that $\{[c_1],\dots,[c_{2g}]\}$ forms a basis for $H_1(\Sigma_g;\mathbb{F}_2)$; otherwise the above paragraph shows that $\{T_{c_1},\dots,T_{c_{2g}}\}$ does not generate $\mathrm{Mod}(\Sigma_g)[\hat{\phi}]$. Since $\{[c_1],\dots,[c_{2g}]\}$ is linearly independent, one can construct a 2-spin structure q for which each $[c_i]$ is admissible, so that the action of $\{T_{c_1},\dots,T_{c_{2g}}\}$ on $H_1(\Sigma_g;\mathbb{F}_2)$ preserves q. On the other hand, for r odd, $\mathrm{Mod}(\Sigma_g)[\hat{\phi}]$ surjects onto $\mathrm{Sp}(2g;\mathbb{Z})$ and hence onto $\mathrm{Sp}(2g;\mathbb{F}_2)$ [27, Lemma 5.4]. It follows that no set of $2g$ Dehn twists can generate in the case of r odd. ■

Acknowledgments. The first author would like to thank Yair Minsky for his support and guidance. The authors would also like to thank Barak Weiss for prompting them to consider monodromy valued in the pronged mapping class group, as well as for alerting them to Boissy's work on prong-marked strata [3]. They are grateful to Ursula Hamenstädt for some enlightening discussions and useful suggestions, and to Dick Hain for his interest in the project and a productive correspondence. Finally, they would like to thank the referee for a very careful reading of this paper and for their helpful comments.

Funding. Large parts of this work were accomplished when the authors were visiting MSRI for the "Holomorphic Differentials in Mathematics and Physics" program in the fall of 2019, and both authors would like to thank the venue for its hospitality, excellent working environment, and generous travel support. The first author gratefully acknowledges financial support from NSF grants DMS 1107452, 1107263, and 1107367 "RNMS: Geometric Structures and Representation Varieties" (the GEAR Network) as well as from NSF grant DMS-161087. AC is supported by NSF Award No. DGE-1122492. NS is supported by NSF Awards No. DMS-1703181, DMS-2003891, and DMS-2153879.

References

[1] Bainbridge, M., Smillie, J., Weiss, B.: Horocycle dynamics: new invariants and eigenform loci in the stratum H(1,1). arXiv:1603.00808 (2016)

[2] Boissy, C.: Classification of Rauzy classes in the moduli space of abelian and quadratic differentials. Discrete Contin. Dynam. Systems **32**, 3433–3457 (2012) Zbl 1268.37039 MR 2945824

[3] Boissy, C.: Labeled Rauzy classes and framed translation surfaces. Ann. Inst. Fourier (Grenoble) **65**, 905–932 (2015) Zbl 1332.37030 MR 3449170

[4] Calderon, A.: Connected components of strata of Abelian differentials over Teichmüller space. Comment. Math. Helv. **95**, 361–420 (2020) Zbl 1447.57018 MR 4115287

[5] Calderon, A., Salter, N.: Higher spin mapping class groups and strata of abelian differentials over Teichmüller space. Adv. Math. **389**, art. 107926, 56 pp. (2021) Zbl 1479.30027 MR 4289049

[6] Calderon, A., Salter, N.: Relative homological representations of framed mapping class groups. Bull. London Math. Soc. **53**, 204–219 (2021) Zbl 07367029 MR 4224523

[7] Chillingworth, D. R. J.: Winding numbers on surfaces. I. Math. Ann. **196**, 218–249 (1972) Zbl 0221.57001 MR 300304

[8] Church, T., Ershov, M., Putman, A.: On finite generation of the Johnson filtrations. J. Eur. Math. Soc. **24**, 2875–2914 (2022) Zbl 07523092 MR 4416592

[9] Ershov, M., He, S.: On finiteness properties of the Johnson filtrations. Duke Math. J. **167**, 1713–1759 (2018) Zbl 06904638 MR 3813595

[10] Eskin, A., Masur, H., Zorich, A.: Moduli spaces of abelian differentials: the principal boundary, counting problems, and the Siegel–Veech constants. Publ. Math. Inst. Hautes Études Sci. **97**, 61–179 (2003) Zbl 1037.32013 MR 2010740

[11] Farb, B., Margalit, D.: A Primer on Mapping Class Groups. Princeton Math. Ser. 49, Princeton Univ. Press, Princeton, NJ (2012) Zbl 1245.57002 MR 2850125

[12] Hain, R.: Private correspondence (2020)

[13] Hamenstädt, U.: Quotients of the orbifold fundamental group of strata of abelian differentials. Preprint, Univ. Bonn (2019)

[14] Hamenstädt, U.: On the orbifold fundamental group of the odd component of the stratum $\mathcal{H}(2,\dots,2)$. Preprint, Univ. Bonn (2020)

[15] Humphries, S. P., Johnson, D.: A generalization of winding number functions on surfaces. Proc. London Math. Soc. (3) **58**, 366–386 (1989) Zbl 0681.57003 MR 977482

[16] Johnson, D.: The structure of the Torelli group. I. A finite set of generators for $\mathcal{I}$. Ann. of Math. (2) **118**, 423–442 (1983) Zbl 0549.57006 MR 727699

[17] Johnson, D.: The structure of the Torelli group. II. A characterization of the group generated by twists on bounding curves. Topology **24**, 113–126 (1985) Zbl 0571.57009 MR 793178

[18] Johnson, D. L.: Homeomorphisms of a surface which act trivially on homology. Proc. Amer. Math. Soc. **75**, 119–125 (1979) Zbl 0407.57003 MR 529227

[19] Kawazumi, N.: The mapping class group orbits in the framings of compact surfaces. Quart. J. Math. **69**, 1287–1302 (2018) Zbl 1432.57029 MR 4064155

[20] Kontsevich, M., Zorich, A.: Lyapunov exponents and Hodge theory. arXiv:hep-th/9701164 (1997)

[21] Kontsevich, M., Zorich, A.: Connected components of the moduli spaces of Abelian differentials with prescribed singularities. Invent. Math. **153**, 631–678 (2003) Zbl 1087.32010 MR 2000471

[22] Looijenga, E., Mondello, G.: The fine structure of the moduli space of abelian differentials in genus 3. Geom. Dedicata **169**, 109–128 (2014) Zbl 1308.14034 MR 3175239

[23] Matsumoto, M.: A presentation of mapping class groups in terms of Artin groups and geometric monodromy of singularities. Math. Ann. **316**, 401–418 (2000) Zbl 0968.57014 MR 1752777

[24] Portilla Cuadrado, P., Salter, N.: Vanishing cycles, plane curve singularities and framed mapping class groups. Geom. Topol. **25**, 3179–3228 (2021) Zbl 1477.14018 MR 4347314

[25] Putman, A.: The Johnson homomorphism and its kernel. J. Reine Angew. Math. **735**, 109–141 (2018) Zbl 1456.57016 MR 3757472

[26] Randal-Williams, O.: Homology of the moduli spaces and mapping class groups of framed, r-Spin and Pin surfaces. J. Topol. **7**, 155–186 (2014) Zbl 1322.55006 MR 3180616

[27] Salter, N.: Monodromy and vanishing cycles in toric surfaces. Invent. Math. **216**, 153–213 (2019) Zbl 1440.57028 MR 3935040

[28] Walker, K. C.: Connected components of the strata of quadratic differentials over the Teichmüller space. Geom. Dedicata **142**, 47–60 (2009) Zbl 1182.30070 MR 2545455

[29] Walker, K. C.: Quotient groups of the fundamental groups of certain strata of the moduli space of quadratic differentials. Geom. Topol. **14**, 1129–1164 (2010) Zbl 1213.30076 MR 2651550

[30] Yıldız, O.: Generating the mapping class group by two torsion elements. Mediterr. J. Math. **19**, art. 59, 23 pp. (2022) Zbl 07488615 MR 4383346

[31] Zorich, A.: Explicit Jenkins–Strebel representatives of all strata of abelian and quadratic differentials. J. Modern Dynam. **2**, 139–185 (2008) Zbl 1149.30033 MR 2366233