

Stratified braid groups: monodromy

BY NICK SALTER

*Department of Mathematics, 255 Hurley Building, University of Notre Dame, Hurley Hall, Notre Dame, IN 46556, U.S.A.**e-mail: nsalter@nd.edu**(Received 15 February 2024; revised 04 December 2024; accepted 09 December 2024)**Abstract*

The space of monic squarefree complex polynomials has a stratification according to the multiplicities of the critical points. We introduce a method to study these strata by way of the infinite-area translation surface associated to the logarithmic derivative df/f of the polynomial. We determine the monodromy of these strata in the braid group, thus describing which braidings of the roots are possible if the orders of the critical points are required to stay fixed. Mirroring the story for holomorphic differentials on higher-genus surfaces, we find the answer is governed by the framing of the punctured disk induced by the horizontal foliation on the translation surface.

2020 Mathematics Subject Classification: 30F30 (Primary); 20F36 (Secondary)

1. Introduction

Let $\text{UConf}_n(\mathbb{C})$ denote the space of unordered configurations of n distinct points in \mathbb{C} . Equivalently, this is the space of monic squarefree polynomials of degree n , which will be written as $\text{Poly}_n(\mathbb{C})$. From this point of view, $\text{Poly}_n(\mathbb{C})$ carries a natural *equicritical stratification* $\{\text{Poly}_n(\mathbb{C})[\kappa]\}$ indexed by partitions κ of $n - 1$. A polynomial $f \in \text{Poly}_n(\mathbb{C})$ belongs to $\text{Poly}_n(\mathbb{C})[\kappa]$ if and only if the roots of f' form the partition κ of $n - 1$. Thinking of f as a mapping $f : \mathbb{C} \rightarrow \mathbb{C}$, the partition κ describes the multiplicities of the critical (i.e. branch) points.

Main Question. Understand the topology of $\text{Poly}_n(\mathbb{C})[\kappa]$. What is the fundamental group $\mathcal{B}_n[\kappa] := \pi_1(\text{Poly}_n(\mathbb{C})[\kappa])$ (a “stratified braid group”)? Is $\text{Poly}_n(\mathbb{C})[\kappa]$ a $K(\pi, 1)$ space?

The strata $\text{Poly}_n(\mathbb{C})[\kappa]$ admit descriptions as certain discriminant complements, and so the Main Question is reminiscent of the conjecture of Arnol’d–Pham–Thom on the homotopy type of discriminant complements associated to isolated hypersurface singularities. As we will see, $\text{Poly}_n(\mathbb{C})[\kappa]$ is also closely related to a certain stratum Ω_κ of meromorphic differentials on \mathbb{CP}^1 , and so the Main Question is a version of the conjecture posed by Kontsevich–Zorich originally for strata of holomorphic differentials on higher-genus curves [KZ97].

In this paper, we begin this project by answering a natural question about the groups $\mathcal{B}_n[\kappa]$. The inclusion map $\text{Poly}_n(\mathbb{C})[\kappa] \hookrightarrow \text{Poly}_n(\mathbb{C})$ induces a *monodromy homomorphism*

into the braid group $B_n := \pi_1(\text{Poly}_n(\mathbb{C}))$

$$\bar{\rho} : \mathcal{B}_n[\kappa] \rightarrow B_n,$$

and we describe the image $B_n[\kappa] := \bar{\rho}(\mathcal{B}_n[\kappa])$. We find (cf. Section 4) that there is a *crossed homomorphism* $\phi_\kappa : B_n \rightarrow (\mathbb{Z}/r\mathbb{Z})^r$, where $r = \text{gcd}(\kappa)$ is the gcd of the parts of κ , which characterises the monodromy image when $n \geq n_0(r)$ is sufficiently large compared to r . The precise bounds $n_0(r)$ are somewhat intricate and are stated in Section 6, but we note here that if $r = 1$ then $n_0(1) = 2$ (in which case ϕ_κ is trivial and $B_n[\kappa] = B_n$) and if $r = 2$ then $n_0(2) = 8$; in general, $n_0(r) \leq 14r$ in the worst-case scenario.

THEOREM A. *For any $n \geq 2$ and any partition κ of $n - 1$ with $\text{gcd}(\kappa) = r$, there is a containment*

$$B_n[\kappa] \leq \ker(\phi_\kappa).$$

If $n \geq n_0(r)$, then this containment is an equality.

Remark 1.1. Theorem 6.10 gives a simple characterisation of $\ker(\phi_\kappa)$ (and hence $B_n[\kappa]$) in the range $n \geq n_0(r)$: it is the subgroup of B_n consisting of braids admitting representatives where at each overcrossing, there are r strands passing underneath.

Remark 1.2. The containment $B_n[\kappa] \leq \ker(\phi_\kappa)$ is *not* always an equality. For $\kappa = \{n - 1\}$, the fundamental group $\pi_1(\text{Poly}_n(\mathbb{C})[\kappa])$ is cyclic, since the only polynomials in this stratum are of the form $f(z) = (z - z_0)^n - c$, while the corresponding $\ker(\phi_\kappa)$ has finite index in B_n . On the other hand, the braid-theoretic methods underlying the proof of Theorem A are almost certainly not optimised, and it would be interesting to know exactly how small $n_0(r)$ can be taken.

Remark 1.3 (The braided Gauss–Lucas theorem). The classical Gauss–Lucas theorem asserts that the critical points lie inside the convex hull of the roots. There is a refinement $\rho : \mathcal{B}_n[\kappa] \rightarrow B_n[\kappa]$ of $\bar{\rho}$ where one tracks both roots and critical points, and the study of ρ can be viewed as a Gauss–Lucas theorem for *families* of polynomials. The classical theorem implies that every braid in the image of ρ admits a representative where at each time t , the critical points lie in the convex hull of the roots. Our study of the monodromy shows that *this is not sufficient*. Figure 4 in Section 4.2 gives an example of a braid satisfying this convexity condition which is not realisable as the braid of root and critical points of any family of polynomials. Theorem A shows that when $r \geq 2$, there are even certain braidings of the roots alone (e.g. a half-twist) which cannot be realised by polynomial families. We plan to return to a study of the refined monodromy ρ in future work.

Remark 1.4 (A finer stratification). The equicritical stratification $\{\text{Poly}_n[\kappa]\}$ is concerned only with the multiplicities of the critical *points*; it can (and does) happen that distinct critical points lie over the same critical value. Therefore, each equicritical stratum admits a further stratification by the *profile* (multiplicities plus ramification) of the critical values. Let us call this the “profile stratification”.

From one point of view, the profile stratification is perhaps more natural to investigate, being built from more fundamental pieces. On the other hand, the Main Question posed




Fig. 1. A strip decomposition of a differential df/f . The polynomial f has simple zeroes z_1, \dots, z_5 , and critical points w_1, w_2, w_3 with w_2 having multiplicity 2. On the left, the points z_i, w_j are shown in \mathbb{C} along with the prongs of the horizontal foliation. The strip decomposition is shown on the right. Colours on the right indicate gluing instructions, and correspond to the colourings of the prongs on the left. (Note: this figure is best viewed in colour).

above is strictly less interesting. This is because each stratum in the profile stratification is very closely related to a certain Hurwitz space of branched covers of \mathbb{C}^* (they are not exactly the same thing, since a point in a Hurwitz space only describes a polynomial up to automorphisms of the domain). Therefore the fundamental group of a profile stratum can be computed in terms of the fundamental group of the associated Hurwitz space, which is nothing more than a finite-index subgroup of a braid group. At the space level, a profile stratum fibers over the associated Hurwitz space with fiber given by the affine group $\mathbb{C} \rtimes \mathbb{C}^*$; as base and fiber are aspherical, so too is the total space. Thus only by passing to the coarser equicritical stratum do we reach uncharted territory.

From polynomials to translation surfaces. Our method of study is built around a type of uniformisation map. Namely, we associate to $f \in \text{Poly}_n(\mathbb{C})[\kappa]$ its logarithmic derivative df/f . This is a meromorphic differential on \mathbb{CP}^1 , with n simple poles of residue $2\pi i$ at the zeroes of f and an additional simple pole of residue $-2\pi in$ at infinity. Such an object can be viewed as a translation surface - the poles give the surface infinite area, but it nevertheless has a very simple global structure (see, e.g. Figure 1). Let Ω_κ denote the moduli space of meromorphic differentials on \mathbb{CP}^1 with n simple poles of residue $2\pi i$, a simple pole at infinity (necessarily of residue $-2\pi in$), and zeroes of multiplicity specified by κ . Some elementary complex analysis (Lemma 2.1) shows that every such differential is of the form df/f for $f \in \text{Poly}_n(\mathbb{C})[\kappa]$. The assignment $f \mapsto df/f$ therefore gives a classifying map

$$\mu : \text{Poly}_n(\mathbb{C})[\kappa] \rightarrow \Omega_\kappa.$$

It is clear that if $f(z)$ and $g(z)$ are related by an affine change of variables $g(z) = f(az + b)$, then the associated differentials df/f and dg/g determine the same point in Ω_κ . The converse is not much harder, but this identification is fundamental to our approach, and we record it here for good measure.

THEOREM 1.5. *The classifying map μ induces an isomorphism of complex orbifolds*

$$\text{Poly}_n(\mathbb{C})[\kappa]/\text{Aff} \cong \Omega_\kappa.$$

The advantage in studying Ω_κ is that its global structure is much more apparent. The equations defining $\text{Poly}_n(\mathbb{C})[\kappa]$ as a discriminant complement are highly nonlinear, and it is difficult to construct and analyse the behavior of explicit loops inside $\text{Poly}_n(\mathbb{C})[\kappa]$. On the other hand, in Ω_κ the corresponding analysis is elementary via deformations of the associated translation surfaces. Moreover, in Proposition 3.10, we use the combinatorics of the translation surfaces to obtain an explicit finite cell structure on Ω_κ , in principle reducing the study of the topology of Ω_κ to the combinatorics of the “labelling systems” that index the cells.

The meromorphic differential df/f , and more precisely its incarnation as an infinite-area translation surface, plays a fundamental role in our analysis of the monodromy. Theorem A exactly parallels a result in the setting of strata of holomorphic differentials on higher-genus surfaces. Here, the problem is to determine the image of the orbifold fundamental group in the mapping class group of the surface. This was answered in [CS23], where it is shown that the image is essentially characterised by the property that the monodromy must preserve the *framing* of the surface (punctured at the locations of the zeroes) associated to the horizontal vector field specified by the translation surface structure. In the case where the locations of the zeroes are not marked, the monodromy must preserve a certain distillate of the framing known as an *r-spin structure*, c.f. [CS21]. Here, we find the exact same sort of characterisation of the monodromy: the crossed homomorphism ϕ_κ measures a “change in winding number” of arcs relative to the framing of the punctured surface induced from df/f . In both of these settings, the integer r is given as the gcd of the orders of the zeroes of the differential.

Related work. In [CW91], Catanese–Wajnryb study the space of what they call “generic polynomials”, i.e. those possessing the maximal number of critical *values* (but with no constraint on simplicity of the roots). Their main result gives a description of the fundamental group of this space. In the language of the profile stratification discussed above in Remark 1.4, the space of generic polynomials is very closely related to the top-dimensional stratum in the profile stratification, the only difference being that here, one does not insist that 0 be a regular value, so the Hurwitz space in question parametrizes branched covers over \mathbb{C} and not \mathbb{C}^* .

In [TBY⁺20], W. Thurston and collaborators associate a “degree- d -invariant lamination” to a complex polynomial of degree d . The combinatorics of this lamination record the multiplicities of the critical points, and moreover there is a space (“spine”) of such laminations which they probe by means of deformations of the translation surface associated to df/f (although stated in different language). See in particular [TBY⁺20, theorem 3.1].

The space Ω_κ fits into the theory of the “isoresidual fibration” studied by Gendron–Tahar [GT21, GT22]. They consider the map from the space of meromorphic differentials with prescribed zero and pole orders to the vector space of residues, showing among other things that in the case of a single zero, the map is a fibration away from a hyperplane arrangement. Our Ω_κ is the fiber of the isoresidual map of differentials on \mathbb{C} over the vector of residues $(2\pi i, \dots, 2\pi i)$, where the zeroes have order specified by κ .

In [DM22], Dougherty–McCammond investigate various combinatorial structures induced from polynomial maps. One of their key tools is a pair of transverse singular

foliations on \mathbb{C} with singularities at the zeroes and critical points of f . We obtain an equivalent pair of foliations from the horizontal and vertical foliations of the translation surface structure on \mathbb{C} induced by df/f . See Remark 3.2.

Outline. In Section 2, we establish Theorem 1.5, showing that one can study polynomials in a stratum $\text{Poly}_n(\mathbb{C})[\kappa]$ by instead studying the translation surfaces associated to their logarithmic derivatives. In Section 3, we describe the structure of an individual df/f as a translation surface, as well as the global structure of the stratum Ω_κ . Our main results here are the discussion in Section 3.2 of the “strip decomposition” of df/f , and the global structure theorem Proposition 3.10, which exhibits a cell structure on Ω_κ coming from the combinatorics of this decomposition. The proof of Theorem A is carried out in Sections 4 to 6. In Section 4, we show how the translation surface structure associated to df/f constrains the monodromy image $B_n[\kappa]$, forcing it to preserve *winding numbers* of arcs on the disk. In Section 5, we exhibit certain loops in Ω_κ and analyse their monodromies in B_n . Finally in Section 6, we show that this finite collection of elements is enough to generate the kernel of ϕ_κ (when n is sufficiently large). The key tool here is to relate the winding number crossed homomorphism ϕ_κ to an *a priori* totally different crossed homomorphism Υ_r formulated in terms of a count of “virtual undercrossings” on a braid diagram, and then to establish a *factorisation algorithm* (Lemma 6.9) for expressing the kernel of Υ_r in terms of elements known to lie in the monodromy image.

2. Moduli spaces of polynomials, differentials and translation surfaces

We begin with a discussion of the space Ω_κ , the stratum of translation surfaces associated to the differentials df/f . We construct this here as a moduli space, by taking a quotient of the space of differentials by the relevant automorphism group. The main result of this section is Theorem 1.5, recorded here as Proposition 2.2, which amounts to little more than an unpacking of the definitions, but lays the foundation for what is to follow, as it will allow us to explore the space $\text{Poly}_n(\mathbb{C})[\kappa]$ by instead exploring the space Ω_κ of translation surfaces.

Let $\kappa = \{k_1, \dots, k_p\}$ be a partition of $n - 1$. Here and throughout, we write $|\kappa| = p$ to denote the number of parts of the partition. Let $\mathcal{MD}(\kappa)$ be the set of meromorphic differential forms ω on $\mathbb{C}\mathbb{P}^1$ satisfying the following properties:

- (i) there are exactly p zeroes of ω of orders k_1, \dots, k_p , and each zero lies in $\mathbb{C} \subset \mathbb{C}\mathbb{P}^1$;
- (ii) there are n simple poles each of residue $2\pi i$ contained in \mathbb{C} , and an additional simple pole at ∞ of residue $-2\pi in$.

The following is basic complex analysis; we include the argument for the sake of completeness.

LEMMA 2.1. *Let $\omega \in \mathcal{MD}(\kappa)$ be given. Then there is a unique $f \in \text{Poly}_n(\mathbb{C})[\kappa]$ such that $\omega = df/f$.*

Proof. Let $f \in \text{Poly}_n(\mathbb{C})$ be the polynomial with simple roots at the n poles z_1, \dots, z_n of ω contained in \mathbb{C} . By the theory of partial fractions,

$$\frac{df}{f} = \left(\frac{1}{z - z_1} + \dots + \frac{1}{z - z_n} \right) dz,$$

on \mathbb{C} , showing that $\omega - df/f$ has no poles on \mathbb{C} . By hypothesis, ω and df/f have simple poles at ∞ of equal residue, so that $\omega - df/f$ is moreover holomorphic in a neighbourhood of ∞ . Thus $\omega - df/f$ is a holomorphic differential form on $\mathbb{C}P^1$; the only such form is 0 (see, e.g. [Mir95, exercise IV.1.A]).

Observe that the affine group

$$\text{Aff} = \{\alpha \in \text{Aut}(\mathbb{C}P^1) \mid \alpha(\infty) = \infty\} = \{az + b \mid a \in \mathbb{C}^*, b \in \mathbb{C}\} \cong \mathbb{C}^* \times \mathbb{C}$$

acts via biholomorphisms on $\mathcal{MD}(\kappa)$ on the left via inverse-pullback:

$$\alpha \cdot \omega = (\alpha^{-1})^*(\omega).$$

Likewise, there is a left action of Aff on $\text{Poly}_n(\mathbb{C}) \subset \mathbb{C}^n$ induced from the diagonal action on \mathbb{C}^n .

Note that for each of these actions, Aff is a Lie group acting properly by holomorphic automorphisms with finite stabilisers. The orbit spaces $\text{Poly}_n(\mathbb{C})[\kappa]/\text{Aff}$ and $\mathcal{MD}(\kappa)/\text{Aff}$ therefore carry complex orbifold structures. We observe that $\dim_{\mathbb{C}}(\text{Poly}_n(\mathbb{C})[\kappa]) = |\kappa| + 1$, since the roots of f' move in a configuration space of dimension $|\kappa|$, and the generic antiderivative of f' has n distinct roots. Thus $\dim_{\mathbb{C}}(\text{Poly}_n(\mathbb{C})[\kappa]/\text{Aff}) = |\kappa| - 1$.

We define the κ -stratum of logarithmic derivatives as the second of the orbifolds discussed above:

$$\Omega_{\kappa} := \mathcal{MD}(\kappa)/\text{Aff}.$$

Observe that there is a natural map

$$\begin{aligned} \mu : \text{Poly}_n(\mathbb{C})[\kappa] &\rightarrow \mathcal{MD}(\kappa) \\ f &\mapsto \frac{df}{f}. \end{aligned}$$

PROPOSITION 2.2 (Theorem 1.5). *The map μ is an Aff-equivariant biholomorphism, inducing an isomorphism of complex orbifolds of dimension $|\kappa| - 1$.*

$$\text{Poly}_n(\mathbb{C})[\kappa]_{\text{Aff}} \cong \Omega_{\kappa}.$$

Proof. That μ is a bijection follows immediately from Lemma 2.1, and it is easy to see that this respects the complex structures on the domain and codomain. Equivariance is also easily verified, as

$$(\alpha^{-1})^* \frac{df}{f} = \frac{d(f \circ \alpha^{-1})}{f \circ \alpha^{-1}},$$

and the polynomial $f \circ \alpha^{-1}$ has simple roots at the points $\alpha(z_1), \dots, \alpha(z_n)$, where z_1, \dots, z_n are the roots of f .

An exact sequence. To conclude this section, we study the relationship between the (orbifold) fundamental groups of $\text{Poly}_n(\mathbb{C})[\kappa]$ and its quotient $\text{Poly}_n(\mathbb{C})[\kappa]/\text{Aff} \cong \Omega_{\kappa}$. Following the discussion in [Loo08, introduction], we find that there is an exact sequence

$$\pi_1(\text{Aff}) \rightarrow \pi_1(\text{Poly}_n(\mathbb{C})[\kappa]) \rightarrow \pi_1^{\text{orb}}(\text{Poly}_n(\mathbb{C})[\kappa]/\text{Aff}) \rightarrow \pi_0(\text{Aff}) \rightarrow 1.$$

Recalling that $\pi_1(\text{Aff}) = \mathbb{Z}$ and $\pi_0(\text{Aff}) = 1$, and also recalling that $\pi_1(\text{Poly}_n(\mathbb{C})[\kappa]) := \mathcal{B}_n[\kappa]$, we obtain the exact sequence

$$\mathbb{Z} \rightarrow \mathcal{B}_n[\kappa] \rightarrow \pi_1^{\text{orb}}(\Omega_\kappa) \rightarrow 1. \tag{1}$$

In particular, we emphasise that the projection $\mathcal{B}_n[\kappa] \rightarrow \pi_1^{\text{orb}}(\Omega_\kappa)$ is surjective. It is not hard to show that (1) is in fact short exact, but we do not need this fact here so we will not elaborate.

3. Ω_κ as a space of translation surfaces

The purpose of this section is to explain the structure of a differential df/f when realised as a translation surface. In Section 3.1, we begin with a discussion of some generalities of translation surfaces induced by meromorphic differentials and the induced horizontal foliation. In Section 3.2, we discuss the notion of a *strip decomposition* of the translation surface for df/f and some important related notions (*strips, slits, fixed/free prongs*). This will give a combinatorial decomposition of Ω_κ into cells; in Section 3.3, we discuss the global structure of this decomposition. In the body of this paper, we will only make use of the *constructive* aspects of the theory we establish here (as a technique for exploring the space Ω_κ and computing the monodromy of loops); in later work, we hope to make use of the global structure theory obtained in Proposition 3.10.

3.1. Flat cone metrics and the horizontal foliation.

The integration map

$$z \mapsto \int_{z_0}^z \frac{df}{f}$$

provides a system of holomorphic charts on \mathbb{C} away from the zeroes of f and f' for which the transition functions are translations $z \mapsto z + c$. The horizontal foliation on \mathbb{C} given by lines of constant real part (equivalently determined as the kernel of the real 1-form $dy = \text{Im}(dz)$) pulls back to a singular foliation \mathcal{F} on the domain.

Near a zero ζ_i of order k_i , these charts realise ζ_i as a cone point with cone angle $2\pi(k_i + 1)$. At such a point \mathcal{F} has a prong singularity of order $2k_i + 2$. In the flat coordinates, the prongs alternate between pointing to the right and left and will be referred to as such. We also note that there is a natural cyclic ordering on both the left and right prongs, and each set of prongs carries the structure of a torsor over $\mathbb{Z}/(k_i + 1)\mathbb{Z}$ by measuring the counterclockwise angle from one prong to the other.

The local structure near a simple pole is slightly less well-known, but is equally straightforward. First note that in the case of $\omega = dz/z$, the integration map (i.e. the logarithm) sends the punctured disk $0 < |z| \leq 1$ to the half-infinite strip

$$S^L = \{a + bi \mid a \leq 0, 0 \leq b \leq 2\pi\}$$

with the top and bottom identified via the translation $z \mapsto z + 2\pi i$; likewise, for $c \in \mathbb{C}^*$, the integration map sends a neighbourhood of $c dz/z$ near 0 to the rotated strip cS^L . For a general differential $\omega = g(z)dz/z$ with a simple pole at $z = 0$, the coordinate

$$w(z) = \exp\left(\int_{z_0}^z \omega\right)$$

pulls ω back to dw/w , showing that in general, a neighbourhood of a simple pole of residue c is realised on the translation surface via the rotated strip cS^L (again with opposite edges identified). At such a pole, \mathcal{F} has an “infinite prong singularity”, where the foliation structure is locally given by the set of rays emanating from a point.

Given a translation surface T represented as a finite collection of disjoint polygons $\{P_i\} \subset \mathbb{C}$ with edge identifications, a *cut move* is a subdivision of some P_i into P_i^1, P_i^2 along with the identification of the cut edges. To perform a *paste move*, take distinct polygons P_i, P_j for which there is an edge of P_i identified to an edge of P_j , and translate P_j so that the identified edges coincide. If P_i and P_j overlap only along this edge, then the paste move can be performed by joining P_i and the translate of P_j into a single polygon, inheriting the remaining edge identifications. It is a basic fact in the theory of translation surfaces that T and T' determine the same point in their stratum (in this case, Ω_κ) if and only if they are related by a sequence of cut/paste moves.

The horizontal foliation for df/f . The integration map $\int df/f$ induces a translation surface structure on an n -times punctured plane, for which the horizontal foliation has the local features discussed above. Conversely, any “combinatorially suitable” translation surface structure T on an n -times punctured plane determines a differential $\omega = df/f \in \Omega_\kappa$. Here, by “combinatorially suitable”, we mean the following:

- (i) T has n half-infinite cylindrical strips S_1^L, \dots, S_n^L , each equivalent to S^L via a translation (each extends infinitely far to the left and has height $2\pi i$);
- (ii) T has one half-infinite cylindrical strip S_∞ equivalent to $-nS^L$ via a translation (thus extending infinitely far to the right and of height $2n\pi i$);
- (iii) T has cone points p_1, \dots, p_m of orders k_1, \dots, k_m , where $\kappa = \{k_1, \dots, k_m\}$;
- (iv) the complement of the strips $S_1^L, \dots, S_n^L, S_\infty$ has finite area.

That every such translation surface is induced by a differential $df/f \in \Omega_\kappa$ is immediate: T induces a Riemann surface structure on an n -times punctured plane, equipped with a differential ω on which S_1^L, \dots, S_n^L correspond to simple poles of residue $2\pi i$, S_∞ corresponds to a simple pole of residue $-2n\pi i$, and which has zeroes of multiplicity specified by κ ; i.e. $\omega \in \Omega_\kappa$.

The global structure of the horizontal foliation \mathcal{F} on \mathbb{CP}^1 induced by df/f is extremely simple.

LEMMA 3.1. *Let $f \in \text{Poly}_n(\mathbb{C})[\kappa]$ be given, and let \mathcal{F} be the horizontal singular foliation on \mathbb{CP}^1 induced by df/f . Then \mathcal{F} has the following properties:*

- (i) *with the finitely many exceptions of leaves incident to a critical point of f , every leaf connects a zero of f to ∞ . In particular, \mathcal{F} has no closed leaves;*
- (ii) *let p_j be a critical point of order k_j , corresponding to a cone point of order k_j on the translation surface and inducing a $2k_j + 2$ -pronged singularity of \mathcal{F} . Then the prongs alternate between terminating at a zero of f and at ∞ . In particular, at most one prong at p_j terminates at each zero of f .*

Proof. We first claim that \mathcal{F} has no closed leaves. Integration of df/f along such a leaf would yield a real period of df/f , but the periods of df/f are purely imaginary. If a leaf

does not terminate at a singularity, it must accumulate somewhere on the compact space \mathbb{CP}^1 . Such a nearly-closed leaf can be completed via a short vertical segment into a simple closed curve whose period has positive real part, again a contradiction. Thus every leaf must terminate at both ends at a singularity of \mathcal{F} . At a critical point of order k , \mathcal{F} has a $2k + 2$ -pronged singularity, so that there are finitely many leaves terminating at a critical point as claimed. Integrating df/f along a path terminating at a zero of f has real part tending to $-\infty$, so that at most one end of every leaf can terminate at such a point; likewise at most one end can terminate at ∞ .

This same observation proves the second assertion: when integrating along consecutive prongs, the real part of $\int df/f$ is *monotonic*, so that exactly one prong in each consecutive pair terminates at a zero of f . If two prongs at p_j terminate at the same zero, we consider the bounded region of the plane enclosed by these leaves. By the above, there must be at least one prong originating inside this region which must terminate at ∞ , but it cannot escape the region enclosed by the two prong leaves, showing a contradiction.

Remark 3.2. In [DM22], Dougherty and McCammond study a pair of transverse singular foliations equivalent to those induced by the real and imaginary parts of df/f . Their point of view is somewhat different: they induce \mathcal{F} by pulling back the transverse foliations $|z| = c$ and $\arg(z) = c$ on \mathbb{C}^* under the map $f: \mathbb{C} \rightarrow \mathbb{C}$, but the result as unmeasured foliations is the same. They equip their foliations with measures that are different from the ones coming here from the flat structure, considering instead the measure induced by the Euclidean structure on \mathbb{C}^* . Using this, they are able to obtain a detailed picture of various combinatorial structures associated to the polynomial f . It would be interesting to see if the translation surface perspective has anything to add to the story they pursue.

3.2. Strip decomposition.

A translation surface $T \in \Omega_\kappa$ admits a finite number of combinatorially-determined standard forms which we call a *strip decomposition*. Assign a numbering S_1^L, \dots, S_n^L to the n left-infinite strips of height $2\pi i$, or equivalently a numbering z_1, \dots, z_n of the zeroes of f . The strips S_i^L extend infinitely far to the left by hypothesis. Following the leaves of the horizontal foliation in S_i^L to the right (towards ∞), we observe that there must be at least one leaf terminating at a critical point w_j , for otherwise, this region would close up into a topological cylinder, rendering the translation surface disconnected (except, of course, in the case $n = 1$ with differential dz/z). Choosing one such leaf, we fix an identification $S_i^L \cong S^L$ by identifying w_j with $0 \sim 2\pi i$ in S^L . We remark that we allow for the non-generic possibility that the leaf connecting z_i to w_j passes through one or more additional cone point.

Define $S_i \subset T$ as the continuation to the right of the leaves of the horizontal foliation passing through S_i^L . This is then a bi-infinite strip of height $2\pi i$, possibly containing additional cone points. The boundary of S_i^L is determined by the prong of w_j from which the leaf terminating at z_i emanates. We say that S_i is *bounded* by the cone point w_j , and call the distinguished prong a *fixed prong*. By Lemma 3.1, all of the leaves of the horizontal foliation not terminating at a cone point are contained in some strip S_i , and so this produces a decomposition of T as claimed.

Strips S_i and S_j are said to be *vertically adjacent* if the top right boundary of S_i is identified with the lower right boundary of S_j or vice versa. We will speak of the strips *above* and *below* S_i via this definition.

Strip coordinates. The strip decomposition of T of course depends on various non-canonical choices. To track this, and moreover to understand the global structure of the space Ω_κ , we define

$$\Omega_\kappa^{ord} \rightarrow \Omega_\kappa$$

as the covering space consisting of differentials $df/f \in \Omega_\kappa$ together with labellings z_1, \dots, z_n and w_1, \dots, w_p of the zeroes and critical points of f , respectively. Note that the stabiliser of a *labelled* configuration of two or more points in \mathbb{C} under the affine group is trivial, so that Ω_κ^{ord} is a manifold cover of the orbifold Ω_κ . For the ensuing discussion, we will lift $df/f \in \Omega_\kappa$ to one of its preimages in Ω_κ^{ord} .

Having fixed such data, one can then encode the combinatorial type of a strip decomposition by tracking the prongs of the cone points. A cone point w_j of order k_j has $2k_j + 2$ prongs emanating from it, of which $k_j + 1$ point to the left on the translation surface. Since $\sum_{j=1}^p k_j = n - 1$, there is a total of $n - 1 + p$ left prongs. Of these, n are fixed prongs; we call the remaining $p - 1$ *free prongs*. Generically, the leaf emanating from a free prong is contained in the interior of a unique strip S_i ; exceptionally it may terminate at some other cone point. Given a labeled differential df/f , a choice of strip decomposition yields the following data:

- (1) for each zero z_i of f , a choice of some left prong to bound the strip S_i ;
- (2) an assignment of the remaining $p - 1$ free prongs to one of the strips containing it (generically, a free prong lies in a unique strip; exceptionally it may lie on the boundary between two that are vertically adjacent);
- (3) the $p - 1$ relative periods γ_j of the arcs connecting each free prong to the fixed prong for its strip, each an element of $\mathbb{R} + [0, 2\pi]i$; if two free prongs belong to the same strip, the relative periods must be distinct (so that the cone points do not collide).

Conversely, we can use the relative periods of the free prongs to put a system of coordinates (“strip coordinates”) on Ω_κ^{ord} . A strip coordinate chart is indexed by a *labelling system* which includes the data specified by (1) and (2) above. Without further constraint, the relative periods of (3) do not yet induce a coordinate patch on Ω_κ^{ord} : as two free prongs in the same strip orbit around one another, one will pass through the slit associated to the other and into a different strip. To prevent this, we additionally impose *orderings* on the imaginary parts of the relative periods of the free prongs within a given strip.

Definition 3.3 (Labelling system). Fix a partition κ of $n - 1$ and consider the associated set of left prongs \mathcal{P}_κ of cardinality $n + p - 1$. A labelling system L is a choice of the following data:

- (1) for each $1 \leq i \leq n$, a choice of prong $v \in \mathcal{P}_\kappa$ as the fixed prong for S_i ;
- (2) an assignment of each of the remaining $p - 1$ prongs in \mathcal{P}_κ to some strip S_i ;
- (3) for each strip S_i , a choice of ordering $v_1 \leq \dots \leq v_{m_i}$ of the m_i free prongs assigned to S_i .

Not every labelling system is realised by some $df/f \in \Omega_\kappa^{ord}$, since some choices of labelling systems will cause the translation surface to be *disconnected*. Here we *state* the combinatorial criterion for connectedness only; we prove that this encodes topological connectedness in Lemma 3.6.

Definition 3.4 (Connected labelling system). Let L be a labelling system for some partition $\kappa = \{k_1, \dots, k_p\}$ of $n - 1$. Let Γ_L be the graph whose vertices are the parts k_i of κ , and where k_i and k_j are connected by an edge if there are prongs of k_i and k_j contained in the same strip S_i . Then L is said to be connected if Γ_L is.

Having specified a labelling system, we turn now to the problem of parametrising the relative periods of the free prongs. For $m \geq 1$, define the closed m -simplex via

$$\Delta^m = \{(y_1, \dots, y_m) \in \mathbb{R}^m \mid 0 \leq y_1 \leq \dots \leq y_m \leq 2\pi\}.$$

Definition 3.5 (Strip coordinate domain). Let L be a labelling system of some partition κ of $n - 1$; for $1 \leq i \leq n$, suppose there are m_i free prongs assigned to the strip S_i . The associated strip coordinate domain is the set

$$\Omega_L := \left(\prod_{i=1}^n \mathbb{R}^{m_i} + i\Delta^{m_i} \right) \setminus \mathcal{D} \subset \mathbb{C}^{p-1},$$

where \mathcal{D} is the union of the following sets:

- (1) \mathcal{D}_1 the set of points where $z_{j_1} = z_{j_2}$ for j_1, j_2 assigned to the same strip;
- (2) \mathcal{D}_2 the set of points where $x_{j_1} = x_{j_2} < 0$ and $y_{j_2} = y_{j_1} + 2\pi i$ for j_1, j_2 assigned to the same strip;
- (3) \mathcal{D}_3 the set of points where $x_j = x_k > 0$ and $y_j = 2\pi i, y_k = 0$, where j is assigned to the strip below the strip containing k ;
- (4) \mathcal{D}_4 the set of points where $z_j = 0$ or $z_j = 2\pi i$.

LEMMA 3.6. *Let L be a connected labelling system of the partition κ . Then there is a realisation map*

$$r : \Omega_L \rightarrow \Omega_\kappa^{ord}.$$

The restriction of r to the interior of Ω_L is a biholomorphism onto its image.

Proof. A point $\gamma = (\gamma_1, \dots, \gamma_{p-1}) \in \Omega_L$ determines a translation surface T_γ as follows: assemble n bi-infinite strips S_1, \dots, S_n , and mark $0 \sim 2\pi i \in S_i$ with the prong specified by L . Given a relative period γ_j , the labelling system specifies a free prong in a strip S_i ; place the free prong at $\gamma_j \in S_i$ and introduce a slit running horizontally to the right from γ_j to ∞ . The cyclic ordering on the prongs at each cone point then specifies gluing instructions on the slits as well as on the right halves of the top and bottom boundary components of each S_j (the top and bottom left halves of S_i are identified to each other). The excision of the set \mathcal{D} from Ω_L ensures that after gluing, no pair of free prongs are identified, and that no free prong is placed at the location of a fixed prong.




Fig. 2. Type 1: changing the fixed prong from v_0 to v_2 .

We claim that T_γ is connected if and only if the labelling system L is connected in the sense of Definition 3.4. Note that the vertices of Γ_L are canonically identified with the cone points w_j of T_γ . A first trivial observation is that T_γ will be connected if and only if there is a path connecting each pair of strips S_i, S_j . Suppose that T_γ is connected; we wish to find a path in Γ_L connecting an arbitrary pair of vertices w_j, w_k . Choose prongs at w_j and at w_k ; these live in strips S, S' respectively. If $S = S'$ then w_j and w_k are connected by definition; otherwise, let c be a path in T_γ connecting S to S' , and one can use this to build a corresponding path in Γ_L by moving through a sequence of cone points lying in the sequence of strips passed through by c .

Conversely, suppose that Γ_L is connected. Given a strip S_i , let w_i be the corresponding bounding cone point. Observe that it suffices to show that strips S_a and S_b are path-connected in T_γ if the associated vertices w_a, w_b of Γ_L are either equal or adjacent. If $w_a = w_b = w$ then a path connecting S_a to S_b can be constructed by winding some number of times around w . If w_a and w_b are adjacent, then there is some strip S_c containing a prong of both w_a and w_b ; a path connecting S_a to S_b can be constructed by concatenating a path between neighbourhoods of w_a and w_b in S_c with paths winding around w_a and w_b between S_c and S_a , resp. S_b .

At this point, we have shown that this construction process yields a well-defined map $r : \Omega_L \rightarrow \Omega_\kappa^{ord}$. We next observe that r is holomorphic - this is a simple consequence of the fact that the relative period maps on Ω_κ^{ord} are holomorphic. It remains to show that r is injective on the interior of Ω_L . To see this, observe that translation surfaces T_γ and T_δ determine the same point in Ω_κ^{ord} only if the sets of relative periods between cone points are equal. The sets of relative periods between fixed cone points is a torsor on the group of absolute periods; in this case the absolute periods is just the set $2\pi i\mathbb{Z} \subset \mathbb{C}$. If $\gamma, \delta \in \text{int}(\Omega_L)$ are distinct, then the corresponding relative periods all have imaginary part strictly between 0 and 2π , so that the relative periods of T_γ cannot be obtained from those of T_δ by translation by some absolute period.

3.3. Change of coordinates; a cell structure on Ω_κ^{ord}

We next consider the transition maps between strip coordinate domains with overlapping image. There are two basic transitions to study: (1) changing which of the prongs in S_i is fixed, and (2) pushing the topmost (relative to the ordering) free prong out the top right side of S_i and into the bottom of the strip above (or in reverse, pushing the bottom free prong through the bottom right side). All coordinate changes are compositions of these two, e.g. pushing a free prong out the top left side is equivalent to changing the free prong to the fixed, and pushing the new free prong (formerly the fixed) out the bottom right. The lemmas below record the effects of these moves on strip coordinates; the proofs follow from inspection of Figures 2 and 3.




Fig. 3. Type 2: pushing \$v_3\$ up from \$S_i\$ to \$S_j\$.

LEMMA 3.7 (Type 1: changing the fixed prong). *Let \$L\$ be a connected labelling system for \$\kappa\$. Choose some strip \$S_i\$; let \$v_0\$ denote the fixed prong and let \$v_1 \le \dots \le v_{m_i}\$ denote the free prongs in \$S_i\$ together with their cyclic ordering. Define \$L'\$ as the labelling system obtained from \$L\$ by choosing some \$v_j\$ as the new fixed prong for \$S_i\$, and ordering the free prongs via*

$$v_{j+1} \le \dots \le v_{m_i} \le v_0 \le v_1 \le \dots \le v_{j-1}.$$

The map on relative periods \$t : \Omega_L \to \Omega_{L'}\$ is given by

$$t(\gamma_1, \dots, \gamma_{m_i}) = (\gamma_{j+1} - \gamma_j, \dots, \gamma_{m_i} - \gamma_j, 2\pi i - \gamma_j, \gamma_1 + 2\pi i - \gamma_j, \dots, \gamma_{j-1} + 2\pi i - \gamma_j).$$

LEMMA 3.8 (Type 2: pushing up/down). *Let \$L\$ be a connected labelling system for \$\kappa\$. Let \$S_i\$ be a strip with fixed prong \$v_0\$ and let \$v_1 \le \dots \le v_{m_i}\$ denote the free prongs in \$S_i\$ together with their cyclic ordering. Denote the relative periods by \$\gamma_1, \dots, \gamma_{m_i}\$, and suppose that \$\gamma_{m_i} = x + 2\pi i\$ with \$x > 0\$. Let \$S_j\$ be the strip whose bottom right boundary is identified with the top right boundary of \$S_i\$, and let \$v'_1 \le \dots \le v'_{m_j}\$ denote the free prongs in \$S_j\$.*

Changing the assignment of \$v_{m_i}\$ from \$S_i\$ to \$S_j\$ yields the labelling system \$L'\$ obtained from \$L\$ by reassigning \$v_{m_i}\$ to \$S_j\$ with ordering in \$S_j\$

$$v_{m_i} \le v'_1 \le \dots \le v'_{m_j}$$

and relative period \$x\$.

Conversely, if \$v_1\$ has period \$x \in \mathbb{R}\$ with \$x > 0\$, then we may reassign it to the strip \$S_j\$ whose top right boundary is identified with the bottom right on \$S_i\$, assigning it to the maximal position in \$S_j\$ with period \$x + 2\pi i\$.

LEMMA 3.9. *Let \$L, L'\$ be connected labelling systems for \$\kappa\$, and let*

$$df/f \in r(\Omega_L) \cap r(\Omega_{L'})$$

be a differential (with zeroes and poles labeled) in the image of the strip coordinate domains for both \$L\$ and \$L'\$. Then \$L'\$ can be obtained from \$L\$ by a sequence of moves of type 1 and 2.

Proof. By hypothesis, there are two translation surfaces \$T\$ and \$T'\$ coming from \$\Omega_L\$ and \$\Omega_{L'}\$, respectively, that determine the same point in \$\Omega_{\kappa}^{ord}\$. Thus \$T\$ and \$T'\$ are equivalent via a sequence of cut/paste moves that moreover preserve the labellings of each of the poles

and zeroes of the associated differential df/f . Applying the strip decomposition to T and T' , it follows that each of the corresponding strips are individually cut/paste equivalent. A cut/paste move applied to a given strip corresponds to a move of type 1 on the labelling system. After applying a cut/paste isomorphism taking T to T' as labeled translation surfaces, the only remaining choices in the assignment of a labelling system arises in assigning prongs lying on the boundary of two strips to one or the other; this corresponds to moves of type 2.

Summary. We summarise the results of the section in the following result, describing the global structure of Ω_κ^{ord} obtained by gluing together strip coordinate patches according to moves of types 1 and 2.

PROPOSITION 3.10. *There is a biholomorphism*

$$\Omega_\kappa^{ord} \cong \coprod \Omega_L / \sim,$$

where the union is taken over all connected labelling systems for κ and \sim is the equivalence relation generated by moves of types 1 and 2 as in Lemmas 3.7 and 3.8.

Proof. We claim that the topological space $\coprod \Omega_L / \sim$ is a manifold under the system of coordinates provided by Ω_L . This is not quite immediate from what we have shown - the strip coordinate domains are not open, and their interiors do not quite cover Ω_κ^{ord} , missing points where some free prong has period with imaginary part 0 or $2\pi i$. But such points lie in the interior of the union of two strip coordinate domains, e.g. as shown in Figure 3. By Lemmas 3.7 and 3.8, the transition functions between overlapping $\Omega_L, \Omega_{L'}$ are holomorphic (and indeed affine), and hence $\coprod \Omega_L / \sim$ is a complex manifold.

By Lemma 3.6 and Lemma 3.9, the set of realisation maps $r : \Omega_L \rightarrow \Omega_\kappa^{ord}$ assemble into a biholomorphism.

Remark 3.11. While one can exhibit a deformation retraction showing that an individual set $r(\Omega_L)$ is contractible, it is not the case that all intersections of sets $r(\Omega_L)$ are contractible. Specifically, if v lies on the top right of its strip, and v' lies on the bottom right of the strip above, then there is a three-fold intersection of labelling systems where at most one of v or v' has been moved up or down. This intersection has two components, arising from the different linear orderings on the real parts. To compute the homotopy type of Ω_κ^{ord} as the nerve of a covering, it is therefore necessary to further subdivide the pieces Ω_L (taking into account the various orderings of the real parts) so as to account for this phenomenon.

4. Winding numbers

In this section, we begin our study of the monodromy of strata of polynomials (Theorem A). Our ultimate objective is Lemma 4.7, which asserts that the monodromy image $B_n[\kappa] = \bar{\rho}(\mathcal{B}_n[\kappa])$ lies in the kernel of a certain crossed homomorphism ϕ_r . This will be constructed as a measure of “change of winding number” for arcs on a translation surface; accordingly, we begin with a discussion of the theory of *relative winding number functions*. In Section 4.2, we use the theory of winding number functions to give an example of a braid which satisfies the convexity condition enforced by the Gauss–Lucas theorem, but which nevertheless cannot be realised as the braid of root and critical points of any family of polynomials.

4.1. Winding number functions

To avoid a lengthy digression, we give here an abbreviated account of the theory of winding number functions which will suffice for our purposes; see [CS23, section 2] for a fuller discussion.

Definition 4.1 (Relative winding number function). Let $\mathbb{C}_{n,p}$ denote the surface $\mathbb{C}\mathbb{P}^1$ with three sets of marked points: n points $S_r \subset \mathbb{C}$ which we call the roots, p points $S_c \subset \mathbb{C}$ called the critical points, and ∞ . Let $S = S_r \cup S_c$. We allow for the possibility of tracking only roots, and not critical points, and hence we permit $p = 0$. We further endow $\mathbb{C}_{n,p}$ with a weighting

$$w : S \rightarrow \mathbb{Z}$$

for which $w(z) = -1$ for each root, and $w(z) > 0$ for each critical point. In the context under study, we think of w as the function that assigns to each point its order as a zero or a pole.

Let $\mathcal{A}_{n,p}$ denote the set of isotopy classes of properly-embedded smooth oriented arcs, disjoint from all marked points on their interior, that connect some root in \mathbb{C} to ∞ (in that order, relative to the orientation). A $\mathbb{Z}/r\mathbb{Z}$ relative winding number function is a set map

$$\psi : \mathcal{A}_{n,p} \rightarrow \mathbb{Z}/r\mathbb{Z}$$

that satisfies the *twist-linearity condition*

$$\psi(T_c(a)) = \psi(a) + \langle a, c \rangle ||c||, \tag{2}$$

where $c \subset \mathbb{C}_{n,p}$ is a simple closed curve and $||c||$ is determined by the formula

$$||c|| = 1 + \sum_{z \in \text{int}(c) \cap S} w(z), \tag{3}$$

where the sum runs over the points of S in the interior of c (i.e. the component of $\mathbb{C}\mathbb{P}^1 \setminus c$ not containing ∞). As usual, $\langle a, c \rangle$ denotes the algebraic intersection pairing, relative to the specified orientation on a and the orientation on c for which ∞ lies to the left. When $r = 0$, we call such an object an *integral relative winding number function*.

Example 4.2 (Horizontal winding number function ψ_T). Let $n \geq 2$ be given, let κ be a partition of $n - 1$, and let $T \in \Omega_\kappa$ be a translation surface structure on $\mathbb{C}\mathbb{P}^1$. Such T corresponds to a differential df/f , and we let $\mathbb{C}_{n,p}(T)$ be the surface with the roots and critical points of f marked. Let w be the weighting given by the order of the corresponding pole or zero of df/f .

T endows $\mathbb{C}_{n,p}(T)$ with an integral relative winding number function ψ_T called the *horizontal winding number function*. Let $a \subset \mathbb{C}_{n,p}(T)$ be a properly-embedded smooth oriented arc connecting a pole of df/f to ∞ . We assign the value $\psi_T(a) \in \mathbb{Z}$ as follows: realise a as an arc on the translation surface T not passing through any of the cone points. As a connects a zero of f to the pole at ∞ , it runs from left to right on T , and as it is properly embedded, it can be isotoped so that it follows a leaf of the horizontal foliation outside of some compact region of T . Such a representative carries an integral winding number $wn_T(a) \in \mathbb{Z}$ by measuring the winding of the forward-pointing tangent vector relative to the horizontal vector

field (the winding number is *integral* because of the condition that the arc coincide with a leaf of the horizontal foliation outside of a compact region).

LEMMA 4.3. *The function*

$$\psi_T(a) := wn_T(a)$$

is a well-defined integral relative winding number function on $\mathbb{C}_{n,p}(T)$.

Setting $r = \gcd(\kappa)$, the mod- r reduction

$$\bar{\psi}_T(a) := \psi_T(\tilde{a}) \pmod{r},$$

where $a \in \mathbb{C}_n(T)$ is an arc and $\tilde{a} \in \mathbb{C}_{n,p}(T)$ is an arbitrary lift, is a well-defined $\mathbb{Z}/r\mathbb{Z}$ relative winding number function on $\mathbb{C}_n(T)$.

Proof. To see that ψ_T is well-defined, we must check (i) that $wn_T(a)$ is unchanged by an isotopy of a , and (ii) that $\psi_T(a)$ satisfies the twist-linearity condition (2). To see that $\bar{\psi}_T(a)$ is well-defined, we must further check (iii) that $wn_T(\tilde{a})$ is unchanged mod r by an isotopy of \tilde{a} across a cone point of T .

To establish (i), we recall that a is horizontal except on a compact set. As the winding number of such an arc is integral (and hence discretely-valued), it follows that the winding number is invariant under any compactly-supported isotopy. Under an isotopy with non-compact support, a can wrap around a pole some number of times, potentially altering the winding number. But this is in fact a special case of (ii): winding a around a pole of T is equivalent to applying the Dehn twist around a curve c enclosing this single pole, i.e. for which $||c|| = 0$.

To establish (ii), let $c \subset \mathbb{C}_{n,p}$ be a simple closed curve. It follows from the Poincaré–Hopf theorem that the winding number of c on T is given by $||c||$ as in (3), as the index of the horizontal vector field at a root or critical point is given by w . Applying the Dehn twist T_c to a , we see that

$$wn_T(T_c(a)) = wn_T(a) + \langle a, c \rangle ||c||$$

holds, since at each intersection between a and c , the twist $T_c(a)$ wraps once around c , contributing $\pm ||c||$ to the winding number, the sign determined by the sign of the intersection.

For (iii), we again invoke the Poincaré–Hopf theorem to see that as a curve is isotoped across a zero of index k on a vector field, the winding number changes by k . By hypothesis, the order of each zero is divisible by r . Thus, after reducing mod r , the quantity $wn_T(\tilde{a})$ is independent of the choice of lift \tilde{a} of $a \subset \mathbb{C}_n(T)$ to $\mathbb{C}_{n,p}(T)$.

As $B_{n,p}$ acts on the set $\mathcal{A}_{n,p}$ of arcs, there is an induced action

$$\beta \cdot \psi(a) = \psi(\beta^{-1}a)$$

on the set of relative winding number functions, and hence there is an associated stabiliser subgroup of $B_{n,p}$, which we call the *framed braid group*. In the case where we track roots but not critical points, we call such groups *r-spin braid groups*, by analogy with the theory

of “ r -spin structures” and their associated “ r -spin mapping class groups” in higher genus, cf. [CS21].

Definition 4.4 (Framed braid group $B_{n,p}(\psi)$, spin braid group $B_n(\bar{\psi})$). Let ψ be an integral relative winding number function on $\mathbb{C}_{n,p}$. The associated framed braid group $B_{n,p}(\psi)$ is the subgroup of $B_{n,p}$ stabilizing ψ under the above action on the set of integral relative winding number functions.

Likewise, if $\bar{\psi}$ is a $\mathbb{Z}/r\mathbb{Z}$ relative winding number function on \mathbb{C}_n , the associated r -spin braid group $B_n(\bar{\psi})$ is the stabiliser of $\bar{\psi}$.

In Lemma 4.6, we will see that the monodromy of a stratum $\text{Poly}_n(\mathbb{C})[\kappa]$ is contained in a certain framed braid group. To establish this, we must digress briefly to give a precise construction of the monodromy homomorphism.

Definition 4.5 (Monodromy). Let κ be a partition of $n - 1$ with $|\kappa| = p$ parts, and let $B_{n,p}$ be the subgroup of B_{n+p} preserving the division of the $n + p$ strands into groups of size n and p . Recalling the definition $\mathcal{B}_n[\kappa] := \pi_1(\text{Poly}_n(\mathbb{C})[\kappa])$, the monodromy is a homomorphism

$$\rho : \mathcal{B}_n[\kappa] \rightarrow B_{n,p}$$

constructed as follows. Let $f \in \text{Poly}_n(\mathbb{C})[\kappa]$ be chosen as a basepoint, and let $T = df/f$ be the associated translation surface in Ω_κ . Fix a choice of marking (i.e. homeomorphism) $\mu : \mathbb{C}_{n,p} \rightarrow T$. Let $\beta : [0, 1] \rightarrow \text{Poly}_n(\mathbb{C})[\kappa]$ be a loop based at f , which induces a loop in Ω_κ , which we will also write β ; we write the image of this latter loop as $\beta(t) = T_t$ with $T_0 = T_1 = T$. The family $\{T_t\}$ of translation surfaces over $[0,1]$ is topologically trivial, and hence there is a well-defined isotopy class of identification $f_t : T_0 \rightarrow T_t$ for $t \in [0, 1]$, which induces a propagation $\mu_t : \mathbb{C}_{n,p} \rightarrow T_t$ of the marking map. The monodromy of β is the element

$$\rho(\beta) = \mu_1^{-1} \mu_0 \in \text{Mod}(\mathbb{C}_{n,p}),$$

where $\text{Mod}(\mathbb{C}_{n,p})$ denotes the mapping class group of $\mathbb{C}_{n,p}$. As the marking μ can be enhanced to identify a tangent vector at $\infty \in \mathbb{C}_{n,p}$ with the canonical horizontal direction on translation surfaces in Ω_κ , we can identify $\text{Mod}(\mathbb{C}_{n,p})$ with the mapping class group of the n,p -times punctured disk $D_{n,p}$, i.e. the subgroup $B_{n,p} \leq B_{n+p}$ preserving setwise the roots and critical points.

That $\rho : \mathcal{B}_n[\kappa] \rightarrow B_{n,p}$ is a homomorphism is a consequence of the fact that if β, γ are loops for which the propagated markings at $t = 1$ are denoted $\mu_\beta, \mu_\gamma : \mathbb{C}_{n,p} \rightarrow T$, then $\mu_\gamma \mu_0^{-1} \mu_\beta$ gives a propagation of the marking along the composite path $\beta\gamma$.

Note that ρ is not completely canonical: it depends on a choice of marking $\mu_0 : \mathbb{C}_{n,p} \rightarrow T$ (and in particular depends on a choice of basepoint $T \in \Omega_\kappa$). However, it is easy to see that different choices of marking lead to *conjugate* monodromy homomorphisms.

Note also that we obtain a reduction

$$\bar{\rho} : \mathcal{B}_n[\kappa] \rightarrow B_n$$

by forgetting the braid of the critical points.

Define

$$B_{n,p}[\kappa] := \rho(\mathcal{B}_n[\kappa])$$

and

$$B_n[\kappa] := \bar{\rho}(\mathcal{B}_n[\kappa]).$$

LEMMA 4.6. *Let $T \in \Omega_\kappa$ be a basepoint. Under the monodromy map $\rho : \mathcal{B}_n[\kappa] \rightarrow B_{n,p}$ based at T , there are containments*

$$B_{n,p}[\kappa] \leq B_{n,p}(\psi_T)$$

and

$$B_n[\kappa] \leq B_n(\bar{\psi}_T).$$

Proof. We assume the notation of Definition 4.5, and consider the monodromy of a loop β in $\text{Poly}_n(\mathbb{C})[\kappa]$. As the marking is propagated along the loop, this induces an identification of the sets $\mathcal{A}_{n,p}(T_t)$ of the arcs on the translation surfaces T_t . It follows that β induces a continuously-varying family of winding number functions ψ_{T_t} on the set $\mathcal{A}_{n,p}$ of arcs on the reference surface $\mathbb{C}_{n,p}$. As the set of winding number functions is a discrete set, it follows that all such winding number functions coincide. In particular, $\psi_{T_0} = \psi_{T_1}$, but from the definitions we have $\psi_{T_1} = \beta \cdot \psi_{T_0}$, showing that $B_{n,p}[\kappa] \leq B_{n,p}(\psi_T)$ as claimed.

The containment $B_n[\kappa] \leq B_n(\bar{\psi}_T)$ is a straightforward consequence of the fact that the integral relative winding number function ψ_T on $\mathbb{C}_{n,p}(T)$ descends to the $\mathbb{Z}/r\mathbb{Z}$ relative winding number function $\bar{\psi}_T$ on $\mathbb{C}_n(T)$ under the forgetful map $\mathbb{C}_{n,p}(T) \rightarrow \mathbb{C}_n(T)$.

4.2. *Convexity is not enough: the braided Gauss–Lucas theorem*

To illustrate Lemma 4.6, we give here in Figure 4 an example of a braid in $B_{n,p}$ that admits a “convex representative”, i.e. where the p -stranded braid of critical points lies inside the convex hull of the n -stranded braid of roots for all times t , and yet which does not arise from any loop of polynomials.

4.3. *Mod- r winding numbers as crossed homomorphisms*

From here to the end of the paper, we will concentrate on the monodromy $\bar{\rho}$ of the roots only, leaving a study of the refinement ρ for future work.

Here, we show that the r -spin braid group $B_n(\bar{\psi}_T)$ can be identified with the kernel of a certain crossed homomorphism ϕ_κ , and show that ϕ_κ has a very simple formula; as this ultimately depends only on $r = \text{gcd}(\kappa)$ and not κ itself, in the sequel we will work instead with the equivalent crossed homomorphism ϕ_r with the simple formula.

LEMMA 4.7. *Let $T \in \Omega_\kappa$, and let S_1, \dots, S_n be the strips in a strip decomposition for T . For $i = 1, \dots, n$, let $a_i \subset \mathbb{C}_n$ be an arc corresponding to a horizontal leaf on T contained entirely in S_i . Then the function*

$$\begin{aligned} \phi_\kappa : B_n &\rightarrow (\mathbb{Z}/r\mathbb{Z})^n \\ \beta &\mapsto \sum \bar{\psi}_T(\beta^{-1}a_i)e_i \end{aligned}$$




Fig. 4. A braid which cannot be realised by a family of polynomials in $\text{Poly}_4(\mathbb{C})[1^3]$. The four roots are illustrated in black, and the three simple critical points are coloured (appearing lighter in grayscale). A choice of arcs connecting roots to infinity (depicted here as the entire boundary for visual simplicity) are shown in gray. As the indicated (blue) point orbits the central root, it alters the winding number of the corresponding arc (as can be seen from the twist-linearity formula), and is thus not contained in the framed braid group $B_{4,3}(\psi_T)$. By Lemma 4.6, it follows that this braid cannot be realised by a loop in $\text{Poly}_4(\mathbb{C})[1^3]$.

is equal to the crossed homomorphism

$$\begin{aligned} \phi_r : B_n &\rightarrow (\mathbb{Z}/r\mathbb{Z})^n \\ \sigma_i &\mapsto e_{i+1}, \end{aligned}$$

(where B_n acts on $(\mathbb{Z}/r\mathbb{Z})^n$ on the left via the coordinate-permutation action induced from the quotient $B_n \rightarrow S_n$, $\beta \mapsto \bar{\beta}$).

In particular, there is a containment

$$B_n[\kappa] \leq \ker(\phi_\kappa) = \ker(\phi_r).$$

Proof. To establish that ϕ_κ is a crossed homomorphism, we make the following observation. If a and a' are two arcs on T with the same beginning and end points, then $a \cup a'$ is an oriented closed curve. There are two cusps at the common endpoints, and otherwise $a \cup a'$ is smoothly immersed. By the Poincaré–Hopf theorem, the winding number of $a \cup a'$ (reduced mod r , as usual) counts the total number of poles on T enclosed by $a \cup a'$ (up to a correction factor of 1 coming from the change in winding number arising from smoothing out the cusps). Thus this quantity is invariant under the action of the braid group:

$$\bar{\psi}_T(\beta a) - \bar{\psi}_T(\beta a') = \bar{\psi}_T(a) - \bar{\psi}_T(a').$$

Now given $\alpha, \beta \in B_n$, we use this to compute

$$\phi_\kappa(\alpha\beta) = \sum \bar{\psi}_T(\beta^{-1}\alpha^{-1}a_i)e_i$$




Fig. 5. The standard marking. At left, the reference translation surface T_κ for the stratum $\kappa = \{k_1, \dots, k_p\}$. The top left and bottom left of each strip are identified. There are p blocks of strips, each one corresponding to a given cone point (depicted as the coloured dots in the middle of each strip). Within each block of strips of the same colour/greyscale value, the top right segment of S_i is identified to the bottom right on S_{i+1} , with remaining gluing instructions specified by colour/letter as in Figure 1. All but the bottom-most cone point (red in the figure) have one free prong in the top strip of the block below. The horizontal lines running across each strip equip T_κ with a marking. At right, the corresponding standard marking of the n -punctured disk.

$$\begin{aligned}
 &= \sum \left(\bar{\psi}_T(\beta^{-1}\alpha^{-1}a_i) - \bar{\psi}_T(\beta^{-1}a_{\alpha^{-1}i}) + \bar{\psi}_T(\beta^{-1}a_{\alpha^{-1}i}) \right) e_i \\
 &= \sum \left(\bar{\psi}_T(\alpha^{-1}a_i) - \bar{\psi}_T(a_{\alpha^{-1}i}) + \bar{\psi}_T(\beta^{-1}a_{\alpha^{-1}i}) \right) e_i.
 \end{aligned}$$

Splitting into three vectors, we observe that the first is $\phi_\kappa(\alpha)$, the second is identically zero (each component $\bar{\psi}_T(a_i)$ is zero since a_i is a leaf of the horizontal foliation), and the third is identified as $\alpha \cdot \phi_\kappa(\beta)$. Thus ϕ_κ is a crossed homomorphism as claimed.

To identify ϕ_κ with ϕ_r , it suffices to check equality on the standard generators σ_i . Under the *standard marking* shown in Figure 5 below, we see that σ_i^{-1} takes a_i to a_{i+1} and a_{i+1} to $T_{c_i}a_i$, where c_i is the boundary of the *standard arc* $A_{i,i+1}$ connecting marked points i and $i + 1$ (cf. Definition 6.5 below). By the twist-linearity formula, it follows that

$$\psi_T(\sigma_i^{-1}a_j) = \begin{cases} 0 & j \neq i + 1 \\ 1 & j = i + 1 \end{cases}$$

from which the claim follows.

5. Constructing monodromy elements

In this section, we “fill out” the monodromy image of $\text{Poly}_n(\mathbb{C})[\kappa]$, showing that the image $B_n[\kappa]$ contains the subgroup Γ'_n of “basic $(r + 1)^{st}$ twists”. This group is defined in

Definition 5.1 below; we exhibit some monodromy elements in Lemma 5.3, and after some group theory carried out in Lemma 5.4, we show the containment $\Gamma_n^r \leq B_n[\kappa]$ in Lemma 5.5.

Definition 5.1 (Basic $(r + 1)^{st}$ twist $\Sigma_{a;r}$, subgroup Γ_n^r). For $n \geq a + r$, the basic $(r + 1)^{st}$ twist $\Sigma_{a;r} \in B_n$ is defined to be the element

$$\Sigma_{a;r} = \sigma_a \dots \sigma_{a+r-1}.$$

The basic $(r + 1)^{st}$ twist group is the subgroup

$$\Gamma_n^r \leq B_n$$

generated by the set of basic $(r + 1)^{st}$ twists.

Remark 5.2. Pictorially, the basic twist $\Sigma_{a;r}$ is given by taking the strand in position a and crossing it over the next r strands to the right, and the inverse $\Sigma_{a;r}^{-1}$ is the same but with the strand crossing over r strands to the left. In particular, $\beta \in \Gamma_n^r$ if and only if it admits a diagram for which each overcrossing passes over a multiple of r strands below it.

LEMMA 5.3. Let

$$\kappa = \{k_1, \dots, k_p\}$$

be a partition of $n - 1$. Then $B_n[\kappa]$ contains the elements

$$\Sigma_{1;k_1}, \quad \Sigma_{k_1+1;k_2}, \quad \dots, \quad \Sigma_{k_1+\dots+k_{p-1}+1;k_p}.$$

Proof. Consider the ‘‘standard marking’’ of the translation surface $T_\kappa \in \Omega_\kappa$ shown in Figure 5. In Figure 6, we exhibit loops in Ω_κ based at T_κ . By comparing markings of the surface before and after, we compute their monodromy in B_n to be $\Sigma_{k_1+\dots+k_{i-1}+1;k_i}^{-1}$. Recalling from (1) that the projection $\mathcal{B}_n[\kappa] \rightarrow \pi_1^{ord}(\Omega_\kappa)$ is surjective, we see that we can lift these loops to $\text{Poly}_n(\mathbb{C})[\kappa]$, realising them as elements of the monodromy group $B_n[\kappa]$.

LEMMA 5.4. Let a, b be integers, and let $G \leq B_{a+b+1}$ be the subgroup generated by $\Sigma_{1;a}$ and $\Sigma_{a+1;b}$. Then $\Sigma_{1;\text{gcd}(a,b)} \in G$.

Proof. Observe that $\Sigma_{1;a} \Sigma_{a+1;b} = \Sigma_{1;a+b}$, and that

$$\Sigma_{1;a+b}^j \Sigma_{p;q} \Sigma_{1;a+b}^{-j} = \Sigma_{p+j;q} \tag{4}$$

so long as the indices $p + j, \dots, p + j + q - 1$ lie on the interval $[1, a + b]$. Thus by conjugating, we can shift the first index of any $\Sigma_{p;q}$ to any valid position, and by taking $\Sigma_{p;q}^{-1} \Sigma_{p;r}$ for $q > r$ and conjugating, we obtain $\Sigma_{p;r-q}$ from $\Sigma_{p;q}$ and $\Sigma_{p;r}$. By repeatedly shifting and deleting initial segments in this way, we can perform the Euclidean algorithm on a, b , eventually obtaining $\Sigma_{1;\text{gcd}(a,b)}$.

LEMMA 5.5. For any $n \geq 2$ and any partition κ of $n - 1$, the group $B_n[\kappa]$ contains the elements $\Sigma_{1;r}$ and $\Sigma_{1;n-1}$, and hence every basic $(r + 1)^{st}$ twist $\Sigma_{k;r}$. Thus,

$$\Gamma_n^r \leq B_n[\kappa].$$




Fig. 6. We depict the block of strips on T_κ between $S_{k_1+\dots+k_{i-1}+1}$ and $S_{k_1+\dots+k_i}$. Gluing instructions are indicated with colours and with letters. Reading lexicographically, we first recut the bottom strip $S_{k_1+\dots+k_{i-1}+1}$ so that it is bounded by the same cone point as the rest. Then we push each of the free prongs down one strip. Next, we apply a cut/paste move to reorder the strips, moving each one up one spot. Finally, we recut the bottom strip once again so that it is bounded by the other cone point. Note that in the case of the bottom block, there is no free prong in the bottom strip, in which case we skip the recutting steps, and in the case of the top block, there is no free prong in the top strip, and there is a slightly different picture (omitted). The picture at bottom right depicts the change of marking, i.e. the monodromy of the loop.

Proof. By Lemma 5.3, $B_n[\kappa]$ contains the elements

$$\Sigma_{1;k_1}, \quad \Sigma_{k_1+1;k_2}, \quad \dots, \quad \Sigma_{k_1+\dots+k_{p-1}-1;k_p}.$$

By repeated application of Lemma 5.4, one sees that $\Sigma_{1;r} \in B_n[\kappa]$, and also

$$\Sigma_{1;n-1} = \Sigma_{1;k_1} \Sigma_{k_1+1;k_2} \cdots \Sigma_{k_1+\dots+k_{p-1}-1;k_p} \in B_n[\kappa].$$

By (4), $B_n[\kappa]$ thus contains all $\Sigma_{k;r}$.

COROLLARY 5.6. *Let $n \geq 2$ be given, and let κ be a partition of $n - 1$ for which $r = \gcd(\kappa) = 1$. Then the monodromy map $\bar{\rho} : \mathcal{B}_n[\kappa] \rightarrow B_n$ is surjective.*

Proof. By Lemma 5.5, the image of $\bar{\rho}$ contains all basic $(r + 1)^{st}$ twists $\Sigma_{k;r}$, but for $r = 1$ these are just the standard half-twist generators of B_n .

6. Generating $\ker(\phi_r)$

In the previous two sections, we have seen how the monodromy image $B_n[\kappa]$ is contained in the kernel of a crossed homomorphism ϕ_r , and conversely contains the subgroup Γ_n^r of basic $(r + 1)^{st}$ twists. Here, we complete the circle of containments, showing that when n is sufficiently large compared to r , the kernel of ϕ_r is generated by basic $(r + 1)^{st}$ twists.

We must first specify what is meant by “sufficiently large”. Define

$$n_0(r, d) = \begin{cases} 8 & r = 2 \\ (6 + d)r & r \text{ odd} \\ (12 + d)r & r \geq 4 \text{ even.} \end{cases} \tag{5}$$

THEOREM 6.1. *Let $n \geq 3$ and $r \geq 2$ be given; let $d \in \{0, 1, 2\}$ be the remainder of $n/3$. Then for $n \geq n_0(r, d)$, the kernel of ϕ_r is generated by $\sigma_1 \dots \sigma_r$ and $\sigma_1 \dots \sigma_{n-1}$.*

The material of this section is purely braid-theoretic and does not require any knowledge e.g. of winding number functions. The outline is as follows. In Section 6.1, we discuss a new crossed homomorphism Υ_r , which can be computed graphically given a braid diagram as a count of “virtual undercrossings”; we show in Lemma 6.3 that $\Upsilon_r = \phi_r$. In Section 6.2, we use this graphical reformulation to give an algorithm for factoring an element of $\ker(\phi_r)$ supported on a small number of strands into the group Γ_n^r of basic $(r + 1)^{st}$ twists. Finally in Section 6.3, we exploit the factorisation algorithm to show the equality $\ker(\phi_r) = \Gamma_n^r$, first in Lemma 6.8 on the level of the pure braid group, and finally in Theorem 6.10 in general.

6.1. ϕ_r as a count of virtual undercrossings

Definition 6.2 (Virtual undercrossing map.) Let $r \geq 1$ be given. The virtual undercrossing map is the homomorphism¹ $\Upsilon_r : B_n \rightarrow \text{GL}_{n+1}(\mathbb{Z}/r\mathbb{Z}) \rtimes (\mathbb{Z}/r\mathbb{Z})^{n+1}$ defined as follows. Number the components of $(\mathbb{Z}/r\mathbb{Z})^{n+1}$ from 0 to n , and, for $1 \leq i \leq n - 1$, let $P_i \in \text{GL}_{n+1}(\mathbb{Z}/r\mathbb{Z})$ be the matrix obtained from I_{n+1} by replacing the i^{th} column with $e_{i-1} - e_i + e_{i+1}$. Then define

$$\Upsilon_r(\sigma_i) = (P_i, e_{i+1} - e_i)$$

for $1 \leq i \leq n - 1$. For $\beta \in B_n$, write

$$\Upsilon_r(\beta) = (M(\beta), v(\beta)).$$

¹ That this is indeed a homomorphism is verified by a routine calculation.

As is common to all homomorphisms into semi-direct products, the second factor $v(\beta)$ defines a crossed homomorphism $v : B_n \rightarrow (\mathbb{Z}/r\mathbb{Z})^{n+1}$ under the action of B_n on $(\mathbb{Z}/r\mathbb{Z})^{n+1}$ via M .

Also note that Υ_r defines an action of B_n on $(\mathbb{Z}/r\mathbb{Z})^{n+1}$ via

$$\beta \cdot \vec{x} = M(\beta)\vec{x} + v(\beta). \tag{6}$$

LEMMA 6.3. *Let $f : (\mathbb{Z}/r\mathbb{Z})^n \rightarrow (\mathbb{Z}/r\mathbb{Z})^{n+1}$ be given by*

$$e_i \in (\mathbb{Z}/r\mathbb{Z})^n \mapsto e_i - e_{i-1} \in (\mathbb{Z}/r\mathbb{Z})^{n+1}.$$

Then f induces a map of B_n -modules, where $(\mathbb{Z}/r\mathbb{Z})^n$ carries the standard permutation action of S_n and $(\mathbb{Z}/r\mathbb{Z})^{n+1}$ carries the action via M . Under the induced map on homology,

$$f_*(\phi_r) = v.$$

Moreover, $\ker(\phi_r) = \ker(v)$.

Proof. That f is a map of B_n -modules under the indicated actions is a routine calculation. To see that $f_*(\phi_r) = v$, it suffices to verify this on the standard generators σ_i . To that end, we compute

$$f_*(\phi_r)(\sigma_i) = f(\phi_r(\sigma_i)) = f(e_{i+1}) = e_{i+1} - e_i = v(r).$$

As $v = f_*(\phi_r)$, there is a containment $\ker(\phi_r) \leq \ker(v)$, and as f is readily seen to be an injection, it follows that this containment is an equality.

Virtual undercrossings. There is a graphical description of Υ_r which provides the key tool for expressing the kernel of these crossed homomorphisms in terms of $(r + 1)^{st}$ -twists. Let $\beta \in B_n$ be given. We imagine β (depicted in black) as sitting “on top of” a trivial braid (in blue) with a very large number of strands, where the ends of the blue strands are not fully fixed but are allowed to “slide” horizontally. Given $\vec{x} \in (\mathbb{Z}/r\mathbb{Z})^{n+1}$, we interpret the entries x_0, \dots, x_n as a mod- r count of the number of strands in the bottom (blue) layer positioned in between each pair of adjacent strands of β at the top of the figure. To compute the action of β on \vec{x} via Υ_r , we thread the strands in the bottom layer downwards, subject to the rule that strands in the bottom layer never cross, and that at each crossing of β , the total number of strands crossing under (counting both layers) is $0 \pmod r$.

Figure 7 illustrates this procedure in the case of a single crossing σ_i , and shows that the effect on the vector $\vec{x} = (x_0, \dots, x_n)$ is exactly given by $\sigma_i \cdot \vec{x}$ as in (6). To compute this action for a general braid, we simply repeat this process at each crossing, working from top to bottom. In particular, the value $v(\beta)$ is computed as the output of the virtual undercrossing procedure applying the zero vector at the top of the braid diagram for β . For future reference, we record the following characterisation of $\ker(\phi_r)$.

LEMMA 6.4. *A braid $\beta \in B_n$ lies in $\ker(\phi_r)$ if and only if the virtual undercrossing action for β satisfies $\beta \cdot \vec{0} = \vec{0}$. If β is moreover a pure braid, then $\beta \cdot \vec{x} = \vec{x}$ for $\vec{x} \in (\mathbb{Z}/r\mathbb{Z})^{n+1}$ arbitrary.*




Fig. 7 The virtual undercrossing procedure for the braid σ_i . The blue (lighter) strands are imagined as lying in a layer below the black strands of σ_i . Each blue (lighter) strand is actually composed of a very large number of individual strands; the numbers of such strands mod r are depicted above and below. By convention, all strands in the middle group (labeled x_i at the top) must cross in the same direction as the undercrossing in β . We then split the right strand (labeled x_{i+1} at the top) so that altogether, the number of strands crossing under (both blue/lighter and black) is $0 \pmod r$. As all x_i blue (lighter) strands in the middle group must cross under, and σ_i itself contributes one, we must borrow $-x_i - 1 \pmod r$ from the right strand to satisfy this condition.

Proof. As noted above, applying the virtual undercrossing procedure on a braid β to $\vec{0}$ yields $v(\beta)$. By Lemma 6.3, $v(\beta) = \vec{0}$ if and only if $\phi_r(\beta) = \vec{0}$. If β is any pure braid, then $M(\beta) = I_{n+1}$, and so $\beta \cdot \vec{x} = M(\beta)\vec{x} + v(\beta) = \vec{x} + v(\beta)$. Thus if $\beta \in \ker(\phi_r)$ is pure, $\beta \cdot \vec{x} = \vec{x}$ as claimed.

6.2. The factorisation algorithm

As discussed in the section outline above, the factorisation algorithm in this section gives a method for expressing an element of $\ker(\phi_r)$ in Γ_n^r when it is supported on a small number of strands. Our algorithm will require that the supporting subdisk have a particularly simple form which we call a *standard embedding*; we begin with this definition.

Definition 6.5 (Standard arc). Let D_n denote the disk with n marked points. An embedded arc $\alpha \subset D_n$ with endpoints at distinct marked points is standard if it is contained entirely in the lower half-disk. For each pair of marked points i, j , there is a unique isotopy class of standard arc connecting i and j , which is denoted α_{ij} .

Definition 6.6 (Standard embedding). Let D_n denote the disk with n marked points. An embedding $i: D_s \rightarrow D_n$ sending marked points to marked points is standard if it can be represented as a regular neighbourhood of a union of standard arcs which are disjoint except at endpoints.

An example of a standard embedding is depicted in Figure 8.




Fig. 8. The disk (shown in red) is a standard embedding of D_5 , as it is a regular neighbourhood of the four standard arcs (shown in blue).

LEMMA 6.7 (Factorisation algorithm). *Let $n \geq rs$. Let $i : D_s \rightarrow D_n$ be a standard embedding, and let $i_* : B_s \rightarrow B_n$ denote the corresponding inclusion of braid groups. Then there is a containment $i_*(PB_s) \cap \ker(\phi_r) \leq \Gamma_n^r$.*

Proof. We begin with an important special case, when D_s is the standard disk consisting of the first s points. This will serve to illustrate all of the key ideas of the argument. Then we will discuss the modifications necessary to apply in the general case.

Special case: first s strands. While reading this portion of the argument, the reader is invited to consult the worked example demonstrated in Figure 9. Let $\beta \in PB_s \cap \ker(\phi_r)$ be given. We view this as a braid α on s strands juxtaposed with a trivial braid τ on $n - s \geq (r - 1)s$ strands lying to the right. The key idea is to treat the strands of τ as the virtual strands in the virtual undercrossing procedure. Accordingly, we will depict the strands of α as black, and those of τ as blue, as in our discussion of virtual undercrossings above.

Recall (Remark 5.2) that a basic $(r + 1)^{st}$ twist $\Sigma_{i,r}$ consists of a single strand passing over r strands, so that in order to exhibit β as an element of Γ_n^r , it suffices to factor β so that all crossings have this form. To perform the factorisation, we will isotope the strands of τ , moving them to the left so that each overcrossing in α has 0 (mod r) total strands (black and blue) passing underneath.

In carrying this factorisation out, we will make use of the following operation. Given a braid β , obviously the product $\beta \Sigma_{i,r}$ lies in the same left coset of Γ_n^r as β . Graphically, $\beta \Sigma_{i,r}$ is obtained from β by taking the packet of r consecutive strands from $i + 1$ to $i + r$ and passing them one unit to the left under the i^{th} strand, the latter of which moves over r units to the right. We call this procedure *passing a packet* to the left; evidently there is also the analogous move of passing a packet to the right, corresponding to right-multiplication by $\Sigma_{i,r}^{-1}$. Likewise, we do not change the left coset by passing packets of r strands at the top of the braid.

Before presenting the algorithm, we make one final observation. Suppose we are given a particular braid diagram for β (not just its isotopy class). As usual, we think of the strands of α as black and the strands of τ as blue. At any vertical level where no two strands of β (black




Fig. 9. The factorisation algorithm applied in the case $r = 2$ to σ_1^4 . Reading left to right, we work our way down through the crossings, borrowing blue (lighter) strands so that successive crossings have a total of $r = 2$ strands crossing under. Since $\phi_2(\sigma_1^4) = \bar{0}$, after performing this procedure at all four crossings, the number of blue (lighter) strands in each position is $0 \pmod{2}$, and they can be passed back under in pairs, preserving the property that every overcrossing has an even number of strands passing underneath. Thus the algorithm produces the factorisation $\sigma_1^4 = (\sigma_1\sigma_2)^2(\sigma_2\sigma_3)^2(\sigma_1\sigma_2)^{-1}(\sigma_2\sigma_3)^{-1}$. The numbers indicate the counts of blue (lighter) strands in the indicated positions at the indicated levels; notice that they record the values of $v(\sigma_1^k)$ for $k = 0, \dots, 4$.

or blue) cross, we have a well-defined count of the number of blue strands in between each adjacent pair of black strands, giving us an integer vector \vec{v} with $s + 1$ entries. We call the space between strands i and $i + 1$ of α as the i th position. If the blue strands of τ are isotoped so as to conform to the conventions of the virtual undercrossings procedure (as described in Figure 7), the reduction of $\vec{v} \pmod{r}$ is equal to $v(\alpha')$, where α' is the portion of α from the top down to the specified vertical level.

We now explain the factorisation algorithm. Express α as a product of the standard generators of B_s , and suppose α begins with σ_i^ε with $\varepsilon = \pm 1$. To begin the factorisation, pass a packet of the first r strands of τ to the left; in the case of $\varepsilon = 1$, pass these to the $(i + 1)^{st}$ position, and if $\varepsilon = -1$, pass these to the $(i - 1)^{st}$. Pass $r - 1$ of these under the overcrossing and the remaining strand straight down, exactly as illustrated in Figure 7. We call this process *resolving a crossing*.

Now repeat this procedure for the remaining crossings of α : pass packets of r blue strands from τ to the relevant position, and borrow the necessary number of strands so as to create an undercrossing by r strands. If ever there are r or more blue strands in a single position after resolving a crossing, pass them in multiples of r all the way back to the right.

We must verify that as long as there are at least $(r - 1)s$ blue strands, it is always possible to pass a packet of r strands from the right over to the location of the overcrossing so as




Fig. 10. The factorisation algorithm as applied to a braid under a standard embedding ($r = 2$). As before, the method is to work down from the top, resolving crossings so as to have r strands crossing under each overcrossing. Note that after each crossing is resolved, the count of blue (lighter) strands in each position is given by $\alpha' \cdot \vec{x}$, where $\vec{x} = (1, 1, 0)$ is the count of blue (lighter) strands at the top of the diagram, α' is the initial segment of α to the given level, and the action is given by (6).

to facilitate a borrowing. Borrowing from the right is necessary only when the number of blue strands in consecutive positions is strictly less than the $r - 1$ needed to ensure that r strands pass under the given overcrossing, i.e. there are consecutive strand counts x_i, x_{i+1} for which $x_i + x_{i+1} \leq r - 2$. Since we have passed packets of r strands to the right (i.e. to x_n) whenever possible, each of the remaining $s - 2$ components x_j for $0 \leq j \leq n - 1$ has $x_j \leq r - 1$. Altogether then, in this situation, we have

$$\sum_{i=0}^{n-1} x_i \leq (s - 2)(r - 1) + r - 2 = (r - 1)s - r,$$

so that $x_n \geq r$ as was to be shown.

By Lemma 6.4, at the conclusion of this process, the number of blue strands mod r in each position is equal to the corresponding component of $v(\alpha) = \vec{0}$. As we have methodically passed packets of r blue strands to the right whenever possible, this shows that in fact there are no blue strands in between the black strands of α . In other words, the resulting braid diagram is isotopic to the original juxtaposition of α and τ . On the other hand, we have isotoped the blue strands of τ so that at every overcrossing, there are $0 \pmod r$ strands passing underneath, exhibiting β as a product of $(r + 1)^{st}$ twists as required.

General case: arbitrary standard embedding. The reader is now invited to consult Figure 10. Let $i : D_s \rightarrow D_n$ be a standard embedding, and let $\beta \in i_*(PB_n) \cap \ker(\phi_r)$ be given. As i is standard, we can represent β as a juxtaposition of a braid $i_*(\alpha)$ for $\alpha \in PB_s$ in black on top

of a trivial braid of $n - s$ strands in blue. In the language established above, the only difference between this setting and the special case above is that here we begin the factorisation algorithm with blue strands in arbitrary positions, not with all $n - s$ strands in position s as above. We proceed as before, working our way down from the top, resolving crossings by borrowing blue strands. The analysis above applies verbatim to show that when $n \geq rs$, there are sufficiently many blue strands available to make borrowing possible. It remains to be shown that the blue strands return to their original positions after all of the crossings of α are resolved. Recalling the hypothesis that β be a *pure* braid, this now follows from Lemma 6.4.

6.3. Factoring general braids

In this section, we conclude the proof of Theorem A. The main technical result is Lemma 6.9, which establishes the containment $\ker(\phi_r) \cap PB_n \leq \Gamma_n^r$. From there, the full containment (Theorem 6.10) and the proof of Theorem A are relatively easy.

To begin the analysis of $\ker(\phi_r) \cap PB_n$, we investigate the restriction of ϕ_r to PB_n .

LEMMA 6.8. *The restriction of ϕ_r to PB_n is a genuine homomorphism $\phi_r : PB_n \rightarrow (\mathbb{Z}/r\mathbb{Z})^n$, given on the standard generators A_{ij} of PB_n via*

$$\phi_r(A_{ij}) = e_i + e_j.$$

Proof. Since the action of B_n on $(\mathbb{Z}/r\mathbb{Z})^n$ factors through the quotient $B_n \rightarrow S_n$, it follows that the restriction of ϕ_r to PB_n is a homomorphism. We evaluate

$$\phi_r(A_{12}) = \phi_r(\sigma_1^2) = \phi_r(\sigma_1) + \sigma_1 \cdot \phi_r(\sigma_1) = (1 + (12)) \cdot e_2 = e_1 + e_2.$$

The formula

$$\phi_r(ghg^{-1}) = g \cdot \phi_r(h)$$

for $g \in B_n$ and $h \in PB_n$ is readily seen to hold, from which the expression $\phi_r(A_{ij}) = e_i + e_j$ follows.

LEMMA 6.9. *Let $n \geq 3$ and $r \geq 2$ be given, and let $d \in \{0, 1, 2\}$ be the remainder of $n/3$. Then for $n \geq n_0(r, d)$ (where $n_0(r, d)$ is defined as in (5)), there is a containment $PB_n \cap \ker(\phi_r) \leq \Gamma_n^r$.*

Proof. Let

$$w = \prod_{k=1}^m A_{i_k j_k}^{\varepsilon_k} \in PB_n \cap \ker(\phi_r) \tag{7}$$

be given. By hypothesis,

$$\phi_r(w) = \sum_{k=1}^m \varepsilon_k (e_{i_k} + e_{j_k}) = \vec{0} \in (\mathbb{Z}/r\mathbb{Z})^n.$$

To express $w \in \Gamma_n^r$, we will exploit the factorisation algorithm (Lemma 6.7) to rewrite the initial segment of w as a product of commuting elements of small support which has the same value under ϕ_r , removing initial segments that lie in $\ker(\phi_r)$ whenever possible. This

will take slightly different forms in the regimes $r = 2$, r odd, and $r \geq 4$ even; we begin with the case $r = 2$ since it is the simplest and will serve to illustrate the essential idea.

The case $r = 2$. Assume $r = 2$, and write

$$w = A_{i_1 j_1}^{\varepsilon_1} A_{i_2 j_2}^{\varepsilon_2} w'$$

with w' given as the product defining w in (7) above but starting at $k = 3$.

There are three possibilities for the size of the set $\{i_1, j_1, i_2, j_2\}$. Suppose first that $\{i_1, j_1\} = \{i_2, j_2\} = \{i, j\}$, in which case necessarily $\varepsilon_1 = \varepsilon_2$. Suppose for simplicity $\varepsilon_1 = \varepsilon_2 = 1$; the argument in the other case is analogous. Here, $A_{ij}^2 \in \ker(\phi_2)$. We note that A_{ij}^2 is a pure braid in $\ker(\phi_2)$ under a standard embedding of a disk with two marked points, and as $n \geq 8$ by hypothesis, we can apply the factorisation algorithm Lemma 6.7 to express $A_{ij}^2 \in \Gamma_n^2$. Thus in this case, we can write $w = \gamma w'$ with $\gamma \in \Gamma_n^2$, proceeding in turn to factorize w' .

Suppose next that $i_1 = i_2$ but $j_1 \neq j_2$. In this case, we have that

$$\phi_2(A_{i_1 j_1}^{\pm 1} A_{i_2 j_2}^{\pm 1}) = e_{j_1} + e_{j_2} = \phi_2(A_{j_1 j_2}),$$

so that

$$A_{i_1 j_1}^{\pm 1} A_{i_2 j_2}^{\pm 1} A_{j_1 j_2}^{-1} \in \ker(\phi_2).$$

This element is again in the image of a standard embedding, so that we can apply the factorisation algorithm (Lemma 6.7) to express $A_{i_1 j_1}^{\pm 1} A_{i_2 j_2}^{\pm 1} A_{j_1 j_2}^{-1}$ as an element of Γ_n^2 , the disk is standard and has four marked points, and hence this is possible for $n \geq 8$. Thus, by left-multiplying by an element of Γ_n^2 we can replace the initial segment $A_{i_1 j_1}^{\pm 1} A_{i_2 j_2}^{\pm 1}$ with the initial segment $A_{j_1 j_2}$, which has smaller support. Similar arguments apply to the various cases when $|\{i_1, j_1, i_2, j_2\}| = 3$.

The remaining possibility is that $|\{i_1, j_1, i_2, j_2\}| = 4$. In this case, we replace the initial segment $A_{i_1 j_1}^{\pm 1} A_{i_2 j_2}^{\pm 1}$ with the segment $A_{a,b} A_{c,d}$, where $\{i_1, j_1, i_2, j_2\} = \{a, b, c, d\}$ and $a < b < c < d$. Note in particular that the elements $A_{a,b}, A_{c,d}$ of the initial segment commute, and that the pair of elements $A_{a,d} A_{b,c}$ also commute; we say that the former pair is *un-nested* and the latter *nested*.

We continue in this way, expressing

$$w = \gamma A_{i_p j_p}^{\pm 1} \prod_{k=p+1}^m A_{i_k j_k}^{\varepsilon_k}$$

with γ a product of pairwise un-nested commuting generators of PB_n . The support of the element $A_{i_p j_p}^{\pm 1}$ intersects the support of 0, 1, or 2 of the elements of γ . If it intersects zero, it may be nested with up to one. This can be resolved by pulling these two commuting elements to the front of γ and replacing them as above with their un-nested counterpart. If it intersects one, these two elements can likewise be moved to the front of γ and resolved into one or two basic elements as above. Finally suppose it intersects the support of two, say $A_{i_1 j_1}$ and $A_{i_2 j_2}$; by the non-nestedness hypothesis, $i_1 < j_1 < i_2 < j_2$. For $A_{i_p j_p}^{\pm 1}$ to intersect both, we must have

$$i_1 \leq i_p \leq j_1 \text{ and } i_2 \leq j_p \leq j_2.$$

Thus the product $A_{i_1 j_1} A_{i_2 j_2} A_{i_p j_p}$ is supported on a standardly-embedded disk of up to six elements. As we are only assuming $n \geq 8$ and the factorisation algorithm (Lemma 6.7)

requires $n \geq 2s = 12$ strands to factor an element supported on six strands, we first rewrite $A_{i_1 j_1} A_{i_2 j_2} = A_{i_1 j_2} A_{i_2 j_1}$. If $i_1 < i_p < j_1 < i_2 < j_p < j_2$, we can then rewrite $A_{i_1 j_2} A_{i_p j_p}^{\pm}$ as the un-nested pair $A_{i_1 i_p} A_{j_p j_2}$ which is then un-nested with $A_{j_1 i_2}$. If $i_1 = i_p$, then the initial segment $A_{i_1 j_2} A_{i_1 j_p}$ can be replaced with $A_{j_2 j_p}$; the remaining cases where the total support is five strands can be handled analogously. Finally, if $A_{i_p j_p}^{\pm 1}$ intersects both $A_{i_1 j_1}$ and $A_{i_2 j_2}$ but the total number of strands is four, then this can be rewritten directly via the factorisation algorithm without any need for initial re-writing.

Altogether then, this process converts an arbitrary word $w \in PB_n \cap \ker(\phi_2)$ into a product of pairwise un-nested and commuting generators A_{ij} . Since $\phi_2(w) = 0$ by hypothesis, this implies that each A_{ij} appears an even number of times; these can then be successively removed from w by means of the factorisation algorithm (Lemma 6.7).

The case $r \geq 3$ odd. In broad outline, we proceed in the same way as in the case $r = 2$. Given $w \in PB_n \cap \ker(\phi_r)$ as in (7), we rewrite the initial segment of w as a set of elements with disjoint and small support. Whereas in the case $r = 2$ these elements were the generators A_{ij} of PB_n , here we will use elements $A_{i;a,b,c}$ which we proceed to define.

Let $a, b, c \in \mathbb{Z}/r\mathbb{Z}$ be given. Set $p = (r + 1)/2$, and define

$$A_{i;a,b,c} = A_{i,i+1}^{p(a+b-c)} A_{i,i+2}^{p(a-b+c)} A_{i+1,i+2}^{p(-a+b+c)}.$$

Observe that

$$\phi_r(A_{i;a,b,c}) = ae_i + be_{i+1} + ce_{i+2}.$$

Also note that $A_{i;a,b,c}$ and $A_{j;a',b',c'}$ have disjoint support whenever $|i - j| \geq 3$.

Divide the strands from 1 to n into groups of three; the last group will contain 3, 4, or 5 depending on the value of d , i.e. the remainder of $n \pmod{3}$. We extend the definition of $A_{1+3k;a,b,c}$ to this last group by taking $A_{1+3k;a,b,c,d,e} = A_{1+3k;a,b,c} A_{3+3k,0,d,e}$ in the case $d = 2$ and similarly for $d = 1$; to simplify notation we will tacitly understand that the last $A_{i;a,b,c}$ may be of this form. Suppose we have a partial factorisation

$$w = \gamma A_{i_p j_p}^{\epsilon_p} w',$$

where γ is a product of elements of the form $A_{1+3k;a,b,c}$. Then $A_{i_p j_p}^{\epsilon_p}$ intersects the support of at most two such elements, and altogether the product of these three elements is supported on a standardly-embedded disk with at most $6 + d$ punctures. Pulling these to the front of γ , since we assume $n \geq n_0(r, d) = (6 + d)r$, we apply the factorisation algorithm (Lemma 6.7) and replace this with a product of up to two elements of the form $A_{1+3k;a,b,c}$ with the same ϕ_r -value.

After completing this process, we have factored $w = \gamma$ into a product of elements of the form $A_{1+3k;a,b,c}$. As $\phi_r(w) = \vec{0}$ by hypothesis, it follows that each $\phi_r(A_{1+3k;a,b,c}) = \vec{0}$ as well. Applying the factorisation algorithm to each of these in turn, we express $w = \gamma$ as an element of Γ_n^r .

The case $r \geq 4$ even. In this last case, we combine the methods of the previous two. Again, the objective is to factor initial segments of w into disjoint elements of small support with the same value of ϕ_r . Like in the case of r odd, we partition the strands into groups of three and attempt to factor the initial segment into elements supported on these groups. But unlike this case, there is a parity phenomenon to keep track of, which will require us to *link* two such groups if the parity of ϕ_r on each is odd.

We define $A_{i;a,b,c}$ analogously as above, this time subject to the requirement that $a + b + c$ be even; under this hypothesis, each of the integers $\pm a \pm b \pm c$ is even and so can be divided by two in the exponent. As before, the last $A_{1+3k;a,b,c}$ may actually be supported on up to five strands. Given an initial segment $\gamma \in B_n$ of w , we say that a group of integers $1 + 3k, 2 + 3k, 3 + 3k$ is *even* if the sum of the coefficients of $\phi_r(\gamma)$ on $e_{1+3k}, e_{2+3k}, e_{3+3k}$ is even, and *odd* otherwise. Observe that there is always an even number of odd groups, since each $A_{ij}^{\pm 1}$ changes the parity of either zero or two groups.

We now describe the structure of the initial segment we will construct. We will express $w = \gamma w'$ where γ is a product of $A_{1+3k;a,b,c}$ over all even groups, along with products of the form

$$A_{1+3k;a,b,c}A_{1+3k';a',b',c'}A_{i,j},$$

where the groups starting at $1 + 3k$ and $1 + 3k'$ are odd, and where the first group contains i and the second contains j . We moreover impose the condition that there are no odd groups in between $1 + 3k$ and $1 + 3k'$. In this way, the structure of the supports mimics that in the case $r = 2$: they are disjoint, un-nested, and supported on standardly-embedded disks.

Given a partial factorisation

$$w = \gamma A_{i_p j_p}^{\epsilon_p} w'$$

of this form, we consider the various possibilities for how the support of $A_{i_p j_p}^{\epsilon_p}$ intersects the supports of the elements in γ . This exactly mirrors the analysis carried out in the case $r = 2$, but this time we apply the factorisation algorithm to elements supported on up to $12 + d$ strands, in the case where we need to convert between nested and un-nested factorisations on two pairs of odd groups, one of which contains the exceptional group of $3 + d$ elements.

THEOREM 6.10. *Let $n_0(r, d)$ and d be given as in Lemma 6.9. Then for $n \geq n_0(r, d)$, there is an equality*

$$\ker(\phi_r) = \Gamma_n^r,$$

i.e. $\ker(\phi_r)$ is generated by the finite set of basic $(r + 1)^{st}$ twists.

Proof. The bulk of the work has been carried out above in Lemma 6.9, which establishes the containment

$$PB_n \cap \ker(\phi_r) \leq \Gamma_n^r$$

in the range $n \geq n_0(r, d)$. Conversely, it is easy to verify that

$$\phi_r(\Sigma_{i;r}) = 0,$$

so that

$$\Gamma_n^r \leq \ker(\phi_r).$$

It remains only to show that the images of $\ker(\phi_r)$ and Γ_n^r in S_n coincide; denote these subgroups of S_n by $\overline{\ker(\phi_r)}$ and $\overline{\Gamma_n^r}$, respectively.

The basic $(r + 1)^{st}$ twists that generate Γ_n^r are sent to the $r + 1$ -cycles

$$(1 \dots r + 1), (2 \dots r + 2), \dots, (n - r - 1 \dots n)$$

in S_n . Thus

$$(1 \ r + 1 \ r + 2) = (1 \dots r + 1)^{-1}(2 \dots r + 2) \in \overline{\Gamma_n^r},$$

and hence also

$$(1 \ 2 \ 3) = (2 \dots r + 2)^2(1 \ r + 1 \ r + 2)(2 \dots r + 2)^{-2} \in \overline{\Gamma_n^r}.$$

As also Γ_n^r contains the element $\Sigma_{1;n-1} = \sigma_1 \dots \sigma_{n-1}$, the image $\overline{\Gamma_n^r}$ contains the cyclic permutation $(1 \dots n)$. Conjugating $(1 \ 2 \ 3)$ by this, it follows that $\overline{\Gamma_n^r}$ contains all 3-cycles of the form $(i \ i + 1 \ i + 2)$. This is well known to generate the alternating group A_n . We conclude that

$$A_n \leq \overline{\Gamma_n^r}$$

for all n, r . For r odd, the $r + 1$ -cycles are odd permutations and are even otherwise, from which it follows that

$$\overline{\Gamma_n^r} = \begin{cases} S_n & r \text{ odd} \\ A_n & r \text{ even.} \end{cases}$$

It remains only to show that $\overline{\ker(\phi_r)} \leq A_n$ when r is even. To see this, we recall that ϕ_r can be viewed as the homomorphism

$$\phi_r : B_n \rightarrow S_n \times (\mathbb{Z}/r\mathbb{Z})^n$$

which sends σ_i to the pair $((i \ i + 1), e_{i+1})$. Let $\text{sgn} : S_n \rightarrow \mathbb{Z}/2\mathbb{Z}$ denote the sign homomorphism, and let $s : (\mathbb{Z}/r\mathbb{Z})^n \rightarrow \mathbb{Z}/2\mathbb{Z}$ be the reduction mod 2 of the sum-of-coefficients map. Then

$$\text{sgn} + s : S_n \times (\mathbb{Z}/r\mathbb{Z})^n \rightarrow \mathbb{Z}/2\mathbb{Z}$$

is a surjective homomorphism, and the composition $(\text{sgn} + s) \circ \phi_r$ is identically zero (being zero on each generator of B_n by above), from which it follows that $\overline{\ker(\phi_r)} \leq A_n$ as was to be shown.

Proof of Theorem A. Lemma 4.7 establishes the containment

$$B_n[\kappa] \leq \ker(\phi_r),$$

Lemma 5.3 shows that

$$\Gamma_n^r \leq B_n[\kappa],$$

and Theorem 6.10 shows the equality

$$\Gamma_n^r = \ker(\phi_r)$$

in the range $n \geq n_0(r, d)$.

Acknowledgements. The author would like to thank Tara Brendle and Matt Day for interesting discussions, and Dan Margalit for very helpful feedback and for alerting the author to the work [DM22] of Dougherty–McCammond. Stepan Orevkov, Michael Lönne, Alex Suciu, and Kathryn Lindsey are due thanks for some helpful bibliographical suggestions. Thanks also to the anonymous referee. The author is supported by NSF Award No. DMS-2153879.

REFERENCES

- [CS21] A. CALDERON and N. SALTER. Higher spin mapping class groups and strata of abelian differentials over Teichmüller space. *Adv. Math.* **389**(107926) (2021).
- [CS23] A. CALDERON and N. SALTER. Framed mapping class groups and the monodromy of strata of Abelian differentials. *J. Eur. Math. Soc.* To appear, 2023.
- [CW91] F. CATANESE and B. WAJNRYB. The fundamental group of generic polynomials. *Topology* **30**(4) (1991), 641–651.
- [DM22] M. DOUGHERTY and J. MCCAMMOND. Geometric combinatorics of polynomials I: The case of a single polynomial. *J. Algebra* **607** (part B) (2022), 106–138.
- [GT21] Q. GENDRON and G. TAHAR. Différentielles abéliennes à singularités prescrites. *J. Éc. polytech. Math.* **8** (2021), 1397–1428.
- [GT22] Q. GENDRON and G. TAHAR. Isoresidual fibration and resonance arrangements. *Lett. Math. Phys.* **12**(2) (33) (2022).
- [KZ97] M. KONTSEVICH and A. ZORICH. Lyapunov exponents and Hodge theory. Preprint, <https://arxiv.org/abs/hep-th/9701164> (1997).
- [Loo08] E. LOOIJENGA. Artin groups and the fundamental groups of some moduli spaces. *J. Topol.* **1**(1) (2008), 187–216.
- [Mir95] R. MIRANDA. *Algebraic Curves and Riemann Surfaces*. Grad. Stud. Math. vol. 5 (American Mathematical Society, Providence, RI, 1995).
- [TBY+20] W. P. THURSTON, H. BAIK, G. YAN, J. H. HUBBARD, K. A. LINDSEY, L. TAN and D. P. THURSTON. Degree- d -invariant laminations. In *What's Next?—the Mathematical Legacy of William P. Thurston*. Ann. of Math. Stud. vol. 205 (Princeton University Press, Princeton, NJ, 2020), pages 259–325.