

Siegel–Veech Constants for Cyclic Covers of Generic Translation Surfaces

David Auricino, Aaron Calderon, Carlos Matheus, Nick Salter, and Martin Schmoll

Abstract

We compute the asymptotic number of cylinders, weighted by their area to any non-negative power, on any cyclic branched cover of any generic translation surface in any stratum. Our formulas depend only on topological invariants of the cover and number-theoretic properties of the degree: in particular, the ratio of the related Siegel–Veech constants for the locus of covers and for the base stratum component is independent of the number of branch values. One surprising corollary is that this ratio for $area^3$ Siegel–Veech constants is always equal to the reciprocal of the the degree of the cover. A key ingredient is a classification of the connected components of certain loci of cyclic branched covers.

Contents

1	Introduction	3
I	Preliminaries and setup	8
2	Branched covers of translation surfaces	9
2.1	Translation surfaces	9
2.2	Branched covers and branching vectors	9
2.3	Cylinders	10
3	Covers of strata and monodromy	10
3.1	Strata	11
3.2	Monodromy representations	11
3.3	Homological monodromy: basic structure	12
3.4	Orbits of branching vectors	13
3.5	Orbits with fixed monodromy	13
II	Structure and cardinality of orbits	13
4	Orbit counts (I): the unsplit case	14
4.1	Statement of results	15
4.2	Proofs	16
5	Framings, winding number functions, and quadratic forms	18
5.1	Framings, winding number functions, and quadratic forms	18
5.2	Geometric splittings	19
5.3	The crossed homomorphism $\Theta_{q,\kappa}$	20
6	The homological monodromy group, non-hyperelliptic case	22

7	The ψ invariant	27
8	Orbit counts (II): Split non-hyperelliptic case	30
8.1	Statement of results	30
8.2	Proofs	31
9	The homological monodromy group, hyperelliptic case	33
9.1	The hyperelliptic mapping class group	34
9.2	The homological monodromy group	35
10	Orbit counts (III): Split hyperelliptic case	36
10.1	Encircling branch points: the b invariant	37
10.2	Statement of results	38
10.3	Proofs	39
III	Computing the Siegel–Veech constants	41
11	Siegel–Veech preliminaries	41
11.1	The Siegel–Veech transform	41
11.2	Cusp volumes and configurations	42
12	From integrals to orbit quotients	44
12.1	Siegel–Veech constants of cyclic covers	45
12.2	Explicitly expressing the counting function	46
12.3	Simplifying the counting function	49
12.4	Conclusion of the proof of Theorem 12.5	50
13	Evaluating the orbit quotient function	53
13.1	Evaluation of orbit ratios for powers of 2	54
14	Main formula	55
14.1	area- and area cubed–Siegel–Veech constants	57
A	Symplectic modules	57
B	Action on the homology classes of cylinders	58
B.1	Non-hyperelliptic	59
B.2	Hyperelliptic	62

1 Introduction

It was discovered in [Vee98, EM01] that the growth rate of the number of periodic trajectories on a generic translation surface satisfy exact asymptotics of the form cL^2 , where L is the length of the trajectory, and c is now known as a *Siegel–Veech constant*. Furthermore, [EM01] found these asymptotics also hold for appropriately weighted periodic trajectories and generic surfaces with respect to any $\mathrm{SL}_2(\mathbb{R})$ -invariant measure. The next breakthrough came with [EMZ03], which related Siegel–Veech constants to *computable* volumes of strata. This opened an industry of computing Siegel–Veech constants for many different invariant measures, see [EMZ03, EMS03, BG15, Gou15, AEZ16, CMZ18, Sau18, CMSZ20, DZ20, DGZZ21] for examples. The main result of the present work is Theorem 14.1, in which we derive explicit formulas for the area^s -Siegel–Veech constant for all loci of cyclic branched covers of all components of all strata in terms of the area^s -Siegel–Veech constant of the base stratum component.

To define our terms, consider a translation surface (X, ω) . Let the width of a cylinder $w(\mathit{cyl})$ on (X, ω) refer to its circumference, i.e., the absolute value of the period of a core curve of the cylinder. Let $s \geq 0$, and for any $L > 0$ and any (X, ω) , define the weighted counting function

$$N_{\mathit{area}^s}((X, \omega), L) = \sum_{\substack{\mathit{cyl} \subseteq (X, \omega) \\ w(\mathit{cyl}) < L}} \frac{\mathit{area}^s(\mathit{cyl})}{\mathit{area}^s(X, \omega)}.$$

By [EM18, EMM15], every $\mathrm{SL}_2(\mathbb{R})$ -invariant measure is fully supported on an invariant subvariety \mathcal{M} in the moduli space of translation surfaces. It follows from [EM01] that the limit

$$c_{\mathit{area}^s}(\mathcal{M}) = \lim_{L \rightarrow \infty} \frac{N_{\mathit{area}^s}((X, \omega), L)}{\pi L^2}$$

exists and is the same for almost every (X, ω) in \mathcal{M} with respect to its Lebesgue measure class. This limit is called the area^s -*Siegel–Veech constant* of \mathcal{M} . When $s = 0$, this is the classical cylinder Siegel–Veech constant. We warn the reader that these constants were defined in the same way in [BG15], but defined differently (in a non- $\mathrm{SL}_2(\mathbb{R})$ -invariant way) in [CMZ18, Section 1.1].

If $(\tilde{X}, \tilde{\omega}) \rightarrow (X, \omega)$ is a branched covering of translation surfaces, then the $\mathrm{SL}_2(\mathbb{R})$ -orbit closure of $(\tilde{X}, \tilde{\omega})$ is readily seen to consist of branched covers of surfaces in the $\mathrm{SL}_2(\mathbb{R})$ -orbit closure of (X, ω) . Mirzakhani conjectured that sufficiently large orbit closures should *only* arise from such branched covering constructions [Wri14]; while this has been disproven in its original form [MMW17], its idea seems to carry through in many cases [MW18, Api18, AN20, AW23]. Thus, understanding the properties of branched covering constructions represents significant progress toward understanding all orbit closures.

Another motivation comes from studying the large-genus asymptotics of Siegel–Veech constants; there has recently been a great deal of progress on this problem for sequences of generic translation surfaces [EZ15, CMZ18, CMSZ20, Agg19, Agg20, Agg21]. Here we consider a novel method for constructing sequences of surfaces of increasing genus: by taking cyclic branched covers with the same degree, but an increasing number of branch points, we obtain sequences of orbit closures with growing dimension and growing genus, but which all have *the same* area -Siegel–Veech constant (Corollary B).

Ratios of Siegel–Veech constants. This paper has two main results. The first is to compute the area^s -Siegel–Veech constants for each component of the locus of cyclic branched covers of surfaces in a given stratum component (Theorem 14.1). These computations crucially rely on our second main result, in which we classify connected components of these loci (Theorem D). The precise formulas are quite intricate and will not be stated in the introduction except in the simplest case; instead, we survey some notable corollaries.

We first set some notation. Let $\kappa = (k_1, \dots, k_n)$ be a partition of $2g - 2$ into non-negative integers. Let \mathcal{H} be a component of a stratum of unit-area translation surfaces with (labeled) zeros of order κ ; a zero of order 0 corresponds to a regular marked point. Pick $(X, \omega) \in \mathcal{H}$ and take a degree d cyclic cover branched over the set \mathcal{B} of zeros and regular marked points. For each point $p_i \in \mathcal{B}$, let d_i denote the *local monodromy* at p_i , i.e., the element $d_i \in \mathbb{Z}/d\mathbb{Z}$ classifying the branched cover near p_i . We allow some d_i to be 0, corresponding to an unbranched cover around p_i .

Let $\mathcal{Z} \subset \mathcal{B}$ be the set of points corresponding to zeros of order at least 1; up to reindexing, we may assume that these points correspond to the first $|\mathcal{Z}|$ elements of \mathcal{B} . Recalling that $|\mathcal{B}| = n$, define

$$\delta := (d_1, \dots, d_n) \in (\mathbb{Z}/d\mathbb{Z})^{\mathcal{B}} \quad \text{and} \quad \pi(\delta) = (d_1, \dots, d_{|\mathcal{Z}|}) \in (\mathbb{Z}/d\mathbb{Z})^{\mathcal{Z}}.$$

Denote $\gcd(\delta) = \gcd(d_1, \dots, d_n, d)$.

Let $\mathcal{M}_\delta \rightarrow \mathcal{H}$ be the locus of all degree d cyclic branched covers of surfaces in \mathcal{H} with local monodromy δ , normalized to have unit area. Depending on the base stratum \mathcal{H} and the invariants defined above, \mathcal{M}_δ may not be connected, and when it is not, its components are classified by certain “ ψ or b invariants” defined at the level of the mod-2 homology of X relative to \mathcal{B} (see Definitions 7.3 and 10.3, and compare Theorem D).

One of the most surprising features of our formulas is that they do not depend on the number of branch points, but only on the structure of the base stratum and coarse numerical data about the branching.

Theorem A (Corollary of Theorem 14.1). *Let \mathcal{H} be a connected component of a stratum of genus g translation surfaces (possibly with regular marked points) and let \mathcal{C} be a connected component of \mathcal{M}_δ . Then*

$$\frac{c_{\text{area}^s}(\mathcal{C})}{c_{\text{area}^s}(\mathcal{H})}$$

depends only on g , s , d , $\gcd(\delta)$, the ψ or b invariant of \mathcal{C} (when defined), and the Arf invariant of \mathcal{H} (when defined). In particular, this quotient is independent of the number of branch values $n = |\mathcal{B}|$.

Remark. We emphasize that Theorem 14.1 allows one to compute numerical values for $c_{\text{area}^s}(\mathcal{C})$. The denominator $c_{\text{area}^s}(\mathcal{H})$ can be computed in terms of cylinder Siegel–Veech constants via [BG15], which can in turn be computed recursively using [EMZ03]. A similar reduction was used to compute area-Siegel–Veech constants in [EKZ14, §2.5.2], which relied on a theorem of [Vor05].

The value of the denominator in Theorem A generally *does* depend on n . However, for $s = 1$, adding regular marked points does not change the area-Siegel–Veech constant of the base stratum (since the areas of cylinders add linearly), and consequently, the numerator also does not change.

Corollary B. *Let \mathcal{H} be a connected component of a stratum of genus g translation surfaces (possibly with regular marked points) and let \mathcal{C} be a connected component of \mathcal{M}_δ . Let \mathcal{H}' be the stratum component obtained from \mathcal{H} by forgetting all regular marked points. Then $c_{\text{area}}(\mathcal{C})$ depends only on g , d , $\gcd(\delta)$, the ψ or b invariant of \mathcal{C} (when defined), and the Arf invariant of \mathcal{H}' (when defined). In particular, it is independent of the number of branch values $n = |\mathcal{B}|$.*

We now highlight a particularly striking example of the phenomenon observed in Theorem A; note that the following result does not depend on which stratum component \mathcal{H} or which component of \mathcal{M}_δ is chosen. To our knowledge, area^3 -Siegel–Veech constants have never been specifically considered in the literature.

Theorem C (Corollary of Theorem 14.1). *Let \mathcal{H} be a component of a stratum of unit-area translation surfaces (possibly with regular marked points) and let \mathcal{C} be a connected component of \mathcal{M}_δ . Then*

$$\frac{c_{\text{area}^3}(\mathcal{C})}{c_{\text{area}^3}(\mathcal{H})} = \frac{1}{d}.$$

A (facile) heuristic. Consider a degree d translation covering $\pi : (\tilde{X}, \tilde{\omega}) \rightarrow (X, \omega)$ and suppose that there were some c such that each cylinder on (X, ω) lifts to exactly c cylinders on $(\tilde{X}, \tilde{\omega})$, each of circumference

d/c times its circumference on (X, ω) .¹ We could then estimate

$$\begin{aligned} N_{\text{area}^s}((\tilde{X}, \tilde{\omega}), L) &:= \sum_{\substack{\tilde{cyl} \subseteq (\tilde{X}, \tilde{\omega}) \\ w(\tilde{cyl}) < L}} \frac{\text{area}^s(\tilde{cyl})}{\text{area}^s(\tilde{X}, \tilde{\omega})} \\ &= \sum_{\substack{cyl \subseteq (X, \omega) \\ w(cyl) < Lc/d}} c \cdot \frac{\text{area}^s(cyl)(d/c)^s}{d^s} = c^{1-s} N_{\text{area}^s}((X, \omega), Lc/d). \end{aligned}$$

Normalizing the covering surface to have unit area and using the fact that our counting function shrinks (approximately) linearly with area and grows (approximately) quadratically with length, this yields

$$N_{\text{area}^s}((\tilde{X}, \tilde{\omega}/\sqrt{d}), L) \approx d \cdot N_{\text{area}^s}((\tilde{X}, \tilde{\omega}), L) \approx \frac{c^{3-s}}{d} N_{\text{area}^s}((X, \omega), L).$$

Thus, we would expect that

$$c_{\text{area}^s}(\mathcal{C}) = \frac{c^{3-s}}{d} c_{\text{area}^s}(\mathcal{H}).$$

This heuristic is incredibly naïve, and shockingly, correct when $s = 3$. Our formulas for general s bear some resemblance to this heuristic: compare especially Theorem 12.5.

Components of loci of covers. The reason that the ψ and b invariants appear in Theorem A and Corollary B is because whenever \mathcal{M}_δ is disconnected, they form a family of complete invariants distinguishing its connected components. Observe again that the number n of branch values does not appear in the classification below.

Theorem D. *Let \mathcal{H} be a connected component of a stratum of unit-area translation surfaces with zeros of order κ , possibly with regular marked points. Let $\mathcal{M}_\delta \rightarrow \mathcal{H}$ be the locus of degree d cyclic branched covers with local branching data δ and let \mathcal{H}' be obtained from \mathcal{H} by forgetting all regular marked points.*

1. *If d is odd or $g = 1$, then \mathcal{M}_δ is connected.*
2. *If \mathcal{H}' is nonhyperelliptic and d is even,*
 - (a) *If $\delta \not\equiv \kappa \pmod{2}$, then \mathcal{M}_δ is connected.*
 - (b) *If $\delta \equiv \kappa \pmod{2}$, then \mathcal{M}_δ has two components, classified by the ψ invariant.*
3. *If \mathcal{H}' is hyperelliptic and d is even,*
 - (a) *If $\pi(\delta) \neq 0$, then \mathcal{M}_δ is connected.*
 - (b) *If $\pi(\delta) = 0$, then the components of \mathcal{M}_δ are classified by the b invariant.*
 - *If $\mathcal{H}' = \mathcal{H}^{\text{hyp}}(2g - 2)$, there are g components.*
 - *If $\mathcal{H}' = \mathcal{H}^{\text{hyp}}(g - 1, g - 1)$, there are $g + 1$ components.*

When \mathcal{M}_δ is connected, the formulas in Theorem 14.1 simplify. To give the reader an idea of their form, we include this reduced version below.

Set d_{abs} to be the greatest divisor of d coprime to $\gcd(\delta)$ and set $d_{\text{rel}} = d/d_{\text{abs}}$.

Theorem E. *If \mathcal{M}_δ is connected (cases 1, 2a, and 3a above), then for all $s \geq 0$,*

$$\frac{c_{\text{area}^s}(\mathcal{M}_\delta)}{c_{\text{area}^s}(\mathcal{H})} = \frac{d_{\text{rel}}}{d^{s-1} \Phi_{2g}(d_{\text{rel}})} \left(\sum_{\mathfrak{d} | d_{\text{rel}}} \Phi_{2g-1} \left(\frac{d_{\text{rel}}}{\mathfrak{d}} \right) \frac{\Phi(\mathfrak{d})}{\mathfrak{d}^{4-2g-s}} \right) \cdot \left(\sum_{\mathfrak{d} | d_{\text{abs}}} \left(\frac{\Phi(\mathfrak{d})}{\mathfrak{d}^{3-s}} \right) \right),$$

where Φ_m denotes the Jordan totient function (see §4.1) and Φ is the usual Euler totient function.

¹To make this heuristic slightly more sophisticated, one could imagine that the cylinders of (X, ω) were “well-distributed” with respect to the cover, so that c is an average number of preimages of a cylinder on (X, ω) .

Remark. Even though this will not be used in this work, we note that the formula in Theorem E can be rewritten using a Dirichlet convolution (denoted $*$). Let $Id_k(n) = n^k$ and $1(n) = 1$. Then

$$\frac{c_{\text{area}^s}(\mathcal{M}_\delta)}{c_{\text{area}^s}(\mathcal{H})} = \left(\frac{d_{rel}}{\Phi_{2g}(d_{rel})} \right) ((Id_{2g+s-4} \cdot \Phi) * \Phi_{2g-1})(d_{rel}) \cdot ((Id_{s-3} \cdot \Phi) * 1)(d_{abs}).$$

In [Sch19], the fifth-named author derived this formula for area^s -Siegel–Veech constants for cyclic torus covers branched over *two* points. The formula there differs from ours due to a different choice of normalization. We emphasize that the formula of Theorem E is new even for torus covers. Substituting $g = s = 1$ into Theorem E and using a standard identity for totient functions, we get the following notable corollary.

Corollary F. *Let \mathcal{M}_δ be the locus of degree d cyclic covers branched over a unit area flat torus with $n \geq 2$ marked points, normalized to have unit area. Then*

$$c_{\text{area}}(\mathcal{M}_\delta) = \frac{3}{\pi^2} \sum_{\mathfrak{d} | d_{abs}} \frac{\Phi(\mathfrak{d})}{\mathfrak{d}^2} = \frac{3}{\pi^2} ((Id_{-2} \cdot \Phi) * 1)(d_{abs}),$$

where Φ denotes the Euler totient function.

In [CMZ18], they compute formulas for area -Siegel–Veech constants for arbitrary branched covers of tori with arbitrarily many marked points. There does not appear to be an obvious simplification to arrive at the formula in Corollary F from those of [CMZ18]. We also observe that the setting of this corollary is distinct from the results of [EKZ11] because they consider tori branched over the 2-torsion points, while our branching occurs over free marked points.

Given the importance of area -Siegel–Veech constants, we give the full general formula for them in all cases in Section 14.1. On the other hand, we do not explicitly state the cylinder Siegel–Veech constant for \mathcal{M}_δ in this manuscript because it is not particularly striking. The interested reader can substitute $s = 0$ into the formulas in Theorem 14.1 to recover the values in all cases.

Large genus asymptotics. Since our formulas do not depend on the number of branch points, by adding more and more branch points to our cover we can build sequences of surfaces with growing genus but with the same area-weighted asymptotic number of cylinders.

Here is an example of one such sequence. Fix a stratum component $\mathcal{H}' \subset \mathcal{H}(k_1, \dots, k_n)$, possibly with regular marked points. Let d be odd and fix some $d_1, \dots, d_n \in \mathbb{Z}/d\mathbb{Z}$ such that $\gcd(d_1, \dots, d_n) = 1$. By appending m 0's to the partition (i.e., marking m freely-varying regular points on the translation surfaces in \mathcal{H}') this picks out a stratum component $\mathcal{H} \subset \mathcal{H}(k_1, \dots, k_n, 0^m)$. Fix any local monodromy data about the m regular marked points to complete (d_1, \dots, d_n) to δ , and let $\mathcal{M}_{m,\delta}$ be the resulting locus of degree d cyclic branched covers of surfaces in \mathcal{H} . Both the dimension of $\mathcal{M}_{m,\delta}$ and the genus of the covering surfaces clearly grow as m increases.

Theorem D ensures that $\mathcal{M}_{m,\delta}$ is connected (and hence the ψ and b invariants are not defined) and our assumptions imply that $\gcd(\delta) = 1$. Thus Corollary B states that $c_{\text{area}}(\mathcal{M}_{m,\delta})$ is always the same for any choice of m and any choice of local monodromy data about the regular points.²

This should be compared with what is known about the area -Siegel–Veech constants of strata: any sequence of strata \mathcal{H}_n with genus tending to infinity has $c_{\text{area}}(\mathcal{H}_n)$ tending to $1/2$ [Agg19], but any such sequence of constants is not a constant sequence [EKZ14]. The most similar sequence of orbit closures of which we are aware lies in [Fou18], which considers sequences of strata of quadratic differentials with fixed genus but increasingly many simple poles. Taking the loci of canonical double covers over these strata yields a sequence of orbit closures with dimension and genus tending to infinity for which Fougerson bounds the sum of Lyapunov exponents and thus the area Siegel–Veech constant.

Growth of Lyapunov exponents. These results can also be used to bound the growth of the smallest Lyapunov exponent for the orbit closures $\mathcal{M}_{m,\delta}$. The Eskin–Kontsevich–Zorich formula [EKZ14, Theorem

²One could of course let d be even and relax the assumption on $\gcd(d_i)$: one just needs to choose branching data so that M and the ψ and b invariants agree.

1] asserts that if \mathcal{M} is an orbit closure in a stratum $\mathcal{H}(k_1, \dots, k_n)$ of translation surfaces of genus g , then the sum of the first g Lyapunov exponents of the Kontsevich–Zorich cocycle over \mathcal{M} is

$$1 + \lambda_2(\mathcal{M}) + \dots + \lambda_g(\mathcal{M}) = \frac{1}{12} \sum_{i=1}^n \frac{k_i(k_i + 2)}{(k_i + 1)} + \frac{\pi^2}{3} c_{\text{area}}(\mathcal{M}),$$

where $c_{\text{area}}(\mathcal{M})$ is the usual ($s = 1$) area–Siegel–Veech constant.

Since the combinatorial term is bounded by

$$\frac{1}{12} \sum_{i=1}^n \frac{k_i(k_i + 2)}{(k_i + 1)} \leq \frac{1}{12} \frac{3}{2} (2g - 2) = \frac{(g - 1)}{4},$$

and for each $2 \leq l \leq g$ the sum of Lyapunov exponents is bounded by

$$1 + (l - 1)\lambda_l(\mathcal{M}) \leq 1 + \lambda_2(\mathcal{M}) + \dots + \lambda_g(\mathcal{M}),$$

it follows from Theorem A that, if g_m is the genus of the translation surfaces in $\mathcal{M}_{m,\delta}$, then³ the smallest Lyapunov exponent of $\mathcal{M}_{m,\delta}$ satisfies

$$\limsup_{m \rightarrow \infty} \lambda_{g_m}(\mathcal{M}_{m,\delta}) \leq 1/4.$$

(This should be compared with the discussion in the beginning of [EKZ14, §2.4].)

Remark. There are heuristic reasons to believe that the inequality above is probably not sharp. Indeed, several numerical simulations by the first named author indicate that $\lambda_{g_m}(\mathcal{M}_{m,\delta})$ vanishes for all m sufficiently large. Furthermore, it is likely that one could rigorously establish this numerical prediction by studying the signatures (p, q) of the eigenspaces of the deck group $\mathbb{Z}/d\mathbb{Z}$ acting on the tangent space of $\mathcal{M}_{m,\delta}$: in fact, these signatures can often be computed by direct methods (cf. [McM13] and [AMY19]) or using Chevalley–Weil type formulas from representation theory (cf. [MYZ14]). Moreover, if the restriction of the Hodge inner product to an eigenspace has signature (p, q) with $p > q$, then that component contributes at least $p - q$ zero Lyapunov exponents (cf. [FMZ14]). Similarly, it is also possible that the inequality in Footnote 3 is not sharp.

Notable aspects of the proof. Basic algebraic topology implies that $\mathbb{Z}/d\mathbb{Z}$ cyclic covers of X branched over \mathcal{B} are classified by the cohomology group $H^1(X \setminus \mathcal{B}; \mathbb{Z}/d\mathbb{Z})$. Theorem D is therefore equivalent to understanding the set of orbits of the monodromy group of \mathcal{H} on $H^1(X \setminus \mathcal{B}; \mathbb{Z}/d\mathbb{Z})$.

The mapping class group-valued monodromy of a non-hyperelliptic stratum-component in genus $g \geq 5$ was studied by the second and fourth-named authors in a sequence of papers [Cal20, CS21a, CS23]. In the companion paper [CS21b], the monodromy action on $H^1(X \setminus \mathcal{B}; \mathbb{Z}/d\mathbb{Z})$ was determined, again for $g \geq 5$. In spite of this progress, the results of the present paper require substantial new developments. A first issue is that the results of [Cal20, CS21a, CS23, CS21b] require $g \geq 5$, but there are non-hyperelliptic components in genus three and four to be studied as well. To that end, in Theorem 6.1 we give a new proof of the monodromy calculation of [CS21b] valid for all non-hyperelliptic stratum components in $g \geq 3$.

The second and more substantial issue is that even knowing the monodromy action on $H^1(X \setminus \mathcal{B}; \mathbb{Z}/d\mathbb{Z})$, it remains to understand the set and size of orbits. This turns out to be a subtle problem with some surprising features, e.g., the presence of “split” orbits: in certain numerical regimes, certain branched covers which are equivalent under the action of the full mapping class group are not monodromy-equivalent; ultimately this is explained by the presence of *framings* on the translation surfaces. This is in one sense similar to Kontsevich–Zorich’s classification [KZ03] of connected components of strata, where certain strata are split into path components according to the Arf invariant of the framing, but in our case, *every* non-hyperelliptic

³More generally, if $h : \mathbb{N} \rightarrow \mathbb{N}$ is asymptotic to the identity (i.e., $h(N)/N \rightarrow 1$ as $N \rightarrow \infty$), one also has $\limsup_{m \rightarrow \infty} \lambda_{h(g_m)}(\mathcal{M}_{m,\delta}) \leq 1/4$.

stratum has exactly one split orbit (Proposition 8.1). In the non-hyperelliptic regime, a single mapping class group orbit of branched covers can split into two monodromy orbits, classified by the “ ψ invariant” defined in Section 7. In the hyperelliptic setting, split orbits are governed by the “ b -invariant” of Section 10.

This work of determining monodromy groups, classifying orbits, computing their cardinalities, and computing a second associated numerical invariant (see Section 3.5) occupies a large part of the paper, spanning Sections 5–10 as well as the two short appendices. This Part II of the paper can be largely read independently from the calculation of Siegel–Veech constants carried out in Part III. The reader only interested in Part III can safely take the formulas in Part II as black boxes.

With an understanding of the monodromy orbits and connected components of \mathcal{M}_δ in hand, we compute the area^s-Siegel–Veech constants in Part III. Our main tool is a method for computing Siegel–Veech constants of \mathcal{M}_δ by defining appropriate counting functions on the base stratum \mathcal{H} . This technique is well-established in the setting of (branched) torus covers [EMS03, CMZ18] and [EKZ14, §10], but to our knowledge this paper represents the first generalization to covers of higher genus surfaces. Strata of higher genus surfaces are far more complicated than moduli spaces of punctured tori, so we must consider the structure of cusps as in [EMZ03], making our arguments more involved. We believe that the perspectives and computations used here for addressing cyclic covers will be valuable to those interested in non-cyclic covers of arbitrary translation surfaces.

Organization. The paper is divided into three parts, accompanied by two short appendices. Part I (Sections 2–3) contains introductory material and sets notation. Part II (Sections 4–10) concerns the topological problem of determining monodromy groups and connected components of loci of branched covers, and Part III (Sections 11–14) carries out the computation of Siegel–Veech constants. Appendix A contains some standard results on the classification of orbits of vectors under the action of certain subgroups of the symplectic group over $\mathbb{Z}/d\mathbb{Z}$, and Appendix B provides an alternative proof of the classification of monodromy orbits of homology classes of cylinders. A detailed section-by-section outline is provided at the beginning of each part.

Acknowledgments. The authors are very grateful to Anton Zorich for numerous discussions and careful explanations of material related to Siegel–Veech constants as well as his continued interest throughout the entirety of the project. The project originated from computer experiments using code written by Charles Fougeron that was made publicly available [DRC⁺24]. The authors are grateful to Vincent Delecroix and Julian R uth for assistance with coding. The authors are also grateful to Jayadev Athreya, Vincent Delecroix, Alex Eskin, Michael Magee, Howard Masur and Martin M oller for helpful conversations and interest.

D.A. was partially supported by the National Science Foundation under Award Nos. DMS - 1738381, DMS - 1600360, the Simons Foundation: Collaboration Grant under Award No. 853471, and several PSC-CUNY Grants. A.C. was supported by NSF Award Nos. DMS-2005328 and DMS-2202703. C.M. was partially supported by ANR-23-CE40-0020 MOST. N.S. was supported by NSF Awards No. DMS-2153879 and DMS-2338485. M.S. was partially supported by Simons Foundation: Collaboration Grant under Award No. 318898.

Part I

Preliminaries and setup

Outline of Part I (Sections 2–3).

- In Section 2, we discuss relevant aspects of the theory of cyclic branched covers.
- In Section 3, we recall various kinds of *monodromy* in the context of strata, and define the numerical invariants $|\overline{\Gamma}_d \alpha|$ and $|(\overline{\Gamma}_d \alpha)_a|$ whose computation will occupy us in Part II.

2 Branched covers of translation surfaces

In this brief section we recall the basics of branched covers and set some notation for the rest of the paper.

2.1 Translation surfaces

A *translation surface* (X, ω) is a Riemann surface X of genus at least one carrying a holomorphic Abelian differential ω . Throughout this manuscript, ω will often be suppressed. We permit (X, ω) to have marked points. A *regular closed trajectory* is a closed trajectory that does not pass through any of the singularities or marked points of (X, ω) . A *cylinder* is a maximal homotopy class of regular closed trajectories relative to the singularities and marked points of (X, ω) . For example, a torus with two marked points typically decomposes into two cylinders of parallel closed trajectories. Cylinders will always be open, i.e., homeomorphic to the Cartesian product of a finite open interval with the circle.

There is a canonical area associated to a flat surface by

$$\text{area}(X, \omega) = \frac{i}{2} \int_X \omega \wedge \bar{\omega}.$$

2.2 Branched covers and branching vectors

For the translation surface (X, ω) , let $\mathcal{B} \subset X$ be a distinguished set of points *containing all of the zeros of ω* , possibly along with some ordinary marked points. We are interested in describing the set of $\mathbb{Z}/d\mathbb{Z}$ cyclic branched covers $p : \tilde{X} \rightarrow X$ branched only over points in \mathcal{B} . *A priori*, such a cover is classified by a homomorphism

$$\alpha : \pi_1(X \setminus \mathcal{B}) \rightarrow \mathbb{Z}/d\mathbb{Z},$$

with \tilde{X} connected if and only if α is surjective. Since the target is Abelian, it follows that α is equivalent to a class

$$\alpha \in H^1(X \setminus \mathcal{B}; \mathbb{Z}/d\mathbb{Z}) \cong \text{Hom}(H_1(X \setminus \mathcal{B}; \mathbb{Z}), \mathbb{Z}/d\mathbb{Z}).$$

We call α the *branching vector* of the branched covering.⁴

An element of $(\mathbb{Z}/d\mathbb{Z})^k$ (for $d \geq 0$) is said to be *primitive* if its components generate $\mathbb{Z}/d\mathbb{Z}$. Surjectivity of $\alpha : \pi_1(X \setminus \mathcal{B}) \rightarrow \mathbb{Z}/d\mathbb{Z}$ (and hence connectivity of the associated branched cover) is equivalent to the condition that $\alpha \in H^1(X \setminus \mathcal{B}; \mathbb{Z}/d\mathbb{Z}) \cong (\mathbb{Z}/d\mathbb{Z})^k$ be primitive. Accordingly, we write

$$H_{\text{prim}}^1(X \setminus \mathcal{B}; \mathbb{Z}/d\mathbb{Z}) \subset H^1(X \setminus \mathcal{B}; \mathbb{Z}/d\mathbb{Z})$$

for the subset of primitive elements.

In the course of our arguments, we will need to understand the *local* behavior of a branched cover at a branch point, as well as some associated numerical invariants.

Definition 2.1. Given $\mathcal{B} = \{p_1, \dots, p_n\} \subset X$, for $1 \leq i \leq n$, choose $\gamma_i \subset X$ to be an oriented loop encircling a small neighborhood of p_i . The *local monodromy of α at p_i* is the element

$$d_i = \alpha(\gamma_i) \in \mathbb{Z}/d\mathbb{Z}.$$

We define the quantity $\text{gcd}(d_1, \dots, d_n, d)$ to be the *monodromy factor* of the covering α , and we denote it by $\text{gcd}(\delta(\alpha))$, which we treat as a formal symbol for this quantity for the moment (for the definition of $\delta(\alpha)$, see the beginning of Section 4).

⁴In this manuscript, *coverings* appear in two distinct (but intertwined) settings: there is the *local* phenomenon of branched coverings of a fixed translation surface under discussion here, and the induced *global* covering of a stratum of translation surfaces consisting of the family of all branched covers. We will have occasion to consider “monodromy actions” of the fundamental group of the base on the fiber in each setting. In the interest of clarity, in the branched cover setting, we will speak as much as possible about the *branching vector*, but the reader should be aware of the potential for confusion.

2.3 Cylinders

The branching vector α determines how curves lift to the branched cover. In general, this behavior is quite intricate (see, for example, [CMZ18, §2.3, Eqn. (24)] which describes how cylinders lift for general torus covers with n generic marked points). In our case, the situation is dramatically simpler since the deck group is Abelian.

Proposition 2.2. *Let (X, ω) be a translation surface with distinguished set $\mathcal{B} \subset X$. Suppose that $C_1, \dots, C_m \subset X \setminus \mathcal{B}$ are curves such that $[C_1] = \dots = [C_m]$ in $H_1(X; \mathbb{Z})$. Let α be a branching vector with monodromy factor $\gcd(\delta(\alpha))$. If $\alpha(C_1) = a$, then for each C_i , there exists $b_i(\alpha) \in \mathbb{Z}/d\mathbb{Z}$ such that $\alpha(C_i)$ is $a + b_i(\alpha) \gcd(\delta(\alpha))$.*

When applying this Proposition, the curves C_i will be core curves of cylinders.

Proof. If $[C_i] = [C_1]$ as elements of $H_1(X; \mathbb{Z})$, then in the excision homology group $H_1(X \setminus \mathcal{B}; \mathbb{Z})$,

$$[C_i] - [C_1] \in \ker(H_1(X \setminus \mathcal{B}; \mathbb{Z}) \rightarrow H_1(X; \mathbb{Z})),$$

so that $[C_i] - [C_1]$ can be written as a sum

$$[C_i] - [C_1] = \sum_{p_k \in \mathcal{B}} m_k \gamma_k,$$

for $m_k \in \mathbb{Z}$, with γ_k a loop encircling p_k as in Definition 2.1. Thus, taking $d_k = \alpha(\gamma_k)$ as in Definition 2.1,

$$\alpha(C_i) = \alpha(C_1) + \sum_{p_k \in \mathcal{B}} m_k d_k,$$

and the difference $\sum_{p_k \in \mathcal{B}} m_k d_k$ is a multiple of $\gcd(\delta(\alpha)) \pmod{d}$ by definition. \square

For later use, we observe that (the proof of) Proposition 2.2 shows that if we fix local monodromy data around the punctures, then the monodromy of all of the C_i depends only on the monodromy of one of them.

Lemma 2.3. *Suppose α and β are branching vectors with $\delta(\alpha) = \delta(\beta)$ and let $C_1, \dots, C_m \subset X \setminus \mathcal{B}$ be a set of curves representing the same class in $H_1(X; \mathbb{Z})$. If $\alpha(C_1) = \beta(C_1)$, then $\alpha(C_i) = \beta(C_i)$ for all i .*

The following lemma follows from the basic principles of covering space theory.

Lemma 2.4. *Let $(\tilde{X}, \tilde{\omega}) \rightarrow (X, \omega)$ be a branched translation covering with branching vector α . Suppose that C_1, \dots, C_m are cylinders on (X, ω) each of width w that satisfy the assumptions of Proposition 2.2. Then C_i lifts to*

$$\sharp(C_i) = \gcd(\alpha(C_i), d)$$

cylinders on $(\tilde{X}, \tilde{\omega})$, each of width

$$\frac{d}{\sharp(C_i)} w = \frac{d}{\gcd(\alpha(C_i), d)} w.$$

3 Covers of strata and monodromy

We now shift our perspective to a *global* one, and consider *families* of branched covers of translation surfaces varying in a component of a stratum.

3.1 Strata

A stratum $\mathcal{H}(\kappa)$ of translation surfaces is a moduli space where the orders of the singularities are specified by κ , a partition of $2g - 2$. We permit κ to be a non-negative partition with finitely many zero entries corresponding to marked points at which ω does not have a singularity. Let $\mathcal{H}_{lab}(\kappa)$ denote the cover of $\mathcal{H}(\kappa)$ on which all of the singularities with orders specified by κ are labeled.

Strata are not necessarily connected; a complete classification of their components was given in [KZ03]. In summary, all strata are connected if they have genus two or an odd order zero and are not of the form $\mathcal{H}(g - 1, g - 1)$. The strata $\mathcal{H}(2g - 2)$ and $\mathcal{H}(g - 1, g - 1)$ always have a unique hyperelliptic component $\mathcal{H}^{hyp}(\kappa)$, in which every translation surface in them admits a hyperelliptic involution. For strata with all even order zeros, the surfaces admit a *spin structure* which can have even or odd parity (cf. the discussion of Sections 5 and 7). In this case, if $g \geq 4$, then $\mathcal{H}(\kappa)$ always has two components not consisting entirely of hyperelliptic differentials, classified by spin parity. In genus 3, the even spin component coincides with the hyperelliptic component, and in genus 2, all components are hyperelliptic. Throughout, we will use \mathcal{H} to denote a connected component of a stratum, suppressing the partition κ when it is understood.

We will also consider strata with extra marked (regular) points. Forgetting the regular points associates to any component of $\mathcal{H}(\kappa)$ a component of $\mathcal{H}(\kappa')$, where κ' is the result of deleting all of the 0's from κ . Because our analysis will depend on which component of $\mathcal{H}(\kappa')$ we are considering, we make the following convention:

Convention. We say that a stratum component \mathcal{H} is hyperelliptic if the result of forgetting all regular marked points is a hyperelliptic stratum component.

For all $A > 0$, there is a natural cross section $\mathcal{H}_A(\kappa)$ of $\mathcal{H}(\kappa)$ given by taking those translation surfaces with area A . In particular, the stratum of labeled unit area translation surfaces is denoted $\mathcal{H}_{1,lab}(\kappa)$, and those with area at most one is denoted $\mathcal{H}_{\leq 1,lab}(\kappa)$.

Each stratum is a quasiprojective variety that is finitely covered by a complex manifold, and there are natural “period coordinates” on such covers. Given a translation surface $(X, \omega) \in \mathcal{H}(\kappa)$, let \mathcal{B} denote the collection of all singularities and marked points of (X, ω) . *Period coordinates* locally model a stratum on $H^1(X, \mathcal{B}; \mathbb{C})$ by integrating a basis of $H_1(X, \mathcal{B}; \mathbb{Z})$ against the differential ω . Strata admit a natural measure known as the *Masur-Veech measure* ν , which is given by pulling back Lebesgue measure on $\mathbb{C}^{2g+n-1} \cong H^1(X, \mathcal{B}; \mathbb{C})$ under period coordinates. Coning off induces a finite measure ν_1 on $\mathcal{H}_1(\kappa)$ [Mas82, Vee82].

3.2 Monodromy representations

Let \mathcal{H} be a component of $\mathcal{H}_{lab}(\kappa)$. Taking (orbifold) π_1 , the natural map $\mathcal{H} \rightarrow \mathcal{M}_{g,n}$ induces a (*topological monodromy representation*)

$$\rho_{top} : \pi_1^{orb}(\mathcal{H}) \rightarrow \text{PMod}(X, \mathcal{B}),$$

where $\text{PMod}(X, \mathcal{B})$ denotes the mapping class group of the surface X relative to \mathcal{B} , i.e., the isotopy classes of orientation-preserving diffeomorphisms of X that fix \mathcal{B} pointwise (here $X \in \mathcal{H}$ is an arbitrary basepoint). The map ρ_{top} measures how a chosen marking of (X, \mathcal{B}) changes as (X, \mathcal{B}) is transported around loops in \mathcal{H} , compare [FM12, § 5.6.1]. Let

$$\text{im}(\rho_{top}) =: \Gamma \leq \text{PMod}(X, \mathcal{B})$$

be the *topological monodromy group* of \mathcal{H} . The topological monodromy group of a hyperelliptic component was classically understood to be the hyperelliptic mapping class group, and was recently determined for non-hyperelliptic components in genus $g \geq 5$ by the second and fourth named authors [CS23], where it is shown to be a “framed mapping class group.”

The topological monodromy group acts on various objects associated to a surface. In our setting, the object of interest is the set of connected $\mathbb{Z}/d\mathbb{Z}$ branched covers, i.e., the set $H_{prim}^1(X \setminus \mathcal{B}; \mathbb{Z}/d\mathbb{Z}) \subset H^1(X \setminus \mathcal{B}; \mathbb{Z}/d\mathbb{Z})$ discussed in the previous section. Accordingly, the monodromy group of primary interest in this manuscript will be the image

$$\bar{\Gamma}_d \subseteq \text{Aut}(H^1(X \setminus \mathcal{B}; \mathbb{Z}/d\mathbb{Z})).$$

We call $\bar{\Gamma}_d$ the *mod d homological monodromy group* of the stratum component \mathcal{H} . (We will omit the subscript in the case $d = 0$). Note that we will often be quite lax about reducing mod d , letting elements $A \in \text{Aut}(H^1(X \setminus \mathcal{B}; \mathbb{Z}))$ act on $H_1(X \setminus \mathcal{B}; \mathbb{Z}/d\mathbb{Z})$.

3.3 Homological monodromy: basic structure

Duality and intersection pairings. The relative intersection pairing

$$\langle \cdot, \cdot \rangle : H_1(X \setminus \mathcal{B}; \mathbb{Z}/d\mathbb{Z}) \otimes H_1(X, \mathcal{B}; \mathbb{Z}/d\mathbb{Z}) \rightarrow \mathbb{Z}/d\mathbb{Z}$$

furnishes a canonical isomorphism $H^1(X \setminus \mathcal{B}; \mathbb{Z}/d\mathbb{Z}) \cong H_1(X, \mathcal{B}; \mathbb{Z}/d\mathbb{Z})$, which intertwines the actions of $\text{PMod}(X, \mathcal{B})$ on both sides. Thus, we can think of $\bar{\Gamma}_d$ as acting on $H_1(X, \mathcal{B}; \mathbb{Z}/d\mathbb{Z})$, which will be better suited for our purposes. The reader should be aware that we will tacitly make use of this isomorphism to think of elements of $H_1(X, \mathcal{B}; \mathbb{Z}/d\mathbb{Z})$ as *cohomology* classes via the intersection pairing; we will abuse notation and write $\alpha \in H_1(X, \mathcal{B}; \mathbb{Z}/d\mathbb{Z})$ when it is convenient to do so.

Absolute and relative automorphism groups. As an abelian group,

$$H_1(X, \mathcal{B}; \mathbb{Z}/d\mathbb{Z}) \cong (\mathbb{Z}/d\mathbb{Z})^{2g+n-1},$$

so that $\bar{\Gamma}_d$ is a subgroup of $\text{GL}_{2g+n-1}(\mathbb{Z}/d\mathbb{Z})$. The underlying topological context leads $\bar{\Gamma}_d$ to preserve a great deal of additional structure. Ultimately, a precise determination of $\bar{\Gamma}_d$ will constitute one of the major results of the paper, appearing in Sections 6 and 9 in the non-hyperelliptic and hyperelliptic settings, respectively. Here, we begin the process of describing $\bar{\Gamma}_d$ by taking relative and absolute homology into account. We follow the discussion of [CS21b, Section 4.2].

Let (X, \mathcal{B}) be a closed surface of genus g endowed with a set $\mathcal{B} = \{p_1, \dots, p_n\} \subset X$ of distinguished points. For any $d \geq 0$, the long exact sequence of this pair specializes to the short exact sequence

$$0 \rightarrow H_1(X; \mathbb{Z}/d\mathbb{Z}) \rightarrow H_1(X, \mathcal{B}; \mathbb{Z}/d\mathbb{Z}) \rightarrow \tilde{H}_0(\mathcal{B}; \mathbb{Z}/d\mathbb{Z}) \rightarrow 0.$$

The action of the pure mapping class group $\text{PMod}(X, \mathcal{B})$ on $H_1(X, \mathcal{B}; \mathbb{Z}/d\mathbb{Z})$ preserves the subspace $H_1(X; \mathbb{Z}/d\mathbb{Z})$ and the restriction preserves the intersection pairing $\langle \cdot, \cdot \rangle$. Define

$$\text{PRelAut}(H_1(X, \mathcal{B}; \mathbb{Z}/d\mathbb{Z})) := \text{Hom}(\tilde{H}_0(\mathcal{B}; \mathbb{Z}/d\mathbb{Z}), H_1(X; \mathbb{Z}/d\mathbb{Z})).$$

The group $\text{PAut}(H_1(X, \mathcal{B}; \mathbb{Z}/d\mathbb{Z}))$ is then characterized by the following short exact sequence:

$$1 \rightarrow \text{PRelAut}(H_1(X, \mathcal{B}; \mathbb{Z}/d\mathbb{Z})) \rightarrow \text{PAut}(H_1(X, \mathcal{B}; \mathbb{Z}/d\mathbb{Z})) \rightarrow \text{Sp}(2g, \mathbb{Z}/d\mathbb{Z}) \rightarrow 1, \quad (1)$$

where, here and throughout, $\text{Sp}(2g, \mathbb{Z}/d\mathbb{Z})$ denotes the *symplectic group* of automorphisms of $H_1(X; \mathbb{Z}/d\mathbb{Z})$ preserving the intersection pairing $\langle \cdot, \cdot \rangle$. By the above discussion,

$$\bar{\Gamma}_d \leq \text{PAut}(H_1(X, \mathcal{B}; \mathbb{Z}/d\mathbb{Z})).$$

In the sequel, we will frequently fix a splitting

$$H_1(X, \mathcal{B}; \mathbb{Z}/d\mathbb{Z}) \cong H_1(X; \mathbb{Z}/d\mathbb{Z}) \oplus \tilde{H}_0(\mathcal{B}; \mathbb{Z}/d\mathbb{Z}),$$

e.g. via a “geometric splitting” discussed in Section 5.2. In such coordinates, elements of $\text{PAut}(H_1(X, \mathcal{B}; \mathbb{Z}/d\mathbb{Z}))$ are upper-triangular, and will be written in the form

$$\begin{pmatrix} M & A \\ 0 & I \end{pmatrix},$$

with $M \in \text{Sp}(2g, \mathbb{Z}/d\mathbb{Z})$ and $A \in \text{Hom}(\tilde{H}_0(\mathcal{B}; \mathbb{Z}/d\mathbb{Z}), H_1(X; \mathbb{Z}/d\mathbb{Z}))$; the lower-right block is the identity by the hypothesis that \mathcal{B} is marked pointwise.

3.4 Orbits of branching vectors

In general, if $p : (\tilde{X}, \tilde{\omega}) \rightarrow (X, \omega)$ is a branched covering of translation surfaces specified by the branching vector α , then p induces a covering map

$$\mathcal{P} : \mathcal{M}_\delta(\alpha) \rightarrow \mathcal{H}_{1,lab},$$

where we recall from the introduction that \mathcal{M}_δ is the set of all translation covers of surfaces in $\mathcal{H}_{1,lab}$ with local monodromy given by δ , and we define $\mathcal{M}_\delta(\alpha)$ to be the component of \mathcal{M}_δ containing $(\tilde{X}, \tilde{\omega})$. The locus $\mathcal{M}_\delta(\alpha)$ can equivalently be defined by taking the $\mathrm{SL}_2(\mathbb{R})$ -orbit closure of $(\tilde{X}, \tilde{\omega})$ when (X, ω) is generic in $\mathcal{H}_{1,lab}$ with respect to the Masur–Veech measure.

The degrees of p and \mathcal{P} rarely coincide. From elementary topology we have the following:

Proposition 3.1. *Let $p : (\tilde{X}, \tilde{\omega}) \rightarrow (X, \omega)$ be a branched cyclic translation covering of degree d with branching vector α and set $\mathcal{P} : \mathcal{M}_\delta(\alpha) \rightarrow \mathcal{H}_{1,lab}$ to be the induced covering map of moduli spaces. Then $\deg(\mathcal{P}) = |\bar{\Gamma}_d \alpha|$.*

3.5 Orbits with fixed monodromy

We turn now to the second set of relevance. To compute Siegel–Veech constants, it is necessary to estimate the volume of the locus in $\mathcal{M}_\delta(\alpha)$ consisting of surfaces that have thin cylinders (see Section 11.2). In order to accomplish this, the fiber in $\mathcal{M}_\delta(\alpha)$ over the point $(X, \omega) \in \mathcal{H}_{1,lab}$ must be decomposed according to the sizes of the periods of the thin cylinders in each cover, which are necessarily multiples of the sizes of the periods of the corresponding cylinders on the base surface (X, ω) .

As we saw in Lemma 2.4 above, the period of a cylinder on $(\tilde{X}, \tilde{\omega})$ is determined by the branching vector classifying the cover $(\tilde{X}, \tilde{\omega}) \rightarrow (X, \omega)$. Thus, for a given cylinder $C \in (X, \omega)$ with core curve c , we are interested in the set

$$(\bar{\Gamma}_d \alpha)_{c;a} := \{\beta \in \bar{\Gamma}_d \alpha \mid \beta(c) = a \pmod{d}\}. \quad (2)$$

Clearly,

$$\bigcup_{a=1}^d (\bar{\Gamma}_d \alpha)_{c;a} = \bar{\Gamma}_d \alpha.$$

Ultimately, we will see that the cardinality of the set $(\bar{\Gamma}_d \alpha)_{c;a}$ is independent of the choice of cylinder C . We will see this in two ways in the non-hyperelliptic setting. Directly, the counting formulas Propositions 4.6 and 8.3 for $|(\bar{\Gamma}_d \alpha)_{c;a}|$ are visibly independent of c . In Appendix B, we present a conceptual argument based on the results Propositions B.1 and B.2 which establish the transitivity of the monodromy group action on excision homology classes of cylinders. While this second proof is not logically necessary in the non-hyperelliptic case (the transitivity having been taken care of via the action of the symplectic group), we believe that the added conceptual clarity, as well as the potential independent interest of Propositions B.1 and B.2, merit inclusion in an appendix. In keeping with this, in the sequel we will write

$$|(\bar{\Gamma}_d \alpha)_a|$$

in place of the more precise $|(\bar{\Gamma}_d \alpha)_{c;a}|$.

Remark. A related concept called the “cusp width” of a locus of torus covers appears in both [EKZ14, §10] and [CMZ18, Proof of Theorem 3.1]. In our setting, this corresponds to the order of a , and in computing Siegel–Veech constants we will group covers by a because they will contribute the same integrals. See (24).

Part II

Structure and cardinality of orbits

Outline of Part II (Sections 4–10). The purpose of this Part is to classify and enumerate orbits of branched covers under the monodromy group of a stratum, and to understand the lifting properties of cylinders along these covers. Topologically, this corresponds to the problem of classifying the components of \mathcal{M}_δ . The ultimate objective is to determine the cardinalities of the sets $\bar{\Gamma}_d\alpha$ and $(\bar{\Gamma}_d\alpha)_{c;a}$ defined above in Sections 3.4 and 3.5. This will be accomplished as follows:

- In Section 4, we begin the project of determining $|\bar{\Gamma}_d\alpha|$ and $|(\bar{\Gamma}_d\alpha)_a|$ by formulating the notion of *split* and *unsplit* orbits. Counting orbits in the unsplit case follows from relatively general principles, and we carry out those computations here.

The remaining work is to determine which orbits are (un)split, and to compute the quantities $|\bar{\Gamma}_d\alpha|$ and $|(\bar{\Gamma}_d\alpha)_a|$ in the split case. We proceed as follows:

- In Section 5, we discuss the theory of framings, winding number functions, and quadratic forms. These notions will be used in our description of $\bar{\Gamma}_d$ in the non-hyperelliptic setting in Section 6.
- In Section 6, we determine the monodromy group $\bar{\Gamma}_d$ in the case of a non-hyperelliptic stratum component.
- In Section 7, we define the ψ *invariant* used to classify split orbits in the non-hyperelliptic setting.
- In Section 8, we determine $|\bar{\Gamma}_d\alpha|$ and $|(\bar{\Gamma}_d\alpha)_a|$ for split orbits in the non-hyperelliptic setting.
- In Section 9, we shift our attention to hyperelliptic components, and give a description of the monodromy group $\bar{\Gamma}_d$ in this setting.
- In Section 10, we determine $|\bar{\Gamma}_d\alpha|$ and $|(\bar{\Gamma}_d\alpha)_a|$ for split orbits in the hyperelliptic setting.

In addition to these sections in the main body of the paper, there are two short appendices related to the material in Part II. In Appendix A, we state and sketch the proofs of some lemmas about the symplectic group over $\mathbb{Z}/d\mathbb{Z}$. In Appendix B, we present a conceptual explanation of the phenomenon that the cardinality of $(\bar{\Gamma}_d\alpha)_{c;a}$ is independent of c , as discussed above in Section 3.5.

4 Orbit counts (I): the unsplit case

We now begin the project of determining the cardinalities of the orbits $\bar{\Gamma}_d\alpha$ and their subsets $(\bar{\Gamma}_d\alpha)_{c;a}$ as defined in (2). To formulate the results, we recall that there are two numerical invariants associated to the elements of the distinguished set \mathcal{B} . The first of these is the integers $k_i \geq 0$, the orders of vanishing of the translation surface at $p_i \in \mathcal{B}$. The second is $d_i \in \mathbb{Z}/d\mathbb{Z}$, the local monodromy of the branching at p_i . Recall that this can be computed as $\alpha(\gamma_i) \in \mathbb{Z}/d\mathbb{Z}$, where $\alpha \in H_1(X, \mathcal{B}; \mathbb{Z}/d\mathbb{Z})$ classifies the branched cover and $\gamma_i \subset X \setminus \mathcal{B}$ is a small loop encircling p_i once counterclockwise (as the covering group $\mathbb{Z}/d\mathbb{Z}$ is abelian, this is well-defined independent of a global basepoint). The vector of the elements d_i can also be expressed as $\delta(\alpha) \in \tilde{H}_0(\mathcal{B}; \mathbb{Z}/d\mathbb{Z})$, where δ is the connecting map for the long exact sequence in homology. This explains the terminology $\gcd(\delta(\alpha))$ introduced in Definition 2.1.

4.1 Statement of results

Here we consider the simplest type of orbit: *unsplit orbits*, where the action of $\bar{\Gamma}_d$ is determined by arithmetic data alone. This notion is sensible in both the hyperelliptic and non-hyperelliptic settings, and in fact can be defined for an *arbitrary* subgroup $\Gamma \leq \text{PAut}(H_1(X, \mathcal{B}; \mathbb{Z}/d\mathbb{Z}))$, not just the monodromy group of a stratum. **Accordingly, in this section, we will work in this more general context, except where otherwise specified.**

Definition 4.1 (Unsplit orbit). Let $\alpha \in H_1^{\text{prim}}(X, \mathcal{B}; \mathbb{Z}/d\mathbb{Z})$, and let $\Gamma \leq \text{PAut}(H_1(X, \mathcal{B}; \mathbb{Z}/d\mathbb{Z}))$ be arbitrary. The orbit $\Gamma\alpha$ is *unsplit* if Γ acts transitively on vectors of given δ -value, i.e. if

$$\Gamma\alpha = \{\beta \in H_1^{\text{prim}}(X, \mathcal{B}; \mathbb{Z}/d\mathbb{Z}) \mid \delta(\beta) = \delta(\alpha)\}.$$

We say α is unsplit if its orbit is.

Let $\text{PAut}(H_1(X, \mathcal{B}; \mathbb{Z}/d\mathbb{Z}))[2]$ denote the “level-2 subgroup”, i.e. the kernel of the reduction map

$$\text{PAut}(H_1(X, \mathcal{B}; \mathbb{Z}/d\mathbb{Z})) \rightarrow \text{PAut}(H_1(X, \mathcal{B}; \mathbb{Z}/2\mathbb{Z}))$$

(where the reduction map is trivial if d is odd). In the sequel (Corollary 6.2, Proposition 9.4), we will see that in both the non-hyperelliptic and hyperelliptic settings, the monodromy group $\Gamma = \bar{\Gamma}_d$ contains the level-2 subgroup. Here, we see that for any Γ containing the level-2 subgroup, the question of whether an orbit splits is determined exclusively by mod-2 data.

Proposition 4.2 (Orbits are determined mod 2). *Let $\alpha, \beta \in H_1^{\text{prim}}(X, \mathcal{B}; \mathbb{Z}/d\mathbb{Z})$ be given, and suppose that $\delta(\alpha) = \delta(\beta)$. Let $\Gamma \leq \text{PAut}(H_1(X, \mathcal{B}; \mathbb{Z}/d\mathbb{Z}))$ be a subgroup that contains the level-2 subgroup $\text{PAut}(H_1(X, \mathcal{B}; \mathbb{Z}/d\mathbb{Z}))[2]$. Then $\Gamma\alpha = \Gamma\beta$ if and only if $\bar{\Gamma}_2\bar{\alpha} = \bar{\Gamma}_2\bar{\beta}$, where $\bar{\alpha}, \bar{\beta}$ are the reductions mod 2 of α, β , and $\bar{\Gamma}_2$ denotes the image of Γ under the reduction map*

$$\text{PAut}(H_1(X, \mathcal{B}; \mathbb{Z}/d\mathbb{Z})) \rightarrow \text{PAut}(H_1(X, \mathcal{B}; \mathbb{Z}/2\mathbb{Z})).$$

In particular, this is trivially satisfied if d is odd, and so in this case, every orbit is unsplit.

A key component of Proposition 4.2 is the assertion that the orbits obey a version of the Chinese remainder theorem. We extract this assertion and record it here as a corollary. It will be essential in the sequel, where we will work one prime at a time.

Corollary 4.3 (CRT for orbits). *Let $\alpha \in H_1(X, \mathcal{B}; \mathbb{Z}/d\mathbb{Z})$ be given, and let $\Gamma \leq \text{PAut}(H_1(X, \mathcal{B}; \mathbb{Z}/d\mathbb{Z}))$ be a subgroup containing $\text{PAut}(H_1(X, \mathcal{B}; \mathbb{Z}/d\mathbb{Z}))[2]$. Then under the decomposition*

$$H_1(X, \mathcal{B}; \mathbb{Z}/d\mathbb{Z}) \cong \bigoplus_{p_i \mid d} H_1(X, \mathcal{B}; \mathbb{Z}/p_i^{k_i}\mathbb{Z})$$

of the Chinese remainder theorem, there is an isomorphism

$$\Gamma\alpha \cong \prod_{p_i \mid d} \bar{\Gamma}_{p_i^{k_i}} \bar{\alpha},$$

where $\bar{\alpha}$ denotes the reduction of α mod $p_i^{k_i}$, and likewise $\bar{\Gamma}_{p_i^{k_i}}$ is the reduction of Γ mod $p_i^{k_i}$.

In the numerical formulas to be obtained below, the *Jordan totient function* will play a key role.

Definition 4.4 ($\Phi_n(d)$). The *Jordan totient function* $\Phi_n(d)$ is the number of primitive ($\text{gcd} = 1$) elements of $(\mathbb{Z}/d\mathbb{Z})^n$. It is computed as

$$\Phi_n(d) = d^n \prod_{p \mid d} \left(1 - \frac{1}{p^n}\right).$$

Note that the case $n = 1$ is the classical Euler totient function; in this case, we will write Φ in place of Φ_1 .

Proposition 4.5 (Orbit size, unsplit case). *Let p be prime, and let $\alpha \in H_1^{\text{prim}}(X, \mathcal{B}; \mathbb{Z}/p^k\mathbb{Z})$ be unsplit for the action of $\Gamma \leq \text{PAut}(H_1(X, \mathcal{B}; \mathbb{Z}/p^k\mathbb{Z}))$. Then*

$$|\Gamma\alpha| = \begin{cases} p^{2gk} & \gcd(\delta(\alpha)) = 1 \pmod{p} \\ \Phi_{2g}(p^k) & \gcd(\delta(\alpha)) \neq 1 \pmod{p}. \end{cases}$$

In the case of monodromy groups of translation surfaces, we will be interested in a finer-scale count than just the size of the orbit.

Proposition 4.6 (Orbit statistics, unsplit case). *Let (X, ω) be a translation surface in an arbitrary component of a stratum in genus $g \geq 1$. Let p be prime, and let $\alpha \in H_1^{\text{prim}}(X, \mathcal{B}; \mathbb{Z}/p^k\mathbb{Z})$ be unsplit for the action of its monodromy group $\bar{\Gamma}_{p^k} \leq \text{PAut}(H_1(X, \mathcal{B}; \mathbb{Z}/p^k\mathbb{Z}))$. Let $c \in H_1^{\text{prim}}(X \setminus \mathcal{B}; \mathbb{Z}/p^k\mathbb{Z})$ be the homology class of a cylinder on X , and let $a \in \mathbb{Z}/p^k\mathbb{Z}$ be given. Then*

$$|(\bar{\Gamma}_{p^k}\alpha)_{c;a}| = \begin{cases} p^{k(2g-1)} & \gcd(\delta(\alpha)) = 1 \text{ or } a \neq 0 \pmod{p} \\ \Phi_{2g-1}(p^k) & \gcd(\delta(\alpha)) \neq 1 \text{ and } a = 0 \pmod{p}. \end{cases}$$

Note that this formula is independent of c , as was asserted in Section 3.5.

4.2 Proofs

4.2.1 Orbit classification

Proof of Proposition 4.2. We first establish a version of the Chinese remainder theorem in our setting. Under the decomposition

$$H_1(X, \mathcal{B}; \mathbb{Z}/d\mathbb{Z}) \cong \bigoplus_{p_i|d} H_1(X, \mathcal{B}; \mathbb{Z}/p_i^{k_i}\mathbb{Z}) \quad (3)$$

of the Chinese remainder theorem, we claim there is an isomorphism

$$\Gamma\alpha \cong \prod_{p_i|d} \bar{\Gamma}_{p_i^{k_i}}\bar{\alpha}, \quad (4)$$

where $\bar{\alpha}$ denotes the reduction of $\alpha \pmod{p_i^{k_i}}$. By reducing modulo each $p_i^{k_i}$, there is an injective map from the left-hand side of (4) to the right; it remains to see that this is surjective. Under the isomorphism (3), a given $(\bar{\beta}_{p_1}, \dots, \bar{\beta}_{p_n}) \in \prod \bar{\Gamma}_{p_i^{k_i}}\bar{\alpha}$ corresponds to some $\beta \in H_1(X, \mathcal{B}; \mathbb{Z}/d\mathbb{Z})$; we seek to show that $\beta \in \Gamma\alpha$.

By hypothesis, Γ contains the level-2 subgroup $\text{PAut}(H_1(X, \mathcal{B}; \mathbb{Z}/d\mathbb{Z}))[2]$. Thus, under the Chinese remainder isomorphism $\mathbb{Z}/d\mathbb{Z} \cong \prod_{p_i|d} \mathbb{Z}/p_i^{k_i}\mathbb{Z}$, there is a containment

$$\text{PAut}(H_1(X, \mathcal{B}; \mathbb{Z}/2^k\mathbb{Z}))[2] \times \prod_{p_i|d \text{ odd}} \text{PAut}(H_1(X, \mathcal{B}; \mathbb{Z}/p_i^{k_i}\mathbb{Z})) \leq \Gamma. \quad (5)$$

In particular, this shows that the reduction of $\alpha \pmod{p_i}$ can be altered via one of the factors in (5) while leaving the reductions $\pmod{p_j}$ for $p_i \neq p_j$ unchanged.

By elementary algebra, for odd p_i , the orbits of $\text{PAut}(H_1(X, \mathcal{B}; \mathbb{Z}/p_i^{k_i}\mathbb{Z}))$ acting on $H_1(X, \mathcal{B}; \mathbb{Z}/p_i^{k_i}\mathbb{Z})$ are classified by \gcd and the value $\delta(\alpha) \in \tilde{H}_0(\mathcal{B}; \mathbb{Z}/p_i^{k_i}\mathbb{Z})$. Using (5), we can send α via Γ to some β' that satisfies $\beta' = \beta \pmod{p_i^{k_i}}$ for all odd primes $p_i | d$.

To complete the argument, we must show that if $\bar{\Gamma}_2\bar{\alpha} = \bar{\Gamma}_2\bar{\beta}$, then $\bar{\Gamma}_{2^k}\alpha = \bar{\Gamma}_{2^k}\beta$. Define

$$\delta := \delta(\alpha) = \delta(\beta).$$

By hypothesis, there is $F \in \Gamma$ such that $F\alpha = \beta'$ with $\beta = \beta' \pmod{2}$, and with $\delta(\beta') = \delta$. Relative to a splitting of $H_1(X, \mathcal{B}; \mathbb{Z}/2^k\mathbb{Z})$, write $\beta = x + \delta$ and $\beta' = x' + \delta$.

The projection $H_1(X, \mathcal{B}; \mathbb{Z}/2^k\mathbb{Z}) \rightarrow H_1(X; \mathbb{Z}/2^k\mathbb{Z})$ associated to such a splitting induces a surjection

$$\mathrm{PAut}(H_1(X, \mathcal{B}; \mathbb{Z}/2^k\mathbb{Z}))[2] \rightarrow \mathrm{Sp}(2g, \mathbb{Z}/2^k\mathbb{Z})[2].$$

According to Proposition A.1.2, the group $\mathrm{Sp}(2g, \mathbb{Z}/2^k\mathbb{Z})[2]$ acts transitively on elements of $H_1(X; \mathbb{Z}/2^k\mathbb{Z})$ with the same gcd and mod-2 reduction. Thus, if $\mathrm{gcd}(x) = \mathrm{gcd}(x')$, there exists $M \in \mathrm{Sp}(2g, \mathbb{Z}/2^k\mathbb{Z})[2]$ such that $Mx = x'$. The element $\begin{pmatrix} M & 0 \\ 0 & I \end{pmatrix}$ is contained in $\mathrm{PAut}(H_1(X, \mathcal{B}; \mathbb{Z}/2^k\mathbb{Z}))[2]$ and takes β to β' , showing $\bar{\Gamma}_{2^k}\alpha = \bar{\Gamma}_{2^k}\beta' = \bar{\Gamma}_{2^k}\beta$ as desired.

In general, however, $\mathrm{gcd}(x) \neq \mathrm{gcd}(x')$ - writing

$$p = \mathrm{gcd}(x), \quad p' = \mathrm{gcd}(x'), \quad q = \mathrm{gcd}(\delta),$$

the hypothesis $\mathrm{gcd}(\alpha) = \mathrm{gcd}(\beta)$ only implies that

$$\mathrm{gcd}(p, q) = \mathrm{gcd}(p', q) = 1.$$

We will show that there exists $G \in \mathrm{PAut}(H_1(X, \mathcal{B}; \mathbb{Z}/2^k\mathbb{Z}))[2]$ such that $G\beta = x'' + \delta$ satisfies $\mathrm{gcd}(x'') = \mathrm{gcd}(p, 2)$ and such that $x'' = x \pmod{2}$. Note that $\mathrm{gcd}(p, 2) = \mathrm{gcd}(p', 2)$, so that applying this construction to both β and β' reduces to the case above.

Write $x = pv$ with $v \in H_1(X; \mathbb{Z}/2^k\mathbb{Z})$ primitive, and let $w \in H_1(X; \mathbb{Z}/2^k\mathbb{Z})$ be a primitive vector linearly independent from v . Let $A \in \mathrm{Hom}(\tilde{H}_0(\mathcal{B}; \mathbb{Z}/2^k\mathbb{Z}), H_1(X; \mathbb{Z}/2^k\mathbb{Z}))$ take the primitive vector $\delta/q \in \tilde{H}_0(\mathcal{B}; \mathbb{Z}/2^k\mathbb{Z})$ to w . Then the element $G = \begin{pmatrix} I & 2A \\ 0 & I \end{pmatrix}$ has trivial mod-2 reduction, and

$$G\beta = (x + 2qw) + \delta = (pv + 2qw) + \delta.$$

Thus $x'' = pv + 2qw$ satisfies $\mathrm{gcd}(x'') = \mathrm{gcd}(p, 2)$ (since $\mathrm{gcd}(p, 2q) = \mathrm{gcd}(p, 2)$ and v, w are linearly independent) and $x'' = x \pmod{2}$ as required. \square

4.2.2 Orbit sizes

Proof of Proposition 4.5. By Proposition 4.2, $\bar{\Gamma}_{p^k}\alpha$ consists of all $\beta \in H_1^{\mathrm{prim}}(X, \mathcal{B}; \mathbb{Z}/p^k\mathbb{Z})$ such that $\delta(\beta) = \delta(\alpha) := \delta$. In the coordinates induced by the geometric splitting, $\beta = y + \delta$, with y subject to the constraint that $\mathrm{gcd}(y + \delta) = 1$. If $\mathrm{gcd}(\delta) = 1$, then y is unconstrained, showing $|\bar{\Gamma}_{p^k}\alpha| = (p^k)^{2g}$ as claimed. Otherwise, if $\mathrm{gcd}(\delta) \neq 1$, then $y \in H_1(X; \mathbb{Z}/p^k\mathbb{Z})$ must be primitive, showing $|\bar{\Gamma}_{p^k}\alpha| = \Phi_{2g}(p^k)$. \square

4.2.3 Orbit statistics

Proof of Proposition 4.6. Since α is unsplit, $\bar{\Gamma}_{p^k}\alpha$ consists of all $\beta \in H_1^{\mathrm{prim}}(X, \mathcal{B}; \mathbb{Z}/p^k\mathbb{Z})$ such that $\delta(\beta) = \delta(\alpha) := \delta$. If $\mathrm{gcd}(\delta) = 1$, then every vector of the form $\beta = y + \delta$ is primitive, and so

$$\bar{\Gamma}_{p^k}\alpha = \{y + \delta \mid y \in H_1(X; \mathbb{Z}/p^k\mathbb{Z})\}.$$

As $c \in H_1(X \setminus \mathcal{B}; \mathbb{Z}/p^k\mathbb{Z})$ is the homology class of a cylinder, it is in particular the class of a simple closed curve, and so it is primitive. It follows that each $a \in \mathbb{Z}/p^k\mathbb{Z}$ is realized as the value $\beta(c)$ an equal number of times as $\beta = y + \delta$ ranges over all $y \in H_1(X; \mathbb{Z}/p^k\mathbb{Z})$. There are p^k distinct values and $(p^k)^{2g}$ elements of $H_1(X; \mathbb{Z}/p^k\mathbb{Z})$, showing

$$|(\bar{\Gamma}_{p^k}\alpha)_{c;a}| := |\{\beta \in \bar{\Gamma}_{p^k}\alpha \mid \beta(c) = a \pmod{p^k}\}| = p^{k(2g-1)}$$

as claimed.

We next consider the other case that $\mathrm{gcd}(\delta) \neq 1$. In this case, the gcd condition is not automatically satisfied, so that

$$\bar{\Gamma}_{p^k}\alpha = \{y + \delta \mid y \in H_1^{\mathrm{prim}}(X; \mathbb{Z}/p^k\mathbb{Z})\}.$$

To count the elements of the orbit assuming the value $a \in \mathbb{Z}/p^k\mathbb{Z}$, we observe that $y \in H_1(X; \mathbb{Z}/p^k\mathbb{Z})$ is primitive if and only if its projection to $H_1(X; \mathbb{Z}/p\mathbb{Z})$ is nonzero. Thus,

$$|(\overline{\Gamma}_{p^k}\boldsymbol{\alpha})_{c;a}| = (p^{k-1})^{(2g-1)} |\{y + \delta \mid y \neq 0 \in H_1(X; \mathbb{Z}/p\mathbb{Z}), \beta(c) = a \pmod{p}\}|.$$

If $a \neq 0 \pmod{p}$, then the condition $\beta(c) = a \pmod{p}$ ensures that $y \neq 0$. The set of such y is an affine hyperplane in $H_1(X; \mathbb{Z}/p\mathbb{Z})$, necessarily of cardinality p^{2g-1} . Putting these calculations together,

$$|(\overline{\Gamma}_{p^k}\boldsymbol{\alpha})_{c;a}| = (p^{k-1})^{(2g-1)} p^{2g-1} = p^{k(2g-1)}$$

as claimed.

If instead $a = 0 \pmod{p}$, the set of such *nonzero* vectors y corresponds to the set of nonzero vectors in the hyperplane determined by the condition $\beta(c) = 0$, a set of size $p^{2g-1} - 1$. In this case,

$$|(\overline{\Gamma}_{p^k}\boldsymbol{\alpha})_{c;a}| = (p^{k-1})^{(2g-1)} (p^{2g-1} - 1) = \Phi_{2g-1}(p^k)$$

as claimed. □

5 Framings, winding number functions, and quadratic forms

Basic principles. The object of central importance in Part II is the monodromy group $\overline{\Gamma}_d$ of a stratum component acting on $H_1(X, \mathcal{B}; \mathbb{Z}/d\mathbb{Z})$. It turns out that in the non-hyperelliptic setting, this admits a description as the kernel of a certain *crossed homomorphism* $\Theta_{q,\kappa}$. Recall that a *crossed homomorphism* $f : G \rightarrow A$ from a group G to a G -module A is a function satisfying

$$f(gh) = f(g) + g \cdot f(h).$$

At root, the constraints on the action of $\overline{\Gamma}_d$ on $H_1(X, \mathcal{B}; \mathbb{Z}/d\mathbb{Z})$ all stem from *topological* constraints imposed by the fact that translation surfaces are endowed with canonical horizontal and vertical foliations. The invariance of such structures imposes subtle homological constraints on $\overline{\Gamma}_d$, as expressed by $\Theta_{q,\kappa}$. The goal of this section is to carefully describe these constraints.

5.1 Framings, winding number functions, and quadratic forms

Recall that a *framing* of a smooth manifold X is a trivialization of the tangent bundle TX ; equivalently, a framing is specified by $n = \dim(X)$ everywhere linearly-independent vector fields. Lemma 5.1 (which is obvious) recalls how framings emerge in the context of translation surfaces.

Lemma 5.1. *Let (X, ω) be a translation surface and let $\mathcal{B} \subset X$ be a set of distinguished points containing all zeros of ω . Then $\operatorname{Re}(\omega), \operatorname{Im}(\omega)$ form a pair of everywhere linearly independent 1-forms on $X \setminus \mathcal{B}$. Endowing X with a Riemannian metric, these dually determine a “translation surface framing” of $X \setminus \mathcal{B}$.*

It turns out that *isotopy classes* of framings are classified by a certain discrete invariant known as a *winding number function*.

Definition 5.2 (Winding number function). Let Σ be a surface equipped with a framing ϕ . Let $\mathcal{S}^+(\Sigma)$ denote the set of isotopy classes of oriented simple closed curves on Σ . The framing ϕ determines a *winding number function* which by abuse of notation we write

$$\phi : \mathcal{S}^+(\Sigma) \rightarrow \mathbb{Z},$$

with $\phi(c) \in \mathbb{Z}$ equal to the winding number of the forward-pointing tangent vector of c relative to the framing.

It should be clear that isotopic framings determine the same winding number function. See [CS23, Section 2] or [RW14] for a more thorough discussion of the relationship between these two notions, including the sense in which the converse holds.

In [CS23, Lemma 7.10], it is shown that the presence of the translation surface framing imposes constraints on the mapping class group-valued monodromy of a stratum component. However, it is not *a priori* clear what kinds of constraints the framing imposes on the *homological* monodromy of interest here. This is because the associated winding number function ϕ is *not homological*, in the sense that $\phi(a)$ for a simple closed curve $a \subset \Sigma$ is not determined by $[a] \in H_1(\Sigma; \mathbb{Z})$. It turns out that the reduction mod 2 of ϕ is homological, but defining this correctly takes some care. To begin, recall the notion of a $\mathbb{Z}/2\mathbb{Z}$ -quadratic form.

Definition 5.3 ($\mathbb{Z}/2\mathbb{Z}$ -quadratic form). Let $V = (\mathbb{Z}/2\mathbb{Z})^{2g}$ be a vector space over $\mathbb{Z}/2\mathbb{Z}$ equipped with a symplectic form $\langle \cdot, \cdot \rangle$, i.e., a nondegenerate bilinear form satisfying $\langle x, x \rangle = 0$ for all $x \in V$. A $\mathbb{Z}/2\mathbb{Z}$ -quadratic form refining $\langle \cdot, \cdot \rangle$ is a function $q : V \rightarrow \mathbb{Z}/2\mathbb{Z}$ that satisfies

$$q(x + y) = q(x) + q(y) + \langle x, y \rangle$$

for all $x, y \in V$.

Lemma 5.4 gives a first indication of the homological character of ϕ .

Lemma 5.4. *Let S be a framed surface. Suppose that either S has a single boundary component, or else $S = X \setminus \mathcal{B}$ is obtained from a translation surface for which every zero has even order. Then the function*

$$q : H_1(S; \mathbb{Z}/2\mathbb{Z}) \rightarrow \mathbb{Z}/2\mathbb{Z}$$

given by

$$q(x) = \phi(\xi) + 1 \pmod{2},$$

where $\xi \subset S$ is a simple closed curve with $[\xi] = x$, is a $\mathbb{Z}/2\mathbb{Z}$ -quadratic form refining the intersection pairing. Moreover, q descends to a $\mathbb{Z}/2\mathbb{Z}$ -quadratic form on $H_1(\bar{S}; \mathbb{Z}/2\mathbb{Z})$, where \bar{S} is the closed surface obtained from S by filling in all boundary components and punctures.

Proof. See [CS21b, Remark 4.3]. □

It turns out that even when the translation surface (X, ω) has zeros of odd order, the mod-2 winding number of a simple closed curve is still homological (determined in general by the class in *excision* homology $H_1(X \setminus \mathcal{B}; \mathbb{Z}/2\mathbb{Z})$, see Lemma 5.7). However, to formulate this precisely, we require the notion of a *geometric splitting*, discussed in the next subsection.

5.2 Geometric splittings

Consider the pair (X, \mathcal{B}) of topological spaces. For any coefficient ring R , the associated long exact sequence in homology degenerates to the short exact sequence

$$0 \rightarrow H_1(X; R) \rightarrow H_1(X, \mathcal{B}; R) \xrightarrow{\delta} \tilde{H}_0(\mathcal{B}; R) \rightarrow 0.$$

This gives a *noncanonical* decomposition of $H_1(X, \mathcal{B}; \mathbb{Z}/d\mathbb{Z})$ into its “relative” and “absolute” parts. The purpose of a “geometric splitting” is to give a geometric meaning to this decomposition.

Definition 5.5 (Geometric splitting). A *geometric splitting* of the pair (X, \mathcal{B}) is a choice of subsurface $X^\circ \subset X \setminus \mathcal{B}$, such that the complement is a disk that contains all the points of \mathcal{B} . For any coefficient ring R , a geometric splitting induces a splitting

$$H_1(X, \mathcal{B}; R) \cong H_1(X; R) \oplus \tilde{H}_0(\mathcal{B}; R)$$

via the condition that classes in $\tilde{H}_0(\mathcal{B}; R)$ be supported on $X \setminus X^\circ$; then also $H_1(X; R)$ is canonically identified with $H_1(X^\circ; R)$.

Definition 5.6 (Associated data of a geometric splitting). Let (X, ω) be a translation surface in the stratum $\mathcal{H}(\kappa)$, with $\kappa = (k_1, \dots, k_n)$; let ϕ be the associated framing of $X \setminus \mathcal{B}$. Let $X^\circ \subset X$ be a geometric splitting. The *associated data* of X° is the pair (q, κ) , where q is the quadratic form on $H_1(X^\circ; \mathbb{Z}/2\mathbb{Z})$ induced from the restriction of ϕ to X° in the sense of Lemma 5.4, and

$$\kappa = \sum_{p_i \in \mathcal{B}} k_i p_i \in \tilde{H}_0(\mathcal{B}; \mathbb{Z}/2\mathbb{Z}) \subset H_1(X, \mathcal{B}; \mathbb{Z}/2\mathbb{Z});$$

note that $\sum k_i = 2g - 2 = 0 \pmod{2}$ and so the expression defining κ is sensible. Note also the slight abuse of notation, conflating κ the ordered partition with a divisor.

The lemma below shows how the mod 2 winding number of any simple closed curve can be computed from the associated data of any geometric splitting, thereby refining and extending Lemma 5.4 (which is the case where $\kappa = 0$).

Lemma 5.7. *Let $a \subset X \setminus \mathcal{B}$ be a simple closed curve, and let $X^\circ \subset X \setminus \mathcal{B}$ be a geometric splitting with associated data (q, κ) . Then*

$$\phi(a) = q(\bar{a}) + 1 + \langle a, \kappa \rangle \pmod{2},$$

where we view \bar{a} in the first term on the right as the induced class in $H_1(X; \mathbb{Z}/2\mathbb{Z}) \cong H_1(X^\circ; \mathbb{Z}/2\mathbb{Z})$ (with the isomorphism furnished by the geometric splitting), and in the expression $\langle a, \kappa \rangle$, we view a as an element of $H_1(X \setminus \mathcal{B}; \mathbb{Z}/2\mathbb{Z})$.

Proof. Let $\Delta = X \setminus X^\circ$ be the disk associated to the geometric splitting. Choose a collection of embedded arcs on Δ representing κ , and then choose a representative for a meeting $\partial\Delta$ and κ transversely. To compute $q(\bar{a})$, we must determine a representative simple closed curve that is homologous to a in $H_1(X; \mathbb{Z}/2\mathbb{Z})$ and supported on X° .

We do so by systematically altering a . By construction, $a \cap \Delta$ is a finite collection of disjoint arcs. Choose an *outermost* such arc a_i , in the sense that one component U of $\Delta \setminus a_i$ contains no other arcs. Isotope a (inside X) by pulling a_i across the “empty” region U . The result is a simple closed curve a' that is homologous to a on $H_1(X; \mathbb{Z}/2\mathbb{Z})$, but which will in general have different winding number. By the Poincaré–Hopf theorem, the difference $\phi(a') - \phi(a)$ is equal to the total order of the zeros associated to the points of \mathcal{B} contained in U . Thus $\phi(a') = \phi(a) \pmod{2}$ if and only if the total multiplicity of zeros in U is even. By the definition of κ , the arc a_i intersects κ an even number of times if and only if this multiplicity is even. We conclude that

$$\phi(a') + \langle a', \kappa \rangle = \phi(a) + \langle a, \kappa \rangle \pmod{2}.$$

After performing this procedure on all of the arcs of $a \cap \Delta$, we obtain a curve a'' supported on X° for which $\phi(a'') + \langle a'', \kappa \rangle = \phi(a) + \langle a, \kappa \rangle \pmod{2}$. By the definition of $q(\bar{a})$, and since $\langle a'', \kappa \rangle = 0$ by construction,

$$q(\bar{a}) = q(\bar{a}) + \langle a'', \kappa \rangle = \phi(a'') + 1 + \langle a'', \kappa \rangle = \phi(a) + 1 + \langle a, \kappa \rangle \pmod{2}$$

as was to be shown. □

5.3 The crossed homomorphism $\Theta_{q, \kappa}$

Here we define the crossed homomorphism $\Theta_{q, \kappa}$. Our initial definition requires a choice of geometric splitting, and so we show in Lemma 5.9 that $\Theta_{q, \kappa}$ coincides with a different crossed homomorphism Θ_ϕ introduced in [CS21b]. The latter does not depend on a geometric splitting, implying that the former is independent of this choice as well. It will not be necessary to give a full discussion of the original crossed homomorphism Θ_ϕ ; for the sake of concision, we will refer to [CS21b] at the few places this is required.

Definition 5.8 (The crossed homomorphism $\Theta_{q, \kappa}$). Let

$$H_1(X, \mathcal{B}; \mathbb{Z}/2\mathbb{Z}) \cong H_1(X; \mathbb{Z}/2\mathbb{Z}) \oplus \tilde{H}_0(\mathcal{B}; \mathbb{Z}/2\mathbb{Z})$$

be a splitting of $H_1(X, \mathcal{B}; \mathbb{Z}/2\mathbb{Z})$, and let q be a quadratic form on $H_1(X; \mathbb{Z}/2\mathbb{Z})$ and $\kappa \in \widetilde{H}_0(\mathcal{B}; \mathbb{Z}/2\mathbb{Z})$ be given. Define a function

$$\Theta_{q, \kappa} : \text{PAut}(H_1(X, \mathcal{B}; \mathbb{Z}/2\mathbb{Z})) \rightarrow H^1(X; \mathbb{Z}/2\mathbb{Z})$$

via

$$\Theta_{q, \kappa} \left(\begin{pmatrix} M & A \\ 0 & I \end{pmatrix} \right) (x) = q(Mx) + q(x) + \langle Mx, A\kappa \rangle.$$

A routine verification shows that $\Theta_{q, \kappa}$ is a crossed homomorphism under the obvious action of $\text{PAut}(H_1(X, \mathcal{B}; \mathbb{Z}/2\mathbb{Z}))$ on $H^1(X; \mathbb{Z}/2\mathbb{Z})$.

Below, we show the equivalence of $\Theta_{q, \kappa}$ with the crossed homomorphism Θ_ϕ studied in [CS21b], from which the invariance under the choice of geometric splitting will follow. Unless the reader is interested in carefully reconciling the results of [CS21b] with those given here, the proof of Lemma 5.9 can be safely skipped.

Lemma 5.9. *Let X° define a geometric splitting of (X, \mathcal{B}) with associated data (q, κ) . Then there is an equality*

$$\Theta_\phi = \Theta_{q, \kappa}.$$

In particular, if (q, κ) and (q', κ') are the associated data for geometric splittings X° and X'° of X , it follows that

$$\Theta_{q, \kappa} = \Theta_{q', \kappa'}.$$

Proof. The group $\text{PAut}(H_1(X, \mathcal{B}; \mathbb{Z}/2\mathbb{Z}))$ is a quotient of $\text{Mod}(X, \mathcal{B})$ via the latter's action on relative homology. As $\text{Mod}(X, \mathcal{B})$ is generated by Dehn twists, it suffices to verify that $\Theta_\phi(T_a) = \Theta_{q, \kappa}(T_a)$ for any Dehn twist T_a .

Let $x \in H_1(X; \mathbb{Z}/2\mathbb{Z})$ be given, and let $c \subset X \setminus \mathcal{B}$ be a representative simple closed curve. Then, following [CS21b, Section 4], [CS21b, Lemma 2.4], and Lemma 5.7 for the first, second and fourth equalities, respectively,

$$\begin{aligned} \Theta_\phi(T_a)(x) &= \phi(T_a(x)) - \phi(x) \\ &= \phi(x) + \langle a, x \rangle \phi(a) - \phi(x) \\ &= \langle a, x \rangle \phi(a) \\ &= \langle a, x \rangle (q(a) + 1 + \langle a, \kappa \rangle); \end{aligned}$$

note that we abuse notation and use the symbol a to denote a curve as well as the associated homology classes over \mathbb{Z} and $\mathbb{Z}/2\mathbb{Z}$.

In order to compute $\Theta_{q, \kappa}(T_a)$, we observe that the action on homology of T_a is given by the transvection formula

$$T_a(x) = x + \langle a, x \rangle a.$$

Relative to the splitting

$$H_1(X, \mathcal{B}; \mathbb{Z}/2\mathbb{Z}) \cong H_1(X; \mathbb{Z}/2\mathbb{Z}) \oplus \widetilde{H}_0(\mathcal{B}; \mathbb{Z}/2\mathbb{Z}),$$

we can write

$$T_a = \begin{pmatrix} M & A \\ 0 & I \end{pmatrix},$$

and we see that if $x \in H_1(X; \mathbb{Z}/2\mathbb{Z})$, then $Mx = x + \langle a, x \rangle a$, but

$$A\kappa = T_a(\kappa) - \kappa = \langle a, \kappa \rangle a.$$

Then (recalling we are working mod 2, so that $\langle a, x \rangle^2 = \langle a, x \rangle$),

$$\begin{aligned} \Theta_{q, \kappa}(T_a)(x) &= q(x + \langle a, x \rangle a) + q(x) + \langle x + \langle a, x \rangle a, \langle a, \kappa \rangle a \rangle \\ &= \langle a, x \rangle q(a) + \langle a, x \rangle^2 + \langle a, \kappa \rangle \langle x, a \rangle \\ &= \langle a, x \rangle (q(\bar{a}) + 1 + \langle a, \kappa \rangle) \\ &= \Theta_\phi(T_a)(x). \end{aligned}$$

Since the definition of Θ_ϕ (cf. [CS21b, Section 4]) does not involve a choice of geometric splitting, the final claim follows. \square

6 The homological monodromy group, non-hyperelliptic case

Here we obtain the first major result of Part II – Theorem 6.1, which determines the monodromy action of a non-hyperelliptic stratum component on relative homology. [CS21b, Theorem B] obtains the same result in the range $g \geq 5$ via different methods; here we give a new and independent proof valid in the maximal range $g \geq 3$. The methods of the proof of Theorem 6.1 will not recur in the rest of the paper, and the casual reader can safely take Theorem 6.1 as a black box.

Theorem 6.1. *Let $\mathcal{H} \subset \mathcal{H}_{lab}(\kappa)$ be a non-hyperelliptic component of a stratum of translation surfaces in genus $g \geq 3$. Choose a basepoint $(X, \omega) \in \mathcal{H}$, and let $\mathcal{B} \subset X$ be the full set of distinguished points. Then the homological monodromy group $\bar{\Gamma} \leq \text{PAut}(H_1(X, \mathcal{B}; \mathbb{Z}))$ is computed to be*

$$\bar{\Gamma} = \ker(\Theta_{q, \kappa}),$$

where $\Theta_{q, \kappa}$ is the crossed homomorphism of Definition 5.8.

Recall that $\text{PAut}(H_1(X, \mathcal{B}; \mathbb{Z}/d\mathbb{Z}))[2]$ denotes the kernel of the reduction map

$$\text{PAut}(H_1(X, \mathcal{B}; \mathbb{Z}/d\mathbb{Z})) \rightarrow \text{PAut}(H_1(X, \mathcal{B}; \mathbb{Z}/2\mathbb{Z}))$$

(where the reduction map is trivial if d is odd). As an immediate corollary, we obtain the following result which will be used in the next section.

Corollary 6.2. *For any non-hyperelliptic stratum component and any integer d , there is a containment*

$$\text{PAut}(H_1(X, \mathcal{B}; \mathbb{Z}/d\mathbb{Z}))[2] \leq \bar{\Gamma}_d.$$

Proof of Theorem 6.1. In [CS21b, proof of Proposition 5.2], it is shown that $\bar{\Gamma} \leq \ker(\Theta_\phi) = \ker(\Theta_{q, \kappa})$; this holds in the maximal range $g \geq 3$. To establish the opposite containment, we recall the decomposition

$$1 \rightarrow \text{PRelAut}(H_1(X, \mathcal{B}; \mathbb{Z})) \rightarrow \text{PAut}(H_1(X, \mathcal{B}; \mathbb{Z})) \rightarrow \text{Sp}(2g, \mathbb{Z}) \rightarrow 1$$

of (1). We will show that the projections of $\bar{\Gamma}$ and $\ker(\Theta_{q, \kappa})$ to $\text{Sp}(2g, \mathbb{Z})$ coincide and that their intersections with the kernel $\text{PRelAut}(H_1(X, \mathcal{B}; \mathbb{Z}))$ are equal.

The monodromy action on *absolute* homology $H_1(X; \mathbb{Z})$ is determined by [GR19, Theorem 6.8]. To state the result, we define

$$r = \gcd(\kappa) = \gcd(k_1, \dots, k_n)$$

as the gcd of the zero orders in the stratum. In the case of r odd, the absolute monodromy is given as

$$\bar{\Gamma}^{abs} = \text{Sp}(2g, \mathbb{Z}), \tag{6}$$

and in the case of r even,

$$\bar{\Gamma}^{abs} = \text{Sp}(2g, \mathbb{Z})[q], \tag{7}$$

where q is the quadratic form on $H_1(X; \mathbb{Z}/2\mathbb{Z})$ induced from ϕ via Lemma 5.4, and $\text{Sp}(2g, \mathbb{Z})[q]$ is the preimage in $\text{Sp}(2g, \mathbb{Z})$ of the stabilizer $\text{Sp}(2g, \mathbb{Z}/2\mathbb{Z})[q]$ of q under the action of $\text{Sp}(2g, \mathbb{Z}/2\mathbb{Z})$ on the set of $\mathbb{Z}/2\mathbb{Z}$ -quadratic forms on $H_1(X; \mathbb{Z}/2\mathbb{Z})$. We must show that likewise, the projection of $\ker(\Theta_{q, \kappa})$ to $\text{Sp}(2g, \mathbb{Z})$ is $\text{Sp}(2g, \mathbb{Z})[q]$. But this is clear: in case r is even, the associated data κ is identically zero, so that in this case,

$$\Theta_{q, \kappa} \left(\begin{pmatrix} M & A \\ 0 & I \end{pmatrix} \right) (x) = q(Mx) + q(x),$$

showing that $\begin{pmatrix} M & A \\ 0 & I \end{pmatrix} \in \ker(\Theta_{q,\kappa})$ if and only if $M \in \mathrm{Sp}(2g, \mathbb{Z})[q]$.

It remains to consider the intersections of $\bar{\Gamma}$ and $\ker(\Theta_{q,\kappa})$ with $\mathrm{PRelAut}(H_1(X, \mathcal{B}; \mathbb{Z}))$. In coordinates, this latter group consists of matrices of the form $\begin{pmatrix} I & A \\ 0 & I \end{pmatrix}$, so that $\Theta_{q,\kappa}$ restricts to the (*uncrossed*) homomorphism

$$\begin{aligned} \kappa^* : \mathrm{PRelAut}(H_1(X, \mathcal{B}; \mathbb{Z})) &\rightarrow H^1(X; \mathbb{Z}/2\mathbb{Z}) \\ A &\mapsto (x \mapsto \langle x, A\kappa \rangle). \end{aligned}$$

Thus it suffices to show that there is a containment

$$\ker(\kappa^*) \leq \bar{\Gamma}.$$

This will be accomplished by exhibiting an explicit collection of elements of $\bar{\Gamma} \cap \ker(\kappa^*)$ that constitute a generating set for $\ker(\kappa^*)$. This computation will be greatly streamlined by the fact that $\ker(\kappa^*)$ carries an action of $\bar{\Gamma}^{abs}$: it suffices to exhibit generators for $\ker(\kappa^*)$ not merely as an abelian group, but as a $\bar{\Gamma}^{abs}$ -module.

It will be useful to have a more explicit description of $\mathrm{PRelAut}(H_1(X, \mathcal{B}; \mathbb{Z}))$. Having defined this above as

$$\mathrm{PRelAut}(H_1(X, \mathcal{B}; \mathbb{Z})) := \mathrm{Hom}(\tilde{H}_0(\mathcal{B}; \mathbb{Z}), H_1(X; \mathbb{Z})),$$

we observe that by choosing bases for $\tilde{H}_0(\mathcal{B}; \mathbb{Z}) \cong \mathbb{Z}^{n-1}$ and for $H_1(X; \mathbb{Z}) \cong \mathbb{Z}^{2g}$, there is an identification

$$\mathrm{PRelAut}(H_1(X, \mathcal{B}; \mathbb{Z})) \cong \mathrm{Mat}_{2g, n-1}(\mathbb{Z}). \quad (8)$$

We will not need to choose an explicit basis for $H_1(X; \mathbb{Z})$, but it will be convenient to do so for $\tilde{H}_0(\mathcal{B}; \mathbb{Z})$, and we take the basis

$$\{p_i - p_n \mid 1 \leq i \leq n-1\}. \quad (9)$$

Recall that by convention, the distinguished point $p_i \in \mathcal{B} \subset X$ is a zero of order $k_i \geq 0$, and that $\kappa \in \tilde{H}_0(\mathcal{B}; \mathbb{Z})$ is defined via

$$\kappa = \sum_{p_i \in \mathcal{B}} k_i p_i.$$

The lemma below gives a generating set for $\ker(\kappa^*)$ in terms of the numerical invariants k_i .

Lemma 6.3. *As a $\bar{\Gamma}^{abs}$ -module, $\ker(\kappa^*) \leq \mathrm{Mat}_{2g, n-1}(\mathbb{Z})$ is generated by the following collection of elements:*

1. *For each $1 \leq i \leq n-1$ such that k_i is even, some matrix $U_1(i) \in \mathrm{Mat}_{2g, n-1}(\mathbb{Z})$ with i^{th} column \mathbf{v} for some primitive vector \mathbf{v} , and all other columns identically zero,*
2. *For each $1 \leq i \leq n-1$ such that k_i is odd, some matrix $U_2(i) \in \mathrm{Mat}_{2g, n-1}(\mathbb{Z})$ with i^{th} column $2\mathbf{v}$ for some primitive vector \mathbf{v} , and all other columns identically zero,*
3. *For each pair of columns $1 \leq i < j \leq n-1$ such that k_i and k_j are both odd, some matrix $U_3(i, j) \in \mathrm{Mat}_{2g, n-1}(\mathbb{Z})$ with i^{th} and j^{th} columns \mathbf{v} for some primitive vector \mathbf{v} , and all other columns identically zero.*

Proof. We first establish that, as an abelian group, $\ker(\kappa^*)$ is generated by the collection of *all* elements of types 1, 2, 3, where \mathbf{v} ranges over the set of *all* primitive vectors. Following this, we will show that, as a $\bar{\Gamma}^{abs}$ -module, only one element of each type is necessary.

Recall the definition of $\kappa^* : \mathrm{PRelAut}(H_1(X, \mathcal{B}; \mathbb{Z})) \rightarrow H^1(X; \mathbb{Z}/2\mathbb{Z})$:

$$\kappa^*(A)(x) = \langle x, A\kappa \rangle \pmod{2}.$$

Concretely, viewing A as an element of $\text{Mat}_{2g, n-1}(\mathbb{Z})$, one can realize $\kappa^*(A)$ as the column vector in $(\mathbb{Z}/2\mathbb{Z})^{2g}$ obtained by summing the columns of A for which the corresponding k_i is odd. It is clear from this that $\ker(\kappa^*)$ is generated as an abelian group by the set of all elements of types 1, 2, 3, with \mathbf{v} ranging over the set of all primitive vectors in $H_1(X; \mathbb{Z})$.

We turn now to the second step, bringing the action of $\bar{\Gamma}^{abs}$ into consideration. An element $M \in \bar{\Gamma}^{abs} \leq \text{Sp}(2g, \mathbb{Z})$ acts on $A \in \ker(\kappa^*) \leq \text{Mat}_{2g, n-1}(\mathbb{Z})$ by left multiplication. In the case of r odd, $\bar{\Gamma}^{abs} = \text{Sp}(2g, \mathbb{Z})$ by (6), which clearly acts transitively on primitive vectors, proving the claim in this case.

In the case of r even, $\bar{\Gamma}^{abs} = \text{Sp}(2g, \mathbb{Z})[q]$ by (7). According to Proposition A.1.1, there are exactly two orbits of primitive vectors, classified by the value $q(\mathbf{v}) \in \mathbb{Z}/2\mathbb{Z}$. Thus the $\bar{\Gamma}^{abs}$ -orbit of a primitive vector \mathbf{v} consists of all primitive vectors \mathbf{w} with $q(\mathbf{w}) = q(\mathbf{v})$. Since \mathcal{H} is non-hyperelliptic, $g \geq 3$ and so by Lemma A.2, it is possible to find nonzero elements $\mathbf{v}_1, \mathbf{v}_2 \in H_1(X; \mathbb{Z})$ with any desired combination of values $q(\mathbf{v}_1), q(\mathbf{v}_2), \langle \mathbf{v}_1, \mathbf{v}_2 \rangle$. Thus, given \mathbf{v} , let \mathbf{w} satisfy $q(\mathbf{w}) = q(\mathbf{v})$ and $\langle \mathbf{v}, \mathbf{w} \rangle = q(\mathbf{v}) + 1$. The mod-2 reduction of $\mathbf{v} + \mathbf{w}$ is necessarily nonzero, and so if $\mathbf{v} + \mathbf{w} \in H_1(X; \mathbb{Z})$ is not primitive, its gcd must be odd. Altering the entries of \mathbf{w} by even values, we may therefore moreover select \mathbf{w} such that $\mathbf{v} + \mathbf{w}$ is primitive. Computing $q(\mathbf{v} + \mathbf{w})$ yields

$$q(\mathbf{v} + \mathbf{w}) = q(\mathbf{v}) + q(\mathbf{w}) + \langle \mathbf{v}, \mathbf{w} \rangle = q(\mathbf{v}) + 1.$$

It follows that the $\mathbb{Z}[\bar{\Gamma}^{abs}]$ -span of \mathbf{v} contains primitive vectors with both values of q , so that in the r even case as well, only one generator of type 1 is required (generators of types 2, 3 of course only exist in the r odd case). \square

Having established Lemma 6.3, we continue with the proof of Theorem 6.1. The outline of the remainder of the argument is as follows. Following Lemma 6.3, it suffices to exhibit loops in \mathcal{H} that represent the three types of generators. This will be accomplished by exhibiting translation surfaces with prescribed patterns of cylinders; combinations of shears about these cylinders will be used to realize the desired elements of $\bar{\Gamma}$.

Before proceeding to the exhibition of these explicit elements, we explain the general method. In order to compute the monodromy action of our loops on $H_1(X, \mathcal{B}; \mathbb{Z})$, we first consider the more refined *topological* monodromy, valued in the pure mapping class group $\text{PMod}(X, \mathcal{B})$ (recall that this is the group of isotopy classes of orientation-preserving diffeomorphisms of X that pointwise fix the set \mathcal{B} of marked points). There is a natural map

$$\Psi^{rel} : \text{PMod}(X, \mathcal{B}) \rightarrow \text{PAut}(H_1(X, \mathcal{B}; \mathbb{Z}))$$

that records the action of a mapping class on $H_1(X, \mathcal{B}; \mathbb{Z})$, and a coarsening

$$\Psi : \text{PMod}(X, \mathcal{B}) \rightarrow \text{Sp}(2g, \mathbb{Z})$$

that tracks only the action on $H_1(X; \mathbb{Z})$.

When a mapping class $f \in \text{PMod}(X, \mathcal{B})$ satisfies $\Psi(f) = I$, the refinement $\Psi^{rel}(f)$ is valued in the subgroup $\text{PRelAut}(H_1(X, \mathcal{B}; \mathbb{Z})) \leq \text{PAut}(H_1(X, \mathcal{B}; \mathbb{Z}))$. To attach an explicit matrix $A \in \text{Mat}_{2g, n-1}(\mathbb{Z})$ to this, it is necessary to choose bases both for $H_1(X; \mathbb{Z})$ and for $\tilde{H}_0(\mathcal{B}; \mathbb{Z})$, and further to lift these to a system of $2g$ simple closed curves, resp. $n - 1$ arcs connecting the points of \mathcal{B} . In light of Lemma 6.3, it will not be necessary to specify explicit coordinates on $H_1(X; \mathbb{Z})$.

If α, α' are two distinct arcs connecting points $p_i, p_j \in \mathcal{B}$, then $\alpha - \alpha'$ is a simple closed curve. For $f \in \ker(\Psi)$, it follows that there is an equality of relative homology classes $f_*([\alpha]) - [\alpha] = f_*([\alpha']) - [\alpha']$, and so, having fixed a basis for $H_1(X; \mathbb{Z})$, the matrix $\Psi^{rel}(f) \in \text{Mat}_{2g, n-1}(\mathbb{Z})$ depends only on the choice of basis for $\tilde{H}_0(\mathcal{B}; \mathbb{Z})$, which we have fixed in (9) to be of the form $\{p_i - p_n\}$.

It is important to note that the loops we exhibit will be based at *different* basepoints in \mathcal{H} , leading to the problem of apparently having to “standardize” by transporting everything to a single fixed basepoint. In spite of this apparent difficulty, the flexibility afforded by recalling the $\bar{\Gamma}^{abs}$ -module structure on $\ker(\kappa^*)$ will render this a non-issue. Indeed, the gcd of the entries in a fixed column of some $A \in \text{Mat}_{2g, n-1}(\mathbb{Z})$ is invariant under the action of $\text{Sp}(2g, \mathbb{Z})$. If a loop $\gamma \subset \mathcal{H}$ based at some point $X \in \mathcal{H}$ acts on $H_1(X, \mathcal{B}; \mathbb{Z})$ via some $A \in \text{Mat}_{2g, n-1}(\mathbb{Z})$, a conjugate $\eta\gamma\eta^{-1}$ acts as MA , where $M = \Psi^{abs}(\eta)$ and $\eta \in \pi_1(\mathcal{H})$. Thus, the gcd

of column entries is invariant under conjugation in $\pi_1(\mathcal{H})$, and hence under change of basepoint. According to Lemma 6.3, the generators of type 1, 2 and 3 are characterized entirely in terms of the gcd of the entries of each column, and hence we are free to change basepoint as suits our needs.

Consider first a generator of type 1 as in Lemma 6.3. This is an element of $\text{PRelAut}(H_1(X, \mathcal{B}; \mathbb{Z})) \cong \text{Mat}_{2g, n-1}(\mathbb{Z})$ with a primitive vector \mathbf{v} in some column i such that k_i is even, and all other columns identically zero. Consider Figure 1. The left side of the figure shows a translation surface constructed as




Figure 1: Exhibiting a generator of type 1.

follows: suppose first that $k_i = 2k$ for some $k \geq 1$ (the exceptional case $k_i = 0$ will be addressed at the end of the paragraph), and let κ' be the partition of $2(g-k) - 2$ obtained by deleting k_i . Let Y be a translation surface in $\mathcal{H}_{lab}(\kappa')$, and let Z be a translation surface in $\mathcal{H}_{lab}(2k-2)$. Let C be a cylinder on Y , and let $X \in \mathcal{H}_{lab}(\kappa)$ be constructed as indicated in Figure 1, by joining Y to Z by cutting Y along the core of C and attaching to Z via the slit construction. In the case where $r = \text{gcd}(\kappa)$ is even, by choosing the components of $\mathcal{H}_{lab}(\kappa')$ and $\mathcal{H}_{lab}(2k-2)$ so that the induced Arf invariant on the glued surface matches with the Arf invariant of \mathcal{H} , it is possible to construct such an X in \mathcal{H} . In the exceptional case $\kappa_i = 0$, we simply proceed by taking $X = Y$ with an extra marked point, “splitting” the cylinder C in two - the reasoning encoded in Figure 1 remains valid.

Any such X is endowed with two cylinders C_1, C_2 which bound a subsurface, as indicated on the right side of Figure 1. Letting S_C denote the (right-handed) shear about the cylinder C , we claim that the loop $S_{C_1} S_{C_2}^{-1}$ acts as a generator of type 1 as in Lemma 6.3. To see this, we first note that the core curves c_1, c_2 of C_1, C_2 are homologous as elements of $H_1(X; \mathbb{Z})$. The action of S_{C_i} on $H_1(X; \mathbb{Z})$ is given by the symplectic transvection T_{c_i} which depends only on $[c_i] \in H_1(X; \mathbb{Z})$, and so $S_{C_1} S_{C_2}^{-1}$ acts trivially on $H_1(X; \mathbb{Z})$. It thus acts on $H_1(X, \mathcal{B}; \mathbb{Z})$ as some $A \in \text{Mat}_{2g, n-1}(\mathbb{Z})$. Following the above discussion, we choose a system of arcs α_i connecting each $p_j (1 \leq j \leq n-1)$ to p_n , choosing α_i to cross c_1 once as shown in Figure 1.

We see from this that

$$[S_{C_1} S_{C_2}^{-1}(\alpha_i) - \alpha_i] = [c_1];$$

this is a primitive element of $H_1(X; \mathbb{Z})$. Any other α_j crosses c_1 and c_2 an equal number of times, and thus the homology class is unchanged under $S_{C_1} S_{C_2}^{-1}$. We conclude that $S_{C_1} S_{C_2}^{-1}$ acts on $H_1(X, \mathcal{B}; \mathbb{Z})$ as a generator of type 1 as claimed.

We next consider a generator of type 2 as in Lemma 6.3. This is an element of $\text{PRelAut}(H_1(X, \mathcal{B}; \mathbb{Z})) \cong \text{Mat}_{2g, n-1}(\mathbb{Z})$ with *twice* a primitive vector $2\mathbf{v}$ in some column i such that k_i is odd, and all other columns identically zero.

We begin with the special case where $\kappa' = (k_1, k_2)$ with k_1 (and hence also k_2) odd. The left side of Figure 2 shows a square-tiled surface X in the stratum $\mathcal{H}_{lab}(\kappa')$. Note the four cylinders labeled C_1, \dots, C_4 with associated core curves c_1, \dots, c_4 . The *chain relation* asserts that the product of shears

$$W = (S_{C_1} S_{C_2} S_{C_3})^4$$

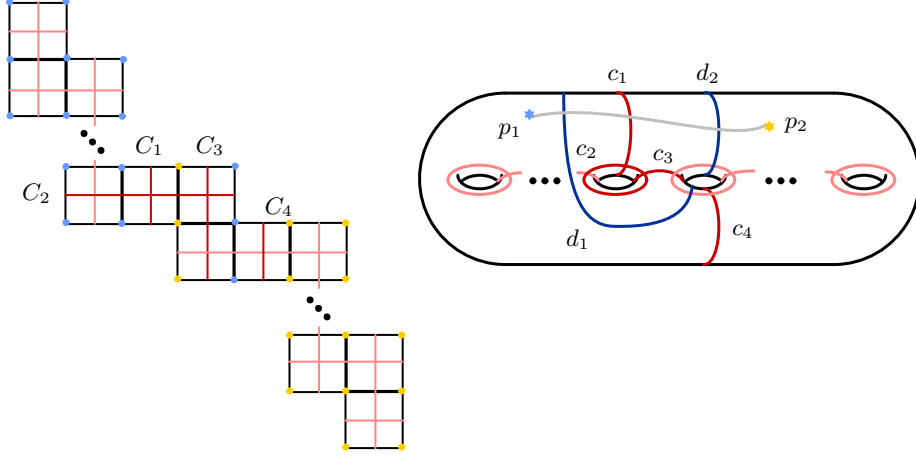


Figure 2: Exhibiting a generator of type 2.

acts as a mapping class via

$$w = (T_{c_1} T_{c_2} T_{c_3})^4 = T_{d_1} T_{d_2},$$

where d_1, d_2 are the curves indicated on the right side of Figure 2. Note that there is an equality

$$[d_1] = [d_2] = [c_4]$$

in $H_1(X; \mathbb{Z})$, and so the loop

$$\gamma = W S_{C_4}^{-2}$$

acts trivially on $H_1(X; \mathbb{Z})$. As in the first case, we can compute the action of γ on $H_1(X, \mathcal{B}; \mathbb{Z})$ by choosing an arc connecting p_1 to $p_2 = p_n$. Doing so shows that the loop acts via the column matrix $A \in \text{Mat}_{2g,1}(\mathbb{Z})$ given by $A = 2[c_4]$.

This exhibits a generator of type 2 in the stratum $\mathcal{H}_{lab}(k_1, k_2)$ with k_1, k_2 both odd. In the general case, we can simply take a small perturbation X' of X which splits the zero p_2 of order k_2 to realize a translation surface with the same configuration of cylinders in any stratum $\mathcal{H}_{lab}(\kappa)$ for $\kappa = \{k_1, \dots\}$ a refinement of κ' . The loop γ exhibiting the generator of type 2 can be transported to a product of shears based at X' . It is clear that one can choose arcs connecting p_n to p_i for $i \geq 2$ that are left invariant by the resulting mapping class, so that the associated $A \in \text{Mat}_{2g,n}(\mathbb{Z})$ has exactly one nonzero column, completing the argument in the general case.

It remains to exhibit a generator of type 3 as in Lemma 6.3. This is an element $A \in \text{Mat}_{2g,n-1}(\mathbb{Z})$ with a primitive vector \mathbf{v} in the i^{th} and j^{th} column, with k_i, k_j both odd, and identically zero elsewhere. We let κ' be the partition obtained from κ by deleting k_i, k_j .




Figure 3: Exhibiting a generator of type 3.

Consider Figure 3. The top left portion of the figure depicts two translation surfaces: $Y \in \mathcal{H}_{lab}(k_i, k_j)$ with k_i, k_j both odd on the left, and $Z \in \mathcal{H}_{lab}(\kappa')$ on the right (we assume here that κ' partitions some *positive* integer; we will deal with the other possibility at the end of the paragraph). As shown, Y and Z have cylinders C_Y and C_Z with the same period. The bottom left of the figure shows the result of cutting Y and Z along these cylinders and re-attaching to make a translation surface $X \in \mathcal{H}_{lab}(\kappa)$ equipped with two cylinders C_1, C_2 with homologous core curves $[c_1] = [c_2] = \mathbf{v}$. The resulting surface X is depicted topologically on the right side of Figure 3. In the exceptional case $\kappa' = (0, \dots, 0)$, we proceed as in Case 1, constructing X from Y by splitting the cylinder C_Y by adding the required number of marked points.

As C_1, C_2 are homologous, the product of shears $S_{C_1} S_{C_2}^{-1}$ acts trivially on $H_1(X; \mathbb{Z})$. Reasoning as in the prior two cases, one computes that $S_{C_1} S_{C_2}^{-1}$ acts on $H_1(X, \mathcal{B}; \mathbb{Z})$ via the matrix $A \in \text{Mat}_{2g, n-1}(\mathbb{Z})$ with the primitive element \mathbf{v} in the columns corresponding to p_i and p_j , and is identically zero elsewhere, as required. \square

7 The ψ invariant

Having computed the monodromy group in the non-hyperelliptic setting, we turn to the question of determining the structure of its orbits on the set of cyclic branched covers (and hence the classification of components of \mathcal{M}_δ). In this section, we define the ψ invariant and establish its basic properties. In the following Section 8, we then show that the ψ invariant classifies split orbits, and we determine the associated cardinalities.

The ψ invariant is closely related to the Arf invariant of a $\mathbb{Z}/2\mathbb{Z}$ -quadratic form. We begin by recalling this notion.

Definition 7.1 (Symplectic basis, Arf invariant). Let $(V, \langle \cdot, \cdot \rangle)$ be a symplectic vector space over a field K . A *symplectic basis* for V is a basis

$$B = \{x_1, y_1, \dots, x_g, y_g\}$$

for which $\langle x_i, y_i \rangle = 1$ and all other pairings among elements of B are zero.

Now specialize to $K = \mathbb{Z}/2\mathbb{Z}$, and let q be a $\mathbb{Z}/2\mathbb{Z}$ -quadratic form refining $\langle \cdot, \cdot \rangle$. The *Arf invariant* $\text{Arf}(q) \in \mathbb{Z}/2\mathbb{Z}$ is defined as

$$\text{Arf}(q) = \sum_{i=1}^g q(x_i)q(y_i),$$

where $B = \{x_i, y_i\}$ is any symplectic basis.

We recall the well-known facts that $\text{Arf}(q)$ is well-defined independently of the choice of B , and that $\text{Arf}(q)$ is the unique invariant of the action of $\text{Sp}(2g, \mathbb{Z}/2\mathbb{Z})$ on the set of $\mathbb{Z}/2\mathbb{Z}$ -quadratic forms [Arf41].

Definition 7.2 (Arf invariant of a framing). Let X be a translation surface for which every zero has even order, and let ϕ be the associated framing of X . Define

$$\text{Arf}(\phi) := \text{Arf}(q),$$

where q is the $\mathbb{Z}/2\mathbb{Z}$ -quadratic form associated to ϕ via Lemma 5.4.

Remark. In the setting of strata of translation surfaces, when the component \mathcal{H} supports a spin structure, the Arf invariant $\text{Arf}(\phi)$ of the translation surface framing coincides with the usual Arf invariant classifying the stratum component in the sense of the Kontsevich-Zorich classification.

We are now ready to define the ψ invariant, a $\mathbb{Z}/2\mathbb{Z}$ -valued invariant of the coset $\delta^{-1}(\kappa) \subset H_1(X, \mathcal{B}; \mathbb{Z}/2\mathbb{Z})$. *A priori*, it depends on a choice of geometric splitting; we show that it is well-defined in Lemma 7.6, which relies on the preliminary results Lemmas 7.4 and 7.5. Finally, in Lemma 7.7, we show that ψ is an invariant of the monodromy action.

Definition 7.3 (The ψ invariant). Fix a geometric splitting X° of X , giving rise to an isomorphism $H_1(X, \mathcal{B}; \mathbb{Z}/2\mathbb{Z}) \cong H_1(X; \mathbb{Z}/2\mathbb{Z}) \oplus \tilde{H}_0(\mathcal{B}; \mathbb{Z}/2\mathbb{Z})$ and associated data q, κ as in Definition 5.6. Given $\alpha = x + \kappa$, we define

$$\psi(\alpha) = \psi(x + \kappa) := q(x) + \text{Arf}(q) \in \mathbb{Z}/2\mathbb{Z}.$$

Recall that the set $Q(H_1(X; \mathbb{Z}/2\mathbb{Z}))$ of quadratic forms $q : H_1(X; \mathbb{Z}/2\mathbb{Z}) \rightarrow \mathbb{Z}/2\mathbb{Z}$ forms a *torsor* over $H^1(X; \mathbb{Z}/2\mathbb{Z})$. In particular, given $q \in Q(H_1(X; \mathbb{Z}/2\mathbb{Z}))$ and $\xi \in H^1(X; \mathbb{Z}/2\mathbb{Z})$, then $q + \xi \in Q(H_1(X; \mathbb{Z}/2\mathbb{Z}))$ is again a quadratic form.

Lemma 7.4. Let $\begin{pmatrix} M & A \\ 0 & I \end{pmatrix} \in \text{PAut}(H_1(X, \mathcal{B}; \mathbb{Z}/2\mathbb{Z}))$ be given, and define

$$\xi = \Theta_{q, \kappa} \left(\begin{pmatrix} M & A \\ 0 & I \end{pmatrix} \right) \in H^1(X; \mathbb{Z}/2\mathbb{Z}).$$

Then $q(A\kappa) = \text{Arf}(q) + \text{Arf}(q + \xi)$; in particular if $\begin{pmatrix} M & A \\ 0 & I \end{pmatrix} \in \ker(\Theta_{q, \kappa})$, then $q(A\kappa) = 0$.

Proof. If $A\kappa = 0$, then $\xi(x) = q(Mx) + q(x)$ and so $(q + \xi)(x) = q(Mx)$. As $M \in \text{Sp}(2g, \mathbb{Z})$, the forms $q(x)$ and $q(Mx)$ have the same Arf invariant, showing the result in this case.

If $A\kappa \neq 0$, then choose a symplectic basis $x_1, y_1, \dots, x_g, y_g$ for $H_1(X; \mathbb{Z}/2\mathbb{Z})$ for which $A\kappa = Mx_1$. As $M \in \text{Sp}(2g, \mathbb{Z}/2\mathbb{Z})$, it preserves the Arf invariant of q :

$$\sum_{i=1}^g q(Mx_i)q(My_i) = \sum_{i=1}^g q(x_i)q(y_i).$$

By the definition of ξ , the condition $A\kappa = Mx_1$ and the symplectic invariance of the intersection pairing,

$$q(Mx) = q(x) + \langle Mx, A\kappa \rangle + \xi(x) = (q + \xi)(x) + \langle x, x_1 \rangle.$$

It follows that

$$\sum_{i=1}^g q(x_i)q(y_i) = \sum_{i=1}^g q(Mx_i)q(My_i) = \sum_{i=1}^g ((q + \xi)(x_i) + \langle x_i, x_1 \rangle)((q + \xi)(y_i) + \langle y_i, x_1 \rangle).$$

In the above, every expression of the form $\langle z, x_1 \rangle$ is zero except $\langle y_1, x_1 \rangle = 1$. Comparing the left and right sides then shows that

$$\text{Arf}(q) = \text{Arf}(q + \xi) + (q + \xi)(x_1).$$

Expanding, recalling that $Mx_1 = A\kappa$,

$$(q + \xi)(x_1) = q(Mx_1) + \langle Mx_1, A\kappa \rangle = q(A\kappa),$$

which completes the proof. \square

Lemma 7.5. For $i \in \{1, 2\}$, let X_i° be geometric splittings with associated data (q_i, κ_i) . Let $f \in \text{Mod}(X, \mathcal{B})$ satisfy $f(X_1^\circ) = X_2^\circ$ and induce the identity on $H_1(X; \mathbb{Z}/2\mathbb{Z})$. Then the quadratic forms q_i on $H_1(X; \mathbb{Z}/2\mathbb{Z})$ are related by the formula

$$q_2(x) = q_1(x) + \Theta_{q_1, \kappa_1}(F)(x) = q_1(x) + \langle x, A\kappa_1 \rangle,$$

where $F = \begin{pmatrix} I & A \\ 0 & I \end{pmatrix} \in \text{PAut}(H_1(X, \mathcal{B}; \mathbb{Z}/2\mathbb{Z}))$ is the automorphism on homology associated to f .

Proof. For $x \in H_1(X; \mathbb{Z}/2\mathbb{Z})$, recall that $q_i(x) = \phi(c_i) + 1$, where c_i is a representative for x supported on X_i° . Given $f \in \text{Mod}(X, \mathcal{B})$ as in the statement of the lemma, we therefore have $q_2(x) = \phi(f(c_1)) + 1$. By Lemma 5.7,

$$q_2(x) = \phi(f(c_1)) + 1 = q_1(x) + \langle f(c_1), \kappa_1 \rangle.$$

It therefore suffices to show $\langle f(c_1), \kappa_1 \rangle = \langle x, A\kappa_1 \rangle$. To see this, we note that the pairing

$$\langle \cdot, \cdot \rangle : H_1(X \setminus \mathcal{B}; \mathbb{Z}/2\mathbb{Z}) \otimes H_1(X, \mathcal{B}; \mathbb{Z}/2\mathbb{Z}) \rightarrow \mathbb{Z}/2\mathbb{Z}$$

intertwines the action of $\text{PMod}(X, \mathcal{B})$ on the factors, so that

$$\langle f(c_1), \kappa_1 \rangle = \langle c_1, f^{-1}(\kappa_1) \rangle.$$

The action of f^{-1} on $H_1(X, \mathcal{B}; \mathbb{Z}/2\mathbb{Z})$ is by F^{-1} ; as we are working mod 2, it follows that $F^{-1} = F = \begin{pmatrix} I & A \\ 0 & I \end{pmatrix}$.

Thus $[f^{-1}(\kappa_1)] = A\kappa_1 + \kappa_1 \in H_1(X, \mathcal{B}; \mathbb{Z}/2\mathbb{Z})$. As c_1 is supported on X_1° and so is disjoint from the representative for κ_1 , it follows that $\langle c_1, A\kappa_1 + \kappa_1 \rangle = \langle c_1, A\kappa_1 \rangle$. This latter expression is now an intersection of *absolute* homology classes, and as $[c_1] = x \in H_1(X; \mathbb{Z}/2\mathbb{Z})$, the claim follows. \square

Lemma 7.6. *The value $\psi(\alpha) \in \mathbb{Z}/2\mathbb{Z}$ is well-defined independent of the choice of geometric splitting.*

Proof. Consider two geometric splittings X_i° for $i \in \{1, 2\}$, along with associated data (q_i, κ_i) . Represent α with $\delta(\alpha) = \kappa$ in the forms

$$\alpha = x_i + \kappa_i,$$

and choose representative simple closed curves $c_i \subset X_i^\circ$ for the classes x_i . Via Lemmas 5.7 and 7.5, and because $\langle x_2, \kappa_2 \rangle = 0$ by construction,

$$\phi(c_2) = q_2(x_2) + 1 = q_1(x_2) + 1 + \langle x_2, \kappa_1 \rangle.$$

Express $\kappa_2 = \kappa_1 + y$ for $y \in H_1(X; \mathbb{Z}/2\mathbb{Z})$. As $\alpha = x_1 + \kappa_1 = x_2 + \kappa_2$, then $x_2 = x_1 + y$, and so

$$\begin{aligned} q_1(x_2) + \langle x_2, \kappa_1 \rangle &= q_1(x_1 + y) + \langle x_2, \kappa_2 + y \rangle \\ &= q_1(x_1) + q_1(y) + \langle x_1, y \rangle + \langle x_1 + y, y \rangle \\ &= q_1(x_1) + q_1(y), \end{aligned}$$

since $\langle x_2, \kappa_2 \rangle = 0$. It therefore suffices to show that

$$q_1(y) = \text{Arf}(q_1) + \text{Arf}(q_2).$$

To see this, let $f \in \text{Mod}(X, \mathcal{B})$ and $F \in \text{PAut}(H_1(X, \mathcal{B}; \mathbb{Z}/2\mathbb{Z}))$ be as in Lemma 7.5. Since $f(X_1^\circ) = X_2^\circ$, it follows that $y = A\kappa_1$. Moreover, $q_2 = F \cdot q_1 = q_1 + \langle \cdot, A\kappa_1 \rangle$ by Lemma 7.5, and so

$$q_1(y) = q_1(A\kappa_1) = \text{Arf}(q_1) + \text{Arf}(q_1 + \Theta_{q_1, \kappa_1}(F)) = \text{Arf}(q_1) + \text{Arf}(q_2)$$

by Lemma 7.4. \square

Having seen that ψ is well-defined independent of choice of geometric splitting, we proceed to show that it is invariant on orbits of $\ker(\Theta_{q, \kappa})$.

Lemma 7.7. *Let $\alpha = x + \kappa \in H_1(X, \mathcal{B}; \mathbb{Z}/2\mathbb{Z})$ and $F \in \ker(\Theta_{q, \kappa})$ be given. Then*

$$\psi(F\alpha) = \psi(\alpha).$$

Proof. Write

$$F = \begin{pmatrix} M & A \\ 0 & I \end{pmatrix};$$

then we must verify that

$$q(Mx + A\kappa) = q(x).$$

As q is a quadratic form,

$$q(Mx + A\kappa) = q(Mx) + q(A\kappa) + \langle Mx, A\kappa \rangle. \quad (10)$$

From Lemma 7.4, $q(A\kappa) = 0$. As $F \in \ker(\Theta_{q,\kappa})$, the equation

$$q(Mx) + q(x) + \langle Mx, A\kappa \rangle = 0$$

holds for all $x \in H_1(X; \mathbb{Z}/2\mathbb{Z})$, and so we can replace $q(Mx) + \langle Mx, A\kappa \rangle$ in (10) with $q(x)$. Altogether, this shows

$$\psi(F\alpha) = q(Mx + A\kappa) + \text{Arf}(q) = q(x) + \text{Arf}(q) = \psi(\alpha)$$

as claimed. \square

8 Orbit counts (II): Split non-hyperelliptic case

Having defined the ψ invariant, we now see how it gives a complete classification of orbits in the non-hyperelliptic setting. Recall that Proposition 4.2 and Corollary 6.2 show that under the Chinese remainder theorem decomposition, the orbit $\bar{\Gamma}_d\alpha$ is split if and only if the mod-2 reduction $\bar{\Gamma}_2\alpha$ is split, and that the Chinese remainder theorem decomposition Corollary 4.3 allows us to consider orbits one prime at a time. Thus in this section, we need only consider the prime $p = 2$.

8.1 Statement of results

Proposition 8.1 ($p = 2$ orbit classification, non-hyperelliptic). *Let (X, ω) be a translation surface in a non-hyperelliptic component of a stratum. Let $\alpha, \beta \in H_1^{\text{prim}}(X, \mathcal{B}; \mathbb{Z}/2^k\mathbb{Z})$ be given, and suppose that $\delta := \delta(\alpha) = \delta(\beta) \in \tilde{H}_0(\mathcal{B}; \mathbb{Z}/2^k\mathbb{Z})$.*

- If $\delta \neq \kappa \pmod{2}$, then every orbit is unsplit: $\bar{\Gamma}_{2^k}\alpha = \bar{\Gamma}_{2^k}\beta$.
- If $\delta = \kappa \pmod{2}$, then $\bar{\Gamma}_{2^k}\alpha = \bar{\Gamma}_{2^k}\beta$ if and only if $\psi(\alpha) = \psi(\beta)$.

Proposition 8.2 ($p = 2$ split orbit sizes, non-hyperelliptic). *Let (X, ω) be a translation surface in a non-hyperelliptic component of a stratum, and let $\alpha \in H_1^{\text{prim}}(X, \mathcal{B}; \mathbb{Z}/2^k\mathbb{Z})$. Assume further that $\delta(\alpha) = \kappa \pmod{2}$, so that the orbit of α is split. Define*

$$\tau(\alpha) = \begin{cases} 1 & \text{if } 2 \mid \text{gcd}(\kappa) \text{ and } \psi(\alpha) = \text{Arf}(\phi) \\ 0 & \text{otherwise.} \end{cases}$$

Then

$$|\bar{\Gamma}_{2^k}\alpha| = 2^{2g(k-1)} \left(2^{g-1}(2^g + (-1)^{\psi(\alpha)}) - \tau(\alpha) \right).$$

Proposition 8.3 ($p = 2$ split orbit statistics, non-hyperelliptic). *Let (X, ω) be a translation surface in a non-hyperelliptic component of a stratum. Let $\bar{\Gamma}_{2^k}\alpha \subset H_1^{\text{prim}}(X, \mathcal{B}; \mathbb{Z}/2^k\mathbb{Z})$ be an orbit. Suppose that $\delta(\alpha) = \kappa$, so that the orbit of α is split. Let $c \in H_1^{\text{prim}}(X \setminus \mathcal{B}; \mathbb{Z}/p^k\mathbb{Z})$ be the homology class of a cylinder on X , and let $a \in \mathbb{Z}/p^k\mathbb{Z}$ be given. Then*

$$|(\bar{\Gamma}_{2^k}\alpha)_{c;a}| = \begin{cases} 2^{(2k-1)g-k}(2^{g-1} + (-1)^{\psi(\alpha)}) & a = 1 \pmod{2} \\ 2^{(2g-1)(k-1)}(2^{2g-2} - \tau(\alpha)) & a = 0 \pmod{2}, \end{cases}$$

where $\tau(\alpha)$ is defined as in Proposition 8.2.

Note that as discussed in Section 3.5 and as in Proposition 4.6, these quantities are independent of c .

8.2 Proofs

Throughout this section, we will tacitly make use of Theorem 6.1, to identify the monodromy group $\bar{\Gamma}_d$ with $\ker(\Theta_{q,\kappa})$. As in Section 5.2, we will also fix a geometric splitting and the induced algebraic splitting.

8.2.1 Orbit classification

Proof of Proposition 8.1. According to Proposition 4.2 and Corollary 6.2, the orbit $\bar{\Gamma}_{2^k}\alpha$ is unsplit if and only if its reduction $\bar{\Gamma}_2\alpha$ is unsplit. Thus it suffices to consider the case $k = 1$ in what follows.

We first consider the case $\delta \neq \kappa$; suppose first that $\delta \neq 0$ as well. Let $y \in H_1(X; \mathbb{Z}/2\mathbb{Z})$ be arbitrary. Since $\delta \neq 0$, we can take $A \in \text{Hom}(\tilde{H}_0(\mathcal{B}; \mathbb{Z}/2\mathbb{Z}), H_1(X; \mathbb{Z}/2\mathbb{Z}))$ such that $A\delta = y - x$. Moreover, since $\delta \neq \kappa$, either δ and κ are linearly independent or else $\kappa = 0$; in either case, we can also stipulate that $A\kappa = 0$. In particular, $\langle z, A\kappa \rangle = 0$ for all $z \in H_1(X; \mathbb{Z}/2\mathbb{Z})$ and so $\begin{pmatrix} I & A \\ 0 & I \end{pmatrix} \in \ker(\Theta_{q,\kappa})$. Now by construction this element takes $x + \delta$ to $x + A\delta + \delta = y + \delta$, and thus $\ker(\Theta_{q,\kappa})$ acts transitively on vectors of the form $x + \delta$, as desired.

In the case $\delta \neq \kappa$ and $\delta = 0$, note first that by the assumption that the cover be connected that x must be nonzero. Take an arbitrary nonzero $y \in H_1(X; \mathbb{Z}/2\mathbb{Z})$; then by Proposition A.1.1 there is an element $M \in \text{Sp}(2g, \mathbb{Z}/2\mathbb{Z})$ such that $Mx = y$. We seek to construct $A \in \text{Hom}(\tilde{H}_0(\mathcal{B}; \mathbb{Z}/2\mathbb{Z}), H_1(X; \mathbb{Z}/2\mathbb{Z}))$ such that

$$q(Mz) + q(z) + \langle Mz, A\kappa \rangle = 0 \quad (11)$$

for all $z \in H_1(X; \mathbb{Z}/2\mathbb{Z})$. Since the symplectic pairing is nondegenerate, (11) completely characterizes $A\kappa$, and so a suitable A can be constructed.

Finally, we consider the case $\delta = \kappa$. Here we seek to show that a primitive $x + \kappa$ can be taken to any primitive $y + \kappa$ with $\psi(x + \kappa) = \psi(y + \kappa)$, i.e., with $q(x) = q(y)$. In the case that both x and y are nonzero, Proposition A.1.3 tells us there is some $M \in \text{Sp}(2g, \mathbb{Z}/2\mathbb{Z})[q]$ taking x to y . Now any element $\begin{pmatrix} M & 0 \\ 0 & I \end{pmatrix}$ for $M \in \text{Sp}(2g, \mathbb{Z}/2\mathbb{Z})[q]$ is an element of $\ker(\Theta_{q,\kappa})$ and takes x to y . If $\kappa = 0$, then this argument completes the proof in that case.

The last thing to prove is that if $\kappa \neq 0$, there is an element of $\ker(\Theta_{q,\kappa})$ taking κ to some $x + \kappa$ for a nonzero $x \in H_1(X; \mathbb{Z}/2\mathbb{Z})$. This can be done by choosing some $M \in \text{Sp}(2g, \mathbb{Z}/2\mathbb{Z})$ that *does not* preserve q . Then (11) characterizes $A\kappa$, and as long as $M \notin \text{Sp}(2g, \mathbb{Z}/2\mathbb{Z})[q]$, then $A\kappa \neq 0$. Then $\begin{pmatrix} M & A \\ 0 & I \end{pmatrix} \in \ker(\Theta_{q,\kappa})$ takes κ to $A\kappa + \kappa$ as required; applying the previous paragraph then allows us to take $A\kappa + \kappa$ to any other $y + \kappa$ so that $q(A\kappa) = q(y)$. □

8.2.2 Orbit sizes

Before turning to the proof of Proposition 8.2, we record a lemma. The following result is standard in the theory of the Arf invariant, and gives a quantitative formulation of the slogan that the Arf invariant is the “democratic invariant,” equal to whichever value in $\mathbb{Z}/2\mathbb{Z}$ is assumed more often by q .

Lemma 8.4. *Let $q : H_1(X; \mathbb{Z}/2\mathbb{Z}) \rightarrow \mathbb{Z}/2\mathbb{Z}$ be a quadratic form. Then for $\varepsilon \in \mathbb{Z}/2\mathbb{Z}$, there are*

$$2^{g-1}(2^g + (-1)^{\text{Arf}(q)+\varepsilon})$$

elements $x \in H_1(X; \mathbb{Z}/2\mathbb{Z})$ satisfying $q(x) = \varepsilon$.

Proof. Induction on g . □

Proof of Proposition 8.2. Fix a geometric splitting X° , inducing an algebraic splitting as in Definition 5.5. By Proposition 8.1, $\bar{\Gamma}_{2^k}\alpha$ consists of all $\beta = y + \kappa \in H_1^{\text{prim}}(X, \mathcal{B}; \mathbb{Z}/2^k\mathbb{Z})$ such that $\psi(\beta) = \psi(\alpha)$. Expressing

$\alpha = x + \kappa$, this latter condition is equivalent to $q(x) = q(y)$. We consider the two possibilities $\gcd(\kappa) = 1$ or $2 \mid \gcd(\kappa)$ in turn.

If $\gcd(\kappa) = 1$, the orbit $\bar{\Gamma}_{2^k} \alpha$ is in bijection with $y \in H_1(X; \mathbb{Z}/2^k \mathbb{Z})$ for which $q(y) = q(x)$. The value $q(x) = q(y)$ is determined by the mod-2 reduction. By Lemma 8.4, there are

$$2^{g-1}(2^g + (-1)^{\text{Arf}(q)+q(x)}) = 2^{g-1}(2^g + (-1)^{\psi(\alpha)})$$

elements $\bar{x} \in H_1(X; \mathbb{Z}/2\mathbb{Z})$ with $q(x) = q(y)$, for a total of

$$2^{2g(k-1)} \left(2^{g-1}(2^g + (-1)^{\psi(\alpha)}) \right) \quad (12)$$

elements $y \in H_1(X; \mathbb{Z}/2^k \mathbb{Z})$ that satisfy $q(y) = q(x)$. Recalling $\tau(\alpha) = 0$ if $\gcd(\kappa) = 1$, this proves the claim in this case.

If $2 \mid \gcd(\kappa)$, then $y \in H_1(X; \mathbb{Z}/2^k \mathbb{Z})$ must additionally be primitive. If $q(y) = 1$, then y is *necessarily* primitive, so that the orbit count is again given by (12). We observe that since $2 \mid \gcd(\kappa)$, the condition $q(x) = 1$ is equivalent to setting $\tau(\alpha) = 0$: indeed, $\tau(\alpha) = 0$ is equivalent to

$$\text{Arf}(\phi) = \psi(\alpha) + 1 = q(x) + \text{Arf}(q) + 1 = q(x) + \text{Arf}(\phi) + 1.$$

In case $q(x) = 0$, we must exclude all nonprimitive $y \in H_1(X; \mathbb{Z}/2^k \mathbb{Z})$ from consideration. Such y is primitive if and only if its reduction mod 2 is nonzero, so that there are

$$2^{2g(k-1)} \left(2^{g-1}(2^g + (-1)^{\psi(\alpha)}) - 1 \right) = 2^{2g(k-1)} \left(2^{g-1}(2^g + (-1)^{\psi(\alpha)}) - \tau(\alpha) \right)$$

primitive $y \in H_1(X; \mathbb{Z}/2^k \mathbb{Z})$ satisfying $q(y) = 0$. □

8.2.3 Orbit statistics

The proof of Proposition 8.3 requires the following preliminary lemma. Let q be a quadratic form on $H_1(X; \mathbb{Z}/2\mathbb{Z})$, and fix $c \in H_1(X; \mathbb{Z}/2\mathbb{Z})$ with $q(c) = 1$. For $\varepsilon, a \in \mathbb{Z}/2\mathbb{Z}$, define

$$N(g, q, \varepsilon, a) = |\{y \in H_1(X; \mathbb{Z}/2\mathbb{Z}) \mid q(y) = \varepsilon, \langle c, y \rangle = a\}|.$$

Lemma 8.5. *For any $c \in H_1(X; \mathbb{Z}/2\mathbb{Z})$ and any quadratic form with $q(c) = 1$,*

$$N(g, q, \varepsilon, a) = \begin{cases} 2^{2g-2} & a = 0 \\ 2^{g-1}(2^{g-1} + (-1)^{\varepsilon + \text{Arf}(q)}) & a = 1. \end{cases}$$

Proof. We first remark that since $\text{Sp}(2g, \mathbb{Z}/2\mathbb{Z})$ acts transitively on the set of quadratic forms of given Arf invariant, we are free to work with whichever particular q is most convenient. Likewise, by Proposition A.1.3, the stabilizer $\text{Sp}(2g, \mathbb{Z}/2\mathbb{Z})[q]$ acts transitively on the set of elements $c \in H_1(X; \mathbb{Z}/2\mathbb{Z})$ satisfying $q(c) = 1$ (note that all such c are nonzero), and so we are free to choose a convenient such c .

We fix such choices as follows. Let $x_1, y_1, \dots, x_g, y_g$ be a symplectic basis for $H_1(X; \mathbb{Z}/2\mathbb{Z})$, and define q by the conditions that $q(x_1) = 1$ and $q(x_i) = q(y_i) = 0$ for $i > 1$; then $q(y_1) = \text{Arf}(q)$. We choose $c = x_1$. Having made such choices, the proof proceeds by induction on g .

In the base case $g = 1$, we observe that $\langle x_1, 0 \rangle = \langle x_1, x_1 \rangle = 0$, and $\langle x_1, y_1 \rangle = \langle x_1, x_1 + y_1 \rangle = 1$. The first pair of equalities shows that

$$N(1, q, \varepsilon, 0) = 1 = 2^{2g-2}.$$

As $q(x_1 + y_1) = q(x_1) + q(y_1) + \langle x_1, y_1 \rangle = q(y_1)$, it follows that $N(1, q, \varepsilon, 1)$ is either 2 or 0 according to whether $\varepsilon = \text{Arf}(q)$ or not, establishing the formula in the case $g = 1, a = 1$.

Now consider $N(g+1, q, \varepsilon, a)$ for $g \geq 1$. Every vector in $(\mathbb{Z}/2\mathbb{Z})^{2g+2}$ is of the form $v + w$ for $v \in (\mathbb{Z}/2\mathbb{Z})^{2g}$ and $w \in \{0, x_{g+1}, y_{g+1}, x_{g+1} + y_{g+1}\}$. Note that $\langle x_1, v + w \rangle = \langle x_1, v \rangle$ and that $q(v + w) = q(v) + q(w)$. Of the four possibilities for w , only $x_{g+1} + y_{g+1}$ has $q = 1$. We therefore deduce the recursion

$$N(g+1, q, \varepsilon, a) = 3N(g, q, \varepsilon, a) + N(g, q, \varepsilon + 1, a),$$

A simple inductive argument then shows that the claimed formulas hold. □

Proof of Proposition 8.3. According to Proposition 8.1, $\bar{\Gamma}_{2^k}\alpha$ consists of all $\beta \in H_1^{\text{prim}}(X, \mathcal{B}; \mathbb{Z}/2^k\mathbb{Z})$ such that

1. $\beta = y + \kappa$ for $y \in H_1(X; \mathbb{Z}/2^k\mathbb{Z})$,
2. $\psi(\beta) = \psi(\alpha) \in \mathbb{Z}/2\mathbb{Z}$.

Here we must make a further distinction between $\gcd(\kappa) = 1$ and $2 \mid \gcd(\kappa)$. In the case $\gcd(\kappa) = 1$, the condition $\gcd(\beta) = 1$ is automatically satisfied. Choose a geometric splitting under which the cylinder with class c is contained in the subsurface X° , and let (q, κ) be the associated data. As c is the class of a cylinder, necessarily $\phi(c) = 0$, i.e., $q(c) = 1$.

We seek to count $y \in H_1(X; \mathbb{Z}/2^k\mathbb{Z})$ such that $\psi(y + \kappa) = q(y) + \text{Arf}(q) = \psi(\alpha)$ and such that

$$\langle c, y + \kappa \rangle = \langle c, y \rangle = a \pmod{2^k}.$$

As $q(y)$ depends only on the class of y in $\mathbb{Z}/2\mathbb{Z}$, it suffices to work with $\mathbb{Z}/2\mathbb{Z}$ coefficients; each such \bar{y} then corresponds to $2^{(k-1)(2g-1)}$ such $y \in H_1(X; \mathbb{Z}/2^k\mathbb{Z})$ with the value $\langle c, y \rangle$ fixed. Applying Lemma 8.5, this quantity is given as

$$\begin{aligned} & |\{y \in H_1(X; \mathbb{Z}/2^k\mathbb{Z}) \mid q(y) + \text{Arf}(q) = \psi(\alpha), \langle c, y \rangle = a\}| \\ &= 2^{(k-1)(2g-1)} N(g, q, \text{Arf}(q) + \psi(\alpha), a). \end{aligned}$$

This completes the argument in the case $\gcd(\kappa) = 1$.

It remains to consider the case $2 \mid \gcd(\kappa)$. Recall from Definition 7.2 that in this case, the Arf invariant $\text{Arf}(\phi) = \text{Arf}(q)$ is globally well-defined independent of the choice of geometric splitting. In this case, the condition $\gcd(\beta) = 1$ will be satisfied if and only if $\beta = y + \kappa$ with $y \in H_1(X; \mathbb{Z}/2^k\mathbb{Z})$ primitive, and this latter condition holds if and only if the reduction $y \pmod{2}$ is nonzero. As in the previous case, we can reduce mod 2:

$$\begin{aligned} & |\{y \in H_1(X; \mathbb{Z}/2^k\mathbb{Z}) \mid y \text{ primitive}, q(y) + \text{Arf}(q) = \psi(\alpha), \langle c, y \rangle = a\}| \\ &= 2^{(k-1)(2g-1)} |\{y \in H_1(X; \mathbb{Z}/2\mathbb{Z}) \mid y \neq 0, q(y) + \text{Arf}(q) = \psi(\alpha), \langle c, y \rangle = a\}|. \end{aligned}$$

If $a = 1$ then the condition $y \neq 0$ is redundant, and we can count as before:

$$\begin{aligned} & |\{y \in H_1(X; \mathbb{Z}/2^k\mathbb{Z}) \mid y \text{ prim.}, q(y) + \text{Arf}(q) = \psi(\alpha), \langle c, y \rangle = 1\}| \\ &= 2^{(k-1)(2g-1)} N(g, q, \text{Arf}(q) + \psi(\alpha), 1) \\ &= 2^{(2k-1)g-k} (2^{g-1} + (-1)^{\psi(\alpha)}) \end{aligned}$$

as claimed.

If $a = 0$ and $\text{Arf}(q) + \psi(\alpha) = 0$, then we must exclude $y = 0$ (and all its preimages before reducing mod 2) from the count. Recall that $\tau(\alpha) = 1$ if and only if $\gcd(\kappa) > 1$ and $\psi(\alpha) = \text{Arf}(q)$, so that

$$\begin{aligned} & |\{y \in H_1(X; \mathbb{Z}/2^k\mathbb{Z}) \mid y \text{ prim.}, q(y) + \text{Arf}(q) = \psi(\alpha), \langle c, y \rangle = 0\}| \\ &= 2^{(k-1)(2g-1)} (N(g, q, \text{Arf}(q) + \psi(\alpha), 0) - \tau(\alpha)) \\ &= 2^{(k-1)(2g-1)} (2^{2g-2} - \tau(\alpha)). \end{aligned}$$

□

9 The homological monodromy group, hyperelliptic case

In these last two sections of Part II, we turn our attention to the hyperelliptic setting. The goal of this section is to establish Proposition 9.4, which is the counterpart of Corollary 6.2 in the hyperelliptic setting. We begin with a recollection of some of the basics of the *topological* monodromy of a hyperelliptic component.




Figure 4: Generators $\{c_i\}$ for hyperelliptic mapping class groups.

9.1 The hyperelliptic mapping class group

We first make some general comments about hyperelliptic translation surfaces; see also [LM14, Cal20], and [KZ03].

Recall [KZ03] that the strata $\mathcal{H}(2g-2)$ and $\mathcal{H}(g-1, g-1)$ each contain a special *hyperelliptic* component, written $\mathcal{H}^{hyp}(2g-2)$ and $\mathcal{H}^{hyp}(g-1, g-1)$, or just as \mathcal{H}^{hyp} when we wish to be agnostic about the number of zeros. These consist of differentials ω on hyperelliptic Riemann surfaces arising as the global square root of a quadratic differential on \mathbb{CP}^1 with a unique zero and either $2g+1$ or $2g+2$ simple poles, respectively. For any $(X, \omega) \in \mathcal{H}^{hyp}(2g-2)$, the hyperelliptic involution ι fixes the zero, and for any $(X, \omega) \in \mathcal{H}^{hyp}(g-1, g-1)$ it interchanges the zeros [KZ03, Remark 3].

Definition 9.1. Set Br to be the number of simple poles of q , i.e., $Br = 2g+1$ if $\mathcal{H} = \mathcal{H}^{hyp}(2g-2)$ and $Br = 2g+2$ if $\mathcal{H} = \mathcal{H}^{hyp}(g-1, g-1)$.

So long as the quadratic differential q is nonsingular at infinity, it admits the explicit expression

$$q = \frac{(z - z_0)^{2g-2}}{(z - z_1) \dots (z - z_{Br})} dz^2$$

and we can take the positions of the singularities $\mathcal{S} = \{z_0, \dots, z_{Br}\}$ of q as moduli for \mathcal{H}^{hyp} . This implies that the component \mathcal{H}^{hyp} is a $K(G, 1)$ for (a finite extension of) the spherical braid group on $Br+1$ strands, with one strand corresponding to the zero of q and the other to its poles [LM14, Section 1.4]. Via the Birman–Hilden theory of symmetric mapping class groups (cf. [MW21] or [FM12, § 9.4]), this spherical braid group can be identified with the centralizer of ι in the mapping class group of X relative to the zero(es) \mathcal{Z} of ω . We therefore arrive at the following:

Lemma 9.2. *The topological monodromy group of a hyperelliptic stratum component where the zeros are unlabeled is the centralizer, denoted $\text{SMod}(X, \mathcal{Z})$, of the hyperelliptic involution ι .*

As a consequence, the topological monodromy group of the cover on which the zeros are labeled is exactly

$$\text{PSMod}(X, \mathcal{Z}) := \text{PMod}(X, \mathcal{Z}) \cap \text{SMod}(X, \mathcal{Z}),$$

where the pure and symmetric mapping class groups are both thought of as subgroups of $\text{Mod}(X, \mathcal{Z})$. Note that when X has two zeros, ι interchanges them and so is not an element of $\text{PMod}(X, \mathcal{Z})$, hence we cannot describe $\text{PSMod}(X, \mathcal{Z})$ as its centralizer.

For both $\mathcal{H}^{hyp}(2g-2)$ and $\mathcal{H}^{hyp}(g-1, g-1)$, fix the curves c_i as shown in Figure 4; then the hyperelliptic mapping class group $\text{PSMod}(X, \mathcal{Z})$ corresponding to the hyperelliptic involution that setwise fixes each c_i is generated by Dehn twists about the c_i 's [FM12, § 9.4]. While we will not explicitly use this generating set, the c_i are a useful set of curves, as we will see in the proof of Proposition 9.4 below.

9.1.1 Ordinary marked points

In our setting, we consider strata of differentials on which some number (possibly zero) of ordinary marked points are distinguished. Let \mathcal{P} denote this set of ordinary marked points, so that the full set \mathcal{B} of distinguished points is given by

$$\mathcal{B} = \mathcal{Z} \sqcup \mathcal{P},$$

with \mathcal{Z} consisting of the zeros, a set of size at most two (and at least one, in the regime $g \geq 2$).

For any partition κ of $2g - 2$, the stratum $\mathcal{H}_{lab}(\kappa, 0^m)$ is a fiber bundle over the stratum $\mathcal{H}_{lab}(\kappa)$, with fiber $\text{PConf}_m(X \setminus \mathcal{B}_\kappa)$, where $\text{PConf}_m(\cdot)$ denotes the space of ordered configurations of m points on a space, and \mathcal{B}_κ is the distinguished set for the stratum $\mathcal{H}_{lab}(\kappa)$. Note in particular that the components of $\mathcal{H}_{lab}(\kappa)$ and $\mathcal{H}_{lab}(\kappa, 0^m)$ are in natural bijective correspondence.

On the level of monodromy, the topological monodromy groups of components of $\mathcal{H}_{lab}(\kappa)$ and $\mathcal{H}_{lab}(\kappa, 0^m)$ are related by the following, which is an immediate consequence of the *Birman exact sequence* [FM12, Theorem 9.1].

Lemma 9.3. *Let $\mathcal{H} \subset \mathcal{H}_{lab}(\kappa)$ be a component of a stratum and let $\mathcal{H}^+ \subset \mathcal{H}_{lab}(\kappa, 0^m)$ be obtained from \mathcal{H} by marking regular points. Let \mathcal{B}_κ be the distinguished set of points for $X \in \mathcal{H}$, and let $\mathcal{B}^+ = \mathcal{B}_\kappa \sqcup \mathcal{P}$ be the distinguished set for $(X, \omega) \in \mathcal{H}^+$. Let $\Gamma \leq \text{PMod}(X, \mathcal{B}_\kappa)$ and $\Gamma^+ \leq \text{PMod}(X, \mathcal{B}^+)$ be the topological monodromy groups for \mathcal{H} and \mathcal{H}^+ , respectively. Then there is a short exact sequence*

$$1 \rightarrow PB_m(X \setminus \mathcal{B}_\kappa) \rightarrow \Gamma^+ \rightarrow \Gamma \rightarrow 1,$$

where $PB_m(X \setminus \mathcal{B}_\kappa) := \pi_1(\text{PConf}_m(X \setminus \mathcal{B}_\kappa))$ is the surface braid group.

The group $PB_m(X \setminus \mathcal{B}_\kappa)$ is generated by *point-pushing elements*: these are mapping classes associated to choosing a marked point $q \in \mathcal{B}^+ \setminus \mathcal{B}_\kappa$, and moving q around $X \setminus \mathcal{B}_\kappa$ while fixing all other marked points, eventually returning to q .

9.2 The homological monodromy group

We come to the main result of the section, a variant of a result of A'Campo [A'C79]. To state it, recall from Section 6 that $\text{PAut}(H_1(X, \mathcal{B}; \mathbb{Z}))[2]$ is the kernel of the action on mod-2 relative homology.

Relative to the short exact sequence (1), this admits the following decomposition:

$$1 \rightarrow \text{Hom}(\tilde{H}_0(\mathcal{B}; \mathbb{Z}), H_1(X; 2\mathbb{Z})) \rightarrow \text{PAut}(H_1(X, \mathcal{B}; \mathbb{Z}))[2] \rightarrow \text{Sp}(2g, \mathbb{Z}[2]) \rightarrow 1.$$

We will also need to examine the kernel term more closely. The set \mathcal{B} of distinguished points has two constituents: zeros \mathcal{Z} and ordinary marked points \mathcal{P} . This leads to a decomposition

$$\tilde{H}_0(\mathcal{B}; \mathbb{Z}) \cong \tilde{H}_0(\mathcal{Z}; \mathbb{Z}) \oplus H_0(\mathcal{P}; \mathbb{Z}), \quad (13)$$

with the latter summand identified with a subspace of $\tilde{H}_0(\mathcal{B}; \mathbb{Z})$ via the correspondence $[q] \in H_0(\mathcal{P}; \mathbb{Z}) \mapsto [q] - [p_1] \in \tilde{H}_0(\mathcal{B}; \mathbb{Z})$ (here, $p_1 \in \mathcal{Z}$ is a zero).

The next two lemmas produce certain useful subgroups of $\bar{\Gamma}$ in the hyperelliptic setting.

Proposition 9.4. *Let \mathcal{H}_{lab}^{hyp} be a hyperelliptic component of labeled translation surfaces of genus $g \geq 2$, possibly with some marked ordinary points. Then there is a containment*

$$\text{PAut}(H_1(X, \mathcal{B}; \mathbb{Z}))[2] \leq \bar{\Gamma}.$$

Proof. To prove the result, we must show that $\bar{\Gamma}$ surjects onto $\text{Sp}(2g, \mathbb{Z}[2])$, and that $\bar{\Gamma}$ contains $\text{Hom}(\tilde{H}_0(\mathcal{B}; \mathbb{Z}), H_1(X; 2\mathbb{Z}))$.

The surjection of $\bar{\Gamma}$ onto $\text{Sp}(2g, \mathbb{Z}[2])$ follows readily from the literature, cf. [AMY18, Theorem 2.9] or [A'C79]. To establish the second claim, we will exhibit special “point-pushing” monodromy elements that induce the identity on absolute homology, *transvect* a chosen basis element $[x] - [p_1]$ to $[x] - [p_1] + 2y$ for $y \in H_1(X; \mathbb{Z})$ arbitrary, and fix the remaining basis elements of $\tilde{H}_0(\mathcal{B}; \mathbb{Z})$.

To show that every element of $H_1(X; 2\mathbb{Z})$ is achieved as a transvection of $[p_2] - [p_1]$, we will make use of the discussion of Section 9.1. There, we gave a description of the topological monodromy group as lifted from a braid group on \mathbb{CP}^1 . We will use this to directly exhibit point-pushes of z_0 on $\mathbb{CP}^1 \setminus \mathcal{S}$ whose lifts to X generate $H_1(X; 2\mathbb{Z})$.




Figure 5: Lifting a point-push via the hyperelliptic involution.

Choose a sequence of pairwise disjoint simple arcs α_i connecting z_i to z_{i+1} for $i = 1, \dots, 2g + 1$; we can then take the set of every other arc $\alpha_1, \alpha_3, \dots, \alpha_{2g+1}$ as branch cuts for the ramified cover $X \rightarrow \mathbb{CP}^1$. Finally, choose simple arcs ε_i whose interiors are disjoint from each other and all α_j and that connect the basepoint z_0 to α_i . See Figure 5.

The lifts of the α_i form a chain c_1, \dots, c_{2g+1} of simple closed curves on X , that is, a set of curves with intersection given by

$$i(c_i, c_j) = \begin{cases} 1 & \text{if } |i - j| = 1 \\ 0 & \text{otherwise.} \end{cases}$$

In particular, $\{c_1, \dots, c_{2g}\}$ is a basis for $H_1(X; \mathbb{Z})$. Now consider the loop ζ_i on $\mathbb{CP}^1 \setminus \mathcal{S}$ formed by traveling along ε_i , going once around α_i , and back along ε_i . Pushing z_0 along ζ_i has the effect of pushing each of p_1 and p_2 (the zeros of the differential) once around c_i , but in opposite directions. Thus, the lift of ζ_i to $\text{Hom}(\tilde{H}_0(\mathcal{B}; \mathbb{Z}), H_1(X; 2\mathbb{Z}))$ transvects any relative homology class by $2[c_i]$. Since the $\{c_i\}$ generate the integral homology of X and every ζ_i can be realized as a loop of quadratic differentials, we see that $[p_2] - [p_1]$ can be transvected by an arbitrary element of $H_1(X; 2\mathbb{Z})$, completing the proof of the proposition. \square

Lemma 9.5. *Let \mathcal{H}_{lab} be a component of labeled translation surfaces of genus $g \geq 1$ (hyperelliptic or otherwise). With respect to the splitting (13) of $\tilde{H}_0(\mathcal{B}; \mathbb{Z})$, there is a containment*

$$\text{Hom}(H_0(\mathcal{P}; \mathbb{Z}), H_1(X; \mathbb{Z})) \leq \bar{\Gamma}.$$

Proof. Following the discussion in Proposition 9.4, we will exhibit further “point-pushing” transvections that send a chosen basis element $[x] - [p_1]$ to $[x] - [p_1] + y$ for $x \in \mathcal{P}$ and $y \in H_1(X; \mathbb{Z})$ arbitrary, and that fix the remaining basis elements of $\tilde{H}_0(\mathcal{B}; \mathbb{Z})$.

Consider a basis element $[q] - [p_1]$ for $q \in \mathcal{P}$. Since q is an ordinary marked point, it is free to move around X along an arbitrary path that avoids the remaining points of \mathcal{B} . Such a path can trace out an arbitrary homology class $y \in H_1(X; \mathbb{Z})$, and so long as this path stays disjoint from the remaining arcs representing a basis of $\tilde{H}_0(\mathcal{B}; \mathbb{Z})$, it will preserve the associated homology classes. \square

10 Orbit counts (III): Split hyperelliptic case

In this section, we complete our discussion of the orbits of $\bar{\Gamma}_d$ on $H_1(X, \mathcal{B}; \mathbb{Z}/d\mathbb{Z})$ (and hence the path component structure of the corresponding space of covers \mathcal{M}_δ), by considering the case of hyperelliptic strata. As in the non-hyperelliptic setting, Proposition 4.2 and Proposition 9.4 together imply that we may restrict our attention to the case of mod-2 coefficients. We will see that in the hyperelliptic case, orbits are classified by how many branch points of the involution are encircled by a nice representative. We state our results in Section 10.2 after introducing this invariant.

10.1 Encircling branch points: the b invariant

In this subsection, we assume that $g \geq 2$ and so the set \mathcal{Z} of zeros of ω is nonempty. Since Lemma 9.3 allows us to split off the action of the monodromy on regular points, we begin by analyzing the action on homology $\text{rel } \mathcal{Z}$.

In order to state our orbit classification, we first define an invariant which records how many branch points of ι that a curve encircles; see Figure 6. To define the b invariant, let a be a simple arc on X with endpoints on \mathcal{Z} (possibly at the same point) with the property that $\iota(a)$ is disjoint from a . Consider the image \bar{a} of a on the $(Br + 1)$ -punctured sphere $X/\langle\iota\rangle$; by assumption it is a simple arc that begins and ends on the marked point corresponding to \mathcal{Z} . The image \bar{a} therefore separates $X/\langle\iota\rangle$ into two spheres, one of which contains

$$b(a) \in \{1, \dots, \lfloor Br/2 \rfloor\}$$

branch points and the other of which contains $Br - b(a)$ branch points.

Definition 10.1 (b invariant). The b invariant of the simple arc a is the integer $b(a)$ following the discussion above.




Figure 6: An arc disjoint from its image under ι and how it separates the quotient $X/\langle\iota\rangle$. In this example, $b(a) = 3$.

This invariant is not defined on every isotopy class, and its definition requires finding a nice representative in an isotopy class; we deal with this issue just below. All the same, this topological definition of $b(a)$ makes it immediate that it is an invariant of $\text{PSMod}(X, \mathcal{Z})$ orbits.

Lemma 10.2. *Suppose that a is an arc on X with endpoints on \mathcal{Z} and so that $\iota(a)$ is disjoint from a . Then for any $c \in \text{PSMod}(X, \mathcal{Z}) \cdot a$, we have that $b(c) = b(a)$.*

Proof. By our discussion in Section 9.1, $\text{PSMod}(X, \mathcal{Z})$ is isomorphic to a spherical braid group on $Br + 1$ strands (after modding out by $\langle\iota\rangle$ in the case of a single zero). Moreover, this identification is equivariant with respect to the action of $\text{PSMod}(X, \mathcal{Z})$ on arcs on X and the action of the braid group on arcs on $X/\langle\iota\rangle$. The lemma now is a consequence of the fact that “the number of distinguished points that an arc cuts off” is a homeomorphism invariant and hence an invariant of the braid group action. \square

We now extend our definition of $b(a)$ to all simple closed curves, essentially by counting the number of branch points p so that the mod-2 winding number of a about p is nonzero. We will not pursue this exact definition, instead giving a homological formulation that will be more convenient in the sequel.

Encircling as intersection. Recall the definition of the curves c_i from Figure 4. We observe that $\{[c_i] \mid i \leq Br - 1\}$ is a basis for the (mod 2) homology of $X \setminus \mathcal{Z}$; however, it will be convenient to instead consider a different basis given by

$$e_j := \sum_{i \leq j} [c_i].$$

Representatives for these homology classes can be expressed by curves arranged in a star-like pattern; see Figure 7.

We now construct a generalization of the b invariant. Define a map

$$\underline{e} : H_1(X, \mathcal{Z}; \mathbb{Z}/2\mathbb{Z}) \rightarrow (\mathbb{Z}/2\mathbb{Z})^{Br-1}$$

by pairing $\alpha \in H_1(X, \mathcal{Z}; \mathbb{Z}/2\mathbb{Z})$ with each e_j .

Definition 10.3. For $\alpha \in H_1(X, \mathcal{Z}; \mathbb{Z}/2\mathbb{Z})$, set $b(\alpha)$ to be either the number of 0's or the number of 1's in the vector $\underline{e}(\alpha)$, whichever is smaller.

Of course, since we have given a new definition for an old invariant, we must show that it coincides with our original one.




Figure 7: Symmetric curves representing the classes e_j and their intersection with an arc. In this example, $b(a) = 3$ and the subsurface U is shaded.

Lemma 10.4. *Let a be a simple arc on X with endpoints on \mathcal{Z} so that $\iota(a)$ is disjoint from a . Then $b(a) = b([a])$.*

Proof. The proof of this lemma essentially follows from inspection of Figure 7. A similar figure exists for $\mathcal{H}^{hyp}(2g-2)$; we leave its construction to the reader. More explicitly, in either case, consider the (isotopy classes of) curves

$$T_{c_j} T_{c_{j-1}} \dots T_{c_2}(c_1)$$

representing the classes e_j . These curves (have representatives which) are symmetric under (a representative for) the hyperelliptic involution ι , and their quotients are a set of arcs $\varepsilon_1, \dots, \varepsilon_{Br-1}$. Each ε_i has one endpoint at a distinguished branch point p_0 ; enumerate the other branch points as p_1, \dots, p_{Br-1} , so that ε_i has endpoints p_0 and p_i .

Now the image \bar{a} of a on $X/\langle \iota \rangle$ separates the surface into two pieces; consider the subsurface U not containing p_0 . For $i \geq 1$, we observe that $p_i \in U$ if and only if the geometric intersection number of ε_i with \bar{a} is odd. The proof is concluded by noting that for each i , the geometric intersection number of ε_i and \bar{a} is the same as the geometric intersection number of $T_{c_i} T_{c_{i-1}} \dots T_{c_2}(c_1)$ and a , which has the same parity as the (algebraic) intersection number of e_j with $[a]$. \square

10.2 Statement of results

As in Section 9.2, when $g \geq 2$ we may write $\mathcal{B} = \mathcal{Z} \cup \mathcal{P}$ where \mathcal{Z} are zeros and \mathcal{P} are regular marked points and split

$$\tilde{H}_0(\mathcal{B}; \mathbb{Z}/d\mathbb{Z}) \cong \tilde{H}_0(\mathcal{Z}; \mathbb{Z}/d\mathbb{Z}) \oplus H_0(\mathcal{P}; \mathbb{Z}/d\mathbb{Z}).$$

Let π denote the projection to the second factor and $H_{\mathcal{Z}} = \ker(\pi)$.

It turns out that the invariant b defined above is in fact a complete invariant of the $\bar{\Gamma}_{2^k}$ action on $H_{\mathcal{Z}}$. Ultimately, we will establish the following:

Proposition 10.5 ($p = 2$ orbit classification, (hyper)elliptic). *Let \mathcal{H}_{lab}^{hyp} be a (hyper)elliptic stratum component of labeled translation surfaces in genus $g \geq 1$, possibly equipped with some ordinary marked points. Let $\alpha, \beta \in H_1^{prim}(X, \mathcal{B}; \mathbb{Z}/2^k\mathbb{Z})$ be given, and suppose that $\delta(\alpha) = \delta(\beta) \in \tilde{H}_0(\mathcal{B}; \mathbb{Z}/2^k\mathbb{Z})$.*

- If $g = 1$, then there is a single (unsplit) orbit: $\bar{\Gamma}_{2^k} \alpha = \bar{\Gamma}_{2^k} \beta$.

- If $g \geq 2$ and $\pi(\delta(\alpha))$ is nonzero, then there is a single (unsplit) orbit.
- Otherwise, $\bar{\Gamma}_{2^k}\alpha = \bar{\Gamma}_{2^k}\beta$ if and only if $b(\alpha) = b(\beta)$.

Remark. The case $g = 1$ is easy and does not require the technology developed in this section (and π is not even defined when $g = 1$), but since it uses the same argument as the unsplit hyperelliptic case, we have chosen to include it here.

Remark. In the case where $\mathcal{H} = \mathcal{H}^{hyp}(g-1, g-1, 0^{n-2})$ and $\pi(\delta) = 0$, the parity of $b(\alpha)$ determines whether α is actually an absolute class (even) or a strictly relative one (odd). So whereas the monodromy group of $\mathcal{H}^{hyp}(2g-2, 0^{n-1})$ acts on mod 2^k absolute homology with g orbits of primitive vectors, the monodromy group of $\mathcal{H}^{hyp}(g-1, g-1, 0^{n-2})$ acts on the mod 2^k reduction of $H_{\mathcal{Z}}$ with $g+1$ orbits, exactly $\lfloor g/2 \rfloor$ of which are absolute and $\lfloor g/2 \rfloor + 1$ of which are relative.

To deal with the fact that the complement of a subset of $2g+2$ elements of size $g+1$ also has size $g+1$, define the following correction factor: set

$$\sigma(\alpha) = \begin{cases} 1 & \mathcal{H} = \mathcal{H}^{hyp}(g-1, g-1, 0^{n-2}) \text{ and } b(\alpha) = g+1 \\ 0 & \text{otherwise.} \end{cases}$$

Proposition 10.6 ($p = 2$ split orbit sizes, hyperelliptic). *With notation as in Proposition 10.5, suppose that $g \geq 2$ and that $\alpha \in H_1^{prim}(X, \mathcal{B}; \mathbb{Z}/2^k\mathbb{Z})$ satisfies $\pi(\delta(\alpha)) = 0$, so the orbit of α is split and the b invariant is well-defined. Then*

$$|\bar{\Gamma}_{2^k}\alpha| = 2^{2g(k-1)-\sigma(\alpha)} \binom{Br}{b(\alpha)}.$$

Proposition 10.7 ($p = 2$ split orbit statistics, hyperelliptic). *With notation as in Propositions 10.5 and 10.6, let $c \in H_1(X \setminus \mathcal{B}; \mathbb{Z}/2^k\mathbb{Z})$ be the homology class of a cylinder on (X, ω) , and let $a \in \mathbb{Z}/2^k\mathbb{Z}$ be given. Suppose $\pi(\delta(\alpha)) = 0$, so that the orbit of α is split and the b invariant is well-defined. Then*

$$\frac{|\bar{\Gamma}_{2^k}\alpha_{c;a}|}{2^{(2g-1)(k-1)-\sigma(\alpha)}} = \begin{cases} 2 \binom{Br-2}{b(\alpha)-1} & a \equiv 1 \pmod{2} \\ \binom{Br-2}{b(\alpha)} + \binom{Br-2}{b(\alpha)-2} & a \equiv 0 \pmod{2}. \end{cases}$$

Remark. There was an arbitrary choice made in our definition of $b(\alpha)$: it would make just as much sense to take instead $Br - b(\alpha)$. In the topological language of Definition 10.1, this is because if an arc cuts off $b(\alpha)$ branch points on one side, then the other side contains $Br - b(\alpha)$ branch points. We observe that making this substitution in Propositions 10.6 and 10.7 above results in the same formulas by the symmetry of the binomial coefficients.

10.3 Proofs

10.3.1 Orbit classification

We first prove a useful result about realizing homology classes by nice arcs; this allows us to freely pass back and forth between our definitions of the b invariant.

Lemma 10.8. *For any class $\alpha \in H_{\mathcal{Z}}$ with mod-2 coefficients, there is a simple arc a on X with endpoints on \mathcal{Z} so that $\iota(a)$ is disjoint from a and $[a] = \alpha$.*

Proof. By the duality between relative and excision homology, and the fact that the e_j curves span $H_1(X \setminus \mathcal{Z}; \mathbb{Z}/2\mathbb{Z})$, it suffices to build an arc a with specified intersection with each e_j . One can do this explicitly as follows: the classes e_j can be represented by curves as in the proof of Lemma 10.4 that project down to arcs ε_j on $X/\langle \iota \rangle$. The ε_j are arranged in a “star” pattern, in that they all share one common endpoint and are otherwise disjoint. Thus, one can build a simple arc a' on $X/\langle \iota \rangle$ with endpoints the (unique) marked point corresponding to $\mathcal{Z}/\langle \iota \rangle$ that meets any specified subset of the ε_j exactly once and is disjoint from the rest.

See Figure 8. Either choice of lift of this arc to X yields an arc a with (not necessarily distinct) endpoints on \mathcal{Z} . Necessarily a and $\iota(a)$ are disjoint (except possibly at common endpoints) since a projects down to a simple arc on $X/\langle\iota\rangle$. Note finally that by construction, a has the specified intersection with each e_j . \square




Figure 8: Constructing an arc with specified intersection with the ε_j arcs.

Proof of Proposition 10.5. By Proposition 4.2 and Proposition 9.4, we need only consider the mod-2 orbits, not the mod- 2^k ones. If either $g = 1$ or $g \geq 2$ and $\pi(\delta(\alpha)) \neq 0$, then Lemma 9.5 allows us to transvect arbitrarily along the fibers of δ , so $\bar{\Gamma}_2$ acts with unsplit orbits.

So suppose $g \geq 2$ and $\pi(\delta(\alpha)) = 0 \pmod{2}$, i.e., $\alpha \in H_{\mathcal{Z}}$. Fix any $F \in \bar{\Gamma}_2$. Using Lemma 10.8, choose some arc a with mod-2 homology class α so that $\iota(\alpha)$ is disjoint from α , and choose any $f \in \text{PSMod}(X, \mathcal{B})$ that induces the F action on relative homology. Then Lemmas 10.4 and 10.2, respectively, imply that

$$b(F\alpha) = b(f(a)) = b(a) = b(\alpha),$$

showing that b is an invariant of the $\bar{\Gamma}_2$ action.

Now suppose that $\beta \in H_1^{\text{prim}}(X, \mathcal{B}; \mathbb{Z}/2\mathbb{Z})$ with $\delta(\beta) = \delta(\alpha)$. Then $\pi(\delta(\beta)) = 0$, and so applying Lemma 10.8 again, we can represent β by some arc b that is disjoint from $\iota(\beta)$. But now a and b both project to simple arcs \bar{a} and \bar{b} on the quotient $X/\langle\iota\rangle$ that cut off $b(\alpha) = b(\beta)$ branch points on each side. By the change-of-coordinates principle [FM12, Section 1.3], there is thus some homeomorphism of $X/\langle\iota\rangle$ rel the branch points, i.e., some braid, taking \bar{a} to \bar{b} . Lifting that element gives an element in the hyperelliptic mapping class group taking a to b ; hence, its induced action on homology must take α to β . Thus, $\bar{\Gamma}_2$ acts transitively on the set of α with fixed $b(\alpha)$. \square

10.3.2 Orbit sizes

Proof of Proposition 10.6. When $k = 1$, we are simply counting the number of primitive α with fixed $\delta(\alpha)$ and $b(\alpha)$. Using Definition 10.3, the number $b(\alpha)$ is either the number of e_j so that $(\alpha, e_j) \neq 0$ or Br minus that number. To count the number of α with fixed $b(\alpha)$, then, we simply need to count the number of ways to have nonzero pairing with either $b(\alpha)$ or $Br - b(\alpha)$ of the e_j , which is

$$\binom{Br-1}{b(\alpha)} + \binom{Br-1}{Br-b(\alpha)} = \binom{Br-1}{b(\alpha)} + \binom{Br-1}{b(\alpha)-1} = \binom{Br}{b(\alpha)}.$$

In the case when $Br = 2g + 2$ and $b(\alpha) = g + 1$, we observe that we must also multiply by a factor of $1/2$: this is because choosing $g + 1$ elements also picks out $g + 1$ elements in the complement, and if we have nonzero intersection with the chosen set then we have zero intersection with each e_j in the complement. So since $b(\alpha)$ is defined in terms of the minimum of the 0's and 1's in the vector $\underline{e}(\alpha)$, we see in this case that $\binom{Br}{b(\alpha)}$ is an overcount by a factor of 2.

To pass from the mod 2 to the mod 2^k case, we just observe there are $2^{2g(k-1)} \pmod{2^k}$ classes with given mod 2 reduction. \square

10.3.3 Orbit statistics

Proof of Proposition 10.7. By Proposition B.2, $\bar{\Gamma}$ acts transitively on the homology classes of core curves of cylinders. We are therefore free to choose any convenient representative: compare Corollary B.3. Without loss of generality, we will therefore assume that the curve $c = c_1$ (which represents the homology class e_1) is a cylinder.

As in the proof of Proposition 10.6, we compute first for $k = 1$. We want to count the number of ways to have nonzero pairing with either $b(\alpha)$ or $Br - b(\alpha)$ of the e_j , constrained so that the intersection with c has fixed parity. In the case that $\alpha(c) = 0 \pmod{2}$, then among the other $Br - 2$ of the e_j 's, we count that there are

$$\binom{Br - 2}{b(\alpha)} + \binom{Br - 2}{Br - b(\alpha)} = \binom{Br - 2}{b(\alpha)} + \binom{Br - 2}{b(\alpha) - 2}$$

ways for this to happen. If we assume $\alpha(c) = 1 \pmod{2}$, then there are

$$\binom{Br - 2}{b(\alpha) - 1} + \binom{Br - 2}{Br - b(\alpha) - 1} = 2 \binom{Br - 2}{b(\alpha) - 1}$$

possible ways to distribute the remainder of the intersections. As before, if $Br = 2g + 2$ and $b(\alpha) = g + 1$, then we must also introduce an extra factor of $1/2$ due to the fact that $b(\alpha) = Br - b(\alpha)$.

To get the final statement for arbitrary k , we note that the number of mod- 2^k classes that have specified intersection with c and that lie over a given mod-2 class is exactly $2^{(2g-1)(k-1)}$. \square

Remark. Both Propositions 10.6 and 10.7 can also be proven by counting the number of ways to partition the branch points into two sets, perhaps subject to the constraint that the arc ε_1 connects zeros in different parts. We leave it as an exercise to the reader to reprove both statements in this language.

Part III

Computing the Siegel–Veech constants

Outline of Part III (Sections 11-14).

- In Section 11, we collect all of the necessary background on Siegel–Veech transforms and constants.
- In Section 12, we compute the Siegel–Veech constants in the main formula in terms of “orbit quotient functions” under the assumption that these functions satisfy several number theoretic properties.
- In Section 13, we use the results of Part II to derive explicit formulas for the orbit quotient functions in all cases, thereby showing that all of the number theoretic assumptions are satisfied.
- In Section 14, we present the main formula for the ratios of Siegel–Veech constants with a guide to all of the notation involved. This section also combines all of the previous results to prove Theorem 14.1 and its corollaries from the Introduction.

11 Siegel–Veech preliminaries

11.1 The Siegel–Veech transform

We begin in the most general context before specifying to strata. Let \mathcal{M}_A be an $\mathrm{SL}_2(\mathbb{R})$ -invariant subvariety of area A translation surfaces and let ν_A denote the $\mathrm{SL}_2(\mathbb{R})$ -invariant measure on \mathcal{M}_A induced from the affine

measure on (the relevant linear subspaces of) period coordinates (this is unique up to scale). The starting point for the Siegel–Veech transform is a choice of *Siegel–Veech measure*, an $\mathrm{SL}_2(\mathbb{R})$ -equivariant map

$$\mathcal{V} : \mathcal{M}_A \rightarrow \{\text{measures on } \mathbb{R}^2 \text{ with discrete support}\}$$

satisfying a certain list of further axioms (see [EM01, §2] or [Vor05]). For example, if one is interested in counts of cylinders, or counts of cylinders weighted by area^s , one may consider the Siegel–Veech measures

$$\mathcal{V}_{cyl} : (X, \omega) \mapsto \sum_{\text{cylinders}} \delta_{\text{hol}(cyl)} \quad \text{and} \quad \mathcal{V}_{\text{area}^s} : (X, \omega) \mapsto \sum_{\text{cylinders}} \frac{\text{area}^s(cyl)}{\text{area}^s(X, \omega)} \delta_{\text{hol}(cyl)}, \quad (14)$$

respectively, where the sums are taken over all cylinders on (X, ω) . Here δ denotes a point mass. Observe that $\mathcal{V}_{cyl} = \mathcal{V}_{\text{area}^0}$.

Given a function $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ with compact support, the *Siegel–Veech transform* of f with respect to \mathcal{V} is a function $\hat{f} : \mathcal{M}_A \rightarrow \mathbb{R}$ defined by

$$\hat{f}(X, \omega) := \int_{\mathbb{R}^2} f(\mathbf{v}) d\mathcal{V}(X, \omega)(\mathbf{v}).$$

In the case that \mathcal{V} does not come with weights (i.e., is a sum of *unit* point masses), this reduces to a sum of the values of f over the support of $\mathcal{V}(X, \omega)$, which is how the Siegel–Veech transform is most usually written.

The *Siegel–Veech formula* [Vee98] relates the integral of the Siegel–Veech transform of a function over \mathcal{M}_A to the integral of the original function:

$$\frac{1}{\nu_A(\mathcal{M}_A)} \int_{\mathcal{M}_A} \hat{f}(X, \omega) d\nu_A = c_{\mathcal{V}}(\mathcal{M}_A) \int_{\mathbb{R}^2} f(\mathbf{v}) d\text{Leb}_{\mathbb{R}^2}(\mathbf{v}),$$

where $c_{\mathcal{V}}(\mathcal{M}_A)$ is a constant depending only on the choice of Siegel–Veech measure \mathcal{V} and not the function f . This is the *Siegel–Veech constant* associated to \mathcal{V} . The main theorem of [EM01] shows that this constant captures the asymptotics of the corresponding weighted count for ν_A -almost every $(X, \omega) \in \mathcal{M}_A$:

$$c_{\mathcal{V}}(\mathcal{M}_A) = \lim_{r \rightarrow \infty} \frac{1}{\pi r^2} \int_{\mathbb{R}^2} \chi_r(\mathbf{v}) d\mathcal{V}(X, \omega)(\mathbf{v}), \quad (15)$$

where χ_r is the indicator function of the ball of radius r about 0. For example, in the cases of (14), the associated constants give the asymptotic count (for generic surfaces) of cylinders whose circumference is at most r (the *cylinder Siegel–Veech constant* c_{cyl}), and the same count weighted by the area^s of the cylinders (the *area^s-Siegel–Veech constant* c_{area^s}), respectively. In particular, the counting functions from the introduction can be expressed as

$$N_{cyl}((X, \omega), L) = \int_{\mathbb{R}^2} \chi_L(\mathbf{v}) d\mathcal{V}_{cyl}(X, \omega)(\mathbf{v}) \quad \text{and} \quad N_{\text{area}^s}((X, \omega), L) = \int_{\mathbb{R}^2} \chi_L(\mathbf{v}) d\mathcal{V}_{\text{area}^s}(X, \omega)(\mathbf{v}). \quad (16)$$

Linearity. Since Siegel–Veech constants $c_{\mathcal{V}}$ are defined in terms of integrals, they behave nicely under linear combinations. In particular, if $\{\mathcal{V}_m\}$ are different maps on \mathcal{M}_A with respect to which we can perform a Siegel–Veech transform, then

$$c_{\mathcal{V}} = \sum a_m c_{\mathcal{V}_m}, \quad \text{where } \mathcal{V} = \sum a_m \mathcal{V}_m \text{ for } a_m \in \mathbb{R}. \quad (17)$$

11.2 Cusp volumes and configurations

We now specialize to strata and indicate how this discussion generalizes to invariant subvarieties. Fix $\varepsilon > 0$ and let χ_ε be the characteristic function of the disc of radius ε centered at the origin in \mathbb{R}^2 ; then the Siegel–Veech formula states that

$$c_{cyl}(\mathcal{H}_{1,lab}) = \frac{1}{\pi \varepsilon^2} \frac{1}{\nu_1(\mathcal{H}_{1,lab})} \int_{\mathcal{H}_{1,lab}} \hat{\chi}_\varepsilon(X, \omega) d\nu_1. \quad (18)$$

where $\widehat{\chi}_\varepsilon$ is the Siegel–Veech transform with respect to \mathcal{V}_{cyl} . A key insight of [EMZ03] is that while it is not possible to compute this exactly for any particular value of ε , we can instead take a limit where ε tends to zero and consider the asymptotics of the integral.

Define the ε -cylinder cusp $\mathcal{H}_{1,lab}^\varepsilon$ to be the set of all (X, ω) that contain a cylinder whose holonomy has length at most ε . Then so long as ε is small enough, the Siegel–Veech transform $\widehat{\chi}_\varepsilon$ is supported entirely on $\mathcal{H}_{1,lab}^\varepsilon$. Moreover, $\widehat{\chi}_\varepsilon$ is identically 1 on most of the cusp. Following [EMZ03], for fixed $\varepsilon' = O(\varepsilon)$, we define the (ε') -thin part of the ε -cusp

$$\mathcal{H}_{1,lab}^{\varepsilon,thin} \subset \mathcal{H}_{1,lab}^\varepsilon$$

to be the set of surfaces which have a saddle connection of length $\leq \varepsilon'$ that is not homologous to the core curve of the ε -thin cylinder; note that this also rules out having any other non-homologous ε -thin cylinders. The precise value of ε' will not particularly matter, and it will be convenient to vary at will (always at the scale of ε). The (ε') -thick part of the ε -cusp is the complement of the thin part.

One can also define a decomposition of the cusp of arbitrary invariant subvarieties: in this case, the thin part of the ε -cylinder cusp consists of those surfaces which have an ε' -short saddle connection that does not remain parallel to the ε -short cylinders as one deforms in a small neighborhood in \mathcal{M} (compare the notion of “ \mathcal{M} -parallel” in [Wri15]).

As recorded in [EMZ03, Corollary 7.2], work of [MS91] implies that the volume of the cusp is concentrated in the thick part (and a similar result is true for invariant subvarieties, see [AMY13]):

$$\nu_1 \left(\mathcal{H}_{1,lab}^{\varepsilon,thin} \right) = o(\varepsilon^2).$$

In [EMZ03, Lemma 7.3] it is shown that the integral of $\widehat{\chi}_\varepsilon$ over the thin part is also negligible:

$$\int_{\mathcal{H}_{1,lab}^{\varepsilon,thin}} \widehat{\chi}_\varepsilon d\nu_1 = o(\varepsilon^2).$$

This fact also holds for area^s-weighted Siegel–Veech transforms, and corresponding results also hold for arbitrary invariant subvarieties.

Combining these two estimates and the fact that $\widehat{\chi}_\varepsilon$ is identically 1 on the thick part, one gets that

$$\int_{\mathcal{H}_{1,lab}} \widehat{\chi}_\varepsilon d\nu_1 = \int_{\mathcal{H}_{1,lab}^\varepsilon} \widehat{\chi}_\varepsilon d\nu_1 = \int_{\mathcal{H}_{1,lab}^{\varepsilon,thick}} d\nu_1 + o(\varepsilon^2). \quad (19)$$

Plugging (19) back into (18) and taking the limit as ε tends to zero, one recovers [EMZ03, Proposition 3.3]:

$$c(\mathcal{H}_{1,lab}) = \lim_{\varepsilon \rightarrow 0} \frac{1}{\pi \varepsilon^2} \frac{\nu_1(\mathcal{H}_{1,lab}^\varepsilon)}{\nu_1(\mathcal{H}_{1,lab})}.$$

Configurations. Of course, in order for the scheme above to be useful, one needs to be able to compute the volume of the cusp. Moreover, to compute area^s Siegel–Veech constants, one must also be able to integrate other functions over these domains. A large portion of the hard, technical work of [EMZ03] is to give local models for (the thick parts of) cusps that facilitate this integration; see that paper for a much more thorough discussion of the concepts we sketch below.

Any surface (X, ω) in the thick part $\mathcal{H}_{1,lab}^{\varepsilon,thick}$ of the ε -cusp contains a unique set of short homologous cylinders. Cutting out those cylinders, one is left with a subsurface with $(\varepsilon$ -short) geodesic boundary. The remaining boundary components can then be contracted to points to yield a (possibly disconnected) translation surface in the “principal boundary.” While this contraction process is generally not canonical, different choices result in only slightly different surfaces in the boundary, allowing one to model (a large measure subset of) $\mathcal{H}_{1,lab}^{\varepsilon,thick}$ by surfaces in a boundary stratum together with data of how to glue the cylinders back in. Making this picture precise takes a great deal of work and some very involved computations; see [EMZ03, Parts I and II].

What is important for us is that different components of the thick part of the cusp correspond to different boundary strata and different ways of gluing in cylinders. All these data are encapsulated in the notion of a *configuration* \mathcal{C} [EMZ03, §3.2]; we use $\mathcal{H}_{1,lab}^{\varepsilon,thick}(\mathcal{C})$ to denote the component of the thick part of the cusp where the cylinders in the configuration are thin. In particular, we can partition the set of cylinders on (X, ω) by the combinatorial data of their configuration. We can then introduce Siegel–Veech measures

$$\mathcal{V}_{\mathcal{C}} : (X, \omega) \mapsto \sum_{\substack{\text{cylinders in} \\ \text{configuration } \mathcal{C}}} \delta_{\text{hol}(cyl)}$$

as well as the corresponding area^s weighted versions $\mathcal{V}_{\mathcal{C}}^s$ and their respective Siegel–Veech transforms and constants. By definition and (19), given any $f \in C_c(\mathbb{R}^2)$, we have that

$$\int_{\mathcal{H}_{1,lab}^{\varepsilon}} \widehat{f}^{\mathcal{C}} d\nu_1 = \int_{\mathcal{H}_{1,lab}^{\varepsilon,thick}(\mathcal{C})} \widehat{f} d\nu_1 + o(\varepsilon^2),$$

where the Siegel–Veech transforms on the left and right are taken with respect to $\mathcal{V}_{\mathcal{C}}$ and \mathcal{V}_{cyl} , respectively. A similar equation also holds for the area^s versions. In light of linearity (17), we have that

$$c_{cyl}(\mathcal{H}_{1,lab}) = \sum_{\mathcal{C}} c_{\mathcal{C}}(\mathcal{H}_{1,lab}) \quad \text{and} \quad c_{\text{area}^s}(\mathcal{H}_{1,lab}) = \sum_{\mathcal{C}} c_{\mathcal{C}}^s(\mathcal{H}_{1,lab}),$$

a fact that we will use many times below.

12 From integrals to orbit quotients

In this section we prove a preliminary result, Theorem 12.5, which facilitates the computation of the Siegel–Veech constants in Theorem 14.1 under certain number theoretic assumptions. In Section 13 below, we verify that the requisite assumptions are always satisfied using the explicit formulas computed in Part II.

To state the theorem, we define two concepts from elementary number theory for which we are unaware of standard terminology, namely, prime power independent functions and D -GCD dependent functions.

Definition 12.1. The *radical of an integer* $\text{rad}(N)$ is its largest square-free divisor.

Definition 12.2. We say that a multiplicative function $f : \mathbb{N} \rightarrow \mathbb{N}$ is *prime power independent*, if $f(N) = f(\text{rad}(N))$ for all N .

Example. A well known example of a non-trivial prime power independent function is $\Phi(N)/N$.

Definition 12.3. A function $f : \mathbb{N} \rightarrow \mathbb{N}$ is *D -GCD dependent*, if $f(N) = f(\text{gcd}(N, D))$ for all N .

Definition 12.4. Given the group $\overline{\Gamma}_D$ defined in Section 3, a branching vector α , and a number N , define the *orbit quotient function*

$$f_{\overline{\Gamma}_D \alpha}(N) = \frac{|(\overline{\Gamma}_D \alpha)_N|}{|\overline{\Gamma}_D \alpha|}.$$

We abuse notation and write $f_D(N)$ when the group and branching vector are understood.

Recall from the Introduction that the *monodromy cofactor* $d_{abs} \in \mathbb{N}$ is the maximal factor of d coprime to $\text{gcd}(\delta(\alpha))$ and $d_{rel} = d/d_{abs}$. We further factor $d = d'_{abs} d'_{rel} 2^k$ so that $d'_{abs} | d_{abs}$ and $d'_{rel} | d_{rel}$ are odd.

Theorem 12.5. *Let $s \geq 0$. Let $\mathcal{M}_{1,\delta}(\alpha)$ be the component of a locus of branched cyclic covers over \mathcal{H}_1 specified by the branching vector α , normalized to have unit area. Suppose that for all $D|d$,*

- f_D is D -GCD dependent,
- f_D can be expressed as a constant multiplied by a prime power independent function,

- $f_{d'_{abs}}$ is a constant, which we denote by $C_{d'_{abs}}$, and
- If the orbit of α is unsplit and $2|d_{abs}$, then $f_{2^k}(a)$ is a constant function.

With these assumptions, then we have

$$\frac{c_{area^s}(\mathcal{M}_{1,\delta}(\alpha))}{c_{area^s}(\mathcal{H}_1)} = \frac{1}{d} \left(\sum_{\mathfrak{d}|d'_{rel}} f_{d'_{rel}}(\mathfrak{d}) \Phi\left(\frac{d'_{rel}}{\mathfrak{d}}\right) \mathfrak{d}^{3-s} \right) \left(C_{d'_{abs}} \sum_{\mathfrak{d}|d'_{abs}} \Phi\left(\frac{d'_{abs}}{\mathfrak{d}}\right) \mathfrak{d}^{3-s} \right) \left(\sum_{\mathfrak{d}|2^k} f_{2^k}(\mathfrak{d}) \Phi\left(\frac{2^k}{\mathfrak{d}}\right) \mathfrak{d}^{3-s} \right).$$

12.1 Siegel–Veech constants of cyclic covers

If a translation surface does not cover a translation surface of lower genus, then the integrals involved in evaluating its Siegel–Veech constants must be performed directly on the orbit closure that contains the translation surface. However, there is a convenient shortcut one can take when considering covers. Rather than evaluate integrals in the space of the cover, a counting function can be defined to capture the finite combinatorial information of the covering; computing the Siegel–Veech constants for a locus of covers therefore reduces to evaluating an integral on the subvariety containing the *base* translation surface. This concept was implemented in [EMS03]. Later it was used to compute area–Siegel–Veech constants of all square-tiled surfaces in [EKZ14, §10], where the base surface is a once punctured torus. Subsequently, this strategy was used in [CMZ18, Part I], where the base was an n -punctured torus. Here we generalize this technique to base surfaces of arbitrary genus, but restrict to cyclic covers.

Remark. Most of the manipulations in this section work for arbitrary G -covers, replacing the numerical quantities with numbers and sizes of cosets, but in order to simplify the formula one needs to understand the $\pi_1(\mathcal{H}_{1,lab})$ action on $\text{Hom}(\pi_1(X \setminus \mathcal{B}), G)$, which is generally a very hard problem.

Defining and decomposing the Siegel–Veech constant. Following the strategy above, we define a counting function on our base translation surface adapted to the counting problem on the cover. We begin with our locus of (area d) cyclic covers $\mathcal{M}_{d,\delta}(\alpha)$ defined above covering a stratum component $\mathcal{H}_{1,lab}$. From the area ^{s} Siegel–Veech formula for $\mathcal{M}_{d,\delta}(\alpha)$, we can compute the area ^{s} Siegel–Veech constant of $\mathcal{M}_{d,\delta}(\alpha)$ by integrating the $\mathcal{V}_{area^s}^s$ Siegel–Veech transform of the indicator function χ_ε of a ball of radius ε :

$$c_{area^s}(\mathcal{M}_{d,\delta}(\alpha)) = \frac{1}{\pi\varepsilon^2} \frac{1}{\nu_d(\mathcal{M}_{d,\delta}(\alpha))} \int_{\mathcal{M}_{d,\delta}(\alpha)} \widehat{\chi}_\varepsilon^s(\tilde{X}, \tilde{\omega}) d\nu_d.$$

It is standard to compute Siegel–Veech constants for unit area translation surfaces. In [EKZ14, §10], it was observed that for any counting problem with quadratic asymptotics and any scaling factor $\lambda > 0$,

$$N((\tilde{X}, \lambda\tilde{\omega}), L) = N\left((\tilde{X}, \tilde{\omega}), \frac{L}{\lambda}\right) = \frac{1}{\lambda^2} N((\tilde{X}, \tilde{\omega}), L).$$

Therefore, in order to normalize to unit area, we multiply the count by $d = \lambda^2$ because if the base surface (X, ω) has unit area, then the covering surface has area d . This yields

$$c_{area^s}(\mathcal{M}_{1,\delta}(\alpha)) = d \cdot \frac{1}{\pi\varepsilon^2} \frac{1}{\nu_d(\mathcal{M}_{d,\delta}(\alpha))} \int_{\mathcal{M}_{d,\delta}(\alpha)} \widehat{\chi}_\varepsilon^s(\tilde{X}, \tilde{\omega}) d\nu_d.$$

We highlight that the integral is still taken over the locus of covers of area d .

As in (18), the integral above is not computable, but we can compute its asymptotics as ε tends to zero. By definition, the Siegel–Veech transform $\widehat{\chi}_\varepsilon^s$ is supported entirely on the ε -cusp of $\mathcal{M}_{d,\delta}(\alpha)$, and by (19), we can restrict to the thick part of the cusp at the cost of introducing an error term of order $o(\varepsilon^2)$.

Suppose that $(\tilde{X}, \tilde{\omega})$ is in the thick part of the ε -cusp of $\mathcal{M}_{d,\delta}(\alpha)$. The base surface (X, ω) must then lie in the thick part of the ε -cusp for the base stratum $\mathcal{H}_{1,lab}$.⁵ Thus, to every component of the thick part of

⁵The precise thickness constants will differ by a constant factor, but since they are absorbed into the $o(\varepsilon^2)$, we will suppress this discussion.

the cusp of $\mathcal{M}_{d,\delta}(\alpha)$, we can associate a configuration of homologous cylinders \mathcal{C} on the base surface (X, ω) . More generally, we can partition the set of holonomy vectors of cylinders on $(\tilde{X}, \tilde{\omega}) \in \mathcal{M}_{d,\delta}(\alpha)$:

$$V(\tilde{X}, \tilde{\omega}) = \bigcup_{\mathcal{C}} V_{\tilde{\mathcal{C}}}(\tilde{X}, \tilde{\omega}),$$

where $V_{\tilde{\mathcal{C}}}$ denotes the set of holonomies of cylinders on $(\tilde{X}, \tilde{\omega})$ projecting down to a configuration \mathcal{C} on (X, ω) . Note, this is why the index set ranges over configurations on the base surface.

Define area^s Siegel–Veech measures associated to each of these sets

$$\mathcal{V}_{\tilde{\mathcal{C}}}^s : (\tilde{X}, \tilde{\omega}) \mapsto \sum_{\mathbf{v} \in V_{\tilde{\mathcal{C}}}(\tilde{X}, \tilde{\omega})} \left(\sum_{\{\text{hol}(\text{cyl})=\mathbf{v}\}} \frac{\text{area}^s(\text{cyl})}{\text{area}^s(\tilde{X}, \tilde{\omega})} \right) \delta_{\mathbf{v}},$$

and let $\widehat{f}^{s,\tilde{\mathcal{C}}}$ denote the Siegel–Veech transform of a function $f \in C_c(\mathbb{R}^2)$ with respect to $\mathcal{V}_{\tilde{\mathcal{C}}}^s$. Thus, by linearity of the Siegel–Veech transform (17),

$$c_{\text{area}^s}(\mathcal{M}_{1,\delta}(\alpha)) = d \cdot \frac{1}{\pi \varepsilon^2} \frac{1}{\nu_d(\mathcal{M}_{d,\delta}(\alpha))} \sum_{\mathcal{C}} \int_{\mathcal{M}_{d,\delta}(\alpha)} \widehat{(\chi_\varepsilon)}^{s,\tilde{\mathcal{C}}}(\tilde{X}, \tilde{\omega}) d\nu_d. \quad (20)$$

12.2 Explicitly expressing the counting function

Next, we convert each integral in (20) over (the thick part of the cusp of) $\mathcal{M}_{d,\delta}(\alpha)$ to an integral over $\mathcal{H}_{1,lab}$. This introduces a new subtlety: even if (X, ω) is in the ε -cusp of $\mathcal{H}_{1,lab}$, the covering surface $(\tilde{X}, \tilde{\omega})$ may not lie in the ε -cusp of $\mathcal{M}_{d,\delta}(\alpha)$, and even if it does, the preimages of homologous ε -thin cylinders on (X, ω) can have different widths on $(\tilde{X}, \tilde{\omega})$.

Recall that the fibers of $\mathcal{M}_{d,\delta}(\alpha) \rightarrow \mathcal{H}_{1,lab}$ are in bijection with the monodromy orbit $\bar{\Gamma}_d(\alpha)$, so that a particular cover $(\tilde{X}, \tilde{\omega}) \rightarrow (X, \omega)$ corresponds with a branching vector β . Suppose a cylinder C on (X, ω) has holonomy vector \mathbf{v} . Set w_C to be the order of $\beta(C) \in \mathbb{Z}/d\mathbb{Z}$. We recall from Lemma 2.4 that the preimage of C in $(\tilde{X}, \tilde{\omega})$ consists of

$$\#(C) := d/w_C = \gcd(\beta(C), d)$$

many cylinders, each of period $w_C \mathbf{v}$. These quantities are locally constant as we deform $(\tilde{X}, \tilde{\omega})$, but implicitly they required a marking of the cylinder C and an identification of the fiber over (X, ω) with (co)homology, and are thus not well-defined once we deform outside a small neighborhood.

All the same, given a configuration \mathcal{C} of homologous cylinders on (X, ω) and a branching vector $\beta \in H_1(X \setminus \mathcal{B}; \mathbb{Z}/d\mathbb{Z})$, we define a function $\Psi_{s,\varepsilon,\mathcal{C},\beta} : \mathbb{R}^2 \rightarrow \mathbb{R}$ via

$$\Psi_{s,\varepsilon,\mathcal{C},\beta}(\mathbf{v}) = \frac{1}{m(\mathcal{C})} \sum_{C \in \mathcal{C}} (\#C)^{1-s} \chi_\varepsilon(w_C \mathbf{v}),$$

where $m(\mathcal{C})$ is the number of cylinders in \mathcal{C} and χ_ε is the indicator function of a ball of radius ε . Again, since this function implicitly depends on a local choice of marking of the cylinders of \mathcal{C} , it is not well-defined after one moves far away in $\mathcal{H}_{1,lab}$. However, the *sum* over β of all of these functions is monodromy invariant, and thus well-defined.

The purpose of defining these functions is that they capture exactly how many ε -thin cylinders on the different covers $(\tilde{X}, \tilde{\omega}) \rightarrow (X, \omega)$ the ε -thin cylinders on (X, ω) lift to. This will allow us to transform the integrals of (20) into integrals involving $\Psi_{s,\varepsilon,\mathcal{C},\beta}$ on the base space $\mathcal{H}_{1,lab}$.

Proposition 12.6. *For any component of any stratum of holomorphic abelian differentials, and for any $s \geq 0$, $\varepsilon > 0$, and configuration \mathcal{C} ,*

$$\int_{\mathcal{M}_{d,\delta}(\alpha)} \widehat{(\chi_\varepsilon)}^{s,\tilde{\mathcal{C}}}(\tilde{X}, \tilde{\omega}) d\nu_d = \int_{\mathcal{H}_{1,lab}} \sum_{\beta \in \bar{\Gamma}_d(\alpha)} \widehat{\Psi}_{s,\varepsilon,\mathcal{C},\beta}(X, \omega) d\nu_1 + o(\varepsilon^2),$$

where the Siegel–Veech transform on the left-hand side is taken with respect to \mathcal{V}_C^s , and the Siegel–Veech transforms on the right-hand side are taken with respect to \mathcal{V}_C^s .

The key point of the proof is to expand out the left-hand integral using period coordinates, manipulate some of the interior terms to separate variables, then collapse the integral back down. The next lemma is crucial for the middle step and is reminiscent of one of the key steps in [CMZ18, Proof of Theorem 3.1].

Lemma 12.7. *For any $L > 0$, define*

$$\Delta_{\leq L} := \left\{ (h_1, \dots, h_m) \in \mathbb{R}_+^m \mid \sum h_i \leq L \right\}.$$

Then for any $s \geq 0$, the integral $\int_{\Delta_{\leq L}} h_i^s dh_1 \dots dh_m$ is independent of i .

Proof. This follows from the fact that the Jacobian of a permutation matrix is 1. \square

Proof of Proposition 12.6. By definition, the function $(\widehat{\chi_\varepsilon})^{s, \tilde{\mathcal{C}}}(\tilde{X}, \tilde{\omega})$ is only nonzero for those $(\tilde{X}, \tilde{\omega})$ such that the base surface (X, ω) has a collection of ε -thin homologous cylinders in configuration \mathcal{C} . Thus, after restricting to thick parts (which introduces an error term of $o(\varepsilon^2)$ by (19)), it is supported only on those components of $\mathcal{M}_{d, \delta}(\alpha)^{\varepsilon, \text{thick}}$ covering the corresponding thick part of the cusp $\mathcal{H}_{1, \text{lab}}^{\varepsilon, \text{thick}}(\mathcal{C})$.

For $(\tilde{X}, \tilde{\omega})$ in such a thick part we can write out the Siegel–Veech transform explicitly. Fix a marking of the base surface, hence an identification of $(\tilde{X}, \tilde{\omega})$ with a cohomology class β . Then if we let h_C denote the height of cylinder $C \in \mathcal{C}$ on (X, ω) and \mathbf{v} denote its holonomy, we write

$$\begin{aligned} (\widehat{\chi_\varepsilon})^{s, \tilde{\mathcal{C}}}(\tilde{X}, \tilde{\omega}) &= \sum_{\tilde{C} \rightarrow C \in \mathcal{C}} \frac{\text{area}^s(\tilde{C})}{\text{area}^s(\tilde{X}, \tilde{\omega})} \chi_\varepsilon(\text{hol}(\tilde{C})) \\ &= \sum_{C \in \mathcal{C}} \#(C) \frac{h_C^s w_C^s |\mathbf{v}|^s}{d^s \text{area}^s(X, \omega)} \chi_\varepsilon(w_C \mathbf{v}) \\ &= \sum_{C \in \mathcal{C}} \#(C)^{1-s} \frac{h_C^s |\mathbf{v}|^s}{\text{area}^s(X, \omega)} \chi_\varepsilon(w_C \mathbf{v}), \end{aligned}$$

where the first sum is the sum over all cylinders \tilde{C} of $(\tilde{X}, \tilde{\omega})$ covering some cylinder C of \mathcal{C} .

Since the covering map $\mathcal{P} : \mathcal{M}_{d, \delta}(\alpha) \rightarrow \mathcal{H}_{1, \text{lab}}$ pushes Masur–Veech measure to Masur–Veech measure, we can push our integral down by (restricting to the thick part of the cusp and) summing over fibers:

$$\begin{aligned} \int_{\mathcal{M}_{d, \delta}(\alpha)} (\widehat{\chi_\varepsilon})^{s, \tilde{\mathcal{C}}}(\tilde{X}, \tilde{\omega}) d\nu_d &= \int_{\mathcal{H}_{1, \text{lab}}^{\varepsilon, \text{thick}}(\mathcal{C})} \sum_{\mathcal{P}^{-1}(X, \omega)} (\widehat{\chi_\varepsilon})^{s, \tilde{\mathcal{C}}}(\tilde{X}, \tilde{\omega}) d\nu_1 + o(\varepsilon^2) \\ &= \int_{\mathcal{H}_{1, \text{lab}}^{\varepsilon, \text{thick}}(\mathcal{C})} \sum_{\beta \in \bar{\Gamma}_d(\alpha)} \sum_{C \in \mathcal{C}} \#(C)^{1-s} \frac{h_C^s |\mathbf{v}|^s}{\text{area}^s(X, \omega)} \chi_\varepsilon(w_C \mathbf{v}) d\nu_1 + o(\varepsilon^2). \end{aligned} \quad (21)$$

We note once more that individual expressions involving $\#(C)$ and w_C depend on a choice of marking, but since we are summing over all β (equivalently, over the entire fiber of \mathcal{P}) the total expression inside the integral is marking-independent.

The integral in (21) now looks almost exactly like the Siegel–Veech transform of $\sum_{\beta} \Psi_{s, \varepsilon, \mathcal{C}, \beta}$, as desired, except the area^s of cylinder C is intertwined with the combinatorial factors. To deal with this, we compute in period coordinates and use the independence established in Lemma 12.7.

We first recall that by definition, the measure ν_1 is computed by taking the Masur–Veech measure of its cone: $\nu_1(U) = \dim_{\mathbb{R}} \mathcal{H} \cdot \nu\{tU \mid t \in [0, 1]\}$. Below, we will use $\mathcal{H}_{\leq 1, \text{lab}}^{\varepsilon, \text{thick}}(\mathcal{C})$ to denote the cone on $\mathcal{H}_{1, \text{lab}}^{\varepsilon, \text{thick}}(\mathcal{C})$, so by the volume of a cone with unit height, we have

$$\frac{1}{\dim_{\mathbb{R}} \mathcal{H}} \int_{\mathcal{H}_{\leq 1, \text{lab}}^{\varepsilon, \text{thick}}(\mathcal{C})} f(X, \omega) d\nu_1 = \int_{\mathcal{H}_{\leq 1, \text{lab}}^{\varepsilon, \text{thick}}(\mathcal{C})} f(X, \omega) d\nu$$

for every area-independent function f (in particular, for the integrand in (21)).

By definition, every surface $(X, \omega) \in \mathcal{H}_{1,lab}^{\varepsilon,thick}(\mathcal{C})$ has a unique collection of ε -thin homologous cylinders in configuration \mathcal{C} . We may therefore choose a finite collection of period coordinate charts covering the cone $\mathcal{H}_{\leq 1,lab}^{\varepsilon,thick}(\mathcal{C})$, each of which corresponds to taking the holonomies of a basis for $H_1(X, \mathcal{B}; \mathbb{Z})$ containing the class of the $C \in \mathcal{C}$ and transversal saddle connections T_C living in each C and connecting singularities in opposite boundaries. Said another way, we can fix a local family of charts $\varphi : \mathbb{C}^N \rightarrow \mathcal{H}_{lab}$ such that $hol(C)$ and $\{hol(T_C)\}_{C \in \mathcal{C}}$ are given by the first $|\mathcal{C}| + 1$ coordinates in each chart.

With respect to each such chart, let $\mathbf{v} = hol(C)$ and decompose each $hol(T_C)$ into a vector t_C parallel to \mathbf{v} (representing twisting about C) and a vector h_C orthogonal to \mathbf{v} (representing the height of C). Since Masur–Veech measure is just Lebesgue measure in period coordinates, we can therefore disintegrate

$$d\nu = \varphi_* \prod_{C \in \mathcal{C}} (dh_C dt_C) d\mathbf{v} d(\text{rest}),$$

where “rest” denotes the (complex) holonomies of the other saddle connections forming a homology basis. Note that when the genus of the base surface is 1 then “rest” is empty.

It will also be important to understand the structure of each chart. Fixing all of the holonomies except for those of T_C specifies a translation structure with boundary on the complement of the cylinders of \mathcal{C} . We can glue in cylinders of arbitrary height and twisting to yield a translation surface in the correct stratum, and so we can extend our chart φ to the product of half-planes $\{h_C > 0, t_C \in \mathbb{R}\}$ with some set U in the orthogonal complement $\text{span}_{\mathbb{C}}(\mathbf{v}, \text{rest})$ (note that increasing t_C too much does a full twist about C , hence the extended map is no longer injective).

The most important part of this discussion for us is how the area behaves. Fixing $(\mathbf{v}, \text{rest}) \in U$ defines a structure of some area A on the complement of \mathcal{C} . The total area is then given by $A + |\mathbf{v}| \sum h_C$, so we get a surface of area ≤ 1 if and only if $\sum h_C \leq (1 - A)/|\mathbf{v}|$.

By passing to the cone and disintegrating in period coordinates as above, the integral in (21) becomes

$$\int_{\mathcal{H}_{1,lab}^{\varepsilon,thick}(\mathcal{C})} \dots d\nu = \dim_{\mathbb{R}} \mathcal{H} \sum_{\text{charts}} \int_U \int_0^{|\mathbf{v}|} \dots \int_0^{|\mathbf{v}|} \underbrace{\int_{\Delta_{\leq (1-A)/|\mathbf{v}|}} \dots \prod_{C \in \mathcal{C}} dh_C \prod_{C \in \mathcal{C}} dt_C d\mathbf{v} d(\text{rest})}_{\text{I}}, \quad (22)$$

(where we have suppressed the integrand for space). Plugging the integrand back in, we notice that most of the terms can be pulled outside the innermost integral.

$$\begin{aligned} \text{I} &= \int_{\Delta_{\leq (1-A)/|\mathbf{v}|}} \sum_{\beta \in \overline{\Gamma}_d(\boldsymbol{\alpha})} \sum_{C \in \mathcal{C}} \#(C)^{1-s} \frac{h_C^s |\mathbf{v}|^s}{\text{area}^s(X, \omega)} \chi_{\varepsilon}(w_C \mathbf{v}) \prod_{C \in \mathcal{C}} dh_C \\ &= \sum_{\beta \in \overline{\Gamma}_d(\boldsymbol{\alpha})} \sum_{C \in \mathcal{C}} \#(C)^{1-s} \frac{|\mathbf{v}|^s}{\text{area}^s(X, \omega)} \chi_{\varepsilon}(w_C \mathbf{v}) \underbrace{\int_{\Delta_{\leq (1-A)/|\mathbf{v}|}} h_C^s \prod_{C \in \mathcal{C}} dh_C}_{\text{II}_C} \end{aligned}$$

Lemma 12.7 ensures that the integrals II_C are independent of our choice of $C \in \mathcal{C}$; let their common value be denoted by II . Then we can factor II out of the sum to yield

$$\begin{aligned} &= \left(\sum_{\beta \in \overline{\Gamma}_d(\boldsymbol{\alpha})} \sum_{C \in \mathcal{C}} \#(C)^{1-s} \chi_{\varepsilon}(w_C \mathbf{v}) \right) \frac{|\mathbf{v}|^s}{\text{area}^s(X, \omega)} \cdot \text{II} \\ &= \left(\sum_{\beta \in \overline{\Gamma}_d(\boldsymbol{\alpha})} \sum_{C \in \mathcal{C}} \#(C)^{1-s} \chi_{\varepsilon}(w_C \mathbf{v}) \right) \cdot \frac{1}{m(\mathcal{C})} \left(\sum_{C \in \mathcal{C}} \text{II}_C \frac{|\mathbf{v}|^s}{\text{area}^s(X, \omega)} \right). \end{aligned}$$

Expanding each Π_C back out, we recognize the area^s Siegel–Veech transform of $\sum_{\beta} \Psi_{s,\varepsilon,C,\beta}$:

$$\begin{aligned} &= \int_{\Delta_{\leq (1-A)/|\mathbf{v}|}} \left(\sum_{\beta \in \overline{\Gamma}_d(\alpha)} \frac{1}{m(\mathcal{C})} \sum_{C \in \mathcal{C}} \#(C)^{1-s} \chi_{\varepsilon}(w_C \mathbf{v}) \right) \cdot \left(\sum_{C \in \mathcal{C}} \frac{h_C^s |\mathbf{v}|^s}{\text{area}^s(X, \omega)} \right) \prod_{C \in \mathcal{C}} dh_C \\ &= \int_{\Delta_{\leq (1-A)/|\mathbf{v}|}} \sum_{\beta \in \overline{\Gamma}_d(\alpha)} \widehat{\Psi}_{s,\varepsilon,C,\beta}(X, \omega) \prod_{C \in \mathcal{C}} dh_C. \end{aligned}$$

Using (21), plugging I back into (22), and collapsing the integral back down, we are left with

$$\int_{\mathcal{M}_{d,\delta}(\alpha)} (\widehat{\chi}_{\varepsilon})^{s,\tilde{\mathcal{C}}}(\tilde{X}, \tilde{\omega}) d\nu_d = \int_{\mathcal{H}_{1,lab}^{\varepsilon,thick}(C)} \sum_{\beta \in \overline{\Gamma}_d(\alpha)} \widehat{\Psi}_{s,\varepsilon,C,\beta}(X, \omega) d\nu_1 + o(\varepsilon^2),$$

which is equal to what we wanted, up to another error of $o(\varepsilon^2)$ coming from the thin part of $\mathcal{H}_{1,lab}$ (19). \square

12.3 Simplifying the counting function

Eventually, we will use the Siegel–Veech transform on $\mathcal{H}_{1,lab}$ to relate the integral from Proposition 12.6 to the integrals of the $\Psi_{s,\varepsilon,C,\beta}$ functions. We record their values below:

Lemma 12.8. *For any $s \geq 0$, $\varepsilon > 0$, configuration \mathcal{C} , and branching vector β ,*

$$\int_{\mathbb{R}^2} \Psi_{s,\varepsilon,C,\beta}(\mathbf{v}) dx dy = \pi \varepsilon^2 \left(\frac{1}{m(\mathcal{C})d^2} \sum_{C \in \mathcal{C}} \gcd(\beta(C), d)^{3-s} \right),$$

where $m(\mathcal{C})$ is the number of cylinders in the configuration \mathcal{C} .

Proof. The function $\Psi_{s,\varepsilon,C,\beta}$ is a sum of indicator functions on balls of radius ε/w_C , each weighted by $\#(C)^{1-s}/m(\mathcal{C})$. The equation follows by recalling that $w_C = d/\#(C)$ and $\#(C) = \gcd(\beta(C), d)$. \square

We now record a lemma that will allow us to further simplify the summation from Lemma 12.8. Recall the definitions of the relevant number theory terms from the beginning of the Section.

Proposition 12.9. *Let $s \geq 0$. Given an integer D , let M be such that $\text{rad}(D) | M$. Let $f : \mathbb{N} \rightarrow \mathbb{Z}$ be a prime power independent function. Then*

$$\sum_{a=1}^D f(\gcd(a, D)) \gcd(a + M, D)^{3-s} = \sum_{\mathfrak{d} | D} f(\mathfrak{d}) \Phi\left(\frac{D}{\mathfrak{d}}\right) \mathfrak{d}^{3-s}.$$

Proof. First observe that from the definition of radical, a prime $p | \text{rad}(M)$ if and only if $p | M$. For each prime p , $p | \gcd(a, D)$ if and only if $p | a$ and $p | D$. Also, if $p | D$, then $p | M$. Hence, if $p | \gcd(a, D)$, then it also divides $a + M$. This combined with $p | D$ implies $p | \gcd(a + M, D)$. If $p | \gcd(a + M, D)$, then $p | a + M$. Since $p | M$, $p | a$ and we conclude $p | \gcd(a, D)$. Therefore, $f(\gcd(a, D)) = f(\gcd(a + M, D))$ because f is prime power independent. This allows us to reindex to get

$$\begin{aligned} \sum_{a=1}^D f(\gcd(a, D)) \gcd(a + M, D)^{3-s} &= \sum_{a=1}^D f(\gcd(a + M, D)) \gcd(a + M, D)^{3-s} \\ &= \sum_{a'=1}^D f(\gcd(a', D)) \gcd(a', D)^{3-s}. \end{aligned}$$

Since $\{a' \mid a' = 1, \dots, D\}$ covers a full set of remainders modulo D , each value of $\gcd(a', D)^{3-s}$ occurs $\Phi(D/\gcd(a', D))$ times to conclude. \square

12.4 Conclusion of the proof of Theorem 12.5

We now have the pieces to put together the proof of the main theorem of this section.

Reduction to a combinatorial formula. For any $\varepsilon > 0$, combining (20) and Proposition 12.6 gives

$$c_{\text{area}^s}(\mathcal{M}_{1,\delta}(\boldsymbol{\alpha})) = d \frac{1}{\pi \varepsilon^2} \frac{1}{\nu_d(\mathcal{M}_{d,\delta}(\boldsymbol{\alpha}))} \sum_{\mathcal{C}} \int_{\mathcal{H}_{1,lab}} \sum_{\beta \in \bar{\Gamma}_d(\boldsymbol{\alpha})} \widehat{\Psi}_{s,\varepsilon,\mathcal{C},\beta}(X, \omega) d\nu_1 + o(\varepsilon^2)$$

Using the $\mathcal{V}_{\mathcal{C}}^s$ Siegel–Veech formula for $\mathcal{H}_{1,lab}$ and Lemma 12.8, respectively,

$$\begin{aligned} &= d \frac{1}{\pi \varepsilon^2} \frac{\nu_1(\mathcal{H}_{1,lab})}{\nu_d(\mathcal{M}_{d,\delta}(\boldsymbol{\alpha}))} \sum_{\mathcal{C}} c_{\mathcal{C}}^s(\mathcal{H}_{1,lab}) \int_{\mathbb{R}^2} \sum_{\beta \in \bar{\Gamma}_d(\boldsymbol{\alpha})} \Psi_{s,\varepsilon,\mathcal{C},\beta}(\mathbf{v}) dx dy + o(\varepsilon^2) \\ &= \frac{1}{d} \frac{\nu_1(\mathcal{H}_{1,lab})}{\nu_d(\mathcal{M}_{d,\delta}(\boldsymbol{\alpha}))} \sum_{\mathcal{C}} c_{\mathcal{C}}^s(\mathcal{H}_{1,lab}) \frac{1}{m(\mathcal{C})} \sum_{\beta \in \bar{\Gamma}_d(\boldsymbol{\alpha})} \sum_{C \in \mathcal{C}} \gcd(\beta(C), d)^{3-s} + o(\varepsilon^2). \end{aligned}$$

At this point, we have an equation that is completely combinatorial, so sending ε to zero, we have

$$c_{\text{area}^s}(\mathcal{M}_{1,\delta}(\boldsymbol{\alpha})) = \frac{1}{d} \frac{\nu_1(\mathcal{H}_{1,lab})}{\nu_d(\mathcal{M}_{d,\delta}(\boldsymbol{\alpha}))} \sum_{\mathcal{C}} c_{\mathcal{C}}^s(\mathcal{H}_{1,lab}) \frac{1}{m(\mathcal{C})} \sum_{\beta \in \bar{\Gamma}_d(\boldsymbol{\alpha})} \sum_{C \in \mathcal{C}} \gcd(\beta(C), d)^{3-s}.$$

Since $\mathcal{P} : \mathcal{M}_{d,\delta}(\boldsymbol{\alpha}) \rightarrow \mathcal{H}_{1,lab}$ is a covering map of degree $|\bar{\Gamma}_d \boldsymbol{\alpha}|$ (by Proposition 3.1) that locally pushes forward ν_d to ν_1 , we can simplify the leading coefficient

$$= \frac{1}{d} \frac{1}{|\bar{\Gamma}_d \boldsymbol{\alpha}|} \sum_{\mathcal{C}} c_{\mathcal{C}}^s(\mathcal{H}_{1,lab}) \frac{1}{m(\mathcal{C})} \underbrace{\sum_{\beta \in \bar{\Gamma}_d(\boldsymbol{\alpha})} \sum_{C \in \mathcal{C}} \gcd(\beta(C), d)^{3-s}}_{(*)}. \quad (23)$$

Recovering the orbit quotient function. Let us now focus on $(*)$ for a fixed configuration \mathcal{C} . Enumerate the cylinders of \mathcal{C} as C_1, \dots, C_m and partition $\bar{\Gamma}_d(\boldsymbol{\alpha})$ depending on the value $a := \beta(C_1)$; label the pieces of the partition as $(\bar{\Gamma}_d \boldsymbol{\alpha})_{C_1;a}$ as in Equation (2). This choice seems to *a priori* matter, but at the very end of the proof we will see that a different labeling of the cylinders results in the same formula. Using this partition,

$$(*) = \sum_{a=1}^d \sum_{\beta \in (\bar{\Gamma}_d \boldsymbol{\alpha})_{C_1;a}} \sum_{i=1}^{m(\mathcal{C})} \gcd(\beta(C_i), d)^{3-s}.$$

Now we recall that specifying $\beta(C_1)$ also specifies the monodromy of all the other C_i (Lemma 2.3). Thus, for fixed a and i the terms of the summand are independent of our choice of $\beta \in (\bar{\Gamma}_d \boldsymbol{\alpha})_{C_1;a}$. Plugging $(*)$ back into (23), switching the orders of summation, and rearranging, we are left with

$$c_{\text{area}^s}(\mathcal{M}_{1,\delta}(\boldsymbol{\alpha})) = \frac{1}{d} \sum_{\mathcal{C}} c_{\mathcal{C}}^s(\mathcal{H}_{1,lab}) \frac{1}{m(\mathcal{C})} \sum_{i=1}^m \underbrace{\sum_{a=1}^d \frac{|(\bar{\Gamma}_d \boldsymbol{\alpha})_a|}{|\bar{\Gamma}_d \boldsymbol{\alpha}|} \gcd(\beta(C_i), d)^{3-s}}_{(**)}. \quad (24)$$

We remark that by the results in Appendix B, the size of $|(\bar{\Gamma}_d \boldsymbol{\alpha})_a|$ does not depend on our choice of C_1 , just on a , which is why we have dropped it from the subscript.

Simplifying ().** At this point, we appeal to Proposition 12.9 and the assumptions of the Theorem to compute (**). Factor $d = d'_{rel} d'_{abs} 2^k$, for $k \geq 0$, such that $d'_{rel} | d_{rel}$, $d'_{abs} | d_{abs}$, and both d'_{rel} and d'_{abs} are odd. By Corollary 4.3 and the multiplicativity of the gcd function,

$$(**) = \prod_{D \in \{d'_{rel}, d'_{abs}, 2^k\}} \left(\sum_{a=1}^D \frac{|\overline{\Gamma}_D \boldsymbol{\alpha}|}{|\overline{\Gamma}_D \boldsymbol{\alpha}|} \gcd(\boldsymbol{\beta}(C_i), D)^{3-s} \right).$$

Each of the terms of the product contains an orbit quotient function $f_D(a)$, for $D \in \{d'_{rel}, d'_{abs}, 2^k\}$ as defined at the beginning of this section. Replacing these orbit quotient functions and using the third assumption of Theorem 12.5 that $f_{d'_{abs}}$ is a constant $C_{d'_{abs}}$, the product can be rewritten as

$$\left(\sum_{a=1}^{d'_{rel}} f_{d'_{rel}}(a) \gcd(\boldsymbol{\beta}(C_i), d'_{rel})^{3-s} \right) \left(C_{d'_{abs}} \sum_{a=1}^{d'_{abs}} \gcd(\boldsymbol{\beta}(C_i), d'_{abs})^{3-s} \right) \left(\sum_{a=1}^{2^k} f_{2^k}(a) \gcd(\boldsymbol{\beta}(C_i), 2^k)^{3-s} \right).$$

We focus on the first term (involving d'_{rel}). Recall from Proposition 2.2 that $\boldsymbol{\beta}(C_i) = a + b_i \gcd(\delta(\boldsymbol{\alpha}))$ for some b_i . Using the first assumption in the statement of the Theorem that $f_{d'_{rel}}$ is d'_{rel} -GCD dependent, we can write

$$\sum_{a=1}^{d'_{rel}} f_{d'_{rel}}(a) \gcd(\boldsymbol{\beta}(C_i), d'_{rel})^{3-s} = \sum_{a=1}^{d'_{rel}} f_{d'_{rel}}(\gcd(a, d'_{rel})) \gcd(a + b_i \gcd(\delta(\boldsymbol{\alpha})), d'_{rel})^{3-s}.$$

By the second assumption of the Theorem, $f_{d'_{rel}}$ can be factored as a constant multiplied by a prime power independent function (compare Lemma 13.1 below). Next we aim to apply Proposition 12.9. To do so, we observe that $\gcd(\delta(\boldsymbol{\alpha})) | d_{rel}$ by definition of d_{rel} , and moreover, $\text{rad}(\gcd(\delta(\boldsymbol{\alpha}))) = \text{rad}(d_{rel})$. Since $d'_{rel} | d_{rel}$, we know $\text{rad}(d'_{rel}) | b_i \gcd(\delta(\boldsymbol{\alpha}))$ and thus the divisibility of the radical assumption in Proposition 12.9 is satisfied for $M = b_i \gcd(\delta(\boldsymbol{\alpha}))$. Finally, we factor out the constant, apply Proposition 12.9, and multiply back by the constant that was factored out to obtain

$$\sum_{a=1}^{d'_{rel}} f_{d'_{rel}}(\gcd(a, d'_{rel})) \gcd(a + b_i \gcd(\delta(\boldsymbol{\alpha})), d'_{rel})^{3-s} = \sum_{\mathfrak{d} | d'_{rel}} f_{d'_{rel}}(\mathfrak{d}) \Phi \left(\frac{d'_{rel}}{\mathfrak{d}} \right) \mathfrak{d}^{3-s}.$$

The second term can be simplified by reindexing to get

$$\sum_{a=1}^{d'_{abs}} \gcd(\boldsymbol{\beta}(C_i), d'_{abs})^{3-s} = \sum_{a=1}^{d'_{abs}} \gcd(a + b_i \gcd(\delta(\boldsymbol{\alpha})), d'_{abs})^{3-s} = \sum_{\mathfrak{d} | d'_{abs}} \Phi \left(\frac{d'_{abs}}{\mathfrak{d}} \right) \mathfrak{d}^{3-s}.$$

For the final term (involving 2^k), there are several cases to consider. If $2 | \gcd(\delta(\boldsymbol{\alpha}))$, then we apply the same argument above for d'_{rel} using Proposition 12.9 to conclude

$$\sum_{a=1}^{2^k} f_{2^k}(a) \gcd(\boldsymbol{\beta}(C_i), 2^k)^{3-s} = \sum_{\mathfrak{d} | 2^k} f_{2^k}(\mathfrak{d}) \Phi \left(\frac{2^k}{\mathfrak{d}} \right) \mathfrak{d}^{3-s}. \quad (25)$$

Therefore, it suffices to consider the case where $2 \nmid d_{rel}$, in which case $2 | d_{abs}$. We split this remaining case into two additional cases depending on whether the orbit is split or unsplit. If it is unsplit, then by the last assumption $f_{2^k}(a)$ is a constant function and we can use the reindexing argument as in the d'_{abs} case. This leaves the case where the orbit is split and $2 | d_{abs}$, which is addressed by the following lemma.

Lemma 12.10. *Suppose that d is even and the $\overline{\Gamma}_d$ orbit of the branching vector $\boldsymbol{\alpha}$ is split. Then for any configuration of disjoint cylinders $\mathcal{C} = \{C_1, \dots, C_m\}$ on $(X, \omega) \in \mathcal{H}$ all representing the same class in absolute homology, the parity of the $\boldsymbol{\alpha}(C_i)$ all agree.*

Proof. By the proof of Proposition 2.2, for all i and j ,

$$\alpha(C_i) - \alpha(C_j) = \sum_{p_k \in \mathcal{B}} m_k \alpha(\gamma_k)$$

where γ_k is some small loop about p_k and $m_k \in \mathbb{Z}$. We need only show that the right-hand sum is even: in order to do this, we will specify the m_k more concretely.

Consider the operation where each cylinder is stretched to an infinite cylinder resulting in m (infinite area) translation surfaces each carrying a collection of zeros and exactly two simple poles corresponding to homologous half-infinite cylinders. Consider one such component W ; re-indexing if necessary, let us assume that its cylinders correspond to C_i and C_j . This surface W demonstrates the homology in X between C_i and C_j , and in particular we get

$$\alpha(C_i) - \alpha(C_j) = \sum_{p_k \in W} \alpha(\gamma_k). \quad (26)$$

Now by Riemann-Roch, the total order of the zeros minus the total order of the poles is even, so since the total order of the poles is 2, then the total order of the zeros is even. In particular, there is an even number of odd order zeros on W .

Now if (X, ω) is non-hyperelliptic and the orbit is split, then $\delta(\alpha) \equiv \kappa \pmod{2}$ by Proposition 8.1. Therefore, the branching over each odd zero is necessarily odd. Combining this with the fact that there are an even number of odd-order zeros on W implies that the sum in (26) must be even, and so the monodromies of C_i and C_j have the same parity.

Consider now the case where (X, ω) lies in a hyperelliptic stratum component and there are no regular marked points. In this setting, we observe that \mathcal{C} must consist of a single cylinder. Indeed, if there were two disjoint homologous cylinders C_1 and C_2 , then since the hyperelliptic involution fixes each cylinder setwise [Lin15, Lemma 2.1] and has genus 0 quotient, this would imply that both of the components of $S \setminus (C_1 \cup C_2)$ have genus 0, and so S would have to have genus 1. Thus, for hyperelliptic strata with regular marked points, the m_k corresponding to loops about the zeros are all 0. The assumption that the orbit is split further implies that the branching over all of the regular marked points is even (by Proposition 10.5), hence these terms do not change the parity of (26). \square

Therefore, if $2 \nmid d_{rel}$, then $2 \nmid \gcd(\delta(\alpha))$ by definition of d_{rel} and by Lemma 12.10, b_i is even. In this case, we apply Proposition 12.9 with $M = b_i \gcd(\delta(\alpha))$, so that $\text{rad}(2^k) = 2|M$, to again obtain the simplification in Equation (25).

Putting this all together, we get that for any configuration \mathcal{C} and all i , we have that

$$(**) = \left(\sum_{\mathfrak{d} | d'_{rel}} f_{d'_{rel}}(\mathfrak{d}) \Phi \left(\frac{d'_{rel}}{\mathfrak{d}} \right) \mathfrak{d}^{3-s} \right) \left(C_{d'_{abs}} \sum_{\mathfrak{d} | d'_{abs}} \Phi \left(\frac{d'_{abs}}{\mathfrak{d}} \right) \mathfrak{d}^{3-s} \right) \left(\sum_{\mathfrak{d} | 2^k} f_{2^k}(\mathfrak{d}) \Phi \left(\frac{2^k}{\mathfrak{d}} \right) \mathfrak{d}^{3-s} \right). \quad (27)$$

Plugging back in. We now observe that (27) is independent of the cylinder C_i : it depends only on the orbit quotient functions f_D and certain number theoretic data. This (together with the results of Appendix B) implies that our initial indexing of the cylinders of \mathcal{C} did not matter, as when we sum over the cylinders of \mathcal{C} , we get the same value.

Moreover, (27) is *also* independent of the initial configuration \mathcal{C} and the number of cylinders in it. Thus, when we plug (**) back into (24), we can pull it out of all of the summations, yielding

$$c_{\text{area}^s}(\mathcal{M}_{1,\delta}(\alpha)) = \frac{1}{d} \sum_{\mathcal{C}} c_{\mathcal{C}}^s(\mathcal{H}_{1,lab}) \frac{1}{m(\mathcal{C})} \sum_{i=1}^m (***) = \frac{1}{d} (***) \sum_{\mathcal{C}} c_{\mathcal{C}}^s(\mathcal{H}_{1,lab}) = \frac{1}{d} (***) c_{\text{area}^s}(\mathcal{H}_1),$$

where the last equality follows from linearity of Siegel–Veech transforms (17) and the fact that labeling points does not change Siegel–Veech constants (this is clear from their interpretations as asymptotic counts of cylinders). This completes the proof of the Theorem.

13 Evaluating the orbit quotient function

Now we use the results of Part II to show that the hypotheses of Theorem 12.5 hold and evaluate the relevant orbit quotient functions. By definition, in the factorization $d = d'_{rel}d'_{abs}2^k$ we have chosen d'_{rel} and d'_{abs} to be odd. Therefore, there is a single formula for each of these quantities for *all* stratum components.

Lemma 13.1. $f_{d'_{rel}}$ is d'_{rel} -GCD dependent and can be factored as a constant times a prime power independent function for all components of all strata. Moreover, the following equation holds:

$$\sum_{\mathfrak{d}|d'_{rel}} f_{d'_{rel}}(\mathfrak{d})\Phi\left(\frac{d'_{rel}}{\mathfrak{d}}\right)\mathfrak{d}^{3-s} = \frac{d'_{rel}}{d'^{s-2}_{rel}\Phi_{2g}(d'_{rel})} \sum_{\mathfrak{d}|d'_{rel}} \Phi_{2g-1}\left(\frac{d'_{rel}}{\mathfrak{d}}\right) \frac{\Phi(\mathfrak{d})}{\mathfrak{d}^{4-2g-s}}.$$

Proof. Let $\mathfrak{d}|d'_{rel}$, and let $p|d'_{rel}$. By Proposition 4.6, if $p|\mathfrak{d}$, then $|\overline{\Gamma}_{p^k}\alpha| = \Phi_{2g-1}(p^k)$, and if $p \nmid \mathfrak{d}$, then $|\overline{\Gamma}_{p^k}\alpha| = p^{k(2g-1)}$. These quantities are multiplicative, in particular,

$$|\overline{\Gamma}_{d'_{rel}}\alpha| = \prod_{p|\mathfrak{d}} p^{k(2g-1)} \prod_{p|\mathfrak{d}} \Phi_{2g-1}(p^k) = \left(\frac{d'_{rel}}{\mathfrak{d}}\right)^{2g-1} \Phi_{2g-1}(\mathfrak{d}).$$

By Proposition 4.5, if $p|d'_{rel}$, then $|\overline{\Gamma}_{p^k}\alpha| = \Phi_{2g}(p^k)$. Hence,

$$f_{d'_{rel}}(\mathfrak{d}) = \frac{|\overline{\Gamma}_{d'_{rel}}\alpha|}{|\overline{\Gamma}_{d'_{rel}}\alpha|} = \left(\frac{d'_{rel}}{\mathfrak{d}}\right)^{2g-1} \frac{\Phi_{2g-1}(\mathfrak{d})}{\Phi_{2g}(d'_{rel})} = \frac{(d'_{rel})^{2g-1}}{\Phi_{2g}(d'_{rel})} \cdot \frac{\Phi_{2g-1}(\mathfrak{d})}{\mathfrak{d}^{2g-1}}.$$

It is immediate that $f_{d'_{rel}}$ is d'_{rel} -GCD dependent. Moreover, we recognize the right-hand side of the equation is a constant depending only on d'_{rel} times $\Phi_{2g-1}(\mathfrak{d})/\mathfrak{d}^{2g-1}$, which is prime power independent by the definition of the Jordan totient function. This facilitates the following simplification:

$$\begin{aligned} \sum_{\mathfrak{d}|d'_{rel}} f_{d'_{rel}}(\mathfrak{d})\Phi\left(\frac{d'_{rel}}{\mathfrak{d}}\right)\mathfrak{d}^{3-s} &= \sum_{\mathfrak{d}|d'_{rel}} \left(\frac{d'_{rel}}{\mathfrak{d}}\right)^{2g-1} \frac{\Phi_{2g-1}(\mathfrak{d})}{\Phi_{2g}(d'_{rel})} \Phi\left(\frac{d'_{rel}}{\mathfrak{d}}\right)\mathfrak{d}^{3-s} \\ &= \frac{(d'_{rel})^{s-2}}{(d'_{rel})^{s-2}d'_{rel}\Phi_{2g}(d'_{rel})} \cdot \left(\sum_{\mathfrak{d}|d'_{rel}} \frac{(d'_{rel})^{2g}\Phi_{2g-1}(\mathfrak{d})}{\mathfrak{d}^{2g-4+s}} \Phi\left(\frac{d'_{rel}}{\mathfrak{d}}\right)\right) \\ &= \frac{d'_{rel}}{(d'_{rel})^{s-2}\Phi_{2g}(d'_{rel})} \cdot \left(\sum_{\mathfrak{d}|d'_{rel}} \left(\frac{d'_{rel}}{\mathfrak{d}}\right)^{2g-4+s} \Phi_{2g-1}(\mathfrak{d})\Phi\left(\frac{d'_{rel}}{\mathfrak{d}}\right)\right), \end{aligned}$$

completing the proof of the lemma. □

Lemma 13.2. The function $f_{d'_{abs}}(\mathfrak{d})$ is constant for all strata and

$$C_{d'_{abs}} \sum_{\mathfrak{d}|d'_{abs}} \Phi\left(\frac{d'_{abs}}{\mathfrak{d}}\right)\mathfrak{d}^{3-s} = \frac{1}{(d'_{abs})^{s-2}} \sum_{\mathfrak{d}|d'_{abs}} \frac{\Phi(\mathfrak{d})}{\mathfrak{d}^{3-s}}.$$

Proof. If $p|d'_{abs}$, then by Propositions 4.5 and 4.6, $|\overline{\Gamma}_{p^k}\alpha| = p^{2gk}$ and for all \mathfrak{d} , $|\overline{\Gamma}_{p^k}\alpha| = p^{k(2g-1)}$, respectively. Therefore,

$$f_{d'_{abs}}(\mathfrak{d}) := \frac{|\overline{\Gamma}_{d'_{abs}}\alpha|}{|\overline{\Gamma}_{d'_{abs}}\alpha|} = \frac{(d'_{abs})^{2g-1}}{(d'_{abs})^{2g}} = \frac{1}{d'_{abs}} = C_{d'_{abs}}.$$

□

The following lemma follows from Propositions 4.5 and 4.6 in exactly the same way as Lemma 13.2 does.

Lemma 13.3. If the orbit of α is unsplit and $2|d_{abs}$, then $f_{2^k}(a)$ is constant.

13.1 Evaluation of orbit ratios for powers of 2

Here we focus on the term

$$\sum_{\mathfrak{d}|2^k} f_{2^k}(\mathfrak{d}) \Phi\left(\frac{2^k}{\mathfrak{d}}\right) \mathfrak{d}^{3-s}.$$

Since we will have to break our proof into cases for when the base stratum is hyperelliptic and when it is not, let us first prove a preliminary (contingent) result that will allow us to get some number-theoretic manipulations out of the way. To condense our expression for the final form of our formula, we define $\xi_2(k) = k + 1$ and for $s \neq 2$, define

$$\xi_s(k) = 2^{k(2-s)+1} + 2^{(2-s)} \frac{1 - 2^{(k-1)(2-s)}}{1 - 2^{2-s}}.$$

Lemma 13.4. *If $f_{2^k}(N)$ can be factored as a constant multiplied by a prime power independent function, then*

$$\sum_{\mathfrak{d}|2^k} f_{2^k}(\mathfrak{d}) \Phi\left(\frac{2^k}{\mathfrak{d}}\right) \mathfrak{d}^{3-s} = 2^{k-1} (f_{2^k}(1) + f_{2^k}(0)\xi_s(k)).$$

Proof. The constant that can be factored out will play no role in this proof and can be safely ignored. The key observation is that prime power independence implies $f_{2^k}(\mathfrak{d}) = f_{2^k}(0)$ if $2|\mathfrak{d}$. Hence,

$$\begin{aligned} \sum_{\mathfrak{d}|2^k} f_{2^k}(\mathfrak{d}) \Phi\left(\frac{2^k}{\mathfrak{d}}\right) \mathfrak{d}^{3-s} &= \sum_{j=0}^k f_{2^k}(2^j) \Phi\left(\frac{2^k}{2^j}\right) (2^j)^{3-s} \\ &= f_{2^k}(1) \Phi(2^k) + \left(\sum_{j=1}^{k-1} f_{2^k}(0) \Phi\left(\frac{2^k}{2^j}\right) (2^j)^{3-s} \right) + f_{2^k}(0) \Phi(1) (2^k)^{3-s} \\ &= 2^{k-1} \left(f_{2^k}(1) + f_{2^k}(0) \left(2^{k(2-s)+1} + \sum_{j=1}^{k-1} (2^{2-s})^j \right) \right). \end{aligned}$$

If $s \neq 2$, then we can simplify using a geometric sum, and if $s = 2$, we can simplify directly. \square

Proposition 13.5. *Let \mathcal{H} be a non-hyperelliptic stratum component and let α be a branching vector. Let $\xi_s(k)$ be as defined above. If the orbit of α is split, then $f_{2^k}(N)$ can be factored as a constant multiplied by a prime power independent function, f_{2^k} is 2^k -GCD dependent, and*

$$\sum_{\mathfrak{d}|2^k} f_{2^k}(\mathfrak{d}) \Phi\left(\frac{2^k}{\mathfrak{d}}\right) \mathfrak{d}^{3-s} = \frac{2^{g-1}(2^{g-1} + (-1)^{\psi(\alpha)}) + (2^{2g-2} - \tau(\alpha))\xi_s(k)}{2^{g-1}(2^g + (-1)^{\psi(\alpha)}) - \tau(\alpha)}.$$

Moreover, if $2|d_{abs}$, then $\tau(\alpha) = 0$.

Proof. From Propositions 8.2 and 8.3, it follows that $f_{2^k}(N)$ is a constant times a prime power independent function. Furthermore, 2^k -GCD dependence follows from Proposition 4.2. If $2|d_{rel}$, then from Lemma 13.4 and Propositions 8.2 and 8.3

$$\begin{aligned} \sum_{\mathfrak{d}|2^k} f_{2^k}(\mathfrak{d}) \Phi\left(\frac{2^k}{\mathfrak{d}}\right) \mathfrak{d}^{3-s} &= 2^{k-1} (f_{2^k}(1) + f_{2^k}(0)\xi_s(k)) \\ &= 2^{k-1} \left(\frac{2^{(2k-1)g-k}(2^{g-1} + (-1)^{\psi(\alpha)})}{2^{2g(k-1)}(2^{g-1}(2^g + (-1)^{\psi(\alpha)}) - \tau(\alpha))} + \frac{2^{(2g-1)(k-1)}(2^{2g-2} - \tau(\alpha))}{2^{2g(k-1)}(2^{g-1}(2^g + (-1)^{\psi(\alpha)}) - \tau(\alpha))} \xi_s(k) \right). \end{aligned}$$

For the last statement, we just note that if $2|d_{abs}$, then $\tau(\alpha) = 0$ because $2|d_{abs}$ implies $\gcd(\delta(\alpha))$ is odd and since the orbit is split, $\delta \equiv \kappa \pmod{2}$ by Proposition 8.1. \square

Proposition 13.6. *Let \mathcal{H}^{hyp} be a hyperelliptic stratum component and let α be a branching vector. Let $\xi_s(k)$ be as defined above. If the orbit of α is split, then $f_{2^k}(N)$ can be factored as a constant multiplied by a prime power independent function, f_{2^k} is 2^k -GCD dependent, and*

$$\sum_{\mathfrak{d}|2^k} f_{2^k}(\mathfrak{d}) \phi\left(\frac{2^k}{\mathfrak{d}}\right) \mathfrak{d}^{3-s} = \xi_s(k) + 2(1 - \xi_s(k)) \frac{\binom{Br-2}{b(\alpha)-1}}{\binom{Br}{b(\alpha)}}.$$

Proof. It follows from Propositions 10.6 and 10.7 that $f_{2^k}(N)$ is a constant times a prime power independent function. As in the non-hyperelliptic case, 2^k -GCD dependence follows from Proposition 4.2. Substituting the formulas derived in Propositions 10.6 and 10.7 into the formula from Lemma 13.4 yields

$$\begin{aligned} \sum_{\mathfrak{d}|2^k} f_{2^k}(\mathfrak{d}) \Phi\left(\frac{2^k}{\mathfrak{d}}\right) \mathfrak{d}^{3-s} &= 2^{k-1} (f_{2^k}(1) + f_{2^k}(0) \xi_s(k)) \\ &= 2^{k-1} \left(\frac{2^{(2g-1)(k-1)-\sigma(\alpha)} \cdot 2^{\binom{Br-2}{b(\alpha)-1}}}{2^{2g(k-1)-\sigma(\alpha)} \binom{Br}{b(\alpha)}} + \frac{2^{(2g-1)(k-1)-\sigma(\alpha)} \cdot \left(\binom{Br-2}{b(\alpha)} + \binom{Br-2}{b(\alpha)-2}\right)}{2^{2g(k-1)-\sigma(\alpha)} \binom{Br}{b(\alpha)}} \xi_s(k) \right) \\ &= \frac{2^{\binom{Br-2}{b(\alpha)-1}}}{\binom{Br}{b(\alpha)}} (1 - \xi_s(k)) + \frac{\left(\binom{Br-2}{b(\alpha)} + 2\binom{Br-2}{b(\alpha)-1} + \binom{Br-2}{b(\alpha)-2}\right)}{\binom{Br}{b(\alpha)}} \xi_s(k). \end{aligned}$$

The final formula follows from the Pascal relation on binomial coefficients. \square

14 Main formula

Notational guide. Let us briefly recall the relevant notation and indicate where in the paper a full discussion can be found.

- \mathcal{H} is a connected component of a stratum $\mathcal{H}(\kappa)$ of translation surfaces, possibly with regular marked points. We use \mathcal{H}_1 to denote its unit-area locus. Recall our convention that \mathcal{H} is hyperelliptic if the stratum component obtained by removing all marked points is hyperelliptic.
- \mathcal{M}_δ is the locus of degree d cyclic branched covers of surfaces with branching data given by δ .
- We factor the degree d into “absolute” and “relative” parts d_{abs} and d_{rel} . If d is even, we further factor $d = d'_{abs} d'_{rel} 2^k$ such that $d'_{abs}|d_{abs}$ and $d'_{rel}|d_{rel}$ are odd.
- The components of \mathcal{M}_δ are classified by the monodromy orbits of the *branching vector* $\alpha \in H^1(X \setminus \mathcal{B}; \mathbb{Z}/d\mathbb{Z})$. We use $\mathcal{M}_\delta(\alpha)$ to denote the component containing the branched covers of surfaces in \mathcal{H} with branching vector α and $\mathcal{M}_{1,\delta}(\alpha)$ to denote the unit-area locus.
- A monodromy orbit of branching vectors is *unsplit* if $\mathcal{M}_\delta(\alpha) = \mathcal{M}_\delta$ and is split otherwise. Whether the orbit of α is split or unsplit is determined by Propositions 4.2, 8.1, and 10.5. In particular, the orbit is split only when d is even.
- The ψ *invariant*, valued in $\mathbb{Z}/2\mathbb{Z}$, is a complete invariant classifying split orbits when \mathcal{H} is non-hyperelliptic. See Section 7.
- The quantity $\tau(\alpha) \in \mathbb{Z}/2\mathbb{Z}$ is defined in Proposition 8.2.
- The quantities Br and $b(\alpha)$ are complete integer invariants classifying split orbits when \mathcal{H} is hyperelliptic. See Definitions 9.1 and 10.1.

- Recall the function from Section 13.1

$$\xi_s(k) = \begin{cases} 2^{k(2-s)+1} + 2^{(2-s)\frac{1-2^{(k-1)(2-s)}}{1-2^{2-s}}} & \text{if } s \neq 2 \\ k + 1 & \text{if } s = 2. \end{cases}$$

- For positive integers A and B , define a function

$$G_{s,g}(A, B) = \frac{A}{\Phi_{2g}(A)} \left(\sum_{\mathfrak{d}|A} \Phi_{2g-1} \left(\frac{A}{\mathfrak{d}} \right) \frac{\Phi(\mathfrak{d})}{\mathfrak{d}^{4-2g-s}} \right) \cdot \left(\sum_{\mathfrak{d}|B} \left(\frac{\Phi(\mathfrak{d})}{\mathfrak{d}^{3-s}} \right) \right).$$

Theorem 14.1. *Let $s \geq 0$. Then $c_{\text{area}^s}(\mathcal{M}_{1,\delta}(\boldsymbol{\alpha}))/c_{\text{area}^s}(\mathcal{H}_1)$ is given by one of the following formulas:*

- If \mathcal{H}_1 is non-hyperelliptic and the orbit of $\boldsymbol{\alpha}$ is split,

$$\frac{2^{k(s-2)}}{d^{s-1}} \cdot \frac{2^{g-1}(2^{g-1} + (-1)^{\psi(\boldsymbol{\alpha})}) + (2^{2g-2} - \tau(\boldsymbol{\alpha}))\xi_s(k)}{2^{g-1}(2^g + (-1)^{\psi(\boldsymbol{\alpha})}) - \tau(\boldsymbol{\alpha})} \cdot G_{s,g}(d'_{rel}, d'_{abs}).$$

- If \mathcal{H}_1 is hyperelliptic and the orbit of $\boldsymbol{\alpha}$ is split,

$$\frac{2^{k(s-2)}}{d^{s-1}} \left(\xi_s(k) + 2(1 - \xi_s(k)) \frac{\binom{Br-2}{b(\boldsymbol{\alpha})-1}}{\binom{Br}{b(\boldsymbol{\alpha})}} \right) \cdot G_{s,g}(d'_{rel}, d'_{abs}).$$

- Otherwise, the orbit of $\boldsymbol{\alpha}$ is unsplit, and the formula is

$$d^{1-s} \cdot G_{s,g}(d_{rel}, d_{abs}).$$

Proof. Having derived every piece of the main formula in the sections above, we combine the results here. Lemmas 13.1, 13.2, and 13.3 along with Propositions 13.5 and 13.6 show that the first two assumptions in Theorem 12.5 are satisfied. Lemma 13.2 verifies the third assumption of Theorem 12.5 holds. Finally, Lemma 13.3 establishes that the fourth assumption of Theorem 12.5 holds. Thus, combining Theorem 12.5 and Lemmas 13.1, 13.2 and 13.3, we have in all cases that the ratio $c_{\text{area}^s}(\mathcal{M}_{d,\delta}(\boldsymbol{\alpha}))/c_{\text{area}^s}(\mathcal{H}_1)$ is given by

$$\begin{aligned} & \frac{1}{d} \left(\sum_{\mathfrak{d}|d'_{rel}} f_{d'_{rel}}(\mathfrak{d}) \Phi \left(\frac{d'_{rel}}{\mathfrak{d}} \right) \mathfrak{d}^{3-s} \right) \left(C_{d'_{abs}} \sum_{\mathfrak{d}|d'_{abs}} \Phi \left(\frac{d'_{abs}}{\mathfrak{d}} \right) \mathfrak{d}^{3-s} \right) \left(\sum_{\mathfrak{d}|2^k} f_{2^k}(\mathfrak{d}) \Phi \left(\frac{2^k}{\mathfrak{d}} \right) \mathfrak{d}^{3-s} \right) \\ &= \frac{1}{d} \left(\frac{d'_{rel}}{(d'_{rel})^{s-2} \Phi_{2g}(d'_{rel})} \sum_{\mathfrak{d}|d'_{rel}} \Phi_{2g-1} \left(\frac{d'_{rel}}{\mathfrak{d}} \right) \frac{\Phi(\mathfrak{d})}{\mathfrak{d}^{4-2g-s}} \right) \left(\frac{1}{(d'_{abs})^{s-2}} \sum_{\mathfrak{d}|d'_{abs}} \left(\frac{\Phi(\mathfrak{d})}{\mathfrak{d}^{3-s}} \right) \right) \\ & \quad \cdot \left(\frac{2^{k(s-2)}}{2^{k(s-2)}} \sum_{\mathfrak{d}|2^k} f_{2^k}(\mathfrak{d}) \Phi \left(\frac{2^k}{\mathfrak{d}} \right) \mathfrak{d}^{3-s} \right). \end{aligned}$$

We have added an extra term to the 2^k term for the convenience of the reader. From this final expression, the reader can readily apply Propositions 13.5 and 13.6 to get the desired set of formulas. \square

We emphasize that this result also holds when \mathcal{H}_1 is a stratum of marked tori (in which case all orbits are unsplit); specializing to the $g = 1$ case yields Corollary F from the introduction.

14.1 area- and area cubed–Siegel–Veech constants

Given the particular importance of area–Siegel–Veech constants in the formula for the sum of the Lyapunov exponents of the Kontsevich–Zorich cocycle [EKZ14], we explicitly state the formula for them for future ease of reference. Observe: $\xi_1(k) = 3 \cdot 2^k - 2$ and

$$G_{1,g}(d_{rel}, d_{abs}) = \frac{d_{rel}}{\Phi_{2g}(d_{rel})} \left(\sum_{\mathfrak{d}|d_{rel}} \Phi_{2g-1} \left(\frac{d_{rel}}{\mathfrak{d}} \right) \frac{\Phi(\mathfrak{d})}{\mathfrak{d}^{3-2g}} \right) \cdot \left(\sum_{\mathfrak{d}|d_{abs}} \left(\frac{\Phi(\mathfrak{d})}{\mathfrak{d}^2} \right) \right).$$

Theorem 14.2. *The ratio $c_{area}(\mathcal{M}_{1,\delta}(\alpha))/c_{area}(\mathcal{H}_1)$ is given by one of the following formulas:*

- If \mathcal{H} is non-hyperelliptic and the orbit of α is split,

$$2^{-k} \cdot \frac{2^{g-1}(2^{g-1} + (-1)^{\psi(\alpha)}) + (2^{2g-2} - \tau(\alpha))(3 \cdot 2^k - 2)}{2^{g-1}(2^g + (-1)^{\psi(\alpha)}) - \tau(\alpha)} \cdot G_{1,g}(d'_{rel}, d'_{abs})$$

- If \mathcal{H} is hyperelliptic and the orbit of α is split, then

$$2^{-k} \left(3 \cdot 2^k - 2 + 6(1 - 2^k) \frac{\binom{Br-2}{b(\alpha)-1}}{\binom{Br}{b(\alpha)}} \right) \cdot G_{1,g}(d'_{rel}, d'_{abs})$$

- Otherwise, the orbit of α is unsplit and the formula is given by $G_{1,g}(d_{rel}, d_{abs})$.

Proof of Theorem C. This follows from the simplification $\xi_3(k) = 1$ and the elementary identities:

$$\Phi_{2g}(d'_{rel}) = \sum_{\mathfrak{d}|d'_{rel}} \Phi_{2g-1} \left(\frac{d'_{rel}}{\mathfrak{d}} \right) \frac{\Phi(\mathfrak{d})}{\mathfrak{d}^{1-2g}} \quad \text{and} \quad d'_{abs} = \sum_{\mathfrak{d}|d'_{abs}} \Phi(\mathfrak{d}).$$

□

A Symplectic modules

Here we collect some basic results on the action of the symplectic group $\mathrm{Sp}(2g, \mathbb{Z}/d\mathbb{Z})$ and certain of its subgroups on $(\mathbb{Z}/d\mathbb{Z})^{2g}$ (equipped with the symplectic form $\langle \cdot, \cdot \rangle$), which do not seem to be readily available in the literature. We allow the possibility that $d = 0$ throughout.

Recall that for d even, there are subgroups $\mathrm{Sp}(2g, \mathbb{Z}/d\mathbb{Z})[2]$ and $\mathrm{Sp}(2g, \mathbb{Z}/d\mathbb{Z})[q]$ defined as follows: $\mathrm{Sp}(2g, \mathbb{Z}/d\mathbb{Z})[2]$ is the kernel of the reduction map $\mathrm{Sp}(2g, \mathbb{Z}/d\mathbb{Z}) \rightarrow \mathrm{Sp}(2g, \mathbb{Z}/2\mathbb{Z})$, and, for a chosen $\mathbb{Z}/2\mathbb{Z}$ quadratic form q on $(\mathbb{Z}/2\mathbb{Z})^{2g}$ (i.e. the pullback of a quadratic form on $(\mathbb{Z}/2\mathbb{Z})^{2g}$ under reduction mod 2), the subgroup $\mathrm{Sp}(2g, \mathbb{Z}/d\mathbb{Z})[q]$ is the stabilizer of q under the action of $\mathrm{Sp}(2g, \mathbb{Z}/d\mathbb{Z})$ on the set of quadratic forms on $(\mathbb{Z}/2\mathbb{Z})^{2g}$.

The common theme in the results presented in Proposition A.1 below is that the orbits of each of these subgroups on $(\mathbb{Z}/2\mathbb{Z})^{2g}$ is “predicted” by the evident invariant.

Proposition A.1. *Fix $g \geq 1$ and $d \geq 0$, and let $x, y \in (\mathbb{Z}/d\mathbb{Z})^{2g}$ be given.*

1. *There exists $A \in \mathrm{Sp}(2g, \mathbb{Z}/d\mathbb{Z})$ for which $Ax = y$ if and only if $\mathrm{gcd}(x) = \mathrm{gcd}(y)$.*
2. *For d even, there exists $B \in \mathrm{Sp}(2g, \mathbb{Z}/d\mathbb{Z})[2]$ for which $Bx = y$ if and only if $\mathrm{gcd}(x) = \mathrm{gcd}(y)$ and $x = y \pmod{2}$.*
3. *For d even, there exists $C \in \mathrm{Sp}(2g, \mathbb{Z}/d\mathbb{Z})[q]$ for which $Cx = y$ if and only if $\mathrm{gcd}(x) = \mathrm{gcd}(y)$ and $q(x) = q(y)$.*

Proof. These results are classical, and we content ourselves with a sketch of the proofs. To begin, we show that it suffices to take $d = 0$ throughout. According to [NS64, Theorem 1], the reduction map $\mathrm{Sp}(2g, \mathbb{Z}) \rightarrow \mathrm{Sp}(2g, \mathbb{Z}/d\mathbb{Z})$ is surjective for all $d \geq 0$. As $\mathrm{Sp}(2g, \mathbb{Z}/d\mathbb{Z})[2]$ is defined as the kernel of the reduction map $\mathrm{Sp}(2g, \mathbb{Z}/d\mathbb{Z}) \rightarrow \mathrm{Sp}(2g, \mathbb{Z}/2\mathbb{Z})$, it follows that reduction mod d induces a surjection $\mathrm{Sp}(2g, \mathbb{Z})[2] \rightarrow \mathrm{Sp}(2g, \mathbb{Z}/d\mathbb{Z})[2]$, and a similar argument shows that $\mathrm{Sp}(2g, \mathbb{Z})[q]$ surjects onto $\mathrm{Sp}(2g, \mathbb{Z}/d\mathbb{Z})[q]$. Recall that an element of $(\mathbb{Z}/d\mathbb{Z})^{2g}$ is *primitive* if its components generate $\mathbb{Z}/d\mathbb{Z}$. We note that every primitive element $\bar{v} \in (\mathbb{Z}/d\mathbb{Z})^{2g}$ lifts to a primitive element $v \in \mathbb{Z}^{2g}$: this can be seen directly by elementary number theory, or else derived from the well-known surjectivity of $\mathrm{SL}_{2g}(\mathbb{Z}) \rightarrow \mathrm{SL}_{2g}(\mathbb{Z}/d\mathbb{Z})$, extending \bar{v} to a unimodular basis of $(\mathbb{Z}/d\mathbb{Z})^{2g}$ (possible since $2g > 1$), and taking an integral lift of the associated matrix. Finally, we note that in case d is even, both the mod-2 reduction and q -values of v and \bar{v} coincide. Taken together, these facts imply that to classify orbits mod d , it suffices to take integral lifts and classify the corresponding orbits in \mathbb{Z}^{2g} .

Next, we reduce to the case $\mathrm{gcd}(x) = 1$ by dividing out the common factor. The key idea in all three parts is an “extension lemma” in the spirit of Witt’s Theorem, which in its most basic form (suitable for (1)) asserts that an arbitrary primitive element of \mathbb{Z}^{2g} can be extended to a symplectic basis. In (2), a stronger extension lemma asserts that given a symplectic basis $\{x_1, y_1, \dots, x_g, y_g\}$ and a vector z such that $z = x_1 \pmod{2}$, then z can be completed to a symplectic basis $\{z = z_1, w_1, \dots, z_g, w_g\}$ such that $z_i = x_i$ and $w_i = y_i \pmod{2}$ for all i . And in (3), the relevant extension lemma asserts that given a $\mathbb{Z}/2\mathbb{Z}$ quadratic form q , a geometric symplectic basis $B = \{x_1, y_1, \dots, x_g, y_g\}$, and a primitive element z satisfying $q(z) = q(x_1)$, then z can be completed to a symplectic basis $\{z = z_1, w_1, \dots, z_g, w_g\}$ such that $q(z_i) = q(x_i)$ and $q(w_i) = q(y_i)$ for $1 \leq i \leq g$.

Granting such extension lemmas, the proof of the theorem follows quickly. For (1), apply the basic extension lemma to x and to y , producing symplectic bases B, B' , and let $A \in \mathrm{GL}(2g, \mathbb{Z})$ be determined by the condition $AB = B'$. Then in fact $A \in \mathrm{Sp}(2g, \mathbb{Z})$, since an element of $\mathrm{GL}(2g, \mathbb{Z})$ is symplectic if and only if it takes any symplectic basis to a symplectic basis. For (2), apply the basic extension lemma to x , producing a symplectic basis B . Then apply the strong extension lemma to y , producing a symplectic basis B' such that every element of B has the same mod-2 reduction as its counterpart in B' . The element of $\mathrm{GL}(2g, \mathbb{Z})$ taking B to B' then is seen to be an element of $\mathrm{Sp}(2g, \mathbb{Z})[2]$. The proof of (3) is similar. \square

The following lemma was invoked in the course of proving Theorem 6.1.

Lemma A.2. *Let q be a quadratic form on $H_1(X; \mathbb{Z}/2\mathbb{Z})$. For $g \geq 3$, any tuple $(a, b, c) \in (\mathbb{Z}/2\mathbb{Z})^3$ is realizable via elements $v, w \in H_1(X; \mathbb{Z}/2\mathbb{Z})$ so that*

$$q(v) = a, \quad q(w) = b, \quad \langle v, w \rangle = c.$$

Proof. Choose a symplectic basis $\{x_1, y_1, \dots, x_g, y_g\}$ for $H_1(X; \mathbb{Z}/2\mathbb{Z})$. Suppose first that $c = 1$. Define a quadratic form q' on $H_1(X; \mathbb{Z}/2\mathbb{Z})$ for which $q'(x_1) = a$ and $q'(y_1) = b$. The values for $q'(x_i)$ and $q'(y_i)$ for $1 < i < g$ can be chosen arbitrarily; then choose $q'(x_g) = 1$ and $q'(y_g)$ to be whichever of 0, 1 makes $\mathrm{Arf}(q') = \mathrm{Arf}(q)$. Thus q and q' lie in the same orbit of $\mathrm{Sp}(2g, \mathbb{Z}/2\mathbb{Z})$ by Proposition A.1, so that $q'(x) = q(Ax)$ for some $A \in \mathrm{Sp}(2g, \mathbb{Z}/2\mathbb{Z})$. Taking then $v = Ax_1$ and $w = Ay_1$ gives v, w with the required properties.

The argument in the case $c = 0$ is similar. This time, choose $q'(x_1) = a$ and $q'(x_2) = b$. As $g \geq 3$, it is possible to choose $q'(x_g) = 1$ and $q'(y_g)$ suitably so that $\mathrm{Arf}(q') = \mathrm{Arf}(q)$, and the rest of the argument follows as above. \square

B Action on the homology classes of cylinders

We return to the setting of Sections 2-10. Recall from Section 3.5 that the crucial quantities

$$|(\bar{\Gamma}_d \alpha)_{c;a}| = |\{\beta \in \bar{\Gamma}_d \alpha \mid \beta(c) = a \pmod{d}\}|$$

are independent of the choice of cylinder C . We emphasize that in the non-hyperelliptic case, this was proven in the above sections. However, the proof of this fact in the hyperelliptic setting relied on the result in this appendix, which we establish below. In this section, we provide a conceptual explanation for this phenomenon, by showing that the monodromy group acts transitively on homology classes of cylinders as elements of $H_1(X \setminus \mathcal{B}; \mathbb{Z})$. For a non-hyperelliptic stratum component in genus $g \geq 5$, this follows from [CS23, Corollary 1.2], which proves the stronger result that the mapping class group monodromy acts transitively on *isotopy classes* of cores of cylinders. A similar strategy will be used in the hyperelliptic setting. However, in this paper we obtain results for *all* stratum components, including the non-hyperelliptic components in genus $g < 5$ to which the results of [CS23] do not apply.

Proposition B.1. *Let \mathcal{H} denote a non-hyperelliptic component of a stratum, and let $(X, \omega) \in \mathcal{H}$ be a basepoint with distinguished set \mathcal{B} . Then the monodromy group $\bar{\Gamma} \leq \text{PAut}(H_1(X \setminus \mathcal{B}; \mathbb{Z}))$ acts transitively on the set of homology classes $[c] \in H_1(X \setminus \mathcal{B}; \mathbb{Z})$ represented by core curves of cylinders on (X, ω) .*

The same result holds in the hyperelliptic setting, but the method of proof will differ.

Proposition B.2. *Let \mathcal{H} denote a hyperelliptic component of a stratum, and let $(X, \omega) \in \mathcal{H}$ be a basepoint with distinguished set \mathcal{B} (possibly including some ordinary marked points). Then the monodromy group $\bar{\Gamma} \leq \text{PAut}(H_1(X \setminus \mathcal{B}; \mathbb{Z}))$ acts transitively on the set of homology classes $[c] \in H_1(X \setminus \mathcal{B}; \mathbb{Z})$ represented by core curves of cylinders on (X, ω) .*

Corollary B.3. *Let \mathcal{H} denote a component of a stratum (hyperelliptic or otherwise), and let $(X, \omega) \in \mathcal{H}$ be a basepoint with distinguished set \mathcal{B} . Let c and $c' \in H_1(X \setminus \mathcal{B}; \mathbb{Z})$ represent the cores of cylinders C and C' , respectively. Then*

$$|\{\beta \in \bar{\Gamma}_d \alpha \mid \beta(c) = a\}| = |\{\beta \in \bar{\Gamma}_d \alpha \mid \beta(c') = a\}|.$$

Proof. By Proposition B.1 or Proposition B.2 as appropriate, the monodromy action of the stratum on $H_1(X \setminus \mathcal{B}; \mathbb{Z}/d\mathbb{Z})$ is transitive on the set of classes $[c]$ represented by core curves of cylinders. Let $f \in \bar{\Gamma}_d$ take c to c' . Then

$$\{\beta \in \bar{\Gamma}_d \alpha \mid \beta(c') = a\} = \{\beta \in \bar{\Gamma}_d \alpha \mid \beta(f(c)) = a\} = f \cdot \{\beta \in \bar{\Gamma}_d \alpha \mid \beta(c) = a\}.$$

In particular, this shows the desired equality of cardinalities. □

B.1 Non-hyperelliptic

Duality. To prove Proposition B.1, it is necessary to discuss some linear-algebraic preliminaries. The core subtlety in what follows is that the monodromy computation Theorem 6.1 determines the action of $\bar{\Gamma}$ on *relative* homology $H_1(X, \mathcal{B}; \mathbb{Z})$. On the other hand, the cylinder classes lie in *excision* homology $H_1(X \setminus \mathcal{B}; \mathbb{Z})$. These spaces are dual to each other via the relative intersection pairing

$$\langle \cdot, \cdot \rangle : H_1(X \setminus \mathcal{B}; \mathbb{Z}) \otimes H_1(X, \mathcal{B}; \mathbb{Z}) \rightarrow \mathbb{Z}$$

discussed in Section 3.3. Accordingly, we can consider $\gamma \in \bar{\Gamma}$ as acting on $H_1(X \setminus \mathcal{B}; \mathbb{Z})$ (denoted γ^*), given by the adjoint formula

$$\langle \gamma^* x, y \rangle = \langle x, \gamma y \rangle.$$

Outline. To prove Proposition B.1, in Definition B.5 we will construct yet another algebraic distillate of the winding number function ϕ induced from the translation surface structure on X , this time, a function $q^* : H_1(X \setminus \mathcal{B}; \mathbb{Z}) \rightarrow \mathbb{Z}/2\mathbb{Z}$ that behaves like a $\mathbb{Z}/2\mathbb{Z}$ quadratic form (note that the intersection pairing $\langle \cdot, \cdot \rangle$ on $H_1(X \setminus \mathcal{B}; \mathbb{Z})$ is *degenerate*). In Lemma B.6, we then show that the adjoint action of $\bar{\Gamma}$ on $H_1(X \setminus \mathcal{B}; \mathbb{Z})$ is characterized by q^* -invariance. Using a geometric interpretation of q^* originally obtained in Lemma 5.7, we then prove Proposition B.1.

Duality and geometric splittings. Let X° be a geometric splitting, inducing an algebraic splitting $H_1(X, \mathcal{B}; \mathbb{Z}) \cong H_1(X; \mathbb{Z}) \oplus \tilde{H}_0(\mathcal{B}; \mathbb{Z})$ as in Definition 5.5. This also induces a splitting

$$H_1(X \setminus \mathcal{B}; \mathbb{Z}) \cong H_1(X; \mathbb{Z}) \oplus \tilde{H}_0(\mathcal{B}; \mathbb{Z}) \quad (28)$$

via the identifications $H_1(X; \mathbb{Z}) \cong H_1(X^\circ; \mathbb{Z}) \leq H_1(X \setminus \mathcal{B}; \mathbb{Z})$ and $\tilde{H}_0(\mathcal{B}; \mathbb{Z}) \cong \ker(i_*)$, where $i_* : H_1(X \setminus \mathcal{B}; \mathbb{Z}) \rightarrow H_1(X; \mathbb{Z})$ is induced by inclusion.

Observe that the subspaces $H_1(X; \mathbb{Z})$ of $H_1(X, \mathcal{B}; \mathbb{Z})$ and $H_1(X \setminus \mathcal{B}; \mathbb{Z})$ are dual to each other under the relative intersection pairing $\langle \cdot, \cdot \rangle$, and that each is isomorphic to $H_1(X^\circ; \mathbb{Z})$. Hence, these subspaces can be canonically identified. This leads to the following subtlety: there are two natural ways to pass from a basis of $H_1(X, \mathcal{B}; \mathbb{Z})$ to a basis for $H_1(X \setminus \mathcal{B}; \mathbb{Z})$: either by taking the dual basis with respect to the intersection pairing, or else by identifying the common subspace $H_1(X; \mathbb{Z})$, and identifying the complementary factors by $\langle \cdot, \cdot \rangle$.

In the sequel we will need to understand the relationship between these. Let $\beta = \{x_1, y_1, \dots, x_g, y_g, a_1, \dots, a_{n-1}\}$ be a basis for $H_1(X, \mathcal{B}; \mathbb{Z})$ such that the restriction $\{x_1, y_1, \dots, x_g, y_g\}$ to $H_1(X; \mathbb{Z})$ is a symplectic basis. Let $\beta^* \subset H_1(X \setminus \mathcal{B}; \mathbb{Z})$ denote the dual basis to β induced by the pairing $\langle \cdot, \cdot \rangle$ on $H_1(X \setminus \mathcal{B}; \mathbb{Z}) \otimes H_1(X, \mathcal{B}; \mathbb{Z})$, and let $\beta' \subset H_1(X \setminus \mathcal{B}; \mathbb{Z})$ denote the basis $\{x_1, y_1, \dots, x_g, y_g, a_1^*, \dots, a_{n-1}^*\}$ given by identifying the common subspace $H_1(X; \mathbb{Z})$ and adjoining the dual basis elements a_i^* to a_i .

Lemma B.4. Let $\gamma = \begin{pmatrix} M & A \\ 0 & I \end{pmatrix} \in \text{PAut}(H_1(X, \mathcal{B}; \mathbb{Z}))$ be given (in coordinates β as discussed above). With β^* and β' as above, the matrix for γ^* in β^* is given by the transpose, while the matrix for γ^* in β' is given by

$$[\gamma^*] = \begin{pmatrix} M^{-1} & 0 \\ A^T & I \end{pmatrix}.$$

Proof. That γ^* is represented in the dual basis by the transpose is a general property of dual bases. The second claim follows from a routine calculation involving the change-of-basis matrix from β^* to β' , keeping in mind that for a symplectic matrix M , the formula $M^{-1} = J^{-1}M^T J$ (where J represents the symplectic form) holds, and that conjugating by J also passes from a basis to its symplectic dual on $H_1(X; \mathbb{Z})$. \square

Definition B.5. In the coordinates on $H_1(X \setminus \mathcal{B}; \mathbb{Z})$ given by the geometric splitting as in (28), the form $q^* : H_1(X \setminus \mathcal{B}; \mathbb{Z}) \rightarrow \mathbb{Z}/2\mathbb{Z}$ is given by

$$q^*(x, b) = q(x) + \langle b, \kappa \rangle,$$

where (q, κ) is the *associated data* of the geometric splitting, cf. Definition 5.6. Note that the expression $q(x)$ exploits the canonical identification of the subspaces $H_1(X; \mathbb{Z})$ of $H_1(X, \mathcal{B}; \mathbb{Z})$ and $H_1(X \setminus \mathcal{B}; \mathbb{Z})$.

Note that by Lemmas 5.4, 5.7, when $c \subset X \setminus \mathcal{B}$ is a simple closed curve with homology class $[c] = (x, b)$ relative to the geometric splitting, there is an equality

$$q^*(x, b) = \phi(c) + 1 \pmod{2}, \quad (29)$$

showing coordinate-independence of q^* . In particular, since $\phi(c) = 0$ for c the core curve of a cylinder, it follows that $q^*(x, b) = 1$ for (x, b) the homology class of any cylinder.

Lemma B.6. Let $\gamma \in \text{PAut}(H_1(X, \mathcal{B}; \mathbb{Z}))$ be given. Relative to a geometric splitting X° of X with associated data (q, κ) , the condition $\Theta_{q, \kappa}(\gamma) = 0$ holds if and only if $\gamma^* \in \text{PAut}(H_1(X \setminus \mathcal{B}; \mathbb{Z}))$ satisfies

$$q^*(\gamma^*(x, b)) = q^*(x, b)$$

for all $(x, b) \in H_1(X \setminus \mathcal{B}; \mathbb{Z})$.

Proof. Recall from Definition 5.8 that if $\gamma = \begin{pmatrix} M & A \\ 0 & I \end{pmatrix}$ in the coordinates given by a geometric splitting, then for all $x \in H_1(X; \mathbb{Z}/2\mathbb{Z})$,

$$\Theta_{q,\kappa}(\gamma)(x) = q(Mx) + q(x) + \langle Mx, A\kappa \rangle.$$

Recalling that q^* is valued in $\mathbb{Z}/2\mathbb{Z}$, we compute the quantity $q^*(\gamma^*(x, b)) + q^*(x, b)$ using Lemma B.4:

$$\begin{aligned} q^*(\gamma^*(x, b)) + q^*(x, b) &= q^* \left(\begin{pmatrix} M & A \\ 0 & I \end{pmatrix}^* (x, b) \right) + q^*(x, b) \\ &= q^*(M^{-1}x, A^T x + b) + q^*(x, b) \\ &= q(M^{-1}x) + \langle A^T x + b, \kappa \rangle + q(x) + \langle b, \kappa \rangle \\ &= q(M^{-1}x) + q(x) + \langle A^T x, \kappa \rangle \\ &= q(M^{-1}x) + q(x) + \langle x, A\kappa \rangle \\ &= \Theta_{q,\kappa}(\gamma)(M^{-1}x). \end{aligned}$$

This calculation shows that $\gamma \in \ker(\Theta_{q,\kappa})$ if and only if γ^* preserves q^* , as desired. \square

Proof of Proposition B.1. Let c and c' denote the core curves of cylinders C and C' on $X \setminus \mathcal{B}$, respectively. Choose a geometric splitting X° with associated data (q, κ) and relative to this, write $[c] = (c_0, c_1)$ and $[c'] = (c'_0, c'_1)$ in the induced coordinates on $H_1(X \setminus \mathcal{B}; \mathbb{Z})$. Since c, c' are core curves of cylinders, each one is *nonseparating* on $X \setminus \mathcal{B}$, so that the elements c_0 and c'_0 of $H_1(X^\circ; \mathbb{Z}) \leq H_1(X \setminus \mathcal{B}; \mathbb{Z})$ are both primitive. By the discussion following Definition B.5 above, $q^*(c_0, c_1) = q^*(c'_0, c'_1) = 1$, since $\phi(C) = \phi(C') = 0$.

The proof will now follow from a slightly more general result, where we only assume the q^* values are *equal*, not necessarily both 1.

Claim. *Let $(a, b), (a', b') \in H_1(X \setminus \mathcal{B}; \mathbb{Z})$ be given such that a and a' are both primitive and such that $q^*(a, b) = q^*(a', b')$. Then there exists $\gamma \in \ker(\Theta_{q,\kappa})$ such that $(a', b') = \gamma^*(a, b)$.*

Proof of Claim: We seek to construct $\gamma = \begin{pmatrix} M & A \\ 0 & I \end{pmatrix}$ such that $a' = M^{-1}a$ and $b' - b = A^T a$ that additionally satisfies the constraint

$$q(M^{-1}x) + q(x) + \langle x, A\kappa \rangle = 0 \tag{30}$$

for all $x \in H_1(X; \mathbb{Z})$.

Assume first that $\kappa = 0$. In this case, the condition $q^*(a, b) = q^*(a', b')$ simplifies to $q(a) = q(a')$. Since a and a' are assumed to be primitive, by Proposition A.1.3, there exists $M \in \text{Sp}(2g, \mathbb{Z})[q]$ such that $M^{-1}a = a'$.

Choosing any A such that $A^T a = b' - b$, it follows that $\gamma = \begin{pmatrix} M & A \\ 0 & I \end{pmatrix}$ satisfies $\gamma^*(a, b) = (a', b')$ and that (30) holds.

Assume then that $\kappa \neq 0$. Choose any $M \in \text{Sp}(2g, \mathbb{Z})$ such that $M^{-1}a = a'$. Note that the function $x \mapsto q(M^{-1}x) + q(x)$ defines an element $\bar{\eta}$ of $H^1(X; \mathbb{Z}/2\mathbb{Z})$; define $\eta \in H^1(X; \mathbb{Z})$ as an integral lift. The matrix A must then satisfy the conditions

$$A^T a = b' - b \quad \text{and} \quad A\kappa = \eta. \tag{31}$$

If such an A exists, then

$$\langle b' - b, \kappa \rangle = \langle A^T a, \kappa \rangle = \langle a, A\kappa \rangle = \langle a, \eta \rangle.$$

Conversely, an argument in elementary linear algebra shows that if $\langle b' - b, \kappa \rangle = \langle a, \eta \rangle$, then A satisfying (31) holds. Recalling the definition of η ,

$$\langle a, \eta \rangle = q(M^{-1}a) + q(a),$$

so that (recalling that we are working mod 2),

$$\langle b' - b, \kappa \rangle + \langle a, \eta \rangle = q(M^{-1}a) + \langle b', \kappa \rangle + q(a) + \langle b, \kappa \rangle = q^*(a', b') + q^*(a, b) = 0$$

by hypothesis. \square

B.2 Hyperelliptic

Proof of Proposition B.2. In the hyperelliptic setting, we will deduce this from the stronger claim that the topological monodromy $\Gamma \leq \text{Mod}(X, \mathcal{B})$ acts transitively on the set of isotopy classes of cores of cylinders. To see this, we write the distinguished set \mathcal{B} as $\mathcal{B} = \mathcal{Z} \sqcup \mathcal{P}$ as before, with \mathcal{Z} the set of zeros and \mathcal{P} a set of $k \geq 0$ ordinary marked points. Recall from Lemma 9.2 and Lemma 9.3 that there is a short exact sequence

$$1 \rightarrow PB_k(X \setminus \mathcal{B}) \rightarrow \Gamma \rightarrow \text{SMod}(X, \mathcal{Z}) \rightarrow 1, \quad (32)$$

where $PB_k(X \setminus \mathcal{B})$ is the surface braid group consisting of point pushes of the set \mathcal{P} on the surface $X \setminus \mathcal{B}$.

Recall from Lemma 9.2 that $\text{SMod}(X, \mathcal{Z})$ is a $\mathbb{Z}/2\mathbb{Z}$ extension of a braid group on $\mathbb{C}\mathbb{P}^1$, where the marked points correspond to the singularities of the associated quadratic differential. This acts *transitively* on the set of isotopy classes of simple arcs on $X/\langle \iota \rangle$, both of whose endpoints lie in the set of distinguished points. We will show that under the projection $X \mapsto X/\langle \iota \rangle$, the core of any cylinder gives a double branched cover of such an arc.

According to [Lin15, Lemma 2.1], every maximal cylinder $C \subset X$ is ι -invariant: $\iota(C) = C$. As ι acts geometrically on the translation surface X by rotation by π , it follows that the core curve c of C is left invariant by ι and gives a double branched cover of a simple arc on $X/\langle \iota \rangle$ connecting two Weierstrass points. Thus any two such core curves c, c' are related by an element of $\text{SMod}(X, \mathcal{Z})$. To complete the argument, we must consider the isotopy class of a core curve of a cylinder on the fully marked surface $X \setminus \mathcal{B}$. But from the description of Γ given in (32), it is clear that Γ acts transitively on the set of isotopy classes of curves $\{\tilde{c}\}$ on $X \setminus \mathcal{B}$ that are isotopic to a fixed curve $c \subset X \setminus \mathcal{Z}$ after forgetting the ordinary points \mathcal{P} . We conclude that Γ acts transitively on the set of core curves c of cylinders on X as desired. \square

References

- [AˆC79] Norbert AˆCampo, *Tresses, monodromie et le groupe symplectique*, Comment. Math. Helv. **54** (1979), no. 2, 318–327. MR 535062
- [AEZ16] Jayadev S. Athreya, Alex Eskin, and Anton Zorich, *Right-angled billiards and volumes of moduli spaces of quadratic differentials on $\mathbb{C}\mathbb{P}^1$* , Ann. Sci. ˆEc. Norm. Supˆer. (4) **49** (2016), no. 6, 1311–1386, With an appendix by Jon Chaika. MR 3592359
- [Agg19] Amol Aggarwal, *Large genus asymptotics for Siegel-Veech constants*, Geom. Funct. Anal. **29** (2019), no. 5, 1295–1324. MR 4025514
- [Agg20] ———, *Large genus asymptotics for volumes of strata of Abelian differentials*, J. Amer. Math. Soc. **33** (2020), no. 4, 941–989, With an appendix by Anton Zorich. MR 4155217
- [Agg21] ———, *Large genus asymptotics for intersection numbers and principal strata volumes of quadratic differentials*, Invent. Math. **226** (2021), no. 3, 897–1010. MR 4337974
- [AMY13] Artur Avila, Carlos Matheus, and Jean-Christophe Yoccoz, *$SL(2, \mathbb{R})$ -invariant probability measures on the moduli spaces of translation surfaces are regular*, Geom. Funct. Anal. **23** (2013), no. 6, 1705–1729. MR 3132901
- [AMY18] ———, *Zorich conjecture for hyperelliptic Rauzy-Veech groups*, Math. Ann. **370** (2018), no. 1-2, 785–809. MR 3747501
- [AMY19] ———, *The Kontsevich-Zorich cocycle over Veech-McMullen family of symmetric translation surfaces*, J. Mod. Dyn. **14** (2019), 21–54. MR 3959355
- [AN20] David Aulicino and Duc-Manh Nguyen, *Rank 2 affine manifolds in genus 3*, J. Differential Geom. **116** (2020), no. 2, 205–280. MR 4168204

- [Api18] Paul Apisa, *GL₂ℝ orbit closures in hyperelliptic components of strata*, Duke Math. J. **167** (2018), no. 4, 679–742. MR 3769676
- [Arf41] Cahit Arf, *Untersuchungen über quadratische formen in körpern der charakteristik 2, I*, J. für die reine und Angew. Math. **183** (1941), 148–167.
- [AW23] Paul Apisa and Alex Wright, *High rank invariant subvarieties*, Ann. of Math. (2) **198** (2023), no. 2, 657–726. MR 4635302
- [BG15] Max Bauer and Elise Goujard, *Geometry of periodic regions on flat surfaces and associated Siegel-Veech constants*, Geom. Dedicata **174** (2015), 203–233. MR 3303050
- [Cal20] Aaron Calderon, *Connected components of strata of Abelian differentials over Teichmüller space*, Comment. Math. Helv. **95** (2020), no. 2, 361–420. MR 4115287
- [CMSZ20] Dawei Chen, Martin Möller, Adrien Sauvaget, and Don Zagier, *Masur-Veech volumes and intersection theory on moduli spaces of Abelian differentials*, Invent. Math. **222** (2020), no. 1, 283–373. MR 4145791
- [CMZ18] Dawei Chen, Martin Möller, and Don Zagier, *Quasimodularity and large genus limits of Siegel-Veech constants*, J. Amer. Math. Soc. **31** (2018), no. 4, 1059–1163. MR 3836563
- [CS21a] Aaron Calderon and Nick Salter, *Higher spin mapping class groups and strata of abelian differentials over Teichmüller space*, Adv. Math. **389** (2021), Paper No. 107926, 56. MR 4289049
- [CS21b] ———, *Relative homological representations of framed mapping class groups*, Bulletin of the London Mathematical Society **53** (2021), no. 1, 204–219.
- [CS23] ———, *Framed mapping class groups and the monodromy of strata of abelian differentials*, J. Eur. Math. Soc. (JEMS) **25** (2023), no. 12, 4719–4790. MR 4662301
- [DGZZ21] Vincent Delecroix, Élise Goujard, Peter Zograf, and Anton Zorich, *Masur-Veech volumes, frequencies of simple closed geodesics, and intersection numbers of moduli spaces of curves*, Duke Math. J. **170** (2021), no. 12, 2633–2718.
- [DRC⁺24] Vincent Delecroix, Julian Rüth, Frédéric Chapoton, Samuel Lelièvre, chriszhang3, Fougeroc, and Diana Davis, *flatsurf/surface-dynamics: 0.5.0*, March 2024.
- [DZ20] Vincent Delecroix and Anton Zorich, *Cries and whispers in wind-tree forests*, What’s next?—the mathematical legacy of William P. Thurston, Ann. of Math. Stud., vol. 205, Princeton Univ. Press, Princeton, NJ, 2020, pp. 83–115.
- [EKZ11] Alex Eskin, Maxim Kontsevich, and Anton Zorich, *Lyapunov spectrum of square-tiled cyclic covers*, J. Mod. Dyn. **5** (2011), no. 2, 319–353. MR 2820564 (2012h:37073)
- [EKZ14] ———, *Sum of Lyapunov exponents of the Hodge bundle with respect to the Teichmüller geodesic flow*, Publ. Math. Inst. Hautes Études Sci. **120** (2014), 207–333. MR 3270590
- [EM01] Alex Eskin and Howard Masur, *Asymptotic formulas on flat surfaces*, Ergodic Theory Dynam. Systems **21** (2001), no. 2, 443–478. MR 1827113
- [EM18] Alex Eskin and Maryam Mirzakhani, *Invariant and stationary measures for the SL(2, ℝ) action on moduli space*, Publ. Math. Inst. Hautes Études Sci. **127** (2018), 95–324. MR 3814652
- [EMM15] Alex Eskin, Maryam Mirzakhani, and Amir Mohammadi, *Isolation, equidistribution, and orbit closures for the SL(2, ℝ) action on moduli space*, Ann. of Math. (2) **182** (2015), no. 2, 673–721. MR 3418528

- [EMS03] Alex Eskin, Howard Masur, and Martin Schmoll, *Billiards in rectangles with barriers*, Duke Math. J. **118** (2003), no. 3, 427–463. MR 1983037 (2004c:37059)
- [EMZ03] Alex Eskin, Howard Masur, and Anton Zorich, *Moduli spaces of abelian differentials: the principal boundary, counting problems, and the Siegel-Veech constants*, Publ. Math. Inst. Hautes Études Sci. (2003), no. 97, 61–179. MR 2010740 (2005b:32029)
- [EZ15] Alex Eskin and Anton Zorich, *Volumes of strata of Abelian differentials and Siegel-Veech constants in large genera*, Arnold Math. J. **1** (2015), no. 4, 481–488. MR 3434506
- [FM12] Benson Farb and Dan Margalit, *A primer on mapping class groups*, Princeton Mathematical Series, vol. 49, Princeton University Press, 2012.
- [FMZ14] Giovanni Forni, Carlos Matheus, and Anton Zorich, *Zero Lyapunov exponents of the Hodge bundle*, Comment. Math. Helv. **89** (2014), no. 2, 489–535. MR 3225454
- [Fou18] Charles Fougerson, *Lyapunov exponents of the Hodge bundle over strata of quadratic differentials with large number of poles*, Math. Res. Lett. **25** (2018), no. 4, 1213–1225. MR 3882161
- [Gou15] Elise Goujard, *Siegel-Veech constants for strata of moduli spaces of quadratic differentials*, Geom. Funct. Anal. **25** (2015), no. 5, 1440–1492. MR 3426059
- [GR19] Rodolfo Gutiérrez-Romo, *Classification of Rauzy-Veech groups: proof of the Zorich conjecture*, Invent. Math. **215** (2019), no. 3, 741–778. MR 3935031
- [KZ03] Maxim Kontsevich and Anton Zorich, *Connected components of the moduli spaces of Abelian differentials with prescribed singularities*, Invent. Math. **153** (2003), no. 3, 631–678. MR MR2000471 (2005b:32030)
- [Lin15] Kathryn Lindsey, *Counting invariant components of hyperelliptic translation surfaces*, Isr. J. Math. **210** (2015), 125–146.
- [LM14] Eduard Looijenga and Gabriele Mondello, *The fine structure of the moduli space of abelian differentials in genus 3*, Geom. Dedicata **169** (2014), 109–128. MR 3175239
- [Mas82] Howard Masur, *Interval exchange transformations and measured foliations*, Ann. of Math. (2) **115** (1982), no. 1, 169–200. MR 644018 (83e:28012)
- [McM13] Curtis T. McMullen, *Braid groups and Hodge theory*, Math. Ann. **355** (2013), no. 3, 893–946. MR 3020148
- [MMW17] Curtis T. McMullen, Ronen E. Mukamel, and Alex Wright, *Cubic curves and totally geodesic subvarieties of moduli space*, Ann. of Math. (2) **185** (2017), no. 3, 957–990. MR 3664815
- [MS91] Howard Masur and John Smillie, *Hausdorff dimension of sets of nonergodic measured foliations*, Ann. of Math. (2) **134** (1991), no. 3, 455–543. MR 1135877
- [MW18] Maryam Mirzakhani and Alex Wright, *Full-rank affine invariant submanifolds*, Duke Math. J. **167** (2018), no. 1, 1–40. MR 3743698
- [MW21] Dan Margalit and Rebecca R. Winarski, *Braids groups and mapping class groups: the Birman-Hilden theory*, Bull. Lond. Math. Soc. **53** (2021), no. 3, 643–659. MR 4275077
- [MYZ14] Carlos Matheus, Jean-Christophe Yoccoz, and David Zmiaikou, *Homology of origamis with symmetries*, Ann. Inst. Fourier (Grenoble) **64** (2014), no. 3, 1131–1176. MR 3330166
- [NS64] M. Newman and J. R. Smart, *Symplectic modular groups*, Acta Arith. **9** (1964), 83–89. MR 162862

- [RW14] Oscar Randal-Williams, *Homology of the moduli spaces and mapping class groups of framed, r -Spin and Pin surfaces*, J. Topol. **7** (2014), no. 1, 155–186. MR 3180616
- [Sau18] Adrien Sauvaget, *Volumes and Siegel-Veech constants of $\mathcal{H}(2G - 2)$ and Hodge integrals*, Geom. Funct. Anal. **28** (2018), no. 6, 1756–1779. MR 3881834
- [Sch19] Martin Schmoll, *Orbit classification and asymptotic constants for d -symmetric covers*, Dynamical systems and random processes, Contemp. Math., vol. 736, Amer. Math. Soc., Providence, RI, 2019, pp. 201–238. MR 4011913
- [Vee82] William A. Veech, *Gauss measures for transformations on the space of interval exchange maps*, Ann. of Math. (2) **115** (1982), no. 1, 201–242. MR 644019 (83g:28036b)
- [Vee98] ———, *Siegel measures*, Ann. of Math. (2) **148** (1998), no. 3, 895–944. MR 1670061
- [Vor05] Yaroslav Vorobets, *Periodic geodesics on generic translation surfaces*, Algebraic and topological dynamics, Contemp. Math., vol. 385, Amer. Math. Soc., Providence, RI, 2005, pp. 205–258. MR 2180238
- [Wri14] Alex Wright, *The field of definition of affine invariant submanifolds of the moduli space of abelian differentials*, Geom. Topol. **18** (2014), no. 3, 1323–1341. MR 3254934
- [Wri15] ———, *Cylinder deformations in orbit closures of translation surfaces*, Geom. Topol. **19** (2015), no. 1, 413–438. MR 3318755