

Linear submanifolds and strata of k -differentials:
invariants and applications

Dissertation
zum Erlangen des Doktorgrades
der Naturwissenschaften

vorgelegt beim Fachbereich 12 Informatik und Mathematik
der Johann Wolfgang Goethe-Universität
Frankfurt am Main

von
Johannes Schwab
aus Offenbach am Main

Frankfurt am Main 2024
(D30)

Vom Fachbereich 12 Informatik und Mathematik der
Johann Wolfgang Goethe-Universität als Dissertation angenommen.

Dekan: Prof. Dr. Bastian von Harrach-Sammet

Gutachter: Prof. Dr. Martin Möller
Prof. Dr. Martin Ulirsch
Dr. Vincent Delecroix

Datum der Disputation: 09.04.2024

Contents

Introduction	v
Chapter I. Realizability of tropical pluri-canonical divisors	1
1. Introduction	1
2. Tropical k -Hodge Bundle	4
3. Moduli space of multi-scale k -differentials	8
4. Reduction to realizability of normalized covers	16
5. The realizability locus	21
6. Examples	31
7. Appendix: Nonemptiness of boundary strata	32
Chapter II. Chern classes of linear submanifolds with applications to spaces of k -differentials and ball quotients	35
1. Introduction	35
2. Logarithmic differential forms and toric varieties	39
3. The closure of linear submanifolds	39
4. Evaluation of tautological classes	47
5. Chern classes of the cotangent bundle via the Euler sequence	50
6. Example: Euler characteristic of the eigenform locus	53
7. Strata of k -differentials	54
8. Ball quotients	65
Chapter III. The multi-scale boundary of the gothic locus	73
1. Introduction	73
2. The gothic locus	76
3. The moduli space of critical values	78
4. The boundary superset	82
5. Non-horizontal divisors	84
6. Non-obvious Teichmüller curves	87
7. Horizontal strata, cylinders, and cusps of Teichmüller curves	89
8. Towards the fundamental class	92
Chapter IV. Visible Lagrangians for Hitchin systems and pillowcase covers	97
1. Introduction	97
2. Symplectic geometry	98
3. Hitchin systems	99
4. Parallelogram-tiled surfaces and pillowcase covers	105
5. Multifold pillowcase covers	107
6. Visible Lagrangians over lines in the Hitchin base	110
Bibliography	117
Zusammenfassung	123

Introduction

1. \mathbb{R} -linear submanifolds

Looking at a rectangular billiard table and starting a ball in one of the corners, you might wonder “will the ball hit a corner ever again?”. For the moment let us assume that your billiard table has integer side lengths. Tracing the trajectory of the ball along all the reflections in the rectangle becomes tedious very soon. Instead of reflecting the ball when it hits the side of the table it is much more convenient to reflect the table instead as in Figure 1. Now the copies of the corners of the table span a lattice in the Euclidean plane. In this setting the question “will the ball hit a corner ever again?” translates to the question “will the ball hit a lattice point ever again?”. The latter question can be answered immediately: The ball will hit a lattice point if and only if the angle between the trajectory and one side of the polygon is a rational multiple of π .

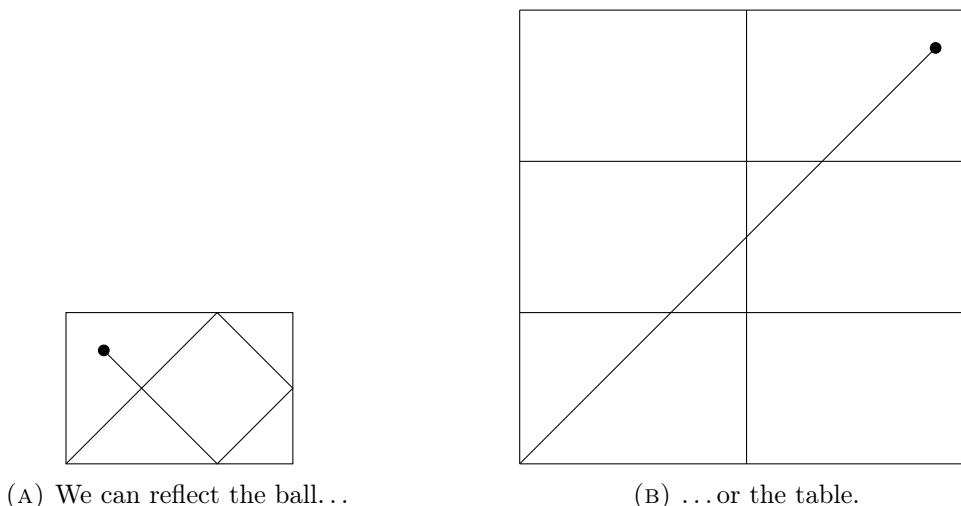


FIGURE 1. A rectangular billiard table

But what will happen if the ball does not hit a corner? In this case the trajectory will obviously be infinite. We can actually say more: The trajectory will be dense in the table as a consequence of the Dirichlet approximation theorem. This motivates the following definition.

DEFINITION 1.1. *A billiard table where each trajectory is either closed or dense is said to have optimal dynamics.*

As we have seen, rectangular billiards with integer (and more generally rational) side lengths have optimal dynamics. For more complicated billiard tables our approach with lattices will not work any more. Instead, we can use the following observation. Flipping the table once produces a table with a different orientation opposed to the original table. Flipping the new table again in the same direction produces a table with the same orientation as the original table. So instead of producing a new table (as we have done to obtain the lattice) we can glue the second table to the first table to produce a manifold. If we think of the original surface as being embedded in the Gaussian plane, the manifold obtained in this way will naturally have the structure of a Riemann surface X . By pulling back the differential form dz from the plane to the surface we obtain a differential form ω on X . This process of obtaining (X, ω) for a polygon is called unfolding. We will see this again in Chapter III. The pair (X, ω) is called a flat surface, as ω induces a flat metric on X .

Conversely, each flat surface (X, ω) can be represented by polygons: if $\gamma_1, \dots, \gamma_n$ is a basis of the relative homology $H_1(X, Z(\omega))$, then the sides of the polygons are given by

$$(1) \quad \int_{\gamma_i} \omega.$$

Let us return to our billiard table. The trajectories of the ball correspond to geodesics in X with respect to the metric given by ω . So the notion of having optimal dynamics may be rephrased in terms of geodesics on (X, ω) . The Hodge bundle $\Omega\mathcal{M}_g \rightarrow \mathcal{M}_g$ over the moduli space of compact Riemann surfaces of genus g is a moduli space for flat surfaces. The group $\mathrm{GL}_2(\mathbb{R})^+$ acts on this moduli space via the action on the polygon representation of the surfaces. Veech [Vee89] [Vee91] observed that a flat surface (X, ω) has optimal dynamics if and only if its $\mathrm{GL}_2(\mathbb{R})^+$ -orbit is closed.

DEFINITION 1.2. *If $(X, \omega) \in \Omega\mathcal{M}_g$ has closed $\mathrm{GL}_2(\mathbb{R})^+$ -orbit then (X, ω) is called Veech surface. The image of $\mathrm{GL}_2(\mathbb{R})^+ \cdot (X, \omega)$ in \mathcal{M}_g is called Teichmüller curve.*

We have seen the simplest example of a Veech surface above: The unfolding of our rectangular billiard table (X, ω) . Since Veech's observation a lot of effort has gone into the classification of Teichmüller curves. As one can easily obtain new Veech surfaces from known ones via covering constructions, one is mainly interested in classifying those that do not arise via covering constructions, so called *primitive* Teichmüller curves. In genus 2, 3, and 4 we know infinitely many primitive Teichmüller curves, discovered by Veech [Vee89], Ward [War98], Bouw-Möller [BM10], McMullen [McM03] [McM06], Calta [Cal04], Vorobets [HS01] and Kenyon-Smillie [KS00]. In each genus greater than four we know only of finitely many primitive Teichmüller curves which all belong to the series discovered by Bouw-Möller, and it is an open question to decide if there exist infinitely many primitive Teichmüller curves in every genus.

Let us fix a genus g and let $\mu = (m_1, \dots, m_n)$ be an integer partition of $2g - 2$. If (X, ω) is a flat surface of genus g we say that ω has *type* μ if ω has precisely n zeros of orders m_1, \dots, m_n . The moduli space of flat surfaces $\Omega\mathcal{M}_g$, also known as the moduli space of abelian differentials, comes with natural stratification by the types of the differentials, and we denote by $\Omega\mathcal{M}_{g,n}(\mu)$ the stratum of differentials of type μ . The integrals in (1) provide local coordinates for the stratum, called *period coordinates*.

DEFINITION 1.3. *A subspace $\Omega\mathcal{H} \subseteq \Omega\mathcal{M}_{g,n}(\mu)$ is called K -linear submanifold if it is cut out by linear equations in period coordinates with coefficients in the field K .*

As the action of $\mathrm{GL}_2(\mathbb{R})^+$ preserves equations with coefficients in \mathbb{R} , any \mathbb{R} -linear submanifolds is the closure of an union of $\mathrm{GL}_2(\mathbb{R})^+$ -orbits. By the ground braking result of Eskin-Mirzakhani-Mohammadi the converse is also true.

THEOREM 1.4 ([EMM15]). *Every $\mathrm{GL}_2(\mathbb{R})^+$ -orbit closure is an \mathbb{R} -linear submanifold.*

An important invariance of a linear submanifold $\Omega\mathcal{H}$ is its *rank* which can be defined as follows. Over $\Omega\mathcal{H}$ we consider the bundle H^1 whose fiber over (X, ω) is $H^1(X, \mathbb{C})$ and the bundle H_{rel}^1 whose fiber over (X, ω) is $H^1(X, Z(\omega), \mathbb{C})$. Let $p : H_{\mathrm{rel}}^1 \rightarrow H^1$ be the natural map. By work of Avila-Eskin-Möller [AEM17] the space $p(T(\Omega\mathcal{H}))$ is symplectic, in particular of even dimension, and we define the rank of $\Omega\mathcal{H}$ as $\frac{1}{2} \dim p(T(\Omega\mathcal{H}))$.

Teichmüller curves (or more precisely the orbit closures of the corresponding Veech surfaces) are \mathbb{R} -linear submanifolds of rank 1. It came as a surprise when McMullen-Mukamel-Wright [MMW17] discovered the first primitive \mathbb{R} -linear submanifold of rank 2, the so-called gothic locus. By now only 6 additional primitive \mathbb{R} -linear submanifolds of rank 2 have been discovered by Eskin-McMullen-Mukamel-Wright [EMMW20]. There is computational evidence for the existence of at least one more such submanifold [DR23]. The existence of a \mathbb{R} -linear submanifold of rank at least 3 is a completely open question.

2. Chern classes of linear submanifolds

To classify mathematical objects (as for example \mathbb{R} -linear submanifolds, which are in fact complex orbifolds) it is often a good idea to compute their invariants. For complex orbifolds an important invariant are the Chern classes. There is a slight problem here: On a

linear submanifold $\Omega\mathcal{H}$ the group \mathbb{C}^\times acts by scaling the differential, so $\Omega\mathcal{H}$ is a trivial \mathbb{C}^\times -bundle and hence all Chern classes are 0. So if we want to expect a useful answer, we should not consider $\Omega\mathcal{H}$ but instead its projectivization $\mathcal{H} := \Omega\mathcal{H}/\mathbb{C}^\times$. In Chapter II, which is joint work with Matteo Costantini and Martin Möller, we prove a formula for the full Chern character of the logarithmic cotangent bundle of a linear submanifold in Theorem II.1.2. This allows in particular to derive a closed formula for the Euler characteristic of a linear submanifold. For a linear submanifold \mathcal{H} , we denote by $\xi_{\mathcal{H}} = c_1(\mathcal{O}(-1))$ the first Chern class of the tautological bundle.

THEOREM 2.1 (Theorem II.1.3). *Let $\mathcal{H} \rightarrow \mathbb{P}\Omega\mathcal{M}_{g,n}(\mu)$ be a projectivized linear submanifold. The orbifold Euler characteristic of \mathcal{H} is given by*

$$\chi(\mathcal{H}) = (-1)^d \sum_{L=0}^d \sum_{\Gamma \in \text{LG}_L(\mathcal{H})} \frac{K_\Gamma^{\mathcal{H}} \cdot N_\Gamma^\top}{|\text{Aut}_{\mathcal{H}}(\Gamma)|} \cdot \prod_{i=0}^{-L} \int_{\mathcal{H}_\Gamma^{[i]}} \xi_{\mathcal{H}_\Gamma^{[i]}}^{d_\Gamma^{[i]}},$$

where the integrals are over the normalization of the closure $\overline{\mathcal{H}} \rightarrow \mathbb{P}\overline{\Xi\mathcal{M}}_{g,n}(\mu)$ inside the moduli space of multi-scale differentials and similar integrals over boundary strata, where

- $\mathcal{H}_\Gamma^{[i]}$ are the linear submanifolds at level i of Γ as defined in Section II.3.5,
- $d_\Gamma^{[i]} := \dim(\mathcal{H}_\Gamma^{[i]})$ is the projectivized dimension,
- $K_\Gamma^{\mathcal{H}}$ is the product of the number of prong-matchings on each edge of Γ that are actually contained in the linear submanifold $\overline{\mathcal{H}}$,
- $\text{Aut}_{\mathcal{H}}(\Gamma)$ is the set of automorphism of the graph Γ whose induced action on a neighborhood of $D_\Gamma^{\mathcal{H}}$ preserves $\overline{\mathcal{H}}$,
- $d := \dim(\mathcal{H})$ is the projectivized dimension.

For most of the notions used in this theorem we refer the reader to Chapter II. The one thing we want to highlight is the fact that the theorem makes use of the moduli space of multi-scale differentials $\mathbb{P}\overline{\Xi\mathcal{M}}_{g,n}(\mu)$. This compactification of the projectivized stratum $\mathbb{P}\Omega\mathcal{M}_{g,n}(\mu)$ was constructed by Bainbridge-Chen-Gendron-Grushevsky-Möller [BCGGM18] [BCGGM19b]. The objects in the boundary roughly consist of nodal curves together with a differential on each irreducible component and a level structure on the irreducible components. For more details see Section I.3. The boundary of this compactification again admits a stratification where the strata are indexed by so-called enhanced level graphs: those are the dual graphs of the underlying nodal curves together with some decoration that records information about the differentials and the level structure. This compactification and enhanced level graphs will appear multiple times hereinafter.

3. The gothic locus

The gothic locus $\Omega G \subseteq \Omega\mathcal{M}_{4,6}(0^3, 2^3)$ is the orbit closure of the unfoldings of all quadrilaterals with angles $(\frac{1}{6}\pi, \frac{1}{6}\pi, \frac{1}{6}\pi, \frac{3}{2}\pi)$. Its name stems from the fact that it contains the surfaces depicted in Figure 2 which resembles the layout of a gothic cathedral. As mentioned above, the gothic locus ΩG was the first known primitive \mathbb{R} -linear subvariety of rank 2. This locus has additional quite surprising properties: It contains a dense subset of primitive Teichmüller curves, and it comes with a natural map to $\mathcal{M}_{1,3}$, and the image of ΩG under this map, the so-called flex locus, is a totally geodesic surface with respect to the Teichmüller metric.

A Teichmüller curve in a stratum of meromorphic differentials is called *obvious* if it arises as the intersection of a covering construction and a condition on the residues. In Chapter III we will analyze the boundary of the closure $\mathbb{P}\overline{\Xi G} := \overline{\mathbb{P}\Omega G} \subseteq \mathbb{P}\overline{\Xi\mathcal{M}}_{4,6}(0^3, 2^3)$. While this analysis is still work in progress, we will present some partial results. As part of this boundary we find an example for a non-obvious Teichmüller curve.

THEOREM 3.1 (Theorem III.1.3). *Let $(X, \omega) \subseteq \Omega\mathcal{M}_{1,6}(-3^2, 2^3)$ be the canonical cover of the 6-differential of type $(-10, -5, 3)$. The differential (X, ω) generates a non-obvious Teichmüller curve. In the chart in Figure 3 this Teichmüller curve is given by the equations*

$$w_i = -w_{i+3} \quad \text{for } i = 1, 2, 3 \quad \text{and} \quad w_1 + w_3 + w_5 = 0.$$




FIGURE 2. The gothic cathedral (opposite sides are identified unless indicated otherwise)

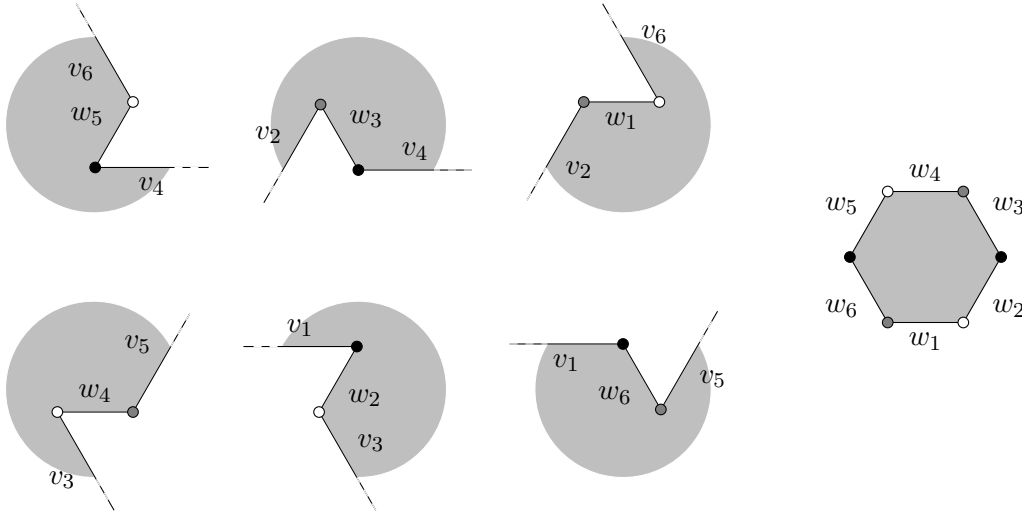


FIGURE 3. A surface of infinite area generating a non-obvious Teichmüller curve in the stratum $\Omega\mathcal{M}_{1,6}(-3^2, 2^3)$

By now we have not succeeded in determining precisely which boundary strata are intersected by the gothic locus, but we have some partial information. For the horizontal boundary strata we prove:

PROPOSITION 3.2 (Proposition III.1.4). *The gothic locus $\mathbb{P}\Xi\overline{G}$ only intersects the horizontal strata listed in Figure 4.*

Recall that the gothic locus itself contains an infinite number of primitive Teichmüller curves. Those curves are not compact. Thus the closure of each of this Teichmüller curves will intersect the boundary of $\mathbb{P}\Xi\overline{G}$ in a number of points, called cusps. Such cusps can only be contained in purely horizontal boundary strata. For the strata corresponding to the enhanced level graphs depicted in Figure 4 we will prove:

PROPOSITION 3.3 (Proposition III.1.6). *The interior of each of the four horizontal strata $D_{\Gamma_1}^G$, $D_{\Gamma_2}^G$, $D_{\Gamma_3}^G$ and $D_{\Gamma_{20}}^G$ contains cusps of a primitive Teichmüller curve contained in the gothic locus ΩG . The interior of the stratum $D_{\Gamma_{19}}^G$ contains cusps of a non-primitive Teichmüller curve.*

One might hope to apply Theorem 2.1 to compute the Euler characteristic of the gothic locus $\mathbb{P}\Omega G$. This theorem can be rephrased in such a way that it is sufficient to know the fundamental class of the image of the gothic locus in $\overline{\mathcal{M}}_{4,6}$. We will outline an approach to the computation of this fundamental class in Chapter III. We are currently short on the necessary computational tools to actually carry out this computation.

4. Strata of k -differentials

Let us consider a generalization of the above setting: Instead of strata of abelian differentials we may also consider strata of k -differentials $\Omega^k\mathcal{M}_{g,n}(\mu)$ parametrizing pairs (X, η) ,




FIGURE 4. The purely horizontal boundary strata in the gothic locus

where η is a section of $\Omega^{\otimes k}(X)$. Here μ is an integer partition of $k(2g - 2)$. As in the abelian case those strata admit a compactification called strata of multi-scale k -differentials $\mathbb{P}\Xi^k \overline{\mathcal{M}}_{g,n}(\mu)$. Work in this direction was done by Bainbridge-Chen-Gendron-Grushevsky-Möller [BCGGM19a] and Costantini-Möller-Zachhuber [CMZ19]. In Chapter II we establish the precise orbifold structure of those spaces. Via a covering construction those spaces are related to linear submanifolds (not necessarily \mathbb{R} -linear), and we give a closed formula for their Euler characteristic in Corollary II.1.5.

We implemented this formula in a Sage package called `diffstrata` which is part of the package `admcycles` [DSZ21]. `Diffstrata` was originally created by Costantini-Möller-Zachhuber [CMZ23] to allow the evaluation of their formula for the Chern classes of strata of abelian differentials [CMZ22]. We extended `diffstrata` to work with all strata of k -differentials. As an example, the Euler characteristic and the Masur-Veech volume of the stratum $\mathbb{P}\Omega^2 \mathcal{M}_{2,2}(-1, 5)$ can be computed with the following commands.

```
sage: from admcycles.diffstrata import Stratum
sage: X = Stratum((-1,5), k=2)
sage: X.euler_characteristic()
-7/15
sage: X.masur_veech_volume()
28/135*pi^4
```

The Euler characteristics of the minimal strata in genus 2 are listed for small k in Table 1. As the above example already shows, the package `diffstrata` can actually do more than only compute the Euler characteristic. For example it can

- list all non-horizontal boundary strata of a stratum of k -differentials,
- compute arbitrary intersection products in the vertical tautological ring (i.e. the ring generated by all non-horizontal strata, ψ - and κ -classes),
- compute the push-forward of classes from the stratum to the moduli space of marked stable curves.

The main limitation `diffstrata` currently has is the fact that it can not work with horizontal strata. This would allow for much more general computations, including pull-backs of arbitrary tautological classes from the moduli space of stable curves to a stratum of k -differentials.

As an application of our computation of the Chern classes we prove that for specific types μ the space $\Omega^k \mathcal{M}_{g,n}(\mu)$ is birational equivalent to a quotient of the complex unit ball.

k	1	3	4	5	6	7	8	9
$\chi(\mathbb{P}\Omega^k \mathcal{M}_{2,1}(2k))$	$-\frac{1}{40}$	$\frac{1}{3}$	$\frac{3}{2}$	$\frac{21}{5}$	9	18	30	51

TABLE 1. Euler characteristics of some minimal strata of k -differentials computed with `diffstrata`

THEOREM 4.1 (Theorem II.1.7). *Suppose that $\mu = (-a_1, \dots, -a_5)$ is a tuple with $a_i \geq 0$ and with the condition*

$$\left(1 - \frac{a_i}{k} - \frac{a_j}{k}\right)^{-1} \in \mathbb{Z} \quad \text{if } a_i + a_k < k$$

for all $i \neq j$. Then there exists a birational contraction morphism $\mathbb{P}\Xi^k \overline{\mathcal{M}}_{0,5}(\mu) \rightarrow \overline{\mathfrak{B}}$ onto a smooth proper DM-stack $\overline{\mathfrak{B}}$ for some ball quotient \mathfrak{B} .

These ball quotients have previously been constructed by Deligne-Mostow [DM86] and Thurston [Thu98] by different methods.

5. The tropical k -Hodge bundle

Similar to the abelian case, the boundary of $\mathbb{P}\Xi^k \overline{\mathcal{M}}_{g,n}(\mu)$ is indexed by enhanced k -level graphs. One might hope to understand the structure of the boundary by studying the graphs themselves as abstract objects. To be able to talk about continuous (un)degeneration of graphs, we might assign to each edge a real length. This leads to the definition of a tropical curve.

DEFINITION 5.1. *A tropical curve is a connected graph with real edge lengths and weights $g : V \rightarrow \mathbb{N}$ assigned to each vertex.*

The level structure of the enhanced level graph can be recorded by assigning integer slopes to the edges, and the zeros of the differential can be recorded by adding legs to the graph. In Chapter I, which is joint work with Felix Röhrlé, we define a tropical k -Hodge bundle $\mathbb{P}\Omega^k M_g^{\text{trop}}$ as, roughly speaking, the space of all tropical curves with legs and integer slopes at the edges that fulfill certain compatibility conditions made to mimic the behaviour of k -differentials. This space is a generalized cone complex, but not equidimensional, see Theorem I.1.1.

The gap between the classical world, that is $\mathbb{P}\Omega^k \mathcal{M}_g$, and the tropical world, that is $\mathbb{P}\Omega^k M_g^{\text{trop}}$, is bridged by a process called tropicalization: There is a continuous tropicalization map $\text{trop}_{\Omega^k} : \mathbb{P}\Omega^k \mathcal{M}_g^{\text{an}} \rightarrow \mathbb{P}\Omega^k M_g^{\text{trop}}$. However, this map is not surjective. The image of trop_{Ω^k} is the *realizability locus*. In Theorem I.1.4 we give a combinatorial criterion to determine for a given element of $\mathbb{P}\Omega^k M_g^{\text{trop}}$ if it is contained in the realizability locus. This allows us to prove that the realizability locus is much nicer behaved than the tropical k -Hodge bundle itself.

THEOREM 5.2 (Theorem I.1.5). *For $k \geq 2$, the realizability locus admits the structure of a generalized cone complex, all of whose maximal cones have dimension $(2+2k)(g-1)-1$. The fiber in the realizability locus over a maximal-dimensional cone in M_g^{trop} is a generalized cone complex, all whose maximal cones have relative dimension $(2k-1)(g-1)$.*

6. Pillowcase covers and visible Lagrangians

Given a covering of Riemann surfaces $f : X \rightarrow \mathbb{P}^1$ that is ramified above at most four points, there is a unique (up to scale) quadratic differential η of type (-1^4) on \mathbb{P}^1 such that the simple poles are supported at the four branch points. We might pull back this differential to X to obtain a quadratic differential $q = f^* \eta$. A quadratic differential (X, q) that arises in this way is called a pillowcase cover. In Chapter IV, which is joint work with Johannes Horn, we study Riemann surfaces X that admit multiple quadratic differentials q_1, \dots, q_n such that

- the vanishing loci of q_1, \dots, q_n are pairwise different,

- the pairs (X, q_i) are pillowcase covers.

We call such a Riemann surface X a *multifold pillowcase cover*. We say that the pillowcase cover $f : X \rightarrow \mathbb{P}^1$ is *uniform* if every fiber consists of ramification points of the same ramification index.

THEOREM 6.1 (Theorem IV.5.2). *For infinitely many genera g there exist multifold uniform pillowcase covers with simple zeros only.*

An example of a multifold pillowcase cover is the Klein quartic. Note that our definition of a multifold pillowcase cover does not require the q_i to be non-isomorphic. Nevertheless, we also provide an example of a multifold pillowcase cover where the quadratic differentials are not all isomorphic.

This has a nice application in the theory of Higgs bundles as follows. For a complex reductive group G consider the moduli space of G -Higgs bundles \mathcal{M}_G with the Hitchin map $\text{Hit} : \mathcal{M}_G \rightarrow \mathcal{B}_G$. A complex Lagrangian $\mathcal{L} \subseteq \mathcal{M}_G$ is called *visible* if the restriction of the Hitchin map factors through a proper subvariety $\mathcal{B}' = \text{Hit}(\mathcal{L}) \subsetneq \mathcal{B}$. For the special case $G = \text{SL}(2, \mathbb{C})$ we will prove:

THEOREM 6.2 (Theorem IV.1.2). *Let $q \in H^0(X, K_X^2)$ be a quadratic differential with simple zeros only. Then there exists a visible Lagrangian*

$$\mathcal{L} \rightarrow \mathcal{B}' = \{tq \mid t \in \mathbb{C}\} \subset \mathcal{B}_{\text{SL}(2, \mathbb{C})}(X)$$

if and only if (X, q) is a pillowcase cover.

Hence our examples of multifold pillowcase covers give examples of Riemann surfaces for which there exist several lines in the $\text{SL}(2, \mathbb{C})$ -Hitchin base $\mathcal{B}_{\text{SL}(2, \mathbb{C})}$ associated to visible Lagrangians.

Notes on collaborations

Chapter I was created in joint work with Felix Röhrle and appeared as a preprint [RS21]. Chapter II was created in joint work with Matteo Costantini and Martin Möller and appeared as a preprint [CMS23]. Sections III.2 and III.3 of Chapter III build on parts of my master's thesis. Chapter IV was created in joint work with Johannes Horn and appeared as a preprint [HS23].

Acknowledgements

I'm grateful to my advisor Martin Möller for introducing me to the fascinating world of flat geometry, his interest in my work, uncountable hours of discussions, his constant support and guidance, and not least his patience in hunting down contradictions (of which I produced a sufficient amount for the next decades to come).

I'm also grateful to my coauthors Matteo Costantini, Johannes Horn, Martin Möller and Felix Röhrle. Being able to do math with so many nice people made the process of writing this thesis even more enjoyable.

I thank Martin Ulirsch for his help in his role as Felix's advisor while we worked on Chapter I. I thank Quentin Gendron for helpful conversations on his joint work with Guillaume Tahar and in particular for sharing the content of the (at the time) forthcoming [GT21b]. I thank Vincent Delecroix and Johannes Schmitt for helpful comments and suggestions while working on the Sage package `diffstrata`, and Johannes Schmitt for his comments on an earlier draft of Chapter II.

Last but not least, I'm grateful to my family, in particular my parents Heike Steidl and Andreas Schwab and my wife Johanna Winkler for all their support and encouragement, without which I wouldn't have been able to study math, let alone work on this thesis.

My research was supported by the LOEWE-Schwerpunkt "Uniformisierte Strukturen in Arithmetik und Geometrie", by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) - project number 456557832 and the Collaborative Research Centre TRR 326 *Geometry and Arithmetic of Uniformized Structures*, project number 444845124.

Realizability of tropical pluri-canonical divisors

1. Introduction

The close analogy between Riemann surfaces and graphs was first described in [Mik05]. Since then many definitions in tropical geometry have been modeled with the aim that tropicalization of algebro-geometric objects produces the corresponding tropical objects. The realizability problem then asks whether a given instance of the tropical notion does indeed arise in this way. For example consider this question for curves with effective divisors. Given an abstract tropical curve Γ (i.e. a vertex weighted metric graph) with an effective divisor D , the realizability problem asks if there exists a smooth proper algebraic curve X with effective divisor \tilde{D} of the same degree and rank as D such that the tropicalization of (X, \tilde{D}) is (Γ, D) (see Section 2.5 below for details on the tropicalization of curves with divisor). This question is very difficult in general. In fact, [Car15] shows that it satisfies a version of Murphy's law that makes a general solution seem unlikely. In this article we restrict our attention to the special case of effective pluri-canonical divisors and give a complete characterization of those tropical objects that are realizable over an algebraically closed base field of characteristic 0.

Let $g \geq 2$ and $k \geq 1$ be integers. In algebraic geometry the k -Hodge bundle $\Omega^k \mathcal{M}_g$ is a moduli space parametrizing pairs (X, η) consisting of a smooth curve X of genus g and a k -differential η , i.e. a global section of the k -th tensor power of the canonical bundle on X . We start our exposition in Section 2 with a review of basic definitions for tropical curves, divisors, and linear equivalence. In Section 2.4 we then construct a tropical counterpart of the projectivized moduli space $\mathbb{P}\Omega^k \mathcal{M}_g$. More precisely, we prove:

THEOREM 1.1. *There exists a generalized cone complex in the sense of [ACP15, Section 2.6] which parametrizes pairs $([\Gamma], D)$ of isomorphism classes of abstract tropical curves Γ of genus g and effective divisor $D \in \text{Div}(\Gamma)$ linearly equivalent to k times the canonical divisor K_Γ . We denote this space by $\mathbb{P}\Omega^k M_g^{\text{trop}}$ and call it the tropical k -Hodge bundle. It is not equidimensional. The dimension of a maximal cone is $(3 + 2k)(g - 1)$.*

Tropicalization of curves with divisor has been described e.g. in [BJ16, Section 6.3] in the following way. Let X be a smooth curve over a non-Archimedean field and let D be an effective divisor on X . Let \mathcal{X} be the stable model of X . Define Γ to be the dual graph of the nodal special fiber endowed with edge lengths obtained from the deformation parameters of the nodes. Furthermore, via the specialization map in the sense of Baker [Bak08, Section 2C] the divisor D gives rise to a divisor on Γ . In [MUW21] the authors gave a description of this procedure as a continuous map between moduli spaces. By restricting this general construction to effective pluri-canonical divisors we obtain a continuous tropicalization map $\text{trop}_{\Omega^k} : \mathbb{P}\Omega^k \mathcal{M}_g^{\text{an}} \rightarrow \mathbb{P}\Omega^k M_g^{\text{trop}}$ in Section 2.5. Throughout, $(\)^{\text{an}}$ denotes analytification in the sense of [Ber90]. The dimension of $\mathbb{P}\Omega^k \mathcal{M}_g$ is at most $(2 + 2k)(g - 1) - 1$ by [BCGGM19a, Theorem 1.1]. Comparing this to the dimension of $\mathbb{P}\Omega^k M_g^{\text{trop}}$ obtained in Theorem 1.1 we see that trop_{Ω^k} cannot be surjective. The realizability problem amounts to describe the image of trop_{Ω^k} , the *realizability locus*, as a subset of $\mathbb{P}\Omega^k M_g^{\text{trop}}$.

As it turns out, the question can be reduced to the realizability problem for so-called *normalized covers*. Recall that the authors of [BCGGM19a] canonically associate to any smooth curve with k -differential an admissible, normalized, cyclic, potentially ramified and disconnected cover $\pi : \widehat{X} \rightarrow X$ with abelian differential ω on \widehat{X} and a deck transformation $\tau : \widehat{X} \rightarrow \widehat{X}$ such that $\omega^k = \pi^* \eta$ and $\tau^* \omega = \zeta \omega$ for a primitive k -th root of unity ζ . Recall that the k -Hodge bundle admits a natural stratification by so-called *types* $\mu =$

$(m_1, \dots, m_n) \in \mathbb{Z}^n$ such that the sum of the m_i is $k(2g - 2)$. The moduli space of multi-scale k -differentials $\mathbb{P}\Xi^k \overline{\mathcal{M}}_{g,n}(\mu)$ was introduced in [CMZ19]. It is a compactification of the projectivized strata $\mathbb{P}\Omega^k \mathcal{M}_g(\mu)$ of the k -Hodge bundle and parametrizes normalized covers in its interior. Accordingly, we define a *tropical normalized cover* in Definition 4.9 to be a tropical Hurwitz cover $\widehat{\Gamma} \rightarrow \Gamma$ in the sense of [CMR16] such that the legs of $\widehat{\Gamma}$ and Γ encode a (pluri-)canonical divisor and additionally we require a deck transformation on $\widehat{\Gamma}$ as well as compatibility conditions mimicking the above. Again we introduce a tropical moduli space in analogy to the algebro-geometric setting.

THEOREM 1.2. *There is a moduli space of tropical normalized covers, denoted $\mathbb{P}\Xi^k M_g^{\text{trop}}$. It carries the structure of a generalized cone complex. The dimension of a maximal cone is $(3 + 2k)(g - 1)$. Furthermore, there is a well-defined, continuous, closed, and proper tropicalization map $\text{trop}_{\Xi^k} : \mathbb{P}\Omega^k \mathcal{M}_g(1, \dots, 1)^{\text{an}} \rightarrow \mathbb{P}\Xi^k M_g^{\text{trop}}$.*

The following corollary will be an easy consequence of the properties of trop_{Ξ^k} . It is the key to reduce our original realizability problem to the realizability of tropical normalized covers (see Corollary 4.13 for the precise statement).

COROLLARY 1.3. *A tropical curve Γ with effective pluri-canonical divisor $D = kK_{\Gamma} + (f)$ is realizable if and only if there exists a realizable tropical normalized cover $\pi : \widehat{\Gamma} \rightarrow \Gamma$ such that the legs of Γ encode D .*

In Section 5 we solve the realizability problem for tropical normalized covers using similar ideas as in [MUW21]. This means that we proceed in two steps.

- (1) For every vertex v in Γ we realize $\pi|_{\pi^{-1}(\{v\})}$ with a normalized cover of smooth curves with meromorphic differentials.
- (2) We glue these parts into a normalized cover of nodal curves which lies in the boundary of $\mathbb{P}\Xi^k \overline{\mathcal{M}}_{g,n}(1, \dots, 1)$ and smoothen these curves.

Observe at this point that for a tropical curve with pluri-canonical divisor $D = kK_{\Gamma} + (f)$ the zero and pole orders of any of the realizations in step (1) are already determined. More precisely, the rational function f gives rise to a canonical *enhanced level graph structure* on Γ (see Definition 3.13 and Lemma 4.1 for details). Consequently, we only need to specify (k -)residues to proceed. Both steps from above impose restrictions on the possible choices. For step (1) these are given by [GT21a], [GT21b] and [GT22a] and lead us to the notions of *illegal vertex* (such a vertex is never realizable) and *inconvenient vertex* (here special care in choosing residues has to be taken). Step (2) is only feasible if the *global residue condition* (see Definition 3.12) as well as the above mentioned compatibilities with π and τ are respected. We will ensure this by assigning residues along τ -orbits of simple closed cycles in $\widehat{\Gamma}$. In contrast to the case $k = 1$ that was treated in [MUW21] not any cycle is sufficient for this purpose. Rather we have to ask for each inconvenient vertex for a corresponding *admissible cycle* (Definition 5.7) or an *independent pair of cycles* (Definition 5.9). Having introduced the necessary notation we state our main result in Theorem 5.11 which roughly says the following.

THEOREM 1.4. *Fix an algebraically closed base field of characteristic 0. Let $g \geq 2$ and fix an integer $k \geq 1$. Let $\pi : \widehat{\Gamma} \rightarrow \Gamma$ be a tropical normalized cover and $D = kK_{\Gamma} + (f)$ be an effective pluri-canonical divisor on Γ . The pair is realizable if and only if the following conditions hold.*

- (i) *There is no illegal vertex in π .*
- (ii) *For every edge \widehat{e} in $\widehat{\Gamma}$ for which $f \circ \pi$ is constant there is an effective cycle in $\widehat{\Gamma}$ through \widehat{e} .*
- (iii) *For every inconvenient vertex v in Γ there is an admissible cycle in $\widehat{\Gamma}$ through one of the preimages \widehat{v} or there is an independent pair of cycles.*

Corollary 1.3 together with Theorem 1.4 provide a complete description of the locus of realizable curves in $\mathbb{P}\Omega^k M_g^{\text{trop}}$. We conclude Section 5 with a result in analogy to [MUW21, Theorem 6.6].

THEOREM 1.5. *For $k \geq 2$, the realizability locus admits the structure of a generalized cone complex, all of whose maximal cones have dimension $(2 + 2k)(g - 1) - 1$. The fiber*

in the realizability locus over a maximal-dimensional cone in M_g^{trop} is a generalized cone complex, all whose maximal cones have relative dimension $(2k - 1)(g - 1)$.

In Section 6 we illustrate the language that we developed throughout Section 5 by applying our theory to give a complete description of the realizability locus over the dumbbell graph for $k = 2$.

- REMARK 1.6. (i) For the general study of realizability of curves with divisors as in [Car15] it is crucial to ask for realizations by divisors of the same rank, i.e. Baker's specialization inequality [Bak08, Corollary 2.11] should be an equality. Without this condition every effective divisor on a tropical curve of genus ≥ 2 would be realizable, simply because tropicalization of curves with divisors is surjective onto the tropical moduli space by [MUW21, Theorem 3.2]. For pluri-canonical divisors, this rank condition is always implicitly included, simply because the (tropical) rank of a (tropical) pluri-canonical divisor is always equal to $(2k - 1)(g - 1)$ by the (tropical) Riemann-Roch theorem (see [GK08] for the tropical Riemann-Roch theorem).
- (ii) For $k = 1$ our Theorem 1.1 contains [LU17, Theorem 4.3 (i) and (ii)] as special case. Furthermore, every tropical normalized cover with $k = 1$ is necessarily the identity and the conditions from Theorem 1.4 reduce to the conditions of [MUW21, Theorem 6.3] (see Remark 5.12 for details). Hence we recover the results of [MUW21].
- (iii) Our construction of $\mathbb{P}\Omega^k M_g^{\text{trop}}$ is a straight-forward generalization of the tropical Hodge bundle introduced in [LU17]. In fact, Theorem 1.1 could have been proved with the same ideas as in [LU17].
- (iv) The techniques involved in the proof of Theorem 1.4 give a very similar criterion to decide which boundary strata of the moduli space of multi-scale k -differentials $\mathbb{P}\Xi^k \overline{\mathcal{M}}_{g,n}(\mu)$ are nonempty, see Appendix 7.
- (v) In Theorem 1.4 we are concerned with finding realizations in the principal stratum $\mu = (1, \dots, 1)$. A slight modification of the ideas from the proof can be used to give a criterion for realizability in any other stratum as well, see Remark 4.2.
- (vi) The reason for reducing the realizability problem to the seemingly more complicated question for normalized covers is subtle. On the classical side, a k -differential is called *primitive* if it is not a power of some k' differential with $k' < k$ and k' dividing k . This property is entirely invisible on the tropical side, i.e. when realizing a tropical curve consisting of a single vertex we may choose to realize it with a primitive or non-primitive differential. This choice has to be fixed in order to proceed and corresponds precisely to choosing a normalized cover.

Very little is known about the topology of the projectivized strata $\mathbb{P}\Omega^k \mathcal{M}_g(\mu)$ of the k -Hodge bundle. We believe that our criterion will be useful for further research in this direction.

In the recent and much-celebrated work [CGP21] the authors computed the top weight cohomology of \mathcal{M}_g from the reduced rational cohomology of the link of M_g^{trop} . The same technique was shortly after applied to compute the top weight cohomology for some instances of the moduli space of abelian varieties \mathcal{A}_g in [BBCMMW21]. In both cases it is vital to identify the tropical moduli space with the boundary complex of the classical moduli space (see e.g. [ACP15] for the case of curves). With our description of the realizability locus in $\mathbb{P}\Omega^k M_g^{\text{trop}}$ we take the first step towards a similar computation of the top weight cohomology of strata of k -differentials.

We want to highlight some work related to this article. Amini-Baker-Brugallé-Rabinoff study in [ABBR15a] and [ABBR15b] the realizability problem for finite harmonic morphisms of tropical curves. Without the extra data of a pluri-canonical divisor, the global obstruction to realizability induced by the global k -residue condition from [BCGGM19a] does not occur. Indeed [ABBR15a, Corollary 1.6] shows that the only obstructions occur locally at the vertices. Furthermore, adding the data of an effective divisor to the problem the condition on the rank of the realization being equal to the rank of the tropical divisor is a non-trivial condition, see [ABBR15b, Section 5].

By [CJP15, Theorem 1.1] every effective divisor class on a chain of loops is realizable by an effective divisor of the same rank. This is not a contradiction to our findings in

Section 6.2 because in this article we consider the much harder problem of realizability of divisors rather than divisor classes.

In some sense Baker-Nicaise use in [BN16] a different framework to discuss tropicalizations of k -differentials. More precisely, they associate to any pluri-canonical form on a curve X a so-called *weight function* on the Berkovich analytification X^{an} . This is related to our divisor-based point of view since the induced divisor of a weight function is again a pluri-canonical divisor by [BN16, Corollary 3.2.5].

2. Tropical k -Hodge Bundle

Fix integers $g \geq 2$ and $k \geq 1$. In this section we will describe a tropical version $\mathbb{P}\Omega^k M_g^{\text{trop}}$ of $\mathbb{P}\Omega^k \mathcal{M}_g$ together with a tropicalization map

$$\text{trop}_{\Omega^k} : \mathbb{P}\Omega^k \mathcal{M}_g^{\text{an}} \longrightarrow \mathbb{P}\Omega^k M_g^{\text{trop}} .$$

The underlying set of the tropical k -Hodge bundle $\mathbb{P}\Omega^k M_g^{\text{trop}}$ parametrizes pairs $([\Gamma], D)$ of isomorphism classes of stable tropical curves Γ of genus g and effective divisors D linearly equivalent to kK_Γ . In the special case of $k = 1$ we recover the description of the tropical Hodge bundle from [LU17, Definition 4.1]. In Section 2.4 we prove Theorem 1.1. To this end we use the moduli space $\text{Div}_{g,d}^{\text{trop}}$ of tropical curves with effective divisor of fixed degree $d = k(2g - 2)$ which was constructed in [MUW21, Definition 2.1] and exhibit $\mathbb{P}\Omega^k M_g^{\text{trop}}$ as a locus in $\text{Div}_{g,d}^{\text{trop}}$. The tropicalization map trop_{Ω^k} is defined in Section 2.5 by restricting the more general tropicalization map from [MUW21, Section 3.1].

We conclude this section by defining the *realizability locus* as the image of trop_{Ω^k} and formally state the realizability problem in Section 2.6.

2.1. Tropical curves.

DEFINITION 2.1. A *graph* is a tuple $G = (V, H, L, \iota, a)$ where

- (i) the finite sets V , H , and L are the vertices, half-edges and legs of the graph respectively,
- (ii) the map $\iota : H \rightarrow H$ is a fixpoint-free involution on the half-edges H that determines the edges of the graph, and
- (iii) the map $a : H \cup L \rightarrow V$ assigns to every half-edge and leg the incident vertex.

For a graph G let $E := \{\{h, h'\} \in H^2 \mid \iota(h) = h'\}$ be the set of unoriented edges. In the following we will often denote a graph simply as 3-tuple (V, E, L) of vertices, edges and legs with the rest of the underlying data remaining implicit. If there are no legs, we abbreviate further and simply write (V, E) .

The *valence* of a vertex $v \in V$ is defined as $\text{val}(v) := |a^{-1}(v)|$.

A *(vertex) weighted graph* is a graph G together with a map $g : V \rightarrow \mathbb{N}$. The weighted graph is called *stable* if for each vertex $v \in V$ the stability condition

$$2g(v) - 2 + \text{val}(v) > 0$$

holds. The *genus* of a weighted graph is defined to be

$$g(G) := b_1(G) + \sum_{v \in V} g(v),$$

where $b_1(G)$ is the first Betti number of G .

A *tropical curve* is a connected weighted metric graph Γ given by the data of a graph G , vertex weights $g : V \rightarrow \mathbb{N}$ and edge lengths $l : E \rightarrow \mathbb{R}_{>0}$. We call $g(v)$ the genus of the vertex v . The topological realization of Γ is the metric space obtained by gluing real intervals $[0, l(e)]$ for every edge and $[0, \infty)$ for every leg according to adjacency in G . Any weighted graph (G', g') giving rise to the same topological realization is referred to as a *model* for Γ . Note that every stable tropical curve has a unique minimal model in the sense of minimal number of edges and vertices. We will usually not distinguish between topological realization and minimal model.

The *genus* of Γ is defined to be

$$g(\Gamma) := b_1(G) + \sum_{v \in V} g(v),$$

The tropical curve Γ is called *stable* if its minimal model is stable.

2.2. Moduli of tropical curves. The following description of the moduli space $M_{g,n}^{\text{trop}}$ of stable tropical curves of genus g with n legs can be found e.g. in [ACP15, Section 4]. Note that in the description of $M_{g,n}^{\text{trop}}$ the legs are usually assumed to be labeled. Let G be a weighted graph and let e be an edge in G . We denote by $G/\{e\}$ the graph that arises from G by contracting e into a single vertex v of weight

$$g(v) = \begin{cases} g(v_1) + g(v_2) & \text{if } e \text{ was connecting } v_1 \text{ and } v_2 \\ g(v_1) + 1 & \text{if } e \text{ was a self-loop at vertex } v_1. \end{cases}$$

Define the category $\mathcal{G}_{g,n}$ with objects being stable weighted graphs of genus g with n legs and morphism are generated by weighted edge contractions $G \rightarrow G/\{e\}$ as well as graph automorphisms respecting the labeling of the legs.

Given $G \in \mathcal{G}_{g,n}$ we associate to it the rational polyhedral cone $\sigma_G := \mathbb{R}_{\geq 0}^{E(G)}$. In fact, this defines a contravariant functor from $\mathcal{G}_{g,n}$ to the category of rational polyhedral cones, where edge contractions are taken to isomorphisms onto faces. The moduli space is now defined as

$$M_{g,n}^{\text{trop}} := \varinjlim_{\mathcal{G}_{g,n}} \sigma_G.$$

Note that the points of $M_{g,n}^{\text{trop}}$ are in one-to-one correspondence with isomorphism classes of tropical curves of genus g with n legs. A topological space arising as colimit over a finite diagram of rational polyhedral cones where all morphisms are isomorphisms onto faces is called *generalized cone complex* in [ACP15, Section 2.6].

2.3. Divisors on tropical curves. Let Γ be a tropical curve without legs. A *divisor* D on Γ is an element of the free abelian group generated by the points in the topological realization of Γ . We denote the abelian group of divisors on Γ by $\text{Div}(\Gamma)$. A divisor $D = \sum a_p p$ is called *effective* if $a_p \geq 0$ for every p . In this case we write $D \geq 0$. The *degree* of D is defined as $\deg(D) := \sum a_p$. The *support* of D is $\text{supp}(D) := \{p \in \Gamma \mid a_p \neq 0\}$. By definition, the support is a finite subset of Γ . One often imagines an effective divisor D as a pile of $D(p) = a_p$ “chips” at every point $p \in \text{supp}(D)$ overing The data of an effective divisor D on a tropical curve Γ without legs is equivalent to a tropical curve $\tilde{\Gamma}$ arising from Γ by attaching $D(p)$ many legs at every $p \in \text{supp}(D)$. From now on all tropical curves are a priori without legs, but given a divisor we will pass to the equivalent representation $\tilde{\Gamma}$ whenever convenient.

A rational function on Γ is a continuous function $f : \Gamma \rightarrow \mathbb{R}$ whose restriction to any edge is piece-wise linear with integer slopes. We denote the set of rational functions on Γ by $\text{Rat}(\Gamma)$. Every $f \in \text{Rat}(\Gamma)$ gives rise to an induced divisor

$$(f) := \sum_{p \in \Gamma} (\text{sum of outgoing slopes of } f \text{ at } p) \cdot p \in \text{Div}(\Gamma).$$

Note that (f) is indeed a finite sum. Two divisors $D, D' \in \text{Div}(\Gamma)$ are *linearly equivalent* if there exists $f \in \text{Rat}(\Gamma)$ such that $D = D' + (f)$. In this case we write $D \sim D'$. Analogously to [LU17, Definition 3.1] we define:

DEFINITION 2.2. Let $D \in \text{Div}(\Gamma)$ be a divisor. We define the *linear system* of D to be

$$|D| := \{D' \in \text{Div}(\Gamma) \mid D' \geq 0 \text{ and } D \sim D'\}.$$

The canonical divisor of a tropical curve Γ without legs is defined as

$$K_\Gamma := \sum_{v \in V} (2g(v) - 2 + \text{val}(v))v.$$

If Γ has legs, then we define K_Γ to be the canonical divisor of the tropical curve arising from Γ by removing the legs. Note that contrary to the classical situation there is a canonical element in the canonical linear system. Furthermore, note that $\deg K_\Gamma = 2g(\Gamma) - 2$. The elements of $|kK_\Gamma|$ are called *pluri-canonical*.

2.4. Tropical k -Hodge bundle. We are now ready to define the central object of this section:

DEFINITION 2.3. Let integers $g \geq 2$ and $k \geq 1$ be given. Define

$$\mathbb{P}\Omega^k M_g^{\text{trop}} := \{([\Gamma], D) \mid [\Gamma] \in M_g^{\text{trop}} \text{ and } D \in |kK_\Gamma|\}.$$

This space is called *tropical k -Hodge bundle*.

Recall [MUW21, Proposition 2.2], where the authors construct a moduli space $\text{Div}_{g,d}^{\text{trop}}$ parametrizing pairs $([\Gamma], D)$ of isomorphism classes of stable tropical curves of genus g and effective divisors of degree $D \in \text{Div}(\Gamma)$. The construction works completely analogous to the one of $M_{g,n}^{\text{trop}}$ given in Section 2.2 above with the only modification being that legs (corresponding to support points of divisors) are now unlabeled. Hence, the colimit involves more automorphisms.

PROOF OF THEOREM 1.1. We identify $\mathbb{P}\Omega^k M_g^{\text{trop}}$ as a subcomplex of (a subdivision of) $\text{Div}_{g,d}^{\text{trop}}$ for $d = k(2g - 2)$ as follows. Let $([\Gamma], D) \in \text{Div}_{g,d}^{\text{trop}}$ and let G be the minimal model for (Γ, D) such that D is supported on the vertices of G . By construction $([\Gamma], D)$ is contained in (a quotient of) the cone $\sigma_G = \mathbb{R}_{\geq 0}^{E(G)}$. We will now describe finitely many rational polyhedral cones in σ_G that contribute to $\mathbb{P}\Omega^k M_g^{\text{trop}}$.

By Definition 2.3 the pair $([\Gamma], D)$ is contained in $\mathbb{P}\Omega^k M_g^{\text{trop}}$ if and only if there exists a rational function f on Γ such that

$$(2) \quad D = kK_\Gamma + (f).$$

Fix an orientation for the edges of G . To specify a rational function f (up to a global additive constant in \mathbb{R}) that satisfies (2) we first need to choose an initial slope $m_e \in \mathbb{Z}$ at the beginning of every edge $e \in E(G)$ subject to the condition that at every vertex $v \in V(G)$

$$(3) \quad D(v) = k(2g(v) - 2 + \text{val}(v)) + \sum_{\text{outward edges at } v} m_e - \sum_{\text{inward edges of } v} m_e$$

holds. By [GK08, Lemma 1.8] there are only finitely many $\{m_e\}_{e \in E(G)}$ subject to (3). For each such choice a linear subspace of σ_G is cut out by the continuity of f . The cones determined this way constitute the entire generalized cone complex structure of $\mathbb{P}\Omega^k M_g^{\text{trop}}$.

For the statement on the dimension recall from [MUW21, Proposition 2.2] that

$$\dim \text{Div}_{g,k(2g-2)}^{\text{trop}} = 3g - 3 + k(2g - 2).$$

This provides an upper bound. This bound is attained by the cone described in Example 2.4. \square

EXAMPLE 2.4. Consider the graph G depicted in Figure 1. It consists of g vertices each of which has one self-loop as well as an incident separating edge joining it to a central chain of $g - 2$ vertices. All vertices have weight 0. This graph is trivalent and hence stable. If G is endowed with edge-lengths, we obtain a tropical curve Γ . The canonical divisor of Γ is the sum over all trivalent vertices, hence kK_Γ has k chips on each vertex. All of these chips can be moved independently onto the bridge edges joining the vertices with self-loops to the rest of the graph. Call the resulting divisor D (see Figure 1 for a picture of D with $k = 3$). The pair (Γ, D) has precisely $(3 + 2k)(g - 1)$ degrees of freedom: g for the length of the self-loops, $2g - 3$ for the lengths of the remaining edges, and $k(2g - 2)$ for the positions of the support points of D along the edges they lie on. Hence, the cone of tropical curves with underlying graph G and divisor D is of maximal dimension in $\mathbb{P}\Omega^k M_g^{\text{trop}}$.

2.5. Tropicalization. Let $g \geq 2$ be an integer and let X be a smooth, proper algebraic curve of genus g over a non-Archimedean field K . Possibly after passing to a finite non-Archimedean field extension $K \subseteq K'$, there is a stable model \mathcal{X} of X over the valuation ring R of K' . The central fiber \mathcal{X}_0 is a nodal curve. Denote the set of irreducible components of \mathcal{X}_0 by $\{C_v\}_{v \in V}$. Let G denote the dual graph of \mathcal{X}_0 , i.e. the set of vertices of G is precisely V and for every node in \mathcal{X}_0 there is one edge in G . Here, the edge corresponding to a node q joins two distinct vertices v and w if $q \in C_v \cap C_w$ and it is a self-loop at vertex v if q is a node of C_v . This graph is vertex weighted by $g(v)$ equal to the genus of




FIGURE 1. Graph G with divisor D defining a maximal cone in $\mathbb{P}\Omega^3 M_g^{\text{trop}}$.

the normalization C_v^ν of C_v . We endow the edge e corresponding to some node $q \in \mathcal{X}_0$ with an edge length in the following way. Write \mathcal{X} locally around q as $xy = f$ for $f \in R$ and let val_R denote the valuation of R . The length of e is defined to be $\text{val}_R(f)$. The resulting metric graph Γ is the tropicalization of X in the sense of curves. The tropicalization map

$$\begin{aligned} \text{trop} : \mathcal{M}_{g,n}^{\text{an}} &\longrightarrow M_{g,n}^{\text{trop}} \\ X &\longmapsto \Gamma \end{aligned}$$

is well-defined (see [Viv13, Lemma-Definition 2.2.7] for independence of the choice of K'), continuous, and surjective by [ACP15, Theorem 1.2.1]. Here $()^{\text{an}}$ denotes analytification in the sense of [Ber90] as before.

If X was endowed with a divisor D then we obtain a divisor on $\text{trop}(X)$ by specialization. [MUW21] presents this extended construction as a map between moduli spaces again. More precisely, for any degree $d \geq 0$ the authors of [MUW21] construct moduli spaces $\text{Div}_{g,d}$ and $\text{Div}_{g,d}^{\text{trop}}$ of pairs of smooth algebraic (resp. stable tropical) curves of genus g together with an effective divisor of degree d and give a tropicalization map $\text{trop}_{g,d} : \text{Div}_{g,d}^{\text{an}} \rightarrow \text{Div}_{g,d}^{\text{trop}}$ in the following way. The curve X can be extended to a semi-stable model \mathcal{X} such that D extends to a divisor \mathcal{D} on \mathcal{X} that does not meet any of the nodes of the special fiber. As before, this might require a base change to a non-Archimedean field extension. The specialization of D to Γ is defined to be the multidegree of $\mathcal{D}_0 := \mathcal{X}_0 \cap \mathcal{D}$, i.e.

$$\text{mdeg}(\mathcal{D}_0) = \sum_{v \in V} \deg(\mathcal{D}_0|_{C_v^\nu}) \cdot [v].$$

For the purposes of this article we simply define

$$\text{trop}_{\Omega^k} : \mathbb{P}\Omega^k \mathcal{M}_g^{\text{an}} \longrightarrow \mathbb{P}\Omega^k M_g^{\text{trop}}$$

to be the restriction of $\text{trop}_{g,k(2g-2)}$.

LEMMA 2.5. *The map trop_{Ω^k} is well-defined, continuous, proper, and closed.*

PROOF. By [Bak08, Lemma 4.20], the specialization of a canonical divisor on a curve X is the canonical divisor on $\text{trop}(X)$. Furthermore, the specialization map is linear and linearly equivalent divisors tropicalize to (tropically) linearly equivalent divisors (see e.g. [BU19, Theorem 4.2]). In particular, trop_{Ω^k} is well-defined. The map is also continuous because $\text{trop}_{g,d}$ is continuous by [MUW21, Theorem 3.2]. Finally, properness and closedness follow from the same properties for $\text{trop}_{g,d}$ (see [MUW21, Section 3.2]) and $\mathbb{P}\Omega^k M_g^{\text{trop}} \subseteq \text{Div}_{g,d}^{\text{trop}}$ being closed. \square

Note that $\mathbb{P}\Omega^k \mathcal{M}_g$ as well as $\mathbb{P}\Omega^k M_g^{\text{trop}}$ admit natural forgetful map to \mathcal{M}_g and M_g^{trop} , respectively. These are compatible with tropicalization maps in the following sense.

PROPOSITION 2.6. *The diagram*

$$\begin{array}{ccc} \mathbb{P}\Omega^k \mathcal{M}_g^{\text{an}} & \xrightarrow{\text{trop}_{\Omega^k}} & \mathbb{P}\Omega^k M_g^{\text{trop}} \\ \downarrow & & \downarrow \\ \mathcal{M}_g^{\text{an}} & \xrightarrow{\text{trop}} & M_g^{\text{trop}} \end{array}$$

commutes, where the vertical arrows are natural forgetful morphisms.

PROOF. This is essentially a modified version of the first part of [ACP15, Theorem 1.2.2] using unlabeled points. Alternatively, one can see this with the explicit descriptions of the two tropicalization maps that were given above. \square

2.6. The realizability problem. We conclude this section with the following observation: by Theorem 1.1, the dimension (of a maximal cone) of $\mathbb{P}\Omega^k M_g^{\text{trop}}$ is $(3+2k)(g-1)$. On the other hand, $\dim \mathbb{P}\Omega^k \mathcal{M}_g \leq (2+2k)(g-1) - 1$ by [BCGGM19a, Theorem 1.1]. By the following argument, this implies that trop_{Ω^k} cannot be surjective.

First note that $\mathbb{P}\Omega^k \mathcal{M}_g$ is a closed substack of $\mathcal{D}iv_{g,d}$. The compactification $\mathcal{D}iv_{g,d} \subseteq \overline{\mathcal{D}iv}_{g,d}$ was identified in [MUW21, Theorem 1.2] as a toroidal embedding of Deligne-Mumford stacks in the sense of [ACP15, Definition 6.1.1]. This means that locally around any geometric point of $\overline{\mathcal{D}iv}_{g,d}$ there exists a so-called *small toric chart* V [ACP15, Definition 6.2.4], i.e. a scheme V and an étale morphism $V \rightarrow \overline{\mathcal{D}iv}_{g,d}$ such that the pull-back V° of $\mathcal{D}iv_{g,d}$ to V is a toroidal embedding $V^\circ \subseteq V$ in the sense of [KKMS73]. In particular, the boundary is without self-intersection. Now consider the pull-back U of $\mathbb{P}\Omega^k \mathcal{M}_g$ to V . The tropicalization $\text{trop}_V(U)$ in the sense of [Uli17] is then a finite rational polyhedral cone complex of dimension $\leq \dim U \leq \dim \mathbb{P}\Omega^k \mathcal{M}_g$ by [Uli15, Theorem 1.1]. Taking the supremum over all small toric charts V we get

$$\dim \text{trop}_{\overline{\mathcal{D}iv}_{g,d}}(\mathbb{P}\Omega^k \mathcal{M}_g) = \sup \dim \text{trop}_V(U) \leq \dim \mathbb{P}\Omega^k \mathcal{M}_g.$$

It remains to argue that this coincides with the dimension of $\text{Im } \text{trop}_{\Omega^k}$. To this end note that by [Uli17, Theorem 1.2] the tropicalization $\text{trop}_{\overline{\mathcal{D}iv}_{g,d}}$ coincides with the retraction map $\rho : \overline{\mathcal{D}iv}_{g,d}^{\text{an}} \rightarrow \mathfrak{S}(\overline{\mathcal{D}iv}_{g,d})$ in the sense of [Thu07]. However, the same holds for trop_{Ω^k} by [MUW21, Theorem 3.2], so the claim follows. Putting everything together we conclude

$$\dim \text{trop}_{\Omega^k}(\mathbb{P}\Omega^k \mathcal{M}_g) \leq (2+2k)(g-1) - 1 < (3+2k)(g-1) = \dim \mathbb{P}\Omega^k M_g^{\text{trop}}.$$

Thus, trop_{Ω^k} cannot be surjective. This motivates the following definition.

DEFINITION 2.7. The image of trop_{Ω^k} is called *realizability locus*. A tropical curve Γ with effective pluri-canonical divisor $D \in |kK_\Gamma|$ is called *realizable* if the pair $([\Gamma], D) \in \mathbb{P}\Omega^k M_g^{\text{trop}}$ is contained in the realizability locus.

The realizability problem asks for a criterion to determine if $([\Gamma], D)$ is in the realizability locus. Corollary 4.13 together with Theorem 5.11 provide its answer.

3. Moduli space of multi-scale k -differentials

Fix integers $g \geq 2$ and $k \geq 1$. Throughout this section we work over the field \mathbb{C} of complex numbers. A tuple $\mu = (m_1, \dots, m_n) \in \mathbb{Z}^n$ such that $\sum m_i = k(2g-2)$ is called a *type*. The stratum of k -differentials of type μ is the subspace $\Omega^k \mathcal{M}_g(\mu)$ of the k -Hodge bundle $\Omega^k \mathcal{M}_g$ parametrizing k -differentials where the zero and pole orders are as prescribed by μ .

For an integer $d \mid k$, taking a global d -th power of a k/d -differential on a curve X gives rise to a k -differential on X . We will often be interested in those k -differentials that are not global powers of k/d -differentials.

DEFINITION 3.1. A k -differential is called *primitive* if it is not a global d -th power of a k/d -differential for some $d > 1$. We denote the union of connected components of $\Omega^k \mathcal{M}_g(\mu)$ parametrizing primitive k -differentials by $\Omega^k \mathcal{M}_g(\mu)^{\text{prim}}$.

For strata of abelian differentials $\Omega \mathcal{M}_g(\mu)$, the authors of [BCGGM19b] constructed a closure $\Xi \overline{\mathcal{M}}_{g,n}(\mu)$, the *moduli space of multi-scale differentials*.¹ Recall that \mathbb{C}^\times acts on $\Omega \mathcal{M}_g(\mu)$ by multiplication on the differential. This action extends to $\Xi \overline{\mathcal{M}}_{g,n}(\mu)$ and the projectivization $\mathbb{P}\Xi \overline{\mathcal{M}}_{g,n}(\mu)$ with respect to this action is a well-behaved compactification

¹ In [BCGGM19b], the marked points of the stratum are labeled. We consider the marked points to be unlabeled, i.e. we consider the quotient of the space in [BCGGM19b] by $\text{Sym}(\mu) = \{\phi \in S_n \mid m_{\phi(i)} = m_i \text{ for } i = 1, \dots, n\}$.

of $\mathbb{P}\Omega\mathcal{M}_g(\mu)$. The construction of $\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ was generalized to strata of primitive k -differentials $\Omega^k\mathcal{M}_g(\mu)^{\text{prim}}$ in [CMZ19].

Let (X, η) be a primitive k -differential in $\Omega^k\mathcal{M}_g(\mu)^{\text{prim}}$. In Section 3.2 we recall that (X, η) admits a canonical cover

$$\pi : (\widehat{X}, \omega) \longrightarrow (X, \eta),$$

where (\widehat{X}, ω) is an abelian differential of some type $\widehat{\mu}$. The canonical cover is unique up to multiplication of ω by a k -th root of unity. Hence, after projectivizing, there is a well-defined map $\mathbb{P}\Omega^k\mathcal{M}_g(\mu)^{\text{prim}} \rightarrow \mathbb{P}\Omega\mathcal{M}_{\widehat{g}}(\widehat{\mu})$ for $\widehat{g}, \widehat{\mu}$ as described in Section 3.2 below. In other words, we can think of $\mathbb{P}\Omega^k\mathcal{M}_g(\mu)^{\text{prim}}$ not only as a space parametrizing primitive k -differentials of type μ , but equivalently as a space parametrizing canonical covers.

The boundary of the compactification constructed in [CMZ19] parametrizes so-called *multi-scale k -differentials*. These are twisted k -differentials with the additional data of an (enhanced) level graph together with some compatibility conditions. Details are outlined in Sections 3.1 and 3.3. The compactification has the following properties.

THEOREM 3.2 ([CMZ19]). *There exists a complex orbifold $\mathbb{P}\Xi^k\overline{\mathcal{M}}_{g,n}(\mu)$, the moduli space of multi-scale k -differentials, with the following properties.²*

- (i) *The space $\mathbb{P}\Xi^k\overline{\mathcal{M}}_{g,n}(\mu)$ is a compactification of $\mathbb{P}\Omega^k\mathcal{M}_g(\mu)^{\text{prim}}$.*
- (ii) *Via the canonical cover construction, the space $\mathbb{P}\Xi^k\overline{\mathcal{M}}_{g,n}(\mu)$ is embedded as a suborbifold in the corresponding stratum $\mathbb{P}\Xi\overline{\mathcal{M}}_{\widehat{g},\widehat{n}}(\widehat{\mu})$ of abelian multi-scale differentials.*

We conclude this section with some more properties of $\mathbb{P}\Xi^k\overline{\mathcal{M}}_{g,n}(\mu)$ in Sections 3.5 and 3.6 which will turn out to introduce major difficulties in solving the realizability problem in Section 5.

3.1. Twisted k -differentials. The underlying curves of the boundary points of the moduli space of multi-scale k -differentials $\mathbb{P}\Xi^k\overline{\mathcal{M}}_{g,n}(\mu)$ will be nodal curves. The k -differentials will degenerate into so-called twisted k -differentials with some additional data and compatibility conditions that we will describe in the following.

For an abelian differential, the residue at a pole is defined as the coefficient in front of z^{-1} in the Laurent expansion around that pole. To define a useful notion of residues for k -differentials, recall from [BCGGM19a, Proposition 3.1] that a k -differential η of order $m = \text{ord}_0 \eta$ may locally be written as

$$(4) \quad \begin{cases} z^m (dz)^k & \text{if } m > -k \text{ or } k \nmid m, \\ \left(\frac{r}{z}\right)^k (dz)^k & \text{if } m = -k, \\ \left(z^{m/k} + \frac{t}{z}\right)^k (dz)^k & \text{if } m < -k \text{ and } k \mid m \end{cases}$$

for some $r \in \mathbb{C}^\times$ and $t \in \mathbb{C}$.

DEFINITION 3.3. For a k -differential η written as in (4), the k -residue of η is defined as

$$(5) \quad \text{Res}_0^k \eta := \begin{cases} 0 & \text{if } m > -k \text{ or } k \nmid m, \\ r^k & \text{if } m = -k, \\ t^k & \text{if } m < -k \text{ and } k \mid m \end{cases}$$

for r and t as above.

DEFINITION 3.4. Let X be a nodal curve and let $\mu = (m_1, \dots, m_n)$ be a type. A *twisted k -differential of type μ* on a stable n -pointed curve (X, \mathbf{s}) is a collection of (possibly meromorphic) k -differentials $\eta = \{\eta_v\}_v$ on the irreducible components X_v of X such that no η_v is identically zero with the following properties.

- (i) **(Vanishing as prescribed)** Each k -differential η_v is holomorphic and nonzero outside the nodes and marked points of X_v . Moreover, if a marked point s_i lies on X_v , then $\text{ord}_{s_i} \eta_v = m_i$.

²Here again, we consider the quotient of the space in [CMZ19] by $\text{Sym}(\mu)$.

- (ii) (**Matching orders**) For any node of X that identifies $q_1 \in X_{v_1}$ and $q_2 \in X_{v_2}$, the vanishing orders satisfy $\text{ord}_{q_1} \eta_{v_1} + \text{ord}_{q_2} \eta_{v_2} = -2k$.
- (iii) (**Matching k -residues condition, MRC**) If at a node of X that identifies $q_1 \in X_{v_1}$ with $q_2 \in X_{v_2}$ the condition $\text{ord}_{q_1} \eta_{v_1} = \text{ord}_{q_2} \eta_{v_2} = -k$ holds, then $\text{Res}_{q_1}^k \eta_{v_1} = (-1)^k \text{Res}_{q_2}^k \eta_{v_2}$.

3.2. Normalized covers. Consider first a smooth curve with k -differential (X, η) . It admits a canonical cover $\pi : (\widehat{X}, \omega) \rightarrow (X, \eta)$ that is unique up to multiplying ω with a k -th root of unity. In particular, the class of (\widehat{X}, ω) in $\mathbb{P}\Omega\mathcal{M}_{\widehat{g}}(\widehat{\mu})$ is unique. The cover \widehat{X} is connected if and only if η is a primitive k -differential. If η is a d -th power of a primitive k/d -differential, then \widehat{X} has k/d isomorphic connected components.

For a twisted k -differential a similar cover can be constructed, but no longer uniquely. Assume that the twisted k -differential η is of type $\mu = (m_1, \dots, m_n)$, let $\widehat{m}_i := (k + m_i) / \gcd(k, m_i) - 1$ and let

$$\widehat{\mu} := \underbrace{(\widehat{m}_1, \dots, \widehat{m}_1)}_{\gcd(k, m_1)} \underbrace{(\widehat{m}_2, \dots, \widehat{m}_2)}_{\gcd(k, m_2)} \dots \underbrace{(\widehat{m}_n, \dots, \widehat{m}_n)}_{\gcd(k, m_n)}.$$

Moreover, let $\widehat{n} := |\widehat{\mu}|$ and $\widehat{g} := \frac{1}{2} \sum_{\widehat{m}_i \in \widehat{\mu}} \widehat{m}_i + 1$.

THEOREM 3.5 ([BCGGM19a]). *For a pointed nodal curve (X, \mathbf{s}) with a twisted k -differential η of type μ , there exists a pointed nodal curve $(\widehat{X}, \widehat{\mathbf{s}})$ with a twisted abelian differential ω of type $\widehat{\mu}$ such that*

- (i) $\pi : \widehat{X} \rightarrow X$ is a cyclic cover of degree k with deck transformation τ ,
- (ii) $\pi^* \eta = \omega^k$,
- (iii) $\tau^* \omega = \zeta \omega$ for a primitive k -th root of unity ζ ,
- (iv) marked points are mapped to marked points, i.e. $\pi(\widehat{\mathbf{s}}) = \mathbf{s}$,
- (v) π is unramified outside of the nodes and marked points of \widehat{X} ,
- (vi) every node or marked point $q \in X_v$ has precisely $\gcd(k, \text{ord}_q \eta_v)$ preimages.

DEFINITION 3.6. We refer to a tuple $(\pi : \widehat{X} \rightarrow X, \mathbf{s}, \omega)$ as above as a *normalized cover* of (X, \mathbf{s}, η) . A normalized cover is called *primitive* if \widehat{X} is connected.

- REMARK 3.7.**
- (1) Condition (vi) in the above theorem is well-defined at nodes because the twisted k -differential η is subject to the matching orders condition of Definition 3.4.
 - (2) If η_v is a d -th power of a primitive k/d -differential, then the irreducible component X_v has precisely d isomorphic preimages.
 - (3) In general, the normalized cover is not unique: While the fibers $\pi|_{\pi^{-1}(v)} : \coprod_{\widehat{v}} X_{\widehat{v}} \rightarrow X_v$ are uniquely determined by the k -differential η_v , there may be a choice how to glue the different fibers along the nodes.

To determine the relation between the residues of the cover and the k -residues of the base curve, let us again consider a normalized cover of smooth curves with differentials $\pi : (\widehat{X}, \omega) \rightarrow (X, \eta)$. Let us fix a pole $p \in X$ of η and let $q \in \pi^{-1}(p)$ be some preimage. If π is ramified at p , we claim that both the k -residue $\text{Res}_p^k(\eta)$ and the residue $\text{Res}_q(\omega)$ vanish. For the k -residue, this is immediate by Definition 3.3, and for the residue this is a consequence of the compatibility with the τ -action as follows. If π is ramified at q , then there is an integer $1 < d \mid k$ such that q is fixed by $\tau^d, \tau^{2d}, \dots, \tau^k$. Because of $\tau^* \omega = \zeta \omega$, this implies that

$$\frac{k}{d} \cdot \text{Res}_q(\omega) = \sum_{i=1}^{k/d} \zeta^{id} \text{Res}_q(\eta) = 0,$$

as the $\frac{k}{d}$ -th roots of unity sum to zero. If the cover is unramified at p , then the k -residues of the twisted k -differential and the residues of the normalized cover are related as follows.

LEMMA 3.8. *If π is unramified at p , then*

$$\text{Res}_p^k(\eta) = (\text{Res}_q(\omega))^k.$$

PROOF. Let $m := \text{ord}_p \eta$ and note that k divides m . Recall from (4) that the k -form η may locally be written as

$$\begin{cases} \left(\frac{r}{z}\right)^k (dz)^k & \text{if } m = -k, \\ \left(z^{m/k} + \frac{t}{z}\right)^k (dz)^k & \text{if } m < -k. \end{cases}$$

In these cases, the k -residues are by definition r^k and t^k , respectively. As π is locally given by $\pi : z \mapsto z$, we get

$$\omega^k = \pi^* \eta = \begin{cases} \left(\frac{r}{z} dz\right)^k & \text{if } m = -k \\ \left(\left(z^{m/k} + \frac{t}{z}\right) dz\right)^k & \text{if } m < -k. \end{cases}$$

Thus the residues of ω at q are r and t , respectively. \square

REMARK 3.9. In general, the k -residue does not coincide with the coefficient in front of z^{-k} in the Laurent expansion around the given pole. Moreover, there is nothing similar to the residue theorem for k -residues.

3.3. Enhanced level graphs. The boundary points of the moduli space of multi-scale differentials $\mathbb{P}\Xi^k \overline{\mathcal{M}}_{g,n}(\mu)$ are normalized covers subject to some conditions on the underlying dual graph of the stable curves. We will recall here the necessary terminology to give the characterization of the boundary points in Section 3.4 below.

Let G be a stable graph. A *full order* on G is an order \succcurlyeq on the vertices V of G that is reflexive, transitive, and such that for any $v_1, v_2 \in V$ at least one of the statements $v_1 \succcurlyeq v_2$ or $v_2 \succcurlyeq v_1$ holds. If $v_1 \succcurlyeq v_2$ and $v_2 \succcurlyeq v_1$, we write $v_2 \asymp v_1$. We call a function $\ell : V \rightarrow \mathbb{Z}_{\leq 0}$ such that $\ell^{-1}(0) \neq \emptyset$ a *level function*. Note that a level function induces a full order on G by setting $v \succcurlyeq w$ whenever $\ell(v) \geq \ell(w)$. A *level graph* (G, ℓ) is a graph G together with a choice of a level function ℓ . When the level function is clear from context, we abuse notation and denote the level graph (G, ℓ) by G as well.

For a given level L we call the subgraph of G that consists of all vertices v with $\ell(v) > L$ along with the edges between them the *graph above level L* of G , and denote it by $G_{>L}$. We similarly define the graph $G_{\geq L}$ *above or at level L* , and the graph $G_{=L}$ *at level L* . An edge $e \in E$ is called *horizontal* if it connects two vertices of the same level, and it is called *vertical* otherwise. Given a vertical edge e , we denote by $v^+(e)$ and $v^-(e)$ the vertex that is its endpoint of higher and lower level, respectively.

Let $\pi : \widehat{G} \rightarrow G$ be a morphism of graphs. By this we mean that π maps vertices to vertices, edges to edges, and legs to legs while respecting edge–vertex and leg–vertex incidences. Assume further that π is surjective on vertices and let \succcurlyeq_G denote a full order on G . We get an induced full order on \widehat{G} by setting $v_1 \succcurlyeq_{\widehat{G}} v_2$ if and only if $\pi(v_1) \succcurlyeq_G \pi(v_2)$. If \succcurlyeq_G was induced by a level function ℓ , then $\succcurlyeq_{\widehat{G}}$ is induced by the *lifted level function* $\widehat{\ell} := \ell \circ \pi$.

In the following, given a twisted k -differential (X, \mathbf{s}, η) and a level graph (G, ℓ) , we will always assume that G is the dual graph of X . We denote by $X_{>L}$ (resp. $X_{\geq L}$ resp. $X_{=L}$) the subcurve whose dual graph is $G_{>L}$ (resp. $G_{\geq L}$ resp. $G_{=L}$).

DEFINITION 3.10. Let $\pi : \widehat{G} \rightarrow G$ be a morphism of graphs. It is called *cover of graphs* if π is surjective on vertices, edges, and legs. Furthermore, it is called *k -cyclic cover of graphs* if there is the additional data of an automorphism τ of \widehat{G} such that $\tau^k = \text{id}$ and π is the quotient map $\widehat{G} \rightarrow \widehat{G}/\tau$.

REMARK 3.11. We would like to stress that morphisms (and covers) of graphs do not contract edges. Also note that in a k -cyclic cover of graphs the order of τ may in fact be $k' < k$ with k' dividing k . We think of a k -cyclic cover of graphs as the dual graphs of a k -cyclic cover of curves. Hence the name k -cyclic.

The following definition is taken from [BCGGM19a].

DEFINITION 3.12. Let $\pi : (\widehat{G}, \widehat{\ell}) \rightarrow (G, \ell)$ be a k -cyclic cover of level graphs. We say that a normalized cover of twisted k -differential $(\pi : \widehat{X} \rightarrow X, \mathbf{s}, \omega)$ is *compatible* with π if it satisfies the following two conditions.

- (iv) **(Partial order)** If a node of \widehat{X} identifies $q_1 \in \widehat{X}_{v_1}$ with $q_2 \in \widehat{X}_{v_2}$, then $v_1 \succcurlyeq v_2$ if and only if $\text{ord}_{q_1} \omega_{v_1} \geq -1$. In particular, $v_1 \asymp v_2$ if and only if $\text{ord}_{q_1} \omega_{v_1} = -1$.
- (v) **(Global residue condition, GRC)** For every level L and every connected component \widehat{Y} of $\widehat{X}_{>L}$ that does not contain a marked point with a prescribed pole the following condition holds: Let q_1, \dots, q_b denote the set of all nodes where \widehat{Y} intersects $\widehat{X}_{=L}$. Then

$$\sum_{j=1}^b \text{Res}_{q_j^-} \omega_{v^-(q_j)} = 0.$$

where $q_j^- \in \widehat{X}_{=L}$ is the point on the irreducible component corresponding to $v^-(q_j) \in \widehat{G}_{=L}$ that is part of the node q_j .

Note that condition (iv) is equivalent to the analogous condition on the induced twisted k -differential η on X : If a node of X identifies $q_1 \in X_{v_1}$ with $q_2 \in X_{v_2}$, then $v_1 \succcurlyeq v_2$ if and only if $\text{ord}_{q_1} \eta_{v_1} \geq -k$, and $v_1 \asymp v_2$ if and only if $\text{ord}_{q_1} \eta_{v_1} = -k$.

Though not strictly necessary at the moment, it will be more convenient later on to consider enhanced level graphs instead of level graphs. Enhanced level graphs additionally carry the data of an integer valued function o which should be thought of as an order at every node and marked point.

DEFINITION 3.13. Let $k \in \mathbb{N}_{\geq 1}$. A k -enhanced level graph $G^+ = (V, H, L, \iota, a, \ell, o)$ is a tuple where (V, H, L, ι, a) is a stable graph, the map $\ell : V \rightarrow \mathbb{Z}_{\leq 0}$ is the *level function* and the so-called *enhancement* $o : H \cup L \rightarrow \mathbb{Z}$ such that the following hold.

- (i) The genus is well-defined, i.e. for all $v \in V$ there is a non-negative integer $g(v)$ such that

$$k(2g(v) - 2) = \sum_{h \in a^{-1}(v)} o(h).$$

We call $\mu(v) := (o(h))_{h \in a^{-1}(v)}$ the *type* of v .

- (ii) The orders at edges match, i.e. for all $h \in H$ we have $o(h) + o(\iota(h)) = -2k$.
- (iii) The orders at the half-edges are compatible with the level function, that is: for all $h \in H$ we have $o(h) \geq o(\iota(h))$ if and only if $\ell(a(h)) \geq \ell(a(\iota(h)))$.

Note that (ii) and (iii) imply that the levels at both ends of an edge are equal if and only if the orders at both ends are $-k$. We call such an edge *horizontal*. Any other edge is called *vertical*.

DEFINITION 3.14. Let $G^+ = (V, H, L, \iota, a, \ell, o)$ be a k -enhanced level graph. A *normalized cover* of G^+ is a triple $(\widehat{G}^+, \pi, \tau)$, where

- (i) $\widehat{G}^+ = (\widehat{V}, \widehat{H}, \widehat{L}, \widehat{\iota}, \widehat{a}, \widehat{\ell}, \widehat{o})$ is a 1-enhanced level graph,
- (ii) $\pi : \widehat{G}^+ \rightarrow G^+$ is a cover of graphs such that
 - (a) π preserves the levels, i.e. $\ell = \widehat{\ell} \circ \pi$,
 - (b) the order at the preimages is the expected one, i.e. for all half-edges and legs $h \in H \cup L$ and all $\widehat{h} \in \pi^{-1}(h)$ it is

$$\widehat{o}(\widehat{h}) + 1 = \frac{o(h) + k}{\text{gcd}(o(h), k)}$$

- (c) the number of preimages is the expected one, i.e. for all half-edges and legs $h \in H \cup L$ we have

$$|\pi^{-1}(h)| = \text{gcd}(o(h), k),$$

- (iii) $\tau : \widehat{G} \rightarrow \widehat{G}$ is a graph automorphism that exhibits π as a k -cyclic cover of graphs.

Note that the genus of each vertex $\widehat{v} \in \widehat{V}$ is an integer by definition of an 1-enhanced level graph.

DEFINITION 3.15. Let $\pi : \widehat{G}^+ \rightarrow G^+$ be a k -cyclic cover of enhanced level graphs. We say that a normalized cover of a twisted k -differential $(\pi : \widehat{X} \rightarrow X, \mathbf{s}, \omega)$ is *compatible* with π if it is compatible with the underlying cover of level graphs $\pi : (\widehat{G}, \widehat{\ell}) \rightarrow (G, \ell)$ and the orders of the differentials coincide with the enhancements.

3.4. Multi-scale k -differentials and the characterization of limit points. The points in the boundary of $\mathbb{P}\Xi^k \overline{\mathcal{M}}_{g,n}(\mu)$ may be described as follows.

DEFINITION 3.16. A *multi-scale k -differential of type μ* on a pointed stable curve (X, \mathbf{s}) consists of the following data.

- (i) A primitive normalized cover of a twisted k -differential $(\pi : \widehat{X} \rightarrow X, \mathbf{s}, \omega)$ of type μ .
- (ii) A compatible k -cyclic cover of enhanced level graphs $\pi : \widehat{G}^+ \rightarrow G^+$.
- (iii) A prong-matching for each node of X joining components on non-equal levels.

A prong-matching roughly represents a choice of gluing the differentials at the nodes of the curve. While it is needed to get a well-behaved compactification, it will be of no importance to us and we will suppress it in the following.

THEOREM 3.17 ([CMZ19]). *The points in the moduli space of multi-scale k -differentials $\mathbb{P}\Xi^k \overline{\mathcal{M}}_{g,n}(\mu)$ are precisely the \mathbb{C}^\times -equivalence classes of multi-scale k -differentials $(\pi : \widehat{X} \rightarrow X, \mathbf{s}, \omega, \pi : \widehat{G}^+ \rightarrow G^+)$ of type μ .*

Note that the tuple $(\pi : \widehat{X} \rightarrow X, \mathbf{s}, \omega, \pi : \widehat{G}^+ \rightarrow G^+)$ is equivalent to the tuple $(\tau \curvearrowright \widehat{X}, \widehat{\mathbf{s}}, \omega, \tau \curvearrowright G^+)$, where $\widehat{\mathbf{s}}$ is the lift of \mathbf{s} to \widehat{X} . We give another version of the same theorem that highlights the possible scaling parameters of one-parameter families approaching the boundary.

Suppose that S is the spectrum of a discrete valuation ring R with residue field \mathbb{C} , whose maximal ideal is generated by t . Let $\widehat{\mathcal{X}}/S$ be a family of semi-stable curves with smooth generic fiber \widehat{X} and special fiber \widehat{X}_0 and such that there is an automorphism τ of degree k on the family $\widehat{\mathcal{X}}/S$. Let ω be a section of the ζ_k -eigenspace (with respect to τ) of $\omega_{\widehat{\mathcal{X}}/S}$ of type $\widehat{\mu} = (\widehat{m}_1, \dots, \widehat{m}_{\widehat{n}})$ whose divisor is given by the sections $\widehat{\mathbf{s}} = (\widehat{s}_1, \dots, \widehat{s}_n)$ with multiplicity \widehat{m}_i . If moreover $(\widehat{\mathcal{X}}/S, \widehat{\mathbf{s}})$ is stable, then the tuple $(\widehat{\mathcal{X}}/S, \tau, \widehat{\mathbf{s}}, \omega)$ is called a *pointed family of stable k -differentials*. (Note that $(X := \widehat{X}/\tau, \eta := (\omega|_{\widehat{X}})/\tau)$ is in fact a k -differential, where $\omega|_{\widehat{X}}$ is the restriction of ω to the generic fiber \widehat{X} .) We define the *scaling factor* $\widehat{\ell}(\widehat{v})$ of a vertex \widehat{v} of the dual graph \widehat{G} of \widehat{X}_0 as the non-positive integer such that the restriction of the meromorphic differential $t^{-\widehat{\ell}(\widehat{v})} \cdot \omega$ to the component $\widehat{X}_{0, \widehat{v}}$ of the special fiber corresponding to \widehat{v} is a well-defined and generically nonzero differential $\omega_{\widehat{v}}$ on $\widehat{X}_{0, \widehat{v}}$. The $\omega_{\widehat{v}}$ are called the *scaling limits* of ω .

THEOREM 3.18 ([CMZ19]). *If $(\widehat{\mathcal{X}}/S, \tau, \widehat{\mathbf{s}}, \omega)$ is as above, then the function $\widehat{\ell}(\widehat{v})$ defines a full order on the dual graph \widehat{G} of the special fiber \widehat{X}_0 and the collection $\omega_{\widehat{v}}$ is a twisted k -differential of type $\widehat{\mu}$ compatible with the level function $\widehat{\ell}$.*

Conversely, suppose that \widehat{X}_0 is a stable \widehat{n} -pointed curve with dual graph \widehat{G} and a degree k automorphism τ . Moreover, suppose that $\omega = \{\omega_{\widehat{v}}\}_{\widehat{v} \in \widehat{V}}$ is a twisted k -differential of type μ in the ζ_k -eigenspace and compatible with a full order on \widehat{G} . Then for every level function $\widehat{\ell} : \widehat{G} \rightarrow \mathbb{Z}$ defining the full order on \widehat{G} and for every τ -invariant assignment of integers $n_{\widehat{e}}$ to horizontal edges there is a stable family $\widehat{\mathcal{X}}/S$ over $S = \text{Spec } \mathbb{C}[[t]]$ with smooth generic fiber and special fiber \widehat{X}_0 that satisfies the following properties.

- (i) *The action of τ extends to a degree k automorphism on $\widehat{\mathcal{X}}/S$.*
- (ii) *There exists a global section ω of the relative dualizing sheaf $\omega_{\widehat{\mathcal{X}}/S}$ whose horizontal divisor $\text{div}_{\text{hor}}(\omega) = \sum_{i=1}^{\widehat{n}} \widehat{m}_i \Sigma_i$ is of type $\widehat{\mu}$ and whose scaling limits are the collection $\{\omega_{\widehat{v}}\}_{\widehat{v} \in \widehat{V}}$. Moreover, the restriction of ω to each fiber is contained in the ζ_k -eigenspace of τ .*
- (iii) *The intersections $\Sigma_i \cap \widehat{X}_0 = \{\widehat{s}_i\}$ are smooth points of the special fiber and ω has a zero of order \widehat{m}_i in \widehat{s}_i .*

(iv) *There exists a positive integer N such that near every node \widehat{e} a local equation for \mathcal{X} is given by*

$$xy = \begin{cases} t^{Nn_{\widehat{e}}} & \text{if } \widehat{e} \text{ is a horizontal edge,} \\ t^{N(\ell(v^+(\widehat{e})) - \ell(v^-(\widehat{e})))} & \text{if } \widehat{e} \text{ is a vertical edge.} \end{cases}$$

PROOF. The first statement is proved in [CMZ19]. Note that the arguments given there hold over any discrete valuation ring with residue field \mathbb{C} .

For the second statement we recall from the proof of [MUW21, Theorem 5.2] that there are no constraints for the plumbing fixtures to be used for plumbing horizontal nodes, whereas for a vertical node corresponding to an edge \widehat{e} the level function $\widehat{\ell}_0$ on the cover used for plumbing has to satisfy the condition

$$(6) \quad (\text{ord}_{v^+(\widehat{e})} \widehat{\omega} + 1) \mid (\widehat{\ell}_0(v^+(\widehat{e})) - \widehat{\ell}_0(v^-(\widehat{e}))).$$

Multiplying the prescribed function $\widehat{\ell}$ by a sufficiently divisible N , the resulting level function $\widehat{\ell}_0 = N \cdot \widehat{\ell}$ satisfies the divisibility property. \square

3.5. Empty primitive strata. The primitive strata of k -differentials $\Omega^k \mathcal{M}_g(\mu)^{\text{prim}}$ are empty for some types μ . To keep the notation concise, we will denote a type $\mu = (m_1, \dots, m_n)$ where multiple m_i agree with exponential notation, e.g. we will denote the type $(0, \dots, 0)$ by (0^n) .

THEOREM 3.19 ([GT21a], [GT21b], [GT22a]). *The primitive stratum $\Omega^k \mathcal{M}_g(\mu)^{\text{prim}}$ (and hence the stratum of multi-scale k -differentials $\mathbb{P}\Xi^k \overline{\mathcal{M}}_{g,n}(\mu)$) is empty if and only if*

- (i) $k = 1$ and $\mu = (-1, m_2, \dots, m_n)$ with $m_2, \dots, m_n \geq 0$,
- (ii) $g = 0$ and $\mu = (m_1, \dots, m_n)$ with $\gcd(m_1, \dots, m_n, k) \neq 1$,
- (iii) $g = 1$ and $\mu = (0^{n-2}, -1, 1)$,
- (iv) $g = 1$, $k \geq 2$ and $\mu = (0^n)$,
- (v) $g = 2$, $k = 2$ and $\mu = (0^{n-1}, 4)$ or $\mu = (0^{n-2}, 1, 3)$.

REMARK 3.20. A stratum $\Omega^k \mathcal{M}_g(\mu)$ may be nonempty even if its primitive part $\Omega^k \mathcal{M}_g(\mu)^{\text{prim}}$ is empty: For each $d \mid k$, the stratum $\Omega^k \mathcal{M}_g(\mu)$ may have nonempty connected components that parametrize d -th powers of primitive k/d -differentials.

3.6. The image of the residue map. In Section 5 when we prove Theorem 1.4, we will first translate a given tropical normalized cover into a normalized cover of enhanced level graphs $\pi : \widehat{G}^+ \rightarrow G^+$ and then try to construct a normalized cover of a twisted k -differential (X, η) that is compatible with π . If this is possible, then the irreducible component (X_v, η_v) is a (possibly meromorphic) k -differential for each $v \in V(G^+)$, whose type $\mu(v)$ is prescribed by the enhancements of G^+ . To construct the twisted k -differential, we will need to fix the k -residues at the poles of η_v . The question which k -residues are valid choices was answered by Gendron-Tahar. We will summarize their results in this section.

DEFINITION 3.21. For a type $\mu = (m_1, \dots, m_n)$, we defined the *reduced type* μ_{red} as the sub-tuple of μ consisting of all nonzero entries. Following [GT22a], we denote this tuple by

$$\mu_{\text{red}} = (a_1, \dots, a_t; -b_1, \dots, -b_p; -c_1, \dots, -c_r; -k^s)$$

where the $a_i > -k$ are the zeros, the $b_i \in k\mathbb{N}_{>1}$ are the poles where the order is greater than k and divisible by k , and the $c_i \in \mathbb{N}_{>k} \setminus (k\mathbb{N})$ are the poles with order not divisible by k . As above, the power $-k^s$ indicates that there are s poles of order k .

Recall that the k -residues at the poles with orders $-c_i$ are zero, while the k -residues at the poles with orders $-k$ cannot be zero. We let

$$\text{Res}_g^k(\mu_{\text{red}}) : \Omega^k \mathcal{M}_g(\mu_{\text{red}})^{\text{prim}} \longrightarrow \mathbb{C}^p \times (\mathbb{C}^\times)^s$$

denote the residue map. For almost all reduced types μ_{red} , the residue map is surjective. In the rest of this section, we will discuss all the cases where the residue map is not surjective. The following proposition lists all those cases where exactly the origin is missing in the image of the residue map.

PROPOSITION 3.22 ([GT21a], [GT21b], [GT22a]). *In the following cases, precisely the origin is missing from the image of the residue map, i.e. the image of the residue map is*

$$\mathrm{Im}(\mathrm{Res}_g^k(\mu_{\mathrm{red}})) = (\mathbb{C}^\times)^p.$$

(i) *If $k = 1$, $g = 0$, $s = 0$ and there exists an index i such that the inequality*

$$a_i > \sum_{j=1}^p b_j - (p+1)$$

holds. (Note that $k = 1$ implies $r = 0$.)

- (ii) *If $k = 2$, $g = 1$ and $\mu_{\mathrm{red}} = (4p; (-4^p))$ or $\mu_{\mathrm{red}} = (2p-1, 2p+1; (-4^p))$ for $p \in \mathbb{N}^\times$.*
 (iii) *If $k \geq 2$, $g = 0$, $\mu_{\mathrm{red}} = (a_1, \dots, a_t; -b_1, \dots, -b_p; -c_1)$ and there is at most one a_i not divisible by k and $\sum_{k|a_i} a_i < kp$.*
 (iv) *If $k \geq 2$, $g = 0$, $\mu_{\mathrm{red}} = (a_1, \dots, a_t; -b_1, \dots, -b_p)$, $p \neq 0$ and none of the following holds:*
 (a) $p = 1$ and $t \geq 3$,
 (b) $p \geq 2$, $t \geq 3$ and there exist at least three a_i not divisible by k ,
 (c) $p \geq 2$, $t \geq 3$ and there exist precisely two a_i not divisible by k and $\sum_{k|a_i} a_i \geq kp$,
 (d) $k = 2$ and $\mu_{\mathrm{red}} = ((2p+b-5)^2; -b, -b-2, (-4^{p-2}))$ or $\mu_{\mathrm{red}} = (2p+b-7, 2p+b-5; (-b^2), (-4^{p-2}))$ for $p \geq 2$ and even $b \geq 4$.

In the cases of the following two propositions not only the origin, but a finite number of \mathbb{C} -lines is missing from the image of the residue map.

PROPOSITION 3.23 ([GT21b], [GT22a]). *For the reduced types μ_{red} in Figure 2, precisely the \mathbb{C} -lines spanned by the vectors w_i are missing from the image of the residue map, i.e.*

$$\mathrm{Im}(\mathrm{Res}_g^k(\mu_{\mathrm{red}})) = (\mathbb{C}^p \times (\mathbb{C}^\times)^s) \setminus \bigcup_i \langle w_i \rangle_{\mathbb{C}}.$$

(Note that if there are multiple poles with the same order in μ_{red} , then the order of the entries of the vectors w_i may not be uniquely determined. In those cases all possible permutations need to be taken into account.)

PROPOSITION 3.24 ([GT21a]). *For $k = 1$, $g = 0$ and $\mu_{\mathrm{red}} = (a_1, \dots, a_t; (-1^s))$ with $s \geq 2$, precisely those \mathbb{C} -lines $\langle (x_1, \dots, x_{s_1}, -y_1, \dots, -y_{s_2}) \rangle_{\mathbb{C}}$ are missing from the image of the residue map for which the $x_i, y_j \in \mathbb{N}$ are pairwise relatively prime and*

$$\sum_{i=1}^{s_1} x_i = \sum_{j=1}^{s_2} y_j \leq \max(a_1, \dots, a_t).$$

Finally, there are some cases where a finite number of at most 2-dimensional subvarieties is missing from the image of the residue map. For $k = 2$, following [GT21b, Définition 1.8] we call three numbers $R_1, R_2, R_3 \in \mathbb{C}^\times$ *triangular*, if there exist square roots r_1, r_2, r_3 of R_1, R_2, R_3 such that $r_1 + r_2 + r_3 = 0$.

PROPOSITION 3.25 ([GT21b]). *For $k = 2$, $g = 0$ and $\mu_{\mathrm{red}} = (a_1, \dots, a_t; (-2^s))$, precisely the following \mathbb{C} -lines are not in the image of the residue map.*

- (i) *For $\mu_{\mathrm{red}} = (2s'-1, 2s'+1; (-2^{2s'+2}))$ with $s' \in \mathbb{N}$ the lines spanned by $(1, \dots, 1, R, R)$ for $R \in \mathbb{C}^\times$ are missing.*
 (ii) *For $\mu_{\mathrm{red}} = ((2s'-1)^2; (-2^{2s'+1}))$ with $s' \in \mathbb{N}^\times$ the lines spanned by $(R_1, R_2, R_3, \dots, R_3)$ for triangular $R_i \in \mathbb{C}^\times$ are missing.*
 (iii) *If precisely two a_i are odd (say a_1 and a_2), the lines spanned by (r_1^2, \dots, r_s^2) for relatively prime $r_i \in \mathbb{N}$ and such that for $S := \sum_i r_i$ either*
 (a) S is odd and $S < \max(a_1, a_2) + 2$ or
 (b) S is even and $S < a_1 + a_2 + 4$
 are missing.

THEOREM 3.26 ([GT21a], [GT21b], [GT22a]). *For $k \geq 1$, the residue map is surjective in all cases not covered by Propositions 3.22, 3.23, 3.24 and 3.25.*

	μ_{red}	w_i
$g=1$	$(2s; (-2^s))$ for $s \in 2\mathbb{N}^\times$	$(1, \dots, 1)$
	$(s-1, s+1; (-2^s))$ for $s \in 2\mathbb{N}^\times$	$(1, \dots, 1)$
$k=2$	$(s-1, s+1; -4; (-2^s))$ for $s \in 2\mathbb{N}^\times$	$(0; 1, \dots, 1)$
	$(2p-1, 2p+1; (-4^p); (-2^2))$ for $p \geq 0$	$(0, \dots, 0; 1, 1)$
	$(s^2; -4; (-2^s))$ for $s \in (\mathbb{N}^\times \setminus 2\mathbb{N})$	$(1; 1, \dots, 1)$
	$((2p-1)^2; (-4^p); -2)$ for $p \geq 1$	$(1, 0, \dots, 0; 1)$
	$((2p+b-5)^2; -b, -b-2, (-4^{p-2}))$ for $p \geq 2$ and even $b \geq 4$	$(1, 1, 0, \dots, 0)$
	$(2p+b-7, 2p+b-5; (-b^2), (-4^{p-2}))$ for $p \geq 2$ and even $b \geq 4$	$(1, 1, 0, \dots, 0)$
$k=3$	$(-1, 4; (-3^3))$	(1^3)
	$(1, 2; (-3^3))$	(1^3)
	$(2, 4; (-3^4))$	$(1^2, -1^2)$
	$(2, 7; (-3^5))$	$(1^4, -1)$
	$(2, 10; (-3^6))$	(1^6)
	$(5, 7; (-3^6))$	(1^6)
$k=4$	$(-1, 5; (-4^3))$	$(1^2, -4)$
	$(3, 5; (-4^4))$	(1^4)
	$(-1, 9; (-4^4))$	(1^4)
	$(3, 13; (-4^6))$	(1^6)
$k=6$	$(-1, 7; (-6^3))$	(1^3)
	$(-1, 13; (-6^4))$	(1^4)
$k \geq 2$	$(-1, 1; (-k^2))$	$(1, (-1)^k)$

FIGURE 2. Reduced types μ_{red} and generators w_i of the \mathbb{C} -lines missing in the image of the residue map.

4. Reduction to realizability of normalized covers

The goal of this section is to reduce the realizability problem for curves with pluri-canonical divisor to the realizability problem for normalized covers, i.e. we want to formally state and prove Corollary 1.3. To do so, we define a notion of tropical normalized cover in Definition 4.9 which should be thought of as a tropical version of Definition 3.6. Essentially, we require a tropical normalized cover to be a cyclic degree k tropical Hurwitz cover $\widehat{\Gamma} \rightarrow \Gamma$ in the sense of [CMR16, Definition 16] with the additional property that the underlying cover of graphs admits the structure of a normalized cover of enhanced level graphs in the sense of Definition 3.14. Of course, the enhancement should be compatible with the divisor marked by the legs of Γ . More precisely, in Lemma 4.1 we give a canonical construction to endow a tropical curve with an effective pluri-canonical divisor with the structure of an enhanced level graph. In the definition of a tropical normalized cover we will then require the structure of normalized cover of enhanced level graphs to coincide on Γ with the output of Lemma 4.1. Once the notion of tropical normalized cover is introduced, we can construct the moduli space of tropical normalized covers $\mathbb{P}\Xi^k M_g^{\text{trop}}$ and the tropicalization map

$$\text{trop}_{\Xi^k} : \mathbb{P}\Omega^k \mathcal{M}_{g,n}^{\text{an}} \longrightarrow \mathbb{P}\Xi^k M_g^{\text{trop}} .$$

Finally, we prove Theorem 1.2 and perform the reduction step.

4.1. From rational functions to k -enhanced level graphs. So far we have considered on one hand tropical curves Γ together with an effective pluri-canonical divisor $D \in |kK_\Gamma|$ and on the other hand k -enhanced level graphs G^+ . These notions are related by

LEMMA 4.1. *Let Γ be a tropical curve and $D = kK_\Gamma + (f) \in |kK_\Gamma|$ an effective pluricanonical divisor. Let G be the minimal graph model of Γ where D is represented with legs. We can associate a natural k -enhanced level graph structure $G^+ = G^+(f)$ on G .*

PROOF. Let G be the minimal model of Γ including $D(p)$ legs for every point $p \in \text{supp } D$. First we endow G with the total order given by f , i.e. for vertices v and w we set $v \preceq w$ if and only if $f(v) \leq f(w)$. Next we define a k -enhancement $o : H \cup L \rightarrow \mathbb{Z}$ compatible with the total order by the following rule

$$(7) \quad o(h) := \begin{cases} 1 & \text{if } h \text{ is a leg} \\ -s(h) - k & \text{if } h \text{ is part of an edge,} \end{cases}$$

where $s(h)$ is the outgoing slope of f on the half-edge h . It is easy to check Definition 3.13. \square

REMARK 4.2. Our choice to represent the effective divisor D by attaching $D(p)$ legs at each point $p \in \Gamma$ and endowing these with o -value 1 amounts to seeking realizations in the principal stratum $\mathbb{P}\Omega^k \mathcal{M}_g(1, \dots, 1)$ in the end. Of course one could also ask for realizability in other strata – the definition in (7) would then have to be adapted.

4.2. Tropical k -cyclic Hurwitz covers. A tropical Hurwitz cover is a harmonic morphism of tropical curves satisfying the local Riemann-Hurwitz conditions. The moduli space of such maps of fixed degree d and ramification profile ξ was introduced in [CMR16]. For our purposes we need the slightly modified notion of tropical k -cyclic Hurwitz covers, which we will introduce now.

DEFINITION 4.3. Let Γ' and Γ be tropical curves. A morphism of metric graphs $\varphi : \Gamma' \rightarrow \Gamma$ is called *morphism of tropical curves* if it maps edges of Γ' linearly to edges of Γ such that the ratio of edge lengths $d_{e'} := \frac{l(\varphi(e'))}{l(e')}$ is an integer for every edge e' of Γ' . In this case the numbers $d_{e'}$ are called *expansion factors*.

DEFINITION 4.4. Let $\varphi : \Gamma' \rightarrow \Gamma$ be a morphism of tropical curves, $p' \in \Gamma'$ and $p = \varphi(p')$. Then φ is called *harmonic at p'* if for every tangent direction $\epsilon \in T_p(\Gamma)$ to p in Γ the value of the *local degree*

$$d_{p'} := \sum_{\substack{e' \in T_{p'}(\Gamma') \\ e' \mapsto \epsilon}} d_{e'}$$

does not depend on ϵ . Here the sum is running over all tangent directions to p' that map to ϵ . A morphism is *harmonic* if it is surjective and harmonic at every $p' \in \Gamma'$. In this case the number $d = \sum_{p' \in f^{-1}(p)} d_{p'}$ is independent of p and is called *degree of φ* .

DEFINITION 4.5. A harmonic morphism $\varphi : \Gamma' \rightarrow \Gamma$ is called *tropical Hurwitz cover* if for every $p' \in \Gamma'$ the *local Riemann-Hurwitz condition* holds, i.e.

$$2 - 2g(p') = d_{p'}(2 - 2g(\varphi(p'))) - \sum_{\substack{h' \text{ half-edge} \\ \text{incident to } p'}} (d_{h'} - 1).$$

We remark that tropical Hurwitz covers are not covers in the sense of topology, i.e. they are not local isomorphisms in general.

DEFINITION 4.6. Let $\pi : \Gamma' \rightarrow \Gamma$ be a tropical Hurwitz cover and let $k \geq 1$ be an integer. We call an automorphism of metric graphs $\tau : \Gamma' \rightarrow \Gamma'$ a *(tropical) deck transformation* if it is an isometry and π is τ -invariant. The data of π together with τ is called *tropical k -cyclic Hurwitz cover* if

- (i) π is of degree k ,
- (ii) the morphism of graphs underlying π is a k -cyclic cover of graphs in the sense of Definition 3.10 (with deck transformation given by the morphism of graphs underlying τ), and
- (iii) $\Gamma = \Gamma'/\tau$.

REMARK 4.7. In a tropical k -cyclic Hurwitz cover the deck transformation τ satisfies necessarily $\tau^k = \text{id}$. This however does not mean that τ is of degree k . Rather we only have that the degree of τ divides k .

We will now construct the moduli space of tropical k -cyclic Hurwitz covers in analogy to the moduli space of tropical Hurwitz covers (see [CMR16, Section 3.2]). The construction follows the same pattern as the construction of $M_{g,n}^{\text{trop}}$ in Section 2.2. Fix a degree k , genera g' and g , and a tuple of ramification profiles $\xi = (\xi_1, \dots, \xi_n)$, i.e. each ξ_i is a partition of d . Consider a tropical k -cyclic Hurwitz cover $(\pi : \Gamma' \rightarrow \Gamma, \tau)$ with the specified parameters. In particular, Γ has to have precisely n legs l_1, \dots, l_n such that the leg l_i has $|\xi_i|$ preimages with expansion factors given by the entries of ξ_i . Then the *combinatorial type* of π is the underlying k -cover of weighted graphs $(\pi : G' \rightarrow G, \tau)$ together with the data of all the expansion factors. Here we denote by G' and G the minimal graph models for Γ' and Γ respectively. We describe a category $\mathcal{J}_{g' \rightarrow g, k}(\xi)$ as follows. Objects are combinatorial types. Morphism are commutative diagrams of the form

$$\begin{array}{ccccc}
 G'_1 & \xrightarrow{f'} & & G'_2 & \\
 \downarrow \pi_1 & \searrow \tau_1 & & \downarrow \tau_2 & \\
 & & G'_1 & \xrightarrow{f'} & G'_2 \\
 & \swarrow \pi_1 & \downarrow \pi_2 & & \swarrow \pi_2 \\
 G_1 & \xrightarrow{f} & & G_2 &
 \end{array}$$

where the maps f' and f are either graph automorphism respecting expansion factors or f is an edge contraction. Note that in case f is an edge contraction, f' is already determined by [CMR16, Proposition 19]. Now associate to each combinatorial type p the rational polyhedral cone $\sigma_p := \mathbb{R}_{\geq 0}^{E(G)}$. This cone parametrizes the set of tropical k -cyclic Hurwitz covers with underlying combinatorial type p (note that edge lengths on G determine edge lengths on G'). Finally the moduli space is defined as

$$H_{g' \rightarrow g, k}^{\text{trop}}(\xi) := \varinjlim_{\mathcal{J}_{g' \rightarrow g, k}(\xi)} \sigma_p.$$

In [CMR16, Definition 25] the authors define a tropicalization which maps the analytification of Hurwitz space to the tropical Hurwitz space. Given a Hurwitz cover $X' \rightarrow X$ defined over a non-Archimedean field, its tropicalization is a tropical Hurwitz cover $\Gamma' \rightarrow \Gamma$ where Γ' is the tropicalization of X' in the sense of curves and Γ the tropicalization of X . More precisely, the Hurwitz space $\overline{\mathcal{H}}_{g' \rightarrow g, k}(\xi)$ comes with two natural forgetful maps,

$$\begin{aligned}
 \text{tgt} : \overline{\mathcal{H}}_{g' \rightarrow g, k}(\xi) &\longrightarrow \overline{\mathcal{M}}_{g, n}, & (X' \rightarrow X) &\longmapsto X \\
 \text{src} : \overline{\mathcal{H}}_{g' \rightarrow g, k}(\xi) &\longrightarrow \overline{\mathcal{M}}_{g', n'}, & (X' \rightarrow X) &\longmapsto X'
 \end{aligned}$$

called *target* and *source map* respectively. We abuse notation and denote the corresponding maps on $H_{g' \rightarrow g, k}^{\text{trop}}(\xi)$ with the same symbols. By [CMR16, Theorem 4] tropicalization (in the sense of [CMR16, Definition 25]) and source (resp. target) map commute. We now define the tropicalization map for the moduli space of k -cyclic Hurwitz covers

$$\text{trop}_H : \mathcal{H}_{g' \rightarrow g, k}(\xi)^{\text{an}} \longrightarrow H_{g' \rightarrow g, k}^{\text{trop}}(\xi)$$

simply by tropicalizing any $X' \rightarrow X$ and adding the induced tropical deck transformation τ to the data. With this definition, compatibility with source and target map remains true.

PROPOSITION 4.8. *Set $n' := \sum |\xi_i|$. Then the following diagram commutes.*

$$\begin{array}{ccc}
 \mathcal{H}_{g' \rightarrow g, k}(\xi)^{\text{an}} & \xrightarrow{\text{tgt}^{\text{an}}} & \mathcal{M}_{g, n}^{\text{an}} \\
 \downarrow \text{src}^{\text{an}} & \searrow \text{trop}_H & \downarrow \text{trop} \\
 & & H_{g' \rightarrow g, k}^{\text{trop}}(\xi) \xrightarrow{\text{tgt}} M_{g, n}^{\text{trop}} \\
 & & \downarrow \text{src} \\
 \mathcal{M}_{g', n'}^{\text{an}} & \xrightarrow{\text{trop}} & M_{g', n'}^{\text{trop}}
 \end{array}$$

PROOF. By definition trop_H commutes with forgetting the deck transformation. The claim then follows from [CMR16, Theorem 4]. \square

4.3. Tropical normalized covers. From now on fix the tuple of ramification profiles to be $\xi = ((k), \dots, (k))$ with $n = k(2g - 2)$ entries.

DEFINITION 4.9. Let $(\pi : \widehat{\Gamma} \rightarrow \Gamma, \tau)$ be a tropical k -cyclic Hurwitz cover with ramification profiles $\xi = ((k), \dots, (k))$. Assume that the effective divisor $D \in \text{Div}(\Gamma)$ given by the legs of Γ is pluri-canonical, i.e. $D = kK_\Gamma + (f)$. Let Γ^+ denote the k -enhanced level graph structure on Γ induced by Lemma 4.1. We say that π is a *tropical normalized cover* if $\widehat{\Gamma}$ admits a structure $\widehat{\Gamma}^+$ of an 1-enhanced level graph such that $\widehat{\Gamma}^+ \rightarrow \Gamma^+$ is a normalized cover of enhanced level graphs in the sense of Definition 3.14.

Define $\mathbb{P}\Xi^k M_g^{\text{trop}}$ as the locus of tropical normalized covers in $H_{g' \rightarrow g, k}^{\text{trop}}((k), \dots, (k))/S_n$ with $n = k(2g - 2)$.

The following lemma motivates that our notion of tropical normalized cover is indeed a tropical analog of Definition 3.6 in the sense that the legs of $\widehat{\Gamma}$ encode a canonical divisor. Note however that all legs of $\widehat{\Gamma}$ necessarily have to carry o -value k by Definition 3.14. Hence, the structure $\widehat{\Gamma}^+$ is not exactly the one constructed in Lemma 4.1 but rather “ k times” the output of Lemma 4.1.

LEMMA 4.10. *Let $\pi : \widehat{\Gamma} \rightarrow \Gamma$ be a tropical normalized cover with divisor $D = kK_\Gamma + (f)$ on the base. Then the legs of $\widehat{\Gamma}$ (neglecting the dilation factor k) mark the canonical divisor*

$$F = K_{\widehat{\Gamma}} + \left(\frac{f \circ \pi}{k} \right).$$

Furthermore, the 1-enhancement of any half-edge which is part of an edge in $\widehat{\Gamma}^+$ coincides with the one constructed in Lemma 4.1 based on F . In particular, the 1-enhanced level graph structure on $\widehat{\Gamma}$ is uniquely determined and thus we may speak of the normalized cover of enhanced level graphs associated to π .

PROOF. Let $p \in \widehat{\Gamma}$. To show the first claim we need to check that the number of legs at p is given by $F(p)$. The claim is clear if p does not carry a leg. Denote the k -enhancement on Γ by o and note that any leg in Γ carries o -value 1. Thus, by definition of normalized cover of enhanced level graphs any leg will have a single preimage in $\widehat{\Gamma}$ and furthermore any point in Γ with at least one leg will have a single preimage under π . In particular, p is the sole preimage of $\pi(p)$ and the local degree at p is k . In the following computation we use the notation $s(h)$ to denote the outgoing slope of f on a half-edge h of Γ and $\widehat{s}(\widehat{h})$ for the outgoing slope of $\frac{f \circ \pi}{k}$ on the half-edge \widehat{h} of $\widehat{\Gamma}$. With this notation we see

$$\begin{aligned} \#\{\text{legs at } p\} &= \#\{\text{legs at } \pi(p)\} \\ &= D(\pi(p)) \\ &= k \left(2g(\pi(p)) - 2 + \text{val}(\pi(p)) \right) + \sum_h s(h). \end{aligned}$$

On the other side, we have by definition $F(p) = 2g(p) - 2 + \text{val}(p) + \sum_{\widehat{h}} \widehat{s}(\widehat{h})$. Since the tropical normalized cover satisfies the local Riemann-Hurwitz condition, we obtain:

$$\begin{aligned} (8) \quad F(p) &= k \left(2g(\pi(p)) - 2 \right) + \sum_{\widehat{h}} (d_{\widehat{h}} - 1) + \text{val}(p) + \sum_{\widehat{h}} \widehat{s}(\widehat{h}) \\ &= k \left(2g(\pi(p)) - 2 \right) + \underbrace{\sum_h \left(\sum_{\widehat{h} \in \pi^{-1}(h)} d_{\widehat{h}} \right)}_{=k} + \sum_{\widehat{h}} \widehat{s}(\widehat{h}) \\ &= k \left(2g(\pi(p)) - 2 + \text{val}(\pi(p)) \right) + \sum_{\widehat{h}} \widehat{s}(\widehat{h}). \end{aligned}$$

The last step of Equation (8) used that π is harmonic. If \widehat{h} maps to h with dilation factor $d_{\widehat{h}}$ then

$$(9) \quad \widehat{s}(\widehat{h}) = \frac{s(h)d_{\widehat{h}}}{k}.$$

Hence, using the harmonic property once more, we see

$$\sum_{\widehat{h}} \widehat{s}(\widehat{h}) = \sum_h \frac{s(h)}{k} \underbrace{\left(\sum_{\widehat{h} \in \pi^{-1}(h)} d_{\widehat{h}} \right)}_{=k} = \sum_h s(h).$$

Plugging this into Equation (8) yields the first claim.

For the second claim we need to verify that the 1-enhancement \widehat{o} on $\widehat{\Gamma}$ which was induced by o does satisfy $\widehat{o}(\widehat{h}) + 1 = -\widehat{s}(\widehat{h})$ for every internal half-edge \widehat{h} of $\widehat{\Gamma}$. By definition of normalized cover of enhanced level graphs we have

$$\widehat{o}(\widehat{h}) + 1 = \frac{o(h) + k}{\gcd(o(h), k)} = -\frac{s(h)}{|\pi^{-1}(h)|}.$$

On the other hand we need to determine $d_{\widehat{h}}$. Once again we use that π is harmonic of degree k , i.e.

$$\sum_{\widehat{h} \in \pi^{-1}(h)} d_{\widehat{h}} = k$$

while at the same time τ is an isometry acting transitively on $\pi^{-1}(h)$. Consequently, $d_{\widehat{h}} = k/|\pi^{-1}(h)|$. Combing this result with Equation (9) we obtain

$$-\widehat{s}(\widehat{h}) = -\frac{s(h)d_{\widehat{h}}}{k} = -\frac{s(h)}{|\pi^{-1}(h)|}.$$

This completes the proof. \square

Note that for $\xi = ((k), \dots, (k))$ the tropicalization trop_H from Section 4.2 is S_n -equivariant. We define tropicalization of normalized covers

$$\text{trop}_{\Xi^k} : \mathbb{P}\Omega^k \mathcal{M}_{g,n}(1, \dots, 1)^{\text{an}} \longrightarrow \mathbb{P}\Xi^k M_g^{\text{trop}}$$

by restricting trop_H/S_n .

LEMMA 4.11. *The map trop_{Ξ^k} is well-defined, continuous, proper, and closed. Furthermore, the following diagram commutes.*

$$(10) \quad \begin{array}{ccccc} \mathbb{P}\Omega^k \mathcal{M}_{g,n}(1, \dots, 1)^{\text{an}} & & \xrightarrow{\text{trop}_{\Omega^k}} & & \mathbb{P}\Omega^k M_g^{\text{trop}} \\ & \searrow \text{trop}_{\Xi^k} & & \text{tgt} & \\ & & \mathbb{P}\Xi^k M_g^{\text{trop}} & & \\ \downarrow \text{src} & & \downarrow \text{src} & & \\ \mathbb{P}\Omega \mathcal{M}_{\widehat{g}}(1, \dots, 1)^{\text{an}} & \xrightarrow{\text{trop}_{\Omega}} & \mathbb{P}\Omega M_{\widehat{g}}^{\text{trop}} & & \end{array}$$

PROOF. To show that trop_{Ξ^k} is well-defined, let $(\widehat{X} \rightarrow X, \mathbf{s}, \omega)$ be a normalized cover of smooth curves defined over a non-Archimedean field. Its tropicalization is in particular a tropical k -cyclic Hurwitz cover $\widehat{\Gamma} \rightarrow \Gamma$. By Proposition 4.8 we know that $\widehat{\Gamma} = \text{trop}_{\Omega}(X', \omega)$ and $\Gamma = \text{trop}_{\Omega^k}(X, \eta)$, where η is the k -differential on X . By well-definedness of trop_{Ω} and trop_{Ω^k} the legs on $\widehat{\Gamma}$ and Γ do indeed represent canonical and pluri-canonical divisors F and D respectively. Finally we need to check that $\widehat{\Gamma} \rightarrow \Gamma$ can be endowed with the structure of a normalized cover of enhanced level graphs $\widehat{\Gamma}^+ \rightarrow \Gamma^+$ such that Γ^+ is induced by D via Lemma 4.1. First of all, there is indeed such a structure $\widehat{G}^+ \rightarrow G^+$ on the underlying k -cover of graphs $\widehat{G} \rightarrow G$ simply because the graphs are dual to the special fiber of the degeneration of $\widehat{X} \rightarrow X$. We check that the enhancements induced by the divisors are the same as \widehat{G}^+ and G^+ . On the cover this was part of the argument of [MUW21, Theorem 1] and hence the claim.

Finally we check that trop_{Ξ^k} is continuous, proper, and closed. The proof of [CMR16, Theorem 1] shows that the retraction map from the analytification of Hurwitz space to its Berkovich skeleton can be extended to a compactification of both of these spaces. Hence the retraction map is continuous, proper, and closed. By [CMR16, Theorem 1] tropicalization of Hurwitz covers behaves on each cone of the tropical moduli space like the retraction map. Thus the properties still hold for tropicalization of Hurwitz covers. It is now easy to check that they remain to hold after taking the S_n -quotient, i.e. for trop_H and continue to hold after restricting to the closed subset $\mathbb{P}\Xi^k M_g^{\text{trop}} \subseteq H_{\hat{g} \rightarrow g, k}^{\text{trop}}((k), \dots, (k))$. This completes the proof. \square

PROOF OF THEOREM 1.2. The claim on trop_{Ξ^k} was proved in Lemma 4.11. We describe the generalized cone complex structure on $\mathbb{P}\Xi^k M_g^{\text{trop}}$. Consider a combinatorial type of a tropical k -cyclic Hurwitz cover, i.e. a normalized cover of enhanced level graphs $p : \widehat{G}^+ \rightarrow G^+$ with dilation factors on every edge of \widehat{G} and with a deck transformation $\tau : \widehat{G} \rightarrow \widehat{G}$. In the proof of Theorem 1.1 we showed that the range of possible choices for edge lengths on G such that it becomes a tropical curve with pluri-canonical divisor is a finite union of rational polyhedral cones. Now note that each such choice determines a unique tropical Hurwitz cover by [CMR16, Lemma 17]. All of these choices give indeed tropical normalized covers, because this property does only depend on the combinatorial type. Hence we obtain a stratification of $\mathbb{P}\Xi^k M_g^{\text{trop}}$ in rational polyhedral cones.

For the statement about the dimension consider again Example 2.4. The graph depicted in Figure 1 can be endowed with a tropical normalized cover by covering every vertex with a single preimage and every edge with as many preimages are necessary to satisfy the definition of a normalized cover of enhanced level graphs. This describes a cone in $\mathbb{P}\Xi^k M_g^{\text{trop}}$ that maps under tgt isomorphically onto the cone from Example 2.4. \square

DEFINITION 4.12. The image of trop_{Ξ^k} in $\mathbb{P}\Xi^k M_g^{\text{trop}}$ is called *locus of realizable covers*.

Note that the locus of realizable covers is of positive codimension in $\mathbb{P}\Xi^k M_g^{\text{trop}}$ by the exact same argument as for the realizability locus in $\mathbb{P}\Omega^k M_g^{\text{trop}}$. Hence, realizability of normalized covers is a nontrivial problem as well. The following is a more precise version of Corollary 1.3 and reduces our original realizability problem to the one for covers.

COROLLARY 4.13. *Let $([\Gamma], D)$ be a pair consisting of an isomorphism class of a tropical curve Γ with an effective pluri-canonical divisor $D = kK_\Gamma + (f)$. The pair is realizable if and only if there exists a realizable tropical normalized cover $\widehat{\Gamma} \rightarrow \Gamma$ with $\text{tgt}([\widehat{\Gamma} \rightarrow \Gamma]) = ([\Gamma], D)$.*

PROOF. This is simply the triangle in Diagram (10) commuting. \square

REMARK 4.14. Note that contrary to the situation of twisted differentials, a tropical curve with pluri-canonical divisor does not admit a unique normalized cover. Indeed, when asking for realizability of a tropical normalized cover where a vertex v is covered with d_v preimages we are asking for a realization by a twisted differential where η_v is precisely a d_v -th power of a primitive k/d_v -differential.

5. The realizability locus

We now turn to the remaining problem of realizability of tropical normalized covers. Let $\pi : \widehat{\Gamma}^+ \rightarrow \Gamma^+$ be a tropical normalized cover with associated enhancements (see Lemma 4.10). This data contains already most of the information of a boundary point of $\mathbb{P}\Xi^k \overline{\mathcal{M}}_{g,n}(1, \dots, 1)$. In fact, only the (k) -residues of the twisted differentials are not yet determined. Realizing the tropical datum amounts to choosing a “valid” combination of (k) -residues. Obstructions arise by the k -residue map being non-surjective for some types (this leads to the notion of inconvenient vertex, see Definition 5.1) and by global compatibility conditions (we tackle this by assigning residues along certain cycles in $\widehat{\Gamma}^+$, see Definitions 5.7 and 5.9). Once notation is established, we state and prove our main theorem (Theorem 5.11).

5.1. Special vertices. As we have seen in Section 3.6, the k -residue map

$$\mathrm{Res}_g^k(\mu_{\mathrm{red}}) : \Omega^k \mathcal{M}_g(\mu_{\mathrm{red}})^{\mathrm{prim}} \longrightarrow \mathbb{C}^p \times (\mathbb{C}^\times)^s$$

is not surjective in general. Vertices with a reduced type for which the residue map is not surjective are a major obstruction to realizability. In the abelian case treated in [MUW21], those vertices were called inconvenient. We will extend the notion of inconvenience to k -differentials and additionally introduce a short list of illegal vertices – a concept which did not appear in the abelian case.

Let us fix some $k \geq 1$. We will formulate the definitions in the language of normalized covers of enhanced level graphs $\pi : \widehat{G}^+ \rightarrow G^+$. Let $v \in V(G^+)$ be a vertex, let $d_v := |\pi^{-1}(v)|$ be the number of preimages and let $k_v := k/d_v$. Recall the notation introduced in Definition 3.21 where we denoted the reduced type of v by

$$\mu_{\mathrm{red}}(v) = (a_1, \dots, a_t; -b_1, \dots, -b_p; -c_1, \dots, -c_r; -k^s).$$

We want to realize $\pi^{-1}(v) \rightarrow v$ as a normalized cover. In particular v has to be realized as a d_v -th power of a primitive k_v -differential of type

$$\mu'_{\mathrm{red}}(v) := \frac{1}{d_v} \cdot \mu_{\mathrm{red}}(v) = (a'_1, \dots, a'_t; -b'_1, \dots, -b'_p; -c'_1, \dots, -c'_r; -k_v^s).$$

For $k_v = 1$ we have $r = 0$, as in this case the c_i would be divisible by $d_v = k$, but the c_i are not divisible by k by definition. Following [MUW21] we call a vertex v inconvenient if the k_v -residue map $\mathrm{Res}_{g(v)}^{k_v}(\mu'_{\mathrm{red}}) : \mathbb{P}\Omega^{k_v} \mathcal{M}_{g(v)}(\mu'_{\mathrm{red}})^{\mathrm{prim}} \rightarrow \mathbb{C}^p \times (\mathbb{C}^\times)^s$ is not surjective. More precisely:

DEFINITION 5.1. A vertex v is called *inconvenient of type I* if μ'_{red} is one of the types in Proposition 3.22 with k_v substituted for k . It is called *inconvenient of type II* if μ'_{red} is one of the types in Propositions 3.23, 3.24 or 3.25 again with k_v instead of k . Summarizing, we call v *inconvenient* if it is inconvenient of type I or II.

Type I inconvenience means that only the origin is missing from the image of the k_v -residue map, whereas type II means that a finite number of lines or at most 2-dimensional subvarieties is missing.

Recall from Theorem 3.19 that for some strata the primitive part $\Omega^k \mathcal{M}_g(\mu)^{\mathrm{prim}}$ is empty. Consequently, vertices that ask to be realized by an element of such an empty primitive part of a stratum are not realizable at all.

DEFINITION 5.2. The vertex v is called *illegal* in the following cases.

- (i) If $\mu'_{\mathrm{red}} = (-1, 1)$.
- (ii) If $k_v = 2$ and $\mu'_{\mathrm{red}} = (1, 3)$.
- (iii) If $g(v) = 0$ and $\mathrm{gcd}(\mu'_{\mathrm{red}}, k_v) \neq 1$.
- (iv) If $k_v \geq 2$ and $\mu'_{\mathrm{red}} = \emptyset$.
- (v) If $k_v = 2$ and $\mu'_{\mathrm{red}} = (4)$.

REMARK 5.3. Being illegal is not an intrinsic property of a vertex. Rather it depends on the context of the given normalized cover. For example a vertex of type $\mu_{\mathrm{red}} = \emptyset$ is not illegal if it is covered by precisely k preimages, i.e. it may be realizable as a k -th power of an abelian differential.

5.2. Special cycles. When assigning (k) -residues to our given enhanced combinatorial data $\pi : \widehat{G}^+ \rightarrow G^+$ we have to ensure some global compatibility conditions. These conditions are compatibility with the deck transformation τ as well as the matching residue condition, global residue conditions, and the residue theorem on \widehat{G}^+ . Compatibility with τ means in particular, that the choice of residues on \widehat{G}^+ determines k -residues on G^+ . Hence we will focus on \widehat{G}^+ .

Let γ be a simple oriented cycle in \widehat{G}^+ and let L_γ denote the lowest level γ passes through. We want to use such a cycle to modify the residues of $\widehat{\Gamma}$ similar to the course of action in the proof of [MUW21, Theorem 6.3]. There the authors chose a complex number $r \in \mathbb{C}^\times$ and added to each half-edge h of \widehat{G}^+ on level L_γ the value r (resp. $-r$) to the residue at the half-edge h if γ leaves (resp. enters) the vertex incident to h along this half-edge. This operation maintains the residue theorem at each vertex, the matching residue

condition, and the global residue condition, i.e. if each of those conditions was true before modifying the residues of \widehat{G}^+ , the conditions still hold for the modified residues.

Recall that compatibility with the deck transformation $\tau : \widehat{G}^+ \rightarrow \widehat{G}^+$ means that the assigned residue at a half-edge $\tau(h)$ has to be ζ times the residue at h . After adding the residues along the cycle γ as described above, this is no longer the case in general. We address this problem by not only considering the cycle γ , but the entire τ -orbit of γ consisting of

$$\gamma_i := \tau_*^i(\gamma) \quad \text{for } i = 0, \dots, k-1.$$

We provide each of the cycles γ_i with the induced orientation and add $\pm\zeta^i r$ to the residues as described above. We will refer to this operation as *assigning the residue r along the orbit of γ* . The total change to the residues under this operation can be easily expressed with the following shorthand notation.

DEFINITION 5.4. Let γ be a simple closed cycle in \widehat{G}^+ with fixed orientation. Given a vertex \widehat{v} in \widehat{G}^+ , let $H'(\widehat{v})$ denote the set of half-edges incident to \widehat{v} with o -value ≤ -1 . Then we define a vector $R_\gamma(\widehat{v}) = (R_\gamma(\widehat{v})_h)_{h \in H'} \in \mathbb{C}^{|H'|}$ as follows. Set it to be 0 if \widehat{v} does not lie on the lowest level that γ passes through and otherwise

$$(11) \quad R_\gamma(\widehat{v})_h := \sum_{i=0}^{k-1} \epsilon_i \zeta^i, \quad \text{where } \epsilon_i := \begin{cases} 1 & \text{if } \tau_*^i \gamma \text{ enters } \widehat{v} \text{ through } h \\ -1 & \text{if } \tau_*^i \gamma \text{ leaves } \widehat{v} \text{ through } h \\ 0 & \text{if } \tau_*^i \gamma \text{ does not pass through } h. \end{cases}$$

Now let v be the image of \widehat{v} under $\pi : \widehat{G}^+ \rightarrow G^+$ and denote μ_{red} the reduced type of v as in Definition 3.21. We define a vector $R_\gamma^k(v) \in \mathbb{C}^{p+s}$ to be 0 if $R_\gamma(\widehat{v})$ is 0 and otherwise $R_\gamma^k(v)_h := (R_\gamma(\widehat{v})_{\widehat{h}})^k$ for $\widehat{h} \in \pi^{-1}(h)$ arbitrary.

When assigning the residue r along the orbit of γ the residues at a vertex \widehat{v} change precisely by adding r times $R_\gamma(\widehat{v})$ to the vector of residues. The residues obtained in this way obviously still preserve the residue theorem, the matching residue condition, and the global residue condition. Additionally, we have maintained compatibility with the deck-transformation. This means that the residues we assigned on \widehat{G}^+ induce well-defined k -residues on G^+ . By Lemma 3.8 the change to the k -residues at a vertex v of G is precisely r^k times $R_\gamma^k(v)$. Observe that a given simple closed cycle γ can only be used to change the k -residues at a vertex v by a \mathbb{C} -multiple of $R_\gamma^k(v)$. This means that some cycles will be more useful for our purpose than others. We illustrate two notable phenomena in Examples 5.6 and 5.8.

REMARK 5.5. The choice of an orientation in Definition 5.4 merely fixes the sign of $R_\gamma(\widehat{v})$. For the rest of this article only the \mathbb{C} -span of $R_\gamma^k(v)$ will be of relevance (see Definitions 5.7 and 5.9 below), hence this choice never really matters.




FIGURE 3. A cover of graphs with the action of the deck-transformation τ .

EXAMPLE 5.6. Consider the cover of enhanced level graphs $\widehat{G}^+ \rightarrow G^+$ for $k = 2$ depicted in Figure 3. Let γ_1 be the simple cycle in \widehat{G}^+ that uses the two topmost edges,

considered with the orientation indicated in the picture. When we compute $R_{\gamma_1}(\hat{v})$ for either of the vertices in \hat{G}^+ we see that the vector is zero. To see this, consider the top left half-edge. Along γ_1 , we add 1. The oriented cycle $\tau_*\gamma_1$ agrees with γ_1 , and thus we add at the same half-edge the value $\zeta = -1$ when distributing the residues for $\tau_*\gamma_1$. In total we have added $1 - 1 = 0$ to the residue. This means that γ_1 cannot be used to assign any nonzero residues at all.

Let γ_2 be the simple cycle in \hat{G}^+ that uses the two outermost edges, again considered with the orientation indicated in the picture. Computing $R_{\gamma_2}(\hat{v})$ for either of the vertices in \hat{G}^+ we see that the vector now is $(-1, -1, 1, 1)$. But this vector is not in the image of the 1-residue map for a vertex of type $(-1^4, 2)$ by Proposition 3.24. These two observations motivate the following definition.

DEFINITION 5.7. Let γ be a simple closed cycle in \hat{G}^+ and let v be a vertex in G^+ . Choose and fix an orientation of γ . We say that γ is *effective* for a half-edge h incident to v if $R_\gamma^k(v)_h$ is nonzero. If v is inconvenient of type I, we say that γ is *admissible for v* if γ is effective for at least one vertical half-edge h incident to v . If v is inconvenient of type II, we say that γ is *admissible for v* if $R_\gamma^{k_v}(v)$ lies in the image of $\text{Res}_g^{k_v}(\mu'_{\text{red}}(v))$.




FIGURE 4. An inconvenient vertex which cannot be redeemed with a single cycle.

EXAMPLE 5.8. Contrary to the abelian case in [MUW21] there are situations, where using only one cycle will not be sufficient to achieve valid residues. To see this, consider the graph in Figure 4. The vertex on lowest level is inconvenient of type II. More specifically, the image of the residue map is missing any tuples where two entries agree while the other two are zero. The depicted cover \hat{G}^+ is the only valid choice – covering a genus 1 vertex of type $(0, 0)$ with only one preimage would be illegal. Notice that each of the four simple cycles in \hat{G}^+ is effective but induces a tuple of k -residues on the base which is not contained in the image of the residue map. Hence none of the available cycles is sufficient to redeem the inconvenient vertex, however using two cycles will work. A converse to this phenomenon is illustrated in Example 5.21: there we have an inconvenient vertex which can only be redeemed with an admissible cycle but no pair of cycles.

DEFINITION 5.9. Let v be a vertex in \hat{G}^+ and let γ and γ' be oriented cycles which are effective for v . We call (γ, γ') an *independent pair for v* if the induced vectors $R_\gamma^{k_v}(v)$ and $R_{\gamma'}^{k_v}(v)$ are not contained in the same linear subspace of the complement of $\text{Res}_g^{k_v}(\mu'_{\text{red}}(v))$.

REMARK 5.10. The upshot behind Definition 5.9 is that linear combinations of the vectors $R_\gamma^{k_v}(v)$ and $R_{\gamma'}^{k_v}(v)$ will generically lie in the image of the k_v -residue map. This is trivially true for inconvenient vertices of type I as the cycles of an independent pair are necessarily admissible, i.e. effective. In fact, the notion of independent pair is not interesting for type I vertices. For inconvenient vertices of type II this is easily seen to be true for all cases where the complement of $\text{Res}_g^{k_v}(\mu'_{\text{red}}(v))$ consists of a finite union of lines or planes, i.e. for all cases except those considered in Proposition 3.25 Part (ii). This is the only case where a finite union of 2 dimensional cones – none of which contains a 2

dimensional plane – is missing. Here a pair of cycles may still be independent even if both R^{kv} -vectors lie in the same irreducible component of the complement of $\text{Im Res}_g^{kv}(\mu'_{\text{red}}(v))$.

5.3. Realizability of covers.

THEOREM 5.11. *Fix an algebraically closed base field of characteristic 0. Let $g \geq 2$ and fix an integer $k \geq 1$. Let $\pi : \widehat{\Gamma}^+ \rightarrow \Gamma^+$ be a tropical normalized cover with enhancements associated by Lemma 4.10. Denote the effective pluri-canonical divisor marked by the legs of Γ by $D = kK_\Gamma + (f)$. Then $(\pi : \widehat{\Gamma}^+ \rightarrow \Gamma, D)$ is realizable if and only if the following conditions hold.*

- (i) *There is no illegal vertex in π .*
- (ii) *For every horizontal edge \widehat{e} in $\widehat{\Gamma}^+$ there is an effective cycle in $\widehat{\Gamma}^+$ through \widehat{e} .*
- (iii) *For every inconvenient vertex v in Γ^+ there is an admissible cycle in $\widehat{\Gamma}^+$ through one of the preimages \widehat{v} or there is an independent pair of cycles.*

REMARK 5.12. Let us explain, how to recover [MUW21, Theorem 6.3] from Theorem 5.11 for $k = 1$. Recall that [MUW21, Theorem 6.3] states the following: the pair (Γ, D) for D as above is realizable if and only if

- (i') For every inconvenient vertex (in the sense of [MUW21, Definition 6.2]) v in Γ there is a simple cycle in Γ through v that does not pass through any node on a level below v .
- (ii') For every horizontal edge e in Γ there is a simple cycle passing through e which does not pass through any node on a level below e .

To see that these conditions are equivalent to ours, note that for $k = 1$ the identity on Γ is the only tropical normalized cover. Now assume $(\text{id} : \Gamma \rightarrow \Gamma, D)$ satisfies the conditions (i), (ii), and (iii) of Theorem 5.11. Every inconvenient vertex in the sense of [MUW21] is inconvenient in our sense as well. Furthermore, every effective or admissible cycle does not pass through any lower level. Hence (Γ, D) satisfies (i') and (ii') as well.

Conversely, suppose (i') and (ii') hold. First note that the only type of illegal vertex for $k = 1$ is $(-1, 1)$ and such a vertex does not admit a simple cycle “at or above level” through the incident horizontal edge. Thus (ii') ensures that there are no illegal vertices. The next observation is that for $k = 1$ a cycle γ is effective for every half-edge at lowest level that γ passes through. In particular, (ii) holds. Furthermore, there are only two kinds of inconvenient vertices: the ones in Proposition 3.22 (i) and the ones in Proposition 3.24. The former is inconvenient in the sense of [MUW21] as well. Hence, (i') provides the necessary effective cycles. The other kind of inconvenient vertex is not an issue for $k = 1$: all simple cycles use precisely two half-edges incident to each vertex they pass through, and thus the residues at ≥ 3 horizontal half-edges may always be chosen sufficiently generic. In other words, the cycles provided by (ii') contain an independent pair.

We split the proof of Theorem 5.11 in three parts. First we prove that the conditions in the theorem are sufficient (resp. necessary) for realizability over the base field \mathbb{C} . Afterwards we generalize the result to arbitrary algebraically closed base fields of characteristic 0.

PROPOSITION 5.13. *The conditions in Theorem 5.11 are sufficient for realizability over the base field \mathbb{C} .*

PROOF. First, we reduce to tropical curves with integer edge lengths. Indeed, the set of tropical curves with rational edge lengths and pluri-canonical divisor is dense in $\mathbb{P}\Xi^k M_g^{\text{trop}}$. Furthermore, the locus in $\mathbb{P}\Xi^k M_g^{\text{trop}}$ described by the conditions in Theorem 5.11 is closed. By Lemma 4.11 trop_{Ξ^k} is continuous and closed. Hence it suffices to consider tropical curves with rational edge lengths. But then again if π is realizable then so is the cover obtained by rescaling the edge lengths with a global constant.

Now assume that the conditions of Theorem 5.11 hold. Ultimately we want to realize π by a normalized cover $\pi : \widehat{X} \rightarrow X$ of smooth curves over a non-Archimedean field with residue field \mathbb{C} , such that the k -differential on the base is of type $(1, \dots, 1)$. To do so, we want to choose for every half-edge in $\widehat{\Gamma}^+$ with o -value ≤ -1 and every half-edge in Γ^+ with o -value $\leq -k$ a (k) -residue (i.e. a complex number) such that all of the following hold.

For every vertex $v \in \Gamma^+$ there exists a smooth curve X_v with meromorphic k -differential η_v realizing v . More precisely, X_v is supposed to be of genus $g(v)$ with distinguished points $z_h \in X_v$ for every half edge h incident to v such that $\text{ord}_{z_h} \eta_v = o(h)$ and $\text{Res}_{z_h}^k \eta_v$ is the value chosen in the beginning and η_v is holomorphic and nonzero outside of $\{z_h\}_h$. Furthermore, η_v is supposed to be the d_v -th power of a primitive k_v -differential of type $\mu'_{\text{red}}(v)$. Yet again further, we require each of the connected components $\widehat{X}_{\widehat{v}}$ of the (uniquely determined) normalized cover of X_v to realize one of the vertices $\widehat{v} \in \pi^{-1}(v)$, again such that orders of the meromorphic abelian differentials $\omega_{\widehat{v}}$ match the o -values on $\widehat{\Gamma}^+$ and the residues coincide with the chosen values from the beginning. Finally, we need to do all of this such that the normalized covers $\coprod \widehat{X}_{\widehat{v}} \rightarrow X_v$ glue into a normalized cover of nodal curves $\widehat{X} \rightarrow X$ with dual graphs given precisely by $\widehat{\Gamma}$ and Γ . Once this is achieved, we obtain the desired normalized cover of smooth curves with deformation parameters corresponding to the edge-lengths of $\widehat{\Gamma}$ and Γ by means of Theorem 3.18.

We note some dependencies among these requirements. Specifying residues on $\widehat{\Gamma}^+$ that satisfy the condition imposed by the residue theorem and that are compatible with the deck transformation τ already determines the k -residues on Γ^+ . This ensures the realizability of v in the above sense if the induced k_v -residues are contained in the image of the k_v -residue map $\text{Res}_{g(v)}^{k_v}(\mu'_{\text{red}}(v))$. In this case realizability of all of the $\widehat{v} \in \pi^{-1}(v)$ is immediate. When it comes to global compatibility note first that compatibility with the level structure is already built into the definition of enhanced level graphs. Beyond that, we only need to ensure MRC and GRC for the cover $\pi : \widehat{\Gamma}^+ \rightarrow \Gamma^+$. To summarize, our goal is to choose residues on $\widehat{\Gamma}^+$ such that:

- For each $\widehat{v} \in \widehat{\Gamma}^+$ the condition imposed by the residue theorem is satisfied. Moreover, MRC and GRC are satisfied.
- Residues on $\widehat{\Gamma}^+$ are compatible with τ , i.e. the residue of $\tau(h)$ is precisely ζ times the residue at h .
- The k -residues which are given on $\pi(h)$ as the k -th power of the residue at h make every vertex of Γ realizable.

Let us now argue that a suitable choice of such residues exists. We start by initializing all residues with 0. Let $\gamma_1, \dots, \gamma_\lambda$ be all the simple cycles in $\widehat{\Gamma}^+$ that exist by assumption, i.e. the cycles containing the horizontal edges and all kinds of inconvenient vertices. For each γ_i we choose and fix an orientation. Note that at this point the first two items of our list of requirements are already satisfied. By construction of the process of “assigning a residue r_i along the τ -orbit of γ_i ”, these conditions continue to hold after doing so.

Let us now pick numbers $r_1, \dots, r_\lambda \in \mathbb{C}$ to be assigned along the orbits of the γ_i such that the third and final condition is met. This amounts to choosing the r_i sufficiently generic such that no undesirable cancellation happens. More precisely, after all the residues have been assigned, the resulting k_v -residues at a vertex $v \in \Gamma^+$ are

$$\sum_{i=1}^{\lambda} r_i^{k_v} R_{\gamma_i}^{k_v}(v)$$

and this has to lie in the image of the k_v -residue map. At every vertex this amounts to avoiding a locus of positive codimension in \mathbb{C}^{p+s} . At the same time, the values that can be achieved using the given γ_i form a vector space V_v . By assumption we have for every deficit in surjectivity of the residue map an admissible cycle or an independent pair of cycles, i.e. V_v is not fully contained in the complement of the image of the residue map. Hence a suitable choice for each r_i is possible and we are done. \square

Let \widehat{X} be a smooth complex curve, and denote by $\text{PD} : H_{\text{dR}}^1(\widehat{X}; \mathbb{R}) \rightarrow H_1(\widehat{X}; \mathbb{R})$ the map given by Poincaré duality. By abuse of notation, we denote the induced map

$$\begin{aligned} \text{PD} : \Omega^1(\widehat{X}) &\longrightarrow H_1(\widehat{X}; \mathbb{C}) \\ \omega &\longmapsto \text{PD}(\text{Re}(\omega)) \oplus i \cdot \text{PD}(\text{Im}(\omega)) \end{aligned}$$

by the same symbol. By naturality of PD there is a commutative diagram

$$(12) \quad \begin{array}{ccc} \Omega^1(\widehat{X}) & \xrightarrow{\text{PD}} & H_1(\widehat{X}; \mathbb{C}) \\ \downarrow \tau^* & & \uparrow \tau_* \\ \Omega^1(\widehat{X}) & \xrightarrow{\text{PD}} & H_1(\widehat{X}; \mathbb{C}). \end{array}$$

PROPOSITION 5.14. *The conditions in Theorem 5.11 are necessary for realizability over the base field \mathbb{C} .*

PROOF. Let $[\widehat{\Gamma} \rightarrow \Gamma] = \text{trop}_{\Xi^k}(\pi : \widehat{X} \rightarrow X, \mathbf{s}, \omega)$ be given. We want to show that the tropical normalized cover satisfies the conditions of Theorem 5.11. By continuity of trop_{Ξ^k} (Lemma 4.11) it suffices to show this for any $\text{trop}_{\Xi^k}(\pi : \widehat{X} \rightarrow X)$ for π taken from a dense subset of $\mathbb{P}\Omega^k \mathcal{M}_{g,n}(1, \dots, 1)^{\text{an}}$. In particular, we may assume that π is defined over a finite extension of $\mathbb{C}(t)$. This allows us to take the equivalent \mathbb{C} -analytic point of view and consider this data as a family of normalized covers $(\pi_t : \widehat{X}_t \rightarrow X_t)_t$ over the punctured unit disc Δ^* . In particular, each \widehat{X}_t and X_t is a smooth curve over \mathbb{C} .

Let $\widehat{X}_0 \rightarrow X_0$ denote the admissible cover obtained as the limit of the family within $\mathbb{P}\Xi^k \overline{\mathcal{M}}_{g,n}(1, \dots, 1)$ for $t \rightarrow 0$. By assumption, the dual graphs of \widehat{X}_0 and X_0 are the underlying unmetrized graphs of $\widehat{\Gamma}$ and Γ respectively. Furthermore, the enhanced level graph structures induced by \widehat{X}_0 and X_0 are precisely $\widehat{\Gamma}^+$ and Γ^+ respectively (see the argument in the proof of Lemma 4.11). We check that the conditions of Theorem 5.11 hold for these.

There cannot be any illegal vertex v in Γ . Otherwise the restriction $\prod_{\widehat{v} \in \pi^{-1}(v)} \widehat{X}_{0, \widehat{v}} \rightarrow X_{0, v}$ of the central fiber would provide an element in a stratum that is empty by Theorem 3.19, a clear contradiction.

Next, we show the existence of effective cycles for horizontal edges and admissible cycles or independent pairs of cycles for all inconvenient vertices. Fix such an edge or vertex and let L denote its level in Γ^+ . Recall that we are only interested in cycles at or above level L . To ensure that any cycle we find during this proof satisfies this condition, we use the following trick. Take the truncated cover $\widehat{X}_{0, \geq L} \rightarrow X_{0, \geq L}$ at or above level L . After restricting to the connected component that contains the edge or vertex under consideration, we obtain a twisted differential from some holomorphic stratum $\mathbb{P}\Xi^k \overline{\mathcal{M}}_{g', n'}(\mu')$. It can be written as limit of a family $(\pi'_t : \widehat{X}'_t \rightarrow X'_t)_t$ of smooth normalized covers. Moreover, any cycle in the dual graph of $\widehat{X}_{0, \geq L}$ is also a cycle in the dual graph of \widehat{X}_0 at or above level L that inherits the property of being effective (resp. admissible). Thus it suffices to find suitable cycles in $\widehat{X}_{0, \geq L}$. Hence we will implicitly work with the family π'_t and assume all our cycles to be at or above level L .

Let ω_t be the abelian differential on \widehat{X}_t and let $\gamma_t := \text{PD}(\omega_t)$. By the commutativity of Diagram (12) we have

$$\gamma_t = \text{PD}(\omega_t) = \tau_* \text{PD}(\tau^* \omega_t) = \tau_* \text{PD}(\zeta \omega_t) = \zeta \tau_* \gamma_t.$$

Repeating the argument with τ being replaced by a power τ^i for $i = 0, \dots, k-1$ and summing the resulting equations we obtain

$$(13) \quad k\gamma_t = \sum_{i=0}^{k-1} \zeta^i \tau_*^i \gamma_t.$$

Now consider a vertex $v \in \Gamma^+$ and let h_1, \dots, h_{p+s} be the half-edges incident to v with α -value $\leq -k$ and divisible by k . Let \widehat{v} be a preimage of v and let $\widehat{h}_1, \dots, \widehat{h}_{p+s}$ be half-edges incident to \widehat{v} such that \widehat{h}_i is a preimage of h_i . Let $\alpha_t^{(1)}, \dots, \alpha_t^{(p+s)}$ be families of simple closed cycles in \widehat{X}_t which get pinched into the corresponding nodes $\widehat{q}_1, \dots, \widehat{q}_{p+s}$. Observe that $\int_{\alpha_t^{(\iota)}} \omega_t$ converges for $t \rightarrow 0$ to the residue $r_{\widehat{q}_\iota}$ of the limiting twisted differential at \widehat{q}_ι . By Poincaré duality and equation (13) this implies

$$(14) \quad \lim_{t \rightarrow 0} \sum_{i=0}^{k-1} \zeta^i \tau_*^i \gamma_t \cap \alpha_t^{(\iota)} = \lim_{t \rightarrow 0} k\gamma_t \cap \alpha_t^{(\iota)} = kr_{\widehat{q}_\iota} \quad \text{for } \iota = 1, \dots, p+s.$$

Next we want to make a consistent choice of a basis of $H_1(\widehat{X}_t; \mathbb{C})$. In general, such a choice is not possible across the entire family. Hence we restrict our family to a real ray in Δ^* , i.e. from now on only consider $t \in (0, 1)$. Recall that there is a surjective map

$$\Phi : H_1(\widehat{X}_0; \mathbb{Z}) \longrightarrow H_1(\widehat{G}; \mathbb{Z})$$

into the homology of the dual graph \widehat{G} of the special fiber \widehat{X}_0 . Let $(\beta_j^{\text{trop}})_{j \in \{1, \dots, b_1(\widehat{G})\}}$ be a basis of $H_1(\widehat{G}; \mathbb{Z})$ consisting of simple cycles and let $\beta'_j \in H_1(\widehat{X}_0; \mathbb{Z})$ be preimages of the β_j^{trop} . Let $J := \{1, \dots, 2g(\widehat{X})\}$. We can complete the $(\beta'_j)_{j \in \{1, \dots, b_1(\widehat{G})\}}$ to a basis $(\beta'_j)_{j \in J}$ of $H_1(\widehat{X}_0; \mathbb{Z})$ in such a way that $\beta'_j \in \ker \Phi$ for $i \in \{b_1(\widehat{G}) + 1, \dots, 2g(\widehat{X}_0)\}$. In other words, the new β'_j have a representative with support in a single irreducible component $\widehat{X}_{0,v}$ of \widehat{X}_0 . Chose cycles β_j on a nearby surface \widehat{X}_t along the real ray such that β_j converges to β'_j for $t \rightarrow 0$. The $(\beta_j)_{j \in J}$ form a basis of $H_1(\widehat{X}_t; \mathbb{Z})$ and $H_1(\widehat{X}_t; \mathbb{C})$ for all t in the real ray.

With our chosen basis we may write $\gamma_t = \sum_{j \in J} c_t^{(j)} \beta_j$ for uniquely determined complex coefficients $c_t^{(j)}$ varying continuously for $t \in (0, 1)$. Equation (14) implies

$$(15) \quad \lim_{t \rightarrow 0} \sum_{i=0}^{k-1} \sum_{j \in J} c_t^{(j)} \zeta^i \tau_*^i \beta_j \cap \alpha_t^{(\iota)} = kr_{\widehat{q}_\iota} \quad \text{for } \iota = 1, \dots, p+s.$$

We claim that for all $\iota = 1, \dots, p+s$ we have

$$\lim_{t \rightarrow 0} \sum_{i=0}^{k-1} \zeta^i \tau_*^i \beta_j \cap \alpha_t^{(\iota)} = \begin{cases} R_{\beta_j^{\text{trop}}}(\widehat{v})_{\widehat{h}_\iota} & \text{for } j = 1, \dots, b_1(\widehat{G}) \\ 0 & \text{otherwise} \end{cases}$$

for $R_{\beta_j^{\text{trop}}}(\widehat{v})_{\widehat{h}_\iota}$ as in (11). For the first claim observe that for $j \in \{1, \dots, b_1(\widehat{G})\}$ the limit $\lim_{t \rightarrow 0} \beta_j \cap \alpha_t^{(\iota)}$ agrees with the coefficient of β_j^{trop} in front of the edge $\widehat{e} \in \widehat{G}$ corresponding to the node $\widehat{q}_\iota \in \widehat{X}_0$. Thus the claim follows by comparing equation (13) with equation (11). Here the sign appearing in (11) is encoded in the intersection product of the implicitly oriented cycles β_j and $\alpha_t^{(\iota)}$. For the second claim observe that for $j \in \{b_1(\widehat{G})+1, \dots, 2g(\widehat{X})\}$ the cycle β_j is chosen such that it does not intersect $\alpha_t^{(\iota)}$. Thus we may rewrite equation (15) as

$$(16) \quad \sum_{j=1}^{b_1(\widehat{G})} c^{(j)} R_{\beta_j^{\text{trop}}}(\widehat{v})_{\widehat{h}_\iota} = kr_{\widehat{q}_\iota} \quad \text{for } \iota = 1, \dots, p+s,$$

where $c^{(j)} := \lim_{t \rightarrow 0} c_t^{(j)}$.

Now assume that h_{ι_0} belongs to a horizontal edge. In this case, $r_{\widehat{q}_{\iota_0}}$ is nonzero. Hence by equation (16), there is a $j_0 \in J$ such that $R_{\beta_{j_0}^{\text{trop}}}(\widehat{v})_{\widehat{h}_{\iota_0}} \neq 0$. In particular $\beta_{j_0}^{\text{trop}}$ is effective for h_{ι_0} .

Now assume that v is inconvenient of type I. In this case, the origin is not in the image of the residue map. Thus there is an ι_0 such that $r_{\widehat{q}_{\iota_0}}$ is nonzero. By the same argument as above, there is an j_0 such that $R_{\beta_{j_0}^{\text{trop}}}(\widehat{v})_{\widehat{h}_{\iota_0}} \neq 0$ by equation (16) and thus $\beta_{j_0}^{\text{trop}}$ is admissible for v .

Now assume that v is inconvenient of type II. In this case, the powers $(r_{\widehat{q}_\iota}^{k_v})_\iota$ are contained in the image of the residue map. Recall that the complement of the image of the residue map is an union of at most 2-dimensional subvarieties $\bigcup_{\sigma'} W'_{\sigma'} \in \mathbb{C}^{p+s}$. Thus the residues $(r_{\widehat{q}_\iota})_\iota$ may not be contained in an union of at most 2-dimensional subvarieties $\bigcup_{\sigma} W_{\sigma} \in \mathbb{C}^{p+s}$, too: Each subvariety W' in the image of the residue map gives rise to multiple subvarieties W corresponding to choices of the k -th root. If there is an admissible cycle β_j^{trop} we are done. So assume that there is none. Moreover, assume for a contradiction that all vectors $(R_{\beta_j^{\text{trop}}}(\widehat{v})_{\widehat{h}_\iota})_\iota$ are contained in a single linear subspace V of a subvariety W_{σ_0} . Then by equation (16), the vector $(r_{\widehat{q}_\iota})_\iota$ is contained in the same linear subspace V . But then the vector $(r_{\widehat{q}_\iota})_\iota$ is not contained in the image of the residue map, which is a contradiction. Thus there is an independent pair of cycles $(\beta_{j_1}^{\text{trop}}, \beta_{j_2}^{\text{trop}})$. \square

So far we have shown the claim to hold over \mathbb{C} . We now generalize to any algebraically closed base field of characteristic 0.

PROPOSITION 5.15. *If Theorem 5.11 is true over the base field \mathbb{C} , it is true over any algebraically closed base field of characteristic 0.*

PROOF. If $K \subseteq L$ is a valued field extension of such fields, then there is a surjective map $(\mathbb{P}\Omega^k \mathcal{M}_g(1, \dots, 1)_L)^{\text{an}} \rightarrow (\mathbb{P}\Omega^k \mathcal{M}_g(1, \dots, 1)_K)^{\text{an}}$ and the following diagram commutes by [Gub13, Proposition 3.7]

$$\begin{array}{ccc} (\mathbb{P}\Omega^k \mathcal{M}_g(1, \dots, 1)_L)^{\text{an}} & & \\ \downarrow & \searrow^{\text{trop}_{\Xi^k}} & \\ (\mathbb{P}\Omega^k \mathcal{M}_g(1, \dots, 1)_K)^{\text{an}} & \xrightarrow{\text{trop}_{\Xi^k}} & \mathbb{P}\Xi^k \mathcal{M}_g^{\text{trop}}. \end{array}$$

Hence, the image of trop_{Ξ^k} does not depend on the base field (whether larger or smaller than \mathbb{C}). \square

5.4. Dimensions. In the abelian case [MUW21, Theorem 6.6] shows that the realizability locus is a pure dimensional generalized cone complex of dimension equal to $4g - 4 = \dim \mathbb{P}\Omega \mathcal{M}_g(1, \dots, 1)$. From [Uli15, Theorem 1.1] we know that the realizability locus for $k \geq 2$ must be a generalized cone complex of dimension $\leq \dim \mathbb{P}\Omega^k \mathcal{M}_g(1, \dots, 1) = (2 + 2k)(g - 1) - 1$ (see the discussion in Section 2.6). Let us now prove Theorem 1.5 from the introduction and show that this bound is in fact attained and all maximal cones have the same dimension. To do so we need two preparational statements. The proof for the following lemma is the same as [MUW21, Lemma 6.8].

LEMMA 5.16. *For every realizable tropical normalized cover $\pi : \widehat{\Gamma}^+ \rightarrow \Gamma^+$ let Γ_0^+ be the level graph obtained by successively contracting edges in Γ^+ that have an $(n + 1)$ -valent genus zero node with $n \geq 1$ marked points at one of its ends. The dimension of the cone in the realizability locus with associated normalized cover π is 1 less than the number of levels plus the number of horizontal edges of Γ_0^+ .*

PROPOSITION 5.17. *Let $k \geq 2$ and let $\pi : \widehat{\Gamma}^+ \rightarrow \Gamma^+$ be a realizable tropical normalized cover. Then π is contained in a cone of dimension $(2 + 2k)(g - 1) - 1$.*

PROOF. Let $c(\pi)$ denote the number of levels minus 1 plus the number of horizontal edges of Γ^+ . As π is realizable, the underlying cover of enhanced level graphs cuts out a boundary stratum $D_\pi \subseteq \mathbb{P}\Xi^k \overline{\mathcal{M}}_{g,n}(1, \dots, 1)$ of codimension $c(\pi)$, i.e. for all multi-scale k -differentials $(\widehat{X} \rightarrow X, \mathbf{s}, \omega, \widehat{G}^+ \rightarrow G^+) \in D_\pi$ the underlying cover of enhanced level graphs $\widehat{G}^+ \rightarrow G^+$ agrees with π , see [CMZ19, Proposition 1.3]. We will prove that the closure \overline{D}_π intersects a boundary stratum of maximal codimension, i.e. that there is a multi-scale k -differential $(\pi' : \widehat{X}' \rightarrow X') \in \overline{D}_\pi$ with $c(\pi') = (2 + 2k)(g - 1) - 1$.

Assume that we have found such a π' . We want to use Lemma 5.16 to see that the tropicalization of π' gives rise to a degeneration of π that spans a cone of the claimed dimension. To apply Lemma 5.16, we need to rearrange any occurring trees of marked points in π' as depicted in Figure 5, where the level structure on the graph on the right may be different depending on the order on the half-edges q_i . This is always possible: First note that all irreducible components of X' need to have genus zero. Second note that the “contracted” vertex in the middle of Figure 5 must be realized by a primitive k -differential. But for a stratum of primitive k -differentials of genus zero, the forgetful map to $\mathcal{M}_{0,n}$ is surjective. Thus the degeneration is contained in the closure.

We will prove the existence of the degeneration π' by induction on $c(\pi)$. Assume that $c(\pi)$ is not maximal, i.e. that $c(\pi) < (2 + 2k)(g - 1) - 1$. Let us prove that \overline{D}_π intersects a boundary stratum of higher codimension. To this end, let for a level L of Γ^+

$$\sigma_L : D_\pi \longrightarrow \left(\prod_{v \in L} \Omega^k \mathcal{M}_{g(v)}(\mu(v)) \right) / \mathbb{C}^\times$$

be the map that cuts out level L . (Note that projectivization of strata of multi-scale differentials is done with respect to the diagonal \mathbb{C}^\times -action, and not with respect to the




FIGURE 5. Rearranging a rational tail

(\mathbb{C}^\times action on each irreducible component.) Since $c(\pi)$ was assumed to not be maximal, there exists a level L such that $\dim \sigma_L(D_\pi) \geq 1$. We need to argue that such a level L can always be degenerated.

If $\sigma_L(D_\pi)$ consists of multiple connected components that may be rescaled independently, then we obtain an additional level by taking the limit of this rescaling, i.e. by distributing the vertices of level L to two levels according to the rescaling. If $\sigma_L(D_\pi)$ does not consist of multiple connected components (or those cannot be rescaled independently), then there must be an irreducible component of positive (projective) dimension in $\sigma_L(D_\pi)$. For ease of notation we assume that $\sigma_L(D_\pi)$ consists of only one such component, i.e.

$$(17) \quad \sigma_L(D_\pi) \subseteq \mathbb{P}\Omega^k \mathcal{M}_{g(v')}(\mu(v')).$$

The map σ_L can be extended to the closure \overline{D}_π using an appropriate moduli space of multi-scale k -differentials as codomain for the extended map. If $\sigma_L(\overline{D}_\pi)$ contains a point from the boundary of $\mathbb{P}\Omega^k \mathcal{M}_{g(v')}(\mu(v'))$, then this point gives rise to a degeneration of the level L . Such a point always exists: The image $\sigma_L(\overline{D}_\pi)$ is a complete variety and the stratum $\mathbb{P}\Omega^k \mathcal{M}_{g(v')}(\mu(v'))$ on the right hand side of (17) does not contain any complete variety by [Gen20, Théorème 1 and Corollaire 2]. This concludes the proof. \square

PROOF OF THEOREM 1.5. The dimension of the realizability locus is bounded from above by the dimension of the domain of the tropicalization map, i.e. by $\dim \mathbb{P}\Xi^k \overline{\mathcal{M}}_{g,n}(1, \dots, 1) = (2 + 2k)(g - 1) - 1$. That this bound is actually obtained and all maximal cones are of the expected dimension follows from the previous proposition. \square

REMARK 5.18. We emphasize that for $k = 1$ the above formula does not give the correct dimension for the maximal cones. This is due to the formula for the dimension of the principal stratum being different in the abelian case.

5.5. Obstructions to realizability. The following are two simple criteria which can be used to recognize non-realizable tropical normalized covers. An application is illustrated in Section 6 below.

COROLLARY 5.19. *Let $\pi : \widehat{\Gamma}^+ \rightarrow \Gamma^+$ be a tropical normalized cover with enhancements associated by Lemma 4.10. Let e be a horizontal edge in Γ^+ and denote by $\pi/\{e\}$ the tropical normalized cover obtain from π by contracting e in the base and every $\widehat{e} \in \pi^{-1}(e)$. If π is realizable, then so is $\pi/\{e\}$.*

PROOF. Let $\widehat{X} \rightarrow X$ be a realization of π . Smoothing of a horizontal edge of $\widehat{X} \rightarrow X$ is always possible, see [CMZ19, Section 3.1], and produces a realization of $\pi/\{e\}$. \square

With the same proof, we get

COROLLARY 5.20. *In the situation of Corollary 5.19 let E' be the set of all edges in Γ^+ connecting two neighboring levels. If π is realizable, then so is π/E' .*

EXAMPLE 5.21. We emphasize that in the situation of Corollary 5.20 only complete level passages may be smoothed. We give an example for this in terms of enhanced level graphs that can easily be imagined as the top most levels of enhanced level graphs associated to a tropical normalized cover by Lemma 4.10. Consider the cover of graphs $\widehat{G}_1^+ \rightarrow G_1^+$ on the left of Figure 6. We assume that the o -value at the top end of all edges is 0. Both vertices on bottom level of G_1^+ are inconvenient. In \widehat{G}_1^+ , there is precisely one effective cycle (up to the τ -action) and this cycle is admissible for both inconvenient




FIGURE 6. A realizable graph and a non-realizable undegeneration

vertices. Therefore, this cover is realizable by Theorem 5.11. After smoothing only some of the edges between top and bottom level, we obtain the cover $\widehat{G}_2^+ \rightarrow G_2^+$ on the right. This is no longer realizable: There is no effective cycle (and hence no admissible cycle or independent pair of cycles) for the inconvenient vertex on bottom level. And besides that, the vertex of genus 2 is illegal.

This example also highlights another aspect. As we have seen in Example 5.8 there are inconvenient vertices of type II for which there is no admissible cycle, but a pair of independent cycles. The inconvenient vertex of genus zero on the left of G_1^+ is inconvenient of type II, but in this case there exists only an admissible cycle and no independent pair of cycles. Thus in fact both situations may occur.

6. Examples

6.1. kK_Γ is always realizable. Consider a stable tropical curve Γ with divisor kK_Γ for some $k \geq 2$. We show that the pair (Γ, kK_Γ) is realizable³. Note that stability implies that every vertex is in the support of the divisor. Hence the construction from Lemma 4.1 produces for every vertex at least one incident leg with enhancement value 1. Thus when constructing a tropical normalized cover $\widehat{\Gamma} \rightarrow \Gamma$, every vertex of Γ has necessarily only a single preimage. Furthermore, all edges in Γ^+ are horizontal. Hence the only possible choice for $\widehat{\Gamma}$ is to replace each edge of Γ with k parallel edges. But now the conditions of Theorem 5.11 are easy to check: the necessary cycles are provided by the parallel edges.

6.2. Dumbbell graph. Let $k = 2$ and consider the *dumbbell graph* Γ consisting of two vertices of genus 0, connected with a bridge edge and having a self-loop at each vertex, see Figure 7. We claim that Figure 9 shows all maximal cones in the realizability locus over the dumbbell graph. Note that each of them is of dimension 5 as was expected by Theorem 1.5.



FIGURE 7. Dumbbell graph.

To verify our claim let us start with the divisor $2K_\Gamma$. We focus on one of the trivalent vertices v : notice that the divisor produces two legs at v . How can we move these to arrive at a combinatorial type with more degrees of freedom? The first option is to move both legs onto the incident self-loop. Necessarily they will be symmetric on the loop. Performing this move on both sides of the graph produces the configuration from Figure 9a. This is realizable by the same argument as in the previous Section 6.1: we may simply cover every

³The authors are very grateful to D. Maclagan who raised this question during a discussion at the 2021 edition of the conference “Effective Methods in Algebraic Geometry”.

vertex by a unique preimage, producing cycles above every horizontal edge. In this case there are no inconvenient vertices.

The second option would be to move a single leg onto the bridge and leave the other one at v . Assume the resulting situation was realizable. Contract the self-loop at v and obtain a vertex of genus 1 and type $(-1, 1)$. This is an illegal vertex and by Corollary 5.19 we obtain a contradiction. Hence, we cannot leave a single leg behind.

So the third and final option is to move both legs onto the bridge. This will leave a part of the graph that is depicted in Figure 8 behind. Note that any cover of this part of the graph must be disconnected. This can be seen directly with Corollary 5.19 again. Alternatively, observe that the vertex has to have two preimages (otherwise it would be illegal). Connecting these two copies with two parallel horizontal edges fails to provide an effective cycle above the (horizontal) self-loop on the base (see Example 5.6). Hence, each of the two instances of the vertex must have a self-loop attached.




FIGURE 8. Part of an enhanced dumbbell graph. This is only realizable when provided with a disconnected cover.

With this observation in mind, we may choose to keep the pair of legs that we just moved onto the bridge together or separate them. In the former case, we arrive at situations from Figure 9b or 9c. In each case there is a unique tropical normalized cover $\widehat{\Gamma}^+$ and it can be checked to satisfy the conditions from Theorem 5.11. Hence, both of these configurations are realizable. The other option does however produce an inconvenient vertex of type $(1, -1; -4)$. This vertex does not have any simple closed cycle above it, which violates the conditions of Theorem 5.11.

Finally, we may push three or four of the legs onto the same self-loop. The first option does not produce a realizable configuration: Similar to the cases discussed above, the leg left behind always produces an inconvenient vertex that is not redeemed by an appropriate cycle. So let us consider the case that all four legs have been pushed to the same self-loop. This gives a realizable configuration if and only if the vertices are pairwise at the same level, as depicted in Figure 9d. We check this by providing a realizable cover $\widehat{\Gamma}^+$ in Figure 10. In fact, taking into consideration what we discussed about Figure 8 this is the only possible normalized cover without illegal vertices. Notice that there are again two inconvenient vertices of type $(-1, 1; -4)$ on the base. But the cover does admit a simple closed cycle above them which can be checked to be effective. We emphasize that this cycle passes through both vertices of type $(-1, 1; -4)$. Thus if those vertices were not on the same level, then there would not exist a simple cycle “at or above level” for the higher of the two. By Theorem 5.11 this would contradict realizability.

7. Appendix: Nonemptiness of boundary strata

While the boundary strata of the moduli space $\mathbb{P}\Xi^k\overline{\mathcal{M}}_{g,n}(\mu)$ of multi-scale k -differentials of type μ are indexed by normalized covers of enhanced level graphs $\pi : \widehat{G}^+ \rightarrow G^+$, not every such cover in fact corresponds to a nonempty boundary stratum D_π . Theorem 5.11 implicitly solves the problem to determine if a boundary stratum D_π is in fact nonempty for the moduli space $\mathbb{P}\Xi^k\overline{\mathcal{M}}_{g,n}(1, \dots, 1)$, that is: Given a cover of enhanced level graphs $\pi : \widehat{G}^+ \rightarrow G^+$ where all legs of G^+ have o -value 1, is there a corresponding normalized cover of twisted k -differentials $\widehat{X} \rightarrow X$ in $\mathbb{P}\Xi^k\overline{\mathcal{M}}_{g,n}(1, \dots, 1)$? In fact our methods can be applied to all strata of k -differentials. We will make this explicit in the slightly more general setting of so-called *generalized strata*.

To motivate the definition of generalized strata, fix a type μ and consider a boundary stratum D_π of $\mathbb{P}\Xi^k\overline{\mathcal{M}}_{g,n}(\mu)$ given by a cover of enhanced level graphs $\pi : \widehat{G}^+ \rightarrow G^+$. Now consider the space B given by the family corresponding to the projection of D_π to




FIGURE 9. Realizability locus over the dumbbell graph. Vertices of Γ^+ connected by dashed lines lie on the same level.




FIGURE 10. A realizable cover for the configuration in Figure 9d.

some level L of G^+ . In general B fails to be a honest stratum for two reasons: The level $G_{=L}^+$ may have several connected components, and the k -residues at the poles connecting $G_{=L}^+$ to higher levels may be restricted by the GRC. In other words, B is a subspace of a product of strata. The definition of a generalized stratum models this kind of spaces. Our definition is a generalization to the setting of k -differentials of the definition given in [CMZ22] for the abelian case.

Let $\Omega_{\mathbf{d}}^k \mathcal{M}_{g,n}(\mu)$ denote the connected components of $\Omega^k \mathcal{M}_{g,n}(\mu)$ which parametrize d -th powers of primitive k/d -differentials. Let

$$\Omega_{\mathbf{d}}^k \mathcal{M}_{\mathbf{g},\mathbf{n}}(\boldsymbol{\mu}) = \prod_{i=1}^{\kappa} \Omega_{d_i}^k \mathcal{M}_{g_i, n_i}(\mu_i)$$

be a disconnected stratum and let

$$\Omega \mathcal{M}_{\hat{\mathbf{g}}, \hat{\mathbf{n}}}(\hat{\boldsymbol{\mu}}) = \prod_{j=1}^{\hat{\kappa}} \Omega \mathcal{M}_{\hat{g}_j, \hat{n}_j}(\hat{\mu}_j)$$

be the product of the strata that contain the canonical covers. Here the bold letters on the left denote the tuples of the corresponding letters on the right, i.e. $\mathbf{d} = (d_1, \dots, d_{\kappa})$. Let $\hat{\boldsymbol{\mu}}_j = (\hat{m}_{j,1}, \dots, \hat{m}_{j,\hat{n}_j})$ and let $H_p := \{(j, l) \mid \hat{m}_{j,l} < -1\}$ be the set of marked non-simple poles in the cover. Let λ be a partition of H_p with parts denoted by $\lambda^{(a)}$ and let $\lambda_{\mathfrak{A}}$ be a

subset of the parts of λ such that $\lambda_{\mathfrak{R}}$ is τ -invariant as a set. Let

$$\mathfrak{R} := \left\{ r = (r_{j,l})_{(j,l) \in H_p} \in \mathbb{C}^{|H_p|} \left| \sum_{(j,l) \in \lambda^{(a)}} r_{j,l} = 0 \text{ for all } \lambda^{(a)} \in \lambda_{\mathfrak{R}} \right. \right\}.$$

Denote by $\Omega\mathcal{M}_{\mathfrak{g},\mathfrak{n}}^{\mathfrak{R}}(\widehat{\boldsymbol{\mu}})$ the subspace with residues in \mathfrak{R} and denote by $\Omega_{\mathbf{d}}^k\mathcal{M}_{\mathfrak{g},\mathfrak{n}}^{\mathfrak{R}}(\boldsymbol{\mu})$ the corresponding subspace of k -differentials. Note that this is well-defined, as we have chosen $\lambda_{\mathfrak{R}}$ to be τ -invariant.

DEFINITION 7.1. We call a stratum of the form $\Omega_{\mathbf{d}}^k\mathcal{M}_{\mathfrak{g},\mathfrak{n}}^{\mathfrak{R}}(\boldsymbol{\mu})$ a *generalized stratum* and denote by $\mathbb{P}\Xi_{\mathbf{d}}^k\overline{\mathcal{M}}_{\mathfrak{g},\mathfrak{n}}^{\mathfrak{R}}(\boldsymbol{\mu})$ the corresponding projectivized generalized stratum of multi-scale k -differentials.

Let $\widehat{G}^+ \rightarrow G^+$ be a cover of enhanced level graphs. Here and in the following we allow for G^+ to be disconnected. Picking up an idea from [CMZ23], we construct the cover of *auxiliary enhanced level graphs* $\widehat{G}_{\infty}^+ \rightarrow G_{\infty}^+$ in the following way. For each $\lambda^{(a)} \in \lambda_{\mathfrak{R}}$, we add a vertex $\widehat{v}_{(a)}$ to \widehat{G}^+ and think of these new vertices as being at level ∞ , i.e. all new vertices are on a new level above all other levels. As the legs of \widehat{G}^+ correspond to the marked points of the stratum with orders $\widehat{m}_{j,l}$, we may think of the legs of \widehat{G}^+ as being indexed by the tuples (j,l) . In particular the leg (j,l) has o -value $\widehat{m}_{j,l}$. For each $(j,l) \in \lambda^{(a)}$, we add an edge to \widehat{G}^+ connecting the leg (j,l) to the new vertex $\widehat{v}_{(a)}$ and we take the o -value at the upper half-edge of the new edge to be $-\widehat{m}_{j,l} - 2k$. If the sum of the o -values of the legs incident to $\widehat{v}_{(a)}$ is odd we add an additional leg with o -value 1 to $\widehat{v}_{(a)}$. Then the genus of $\widehat{v}_{(a)}$ is determined by the o -values of the incident legs. Finally, the action of τ on the new edges, vertices and legs is determined by the action of τ on the legs (j,l) . We call this new graph \widehat{G}_{∞}^+ and we add edges, vertices and legs to G^+ to complete G^+ to the quotient $G_{\infty}^+ := \widehat{G}_{\infty}^+/\tau$. We emphasize that the new vertices added to G_{∞}^+ are never inconvenient, as they do not contain any poles.

We call the marked poles of $\Omega\mathcal{M}_{\mathfrak{g},\mathfrak{n}}^{\mathfrak{R}}(\widehat{\boldsymbol{\mu}})$ that are not contained in $\lambda_{\mathfrak{R}}$ the *free poles*. When we assign residues to the graph \widehat{G}^+ as in the proof of our main theorem, we may alter the residues at free poles at will (while maintaining the residue theorem at each component), as the GRC does not restrict the residues at components containing a free pole in any way, see Definition 3.12. We reflect this in the following definition.

DEFINITION 7.2. A *free pole path* is a simple path in \widehat{G}^+ starting and ending in a (different) free pole. A *generalized cycle* is a free pole path or a simple cycle.

The definitions of an effective (resp. admissible) cycle and of an independent pair of cycles can be adapted for generalized cycles in the obvious way. By applying the methods of our proof to the enveloping stratum of the cover of auxiliary enhanced level graphs $\widehat{G}_{\infty}^+ \rightarrow G_{\infty}^+$, it is not hard to check that the proof of Theorem 1.4 in fact proves

THEOREM 7.3. A *normalized cover of enhanced level graphs* $\pi: \widehat{G}^+ \rightarrow G^+$ corresponds to a *nonempty boundary stratum of a generalized stratum* $\mathbb{P}\Xi_{\mathbf{d}}^k\overline{\mathcal{M}}_{\mathfrak{g},\mathfrak{n}}^{\mathfrak{R}}(\boldsymbol{\mu})$ if and only if the following conditions hold.

- (i) *There is no illegal vertex in π .*
- (ii) *For every horizontal edge \widehat{e} in \widehat{G}^+ there is an effective generalized cycle in \widehat{G}_{∞}^+ through \widehat{e} .*
- (iii) *For every inconvenient vertex v in G^+ there is an admissible generalized cycle in \widehat{G}_{∞}^+ through one of the preimages \widehat{v} or there is an independent pair of generalized cycles.*

REMARK 7.4. For $k = 1$, Theorem 7.3 recovers [CMZ23, Proposition 3.2] in the same way as Theorem 1.4 recovered [MUW21, Theorem 6.3], see Remark 5.12.

Chern classes of linear submanifolds with applications to spaces of k -differentials and ball quotients

1. Introduction

Linear submanifolds are the most interesting and well-studied subvarieties of moduli spaces of abelian differentials $\Omega\mathcal{M}_{g,n}(\mu)$ and their classification seems far from complete at present. They are defined as the normalization of algebraic substacks of $\Omega\mathcal{M}_{g,n}(\mu)$ that are locally a union of linear subspaces in period coordinates. In the holomorphic case, linear submanifolds defined by real linear equations are precisely the closures of $\mathrm{GL}_2^+(\mathbb{R})$ -orbits by the fundamental theorems of Eskin-Mirzakhani-Mohammadi ([EM18], [EMM15]). These orbit closures are automatically algebraic subvarieties by Filip's theorem ([Fil16]). Our results require algebraicity, but they work as well for meromorphic differentials and for subvarieties whose equations are only \mathbb{C} -linear.

Linear submanifolds include

- spaces of quadratic differentials,
- Teichmüller curves,
- eigenform loci and Prym loci,
- the recent sporadic examples from [MMW17] and [EMMW20], but also
- spaces defined by covering constructions, and
- in the meromorphic case, spaces defined by residue conditions.

These examples are \mathbb{R} -linear. Spaces of k -differentials for $k \geq 2$ and in particular the ball quotients in Section 8 are prominent examples that are only \mathbb{C} -linear.

Our primary goal is a formula for the Chern classes of the cotangent bundle of any linear submanifold or rather of its compactification. The Euler characteristic is an intrinsic compactification-independent application. Knowing the Chern classes is a prerequisite for understanding the birational geometry of linear submanifolds, such as computations of the Kodaira dimension, see [CCM22].

This goal was achieved in [CMZ22] for the full projectivized strata of Abelian differentials $\mathbb{P}\Omega\mathcal{M}_{g,n}(\mu)$ themselves, taking the modular smooth normal crossing compactification $\mathbb{P}\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ of multi-scale differentials from [BCGGM19b] as point of departure. In the inextricable zoo of linear manifolds we are not aware of any intrinsic way to construct a smooth compactification with modular properties. Working with the normalization of the closure in some ambient compactification is usually unsuitable for intersection theory computations. Here, however, thanks to the work of Benirschke-Dozier-Grushevsky ([BDG22]) and some minor upgrades we are able to work with this closure.

We now introduce more notation to state the general results and then apply them to specific linear submanifolds. Let $\Omega\mathcal{H} \rightarrow \Omega\mathcal{M}_{g,n}(\mu)$ be a linear submanifold. Let moreover $\mathcal{H} \rightarrow \mathbb{P}\Omega\mathcal{M}_{g,n}(\mu)$ be its projectivization and let $\overline{\mathcal{H}} \rightarrow \mathbb{P}\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ denote the normalization of its closure into the space of multi-scale differentials. The boundary strata D_Γ of $\mathbb{P}\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ are indexed by level graphs Γ as we recall in Section 3.2. By [BDG22, Theorem 1.5] the boundary of $\overline{\mathcal{H}}$ is divisorial and consists two types of divisors: First there are the divisors $D_h^{\mathcal{H}}$ of curves whose level graphs have only horizontal edges (i.e. joining vertices of the same level). Second there are the divisors $D_\Gamma^{\mathcal{H}}$ parameterized by level graphs $\Gamma \in \mathrm{LG}_1(\mathcal{H})$ that have one level below the zero level and no horizontal edges and such that the intersection of $\overline{\mathcal{H}}$ with the interior of the boundary divisor D_Γ is non-empty. Those boundary divisors $D_\Gamma^{\mathcal{H}}$ come with the integer ℓ_Γ , the least common multiple of the prongs κ_e along the edges. We let $\xi = c_1(\mathcal{O}(-1))$ be the first Chern class of the tautological bundle on $\overline{\mathcal{H}}$.

THEOREM 1.1. *The first Chern class of the logarithmic cotangent bundle of a projectivized compactified linear submanifold $\overline{\mathcal{H}}$ is*

$$(18) \quad c_1(\Omega_{\overline{\mathcal{H}}}^1(\log \partial\mathcal{H})) = N \cdot \xi + \sum_{\Gamma \in \text{LG}_1(\mathcal{H})} (N - N_\Gamma^\top) \ell_\Gamma[D_\Gamma^{\mathcal{H}}] \in \text{CH}^1(\overline{\mathcal{H}}),$$

where $N := \dim(\Omega\mathcal{H})$ and where $N_\Gamma^\top := \dim(D_\Gamma^{\mathcal{H},\top}) + 1$ is the dimension of the unprojectivized top level stratum in $D_\Gamma^{\mathcal{H}}$.

To state a formula for the full Chern character we need to recall a procedure that also determines adjacency of boundary strata. It is given by undegeneration maps δ_i that contract all the edges except those that cross from level $-i + 1$ to level $-i$, see Section 3.2. This construction can obviously be generalized so that a larger subset of levels remains. For example the undegeneration map $\delta_i^{\mathbb{C}}$ contracts only the edges crossing from level $-i + 1$ to level $-i$. We can now define for any graph $\Gamma \in \text{LG}_L(\mathcal{H})$ with L levels below zero and without horizontal edges the boundary component $D_\Gamma^{\mathcal{H}}$ of codimension L and the quantity $\ell_\Gamma = \prod_{i=1}^L \ell_{\delta_i(\Gamma)}$.

THEOREM 1.2. *The Chern character of the logarithmic cotangent bundle is*

$$\text{ch}(\Omega_{\overline{\mathcal{H}}}^1(\log \partial\mathcal{H})) = e^\xi \cdot \sum_{L=0}^{N-1} \sum_{\Gamma \in \text{LG}_L(\mathcal{H})} \ell_\Gamma \left(N - N_{\delta_L(\Gamma)}^\top \right) i_{\Gamma*} \left(\prod_{i=1}^L \text{td} \left(\mathcal{N}_{\Gamma/\delta_i^{\mathbb{C}}(\Gamma)}^{\otimes -\ell_{\delta_i(\Gamma)}} \right)^{-1} \right),$$

where $\mathcal{N}_{\Gamma/\delta_i^{\mathbb{C}}(\Gamma)}$ denotes the normal bundle of $D_\Gamma^{\mathcal{H}}$ in $D_{\delta_i^{\mathbb{C}}(\Gamma)}^{\mathcal{H}}$, where td is the Todd class and $i_\Gamma : D_\Gamma^{\mathcal{H}} \hookrightarrow \overline{\mathcal{H}}$ is the inclusion map.

So far the results have been stated to parallel exactly those in [CMZ22]. We start explaining the difference in evaluating this along with the next result, a closed formula for the Euler characteristic.

THEOREM 1.3. *Let $\mathcal{H} \rightarrow \mathbb{P}\Omega\mathcal{M}_{g,n}(\mu)$ be a projectivized linear submanifold. The orbifold Euler characteristic of \mathcal{H} is given by*

$$\chi(\mathcal{H}) = (-1)^d \sum_{L=0}^d \sum_{\Gamma \in \text{LG}_L(\mathcal{H})} \frac{K_\Gamma^{\mathcal{H}} \cdot N_\Gamma^\top}{|\text{Aut}_{\mathcal{H}}(\Gamma)|} \cdot \prod_{i=0}^{-L} \int_{\mathcal{H}_\Gamma^{[i]}} \xi_{\mathcal{H}_\Gamma^{[i]}}^{d_\Gamma^{[i]}},$$

where the integrals are over the normalization of the closure $\overline{\mathcal{H}} \rightarrow \mathbb{P}\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ inside the moduli space of multi-scale differentials and similar integrals over boundary strata, where

- $\mathcal{H}_\Gamma^{[i]}$ are the linear submanifolds at level i of Γ as defined in Section 3.5,
- $d_\Gamma^{[i]} := \dim(\mathcal{H}_\Gamma^{[i]})$ is the projectivized dimension,
- $K_\Gamma^{\mathcal{H}}$ is the product of the number of prong-matchings on each edge of Γ that are actually contained in the linear submanifold $\overline{\mathcal{H}}$,
- $\text{Aut}_{\mathcal{H}}(\Gamma)$ is the set of automorphism of the graph Γ whose induced action on a neighborhood of $D_\Gamma^{\mathcal{H}}$ preserves $\overline{\mathcal{H}}$,
- $d := \dim(\mathcal{H})$ is the projectivized dimension.

The number of *reachable prong matchings* $K_\Gamma^{\mathcal{H}}$ and the number $|\text{Aut}_{\mathcal{H}}(\Gamma)|$ as defined in the theorem are in general non-trivial to determine. Also the description of $\mathcal{H}_\Gamma^{[i]}$ requires specific investigation. For example, for strata of k -differentials, these $\mathcal{H}_\Gamma^{[i]}$ are again some strata of k -differentials, but the markings of the edges have to be counted correctly.

The most important obstacle to evaluate this formula however is to compute the fundamental classes of linear submanifolds, or to use tricks to avoid this. For strata of abelian differentials, this step was provided by the recent advances in relating fundamental classes to Pixton's formula ([HS21], [BHPSS20]). Whenever we have the fundamental classes at our disposal, we can evaluate expressions in the tautological ring, as we briefly summarize in Section 4.

Applications: Teichmüller curves in genus two. As an example where fundamental class considerations can be avoided, we give an alternative quick proof of one of the first computations of Euler characteristics of Teichmüller curves, initially proven in [Bai07], see also [MZ16] for a proof via theta derivatives. We assume familiarity with the notation for linear submanifolds in genus two strata, as recalled in Section 6.

THEOREM 1.4 (Bainbridge). *The Euler characteristic of the Teichmüller curve W_D in the eigenform locus for real multiplication by a non-square discriminant D is $\chi(W_D) = -9\zeta(-1)$ where $\zeta = \zeta_{\mathbb{Q}(\sqrt{D})}$ is the Dedekind zeta function.*

PROOF. The Hilbert modular surface X_D is the disjoint union of the symmetrization of the eigenform locus $E_D \subset \Omega\mathcal{M}_{2,1}(1,1)$, the product locus P_D of reducible Jacobians and the Teichmüller curve W_D . This gives

$$\chi(P_D) + \chi(W_D) + \frac{1}{2}\chi(E_D) = \chi(X_D).$$

Now we apply Theorem 1.3 to E_D . The top- ξ -integral in the $L = 0$ -term of vanishes by Corollary 4.3, since E_D is a linear submanifold with REL non-zero. The codimension-one boundary strata are P_D and W_D . They don't intersect, so there are no codimension-two boundary strata without horizontal nodes and we get

$$(19) \quad \chi(E_D) = -\chi(P_D) - 3\chi(W_D)$$

where the factor 3 stems from the number of prong-matchings. Since Siegel computed $\chi(X_D) = 2\zeta(-1)$ and viewing P_D as the vanishing locus of the product of odd theta functions gives $\chi(P_D) = -5\zeta(-1)$, the theorem follows from the two equations. \square

Strata of k -differentials. The space of quadratic differentials is the cotangent space to moduli space of curves and thus fundamental in Teichmüller dynamics. We give formulas for Chern classes, Euler characteristics and for the intersection theory in these spaces. In fact, our formulas work uniformly for spaces of k -differentials for all $k \geq 1$. Having the quadratic case in mind, we write $\overline{\mathcal{Q}} = \mathbb{P}\Xi^k \overline{\mathcal{M}}_{g,n}(\mu)$ for the space of multi-scale k -differentials defined in [CMZ19], which coincides (up to explicit isotropy groups, see Lemma 7.2) with the compactification as above of the associated linear submanifold obtained via the canonical covering construction.

The formulas in Theorem 1.2 apply to $\overline{\mathcal{Q}}$ viewed as a linear submanifold in some higher genus stratum $\mathcal{M}_{\widehat{g},\widehat{n}}(\widehat{\mu})$. However the fundamental class of these submanifolds is not known, conceivably it is not even a tautological class. The main challenge here is to convert these formulas into formulas that can be evaluated on $\overline{\mathcal{Q}}$ viewed as a submanifold in $\overline{\mathcal{M}}_{g,n}$ where the fundamental class is given by Pixton's formula.

While the boundary strata of the moduli space $\mathbb{P}\Xi \overline{\mathcal{M}}_{g,n}(\mu)$ are indexed by level graphs, the boundary strata of the moduli space of multi-scale k -differentials $\overline{\mathcal{Q}}$ are indexed by coverings of k -level graphs $\pi : \widehat{\Gamma}_{\text{mp}} \rightarrow \Gamma$, where the legs of $\widehat{\Gamma}_{\text{mp}}$ are marked only partially, see Section 7 or also [CMZ19, Section 2] for the definitions of these objects and the labeling conventions of those covers. Each edge $e \in \Gamma$ has an associated k -enhancement κ_e given by $|\text{ord}_e \omega + k|$, where ω is the k -differential on a generic point of the associated boundary stratum D_π . We let $\zeta = c_1(\mathcal{O}(-1))$ be the first Chern class of the tautological bundle on $\overline{\mathcal{Q}}$. Via the canonical cover construction, Theorem 1.3 implies the following formula for the Euler characteristic of strata of k -differentials.

COROLLARY 1.5. *The orbifold Euler characteristic of a projectivized stratum of k -differentials $\mathbb{P}\Omega^k \mathcal{M}_{g,n}(\mu)$ is given by*

$$\chi(\mathbb{P}\Omega^k \mathcal{M}_{g,n}(\mu)) = \left(\frac{-1}{k}\right)^d \sum_{L=0}^d \sum_{(\pi: \widehat{\Gamma}_{\text{mp}} \rightarrow \Gamma) \in \text{LG}_L(\mathcal{Q})} S(\pi) \cdot \frac{N_\pi^\top \cdot \prod_{e \in E(\Gamma)} \kappa_e}{|\text{Aut}(\Gamma)|} \cdot \prod_{i=0}^{-L} \int_{\mathcal{Q}_\pi^{[i]}} \zeta_{\mathcal{Q}_\pi^{[i]}}^{d_\pi^{[i]}}$$

where $S(\pi)$ is the normalized size of a stabilizer of a totally labeled version of the graph $\widehat{\Gamma}_{\text{mp}}$ and $\mathcal{Q}_\pi^{[i]}$ are the strata of k -differentials of D_π at level i .

The full definition of $S(\pi)$ is presented in (65). It equals one for many π , e.g. if all vertices in Γ have only one preimage in $\widehat{\Gamma}_{\text{mp}}$. See Remark 7.6 for values of this combinatorial constant.

k	1	3	4	5	6	7	8	9
$\chi(\mathbb{P}\Omega^k \mathcal{M}_{2,1}(2k))$	$-\frac{1}{40}$	$\frac{1}{3}$	$\frac{3}{2}$	$\frac{21}{5}$	9	18	30	51

TABLE 1. Euler characteristics of some minimal strata of k -differentials

Table 1 gives the Euler characteristics of some strata of quadratic differentials, for more examples and cross-checks see Section 7.5.

All the formulas for evaluations in the tautological ring of strata of k -differentials have been coded in an extension of the sage program `diffstrata` (an extension of `admcycles` by [DSZ21]) that initially had this functionality for abelian differentials only (see [CMZ22], [CMZ23]). See Section 4 for generalities on tautological ring computations and in particular Section 7 for the application to k -differentials. The program `diffstrata` has been used to verify the Hodge-DR-conjecture from [CGHMS22] in low genus. Moreover, `diffstrata` confirms that the values of the tables in [Gou16] can be obtained via intersection theory computations:

PROPOSITION 1.6. *The Conjecture 1.1 in [CMS19] expressing Masur-Veech volumes for strata of quadratic differentials as intersection numbers holds true for strata of projectivized dimension up to six, e.g. $\mathcal{Q}(12) = 5614/6075 \cdot \pi^6$.*

Ball quotients. Deligne-Mostow ([DM86]) and Thurston [Thu98] constructed compactifications of strata of k -differentials on $\mathcal{M}_{0,n}$ for very specific choices of μ and showed that these compactified strata are quotients of the complex $(n-3)$ -ball. These results were celebrated as they give a list of non-arithmetic ball quotients, of which there today are still only finitely many sporadic examples, see [DPP16] and [Der20] for recent progress. The compactifications are given as GIT quotients (in ([DM86]) or in the language of cone manifolds (in [Thu98]) and the proof of the discreteness of the monodromy representation requires delicate arguments for extension of the period at the boundary, resp. surgeries for the cone manifold completion.

As application of our Chern class formulas we give a purely algebraic proof that these compactifications are ball quotients, based on the fact that the equality case in the Bogomolov-Miyaoka-Yau inequality implies a ball quotient structure, see Proposition 8.1. Since this is a proof of concept, we restrict to the case $n=5$, i.e. to quotients of the complex two-ball, and to the condition INT in (20), leaving the analog for Mostow's generalized Σ INT-condition [Mos86] for the reader.

The computation of the hyperbolic volume of these ball quotients had been open for a long time. A solution has been given by McMullen [McM17] and Koziarz-Nguyen [KN18], see also [KM16]. Since computing the hyperbolic volume is equivalent to computing the Euler characteristic by Gauss-Bonnet, our results provide alternative approach to this question, too.

There are only four kinds of boundary divisors of $\overline{\mathcal{Q}}$:

- The divisors Γ_{ij} where two points with $a_i + a_j < k$ collide.
- The divisors L_{ij} where two points with $a_i + a_j > k$ collide.
- The 'horizontal' boundary divisor D_{hor} consisting of all components where two points with $a_i + a_j = k$ collide.
- The 'cherry' boundary divisors ${}_{ij}\Lambda_{kl}$.

THEOREM 1.7. *Suppose that $\mu = (-a_1, \dots, -a_5)$ is a tuple with $a_i \geq 0$ and with the condition*

$$(20) \quad \left(1 - \frac{a_i}{k} - \frac{a_j}{k}\right)^{-1} \in \mathbb{Z} \quad \text{if } a_i + a_j < k \quad (\text{INT})$$

for all $i \neq j$. Then there exists a birational contraction morphism $\overline{\mathcal{Q}} \rightarrow \overline{\mathfrak{B}}$ onto a smooth proper DM-stack $\overline{\mathfrak{B}}$ that contracts precisely all the divisors L_{ij} and ${}_{ij}\Lambda_{kl}$. The target $\overline{\mathfrak{B}}$ satisfies the Bogomolov-Miyaoka-Yau equality for $\Omega_{\overline{\mathfrak{B}}}^1(\log D_{\text{hor}})$.

As a consequence $\mathfrak{B} = \overline{\mathfrak{B}} \setminus D_{\text{hor}}$ is a ball quotient.

The signature of the intersection form on the eigenspace that k -differentials are modeled on has been computed by Veech [Vee93]. The only other case where the signature is $(1, 2)$ are strata in $\mathcal{M}_{1,3}$. As observed by Ghazouani-Pirio in [GP17], (see also [GP20]) there are only few cases where the metric completion of the strata can be a ball quotient. However they also find additional cases where the monodromy of the stratum is discrete. This implies that the period map descends to a map from the compactified stratum to a ball quotient. It would be interesting to investigate if there are more such cases, possibly with non-arithmetic monodromy.

2. Logarithmic differential forms and toric varieties

This section connects the Euler characteristic to integrals of characteristic classes of the sheaf of logarithmic differential forms. We work on a possibly singular but normal and irreducible variety $\overline{\mathcal{H}}$ of dimension d , whose singularities are toric and contained in some boundary divisor $\partial\mathcal{H}$. We are interested in the Euler characteristic of a (Zariski) open subvariety \mathcal{H} with divisorial complement, such that the inclusion $\mathcal{H} \hookrightarrow \overline{\mathcal{H}}$ is a toroidal embedding. In particular the boundary divisor $\partial\mathcal{H} = \overline{\mathcal{H}} \setminus \mathcal{H}$ is locally on open subsets U_α a torus-invariant divisor.

In this situation we define locally $\Omega_{U_\alpha}^1(\log)$ to be the sheaf of $(\mathbb{C}^*)^d$ -invariant meromorphic differential forms. These glue to sheaf $\Omega_{\overline{\mathcal{H}}}^1(\log \partial\mathcal{H})$, that is called *logarithmic differential sheaf*. This terminology is justified by the following idea from [Mum77, Section 4], the details and definitions being given in [KKMS73]. For any 'allowable' smooth modification $p: \overline{W} \rightarrow \overline{\mathcal{H}}$ that maps a normal crossing boundary divisor $\partial W \subset \overline{W}$ onto $\partial\mathcal{H}$ we have $p^*\Omega_{\overline{\mathcal{H}}}^1(\log \partial\mathcal{H}) = \Omega_{\overline{W}}^1(\log \partial W)$ for the usual definition of the logarithmic sheaf on \overline{W} . Moreover, such an 'allowable' smooth modification always exists.

PROPOSITION 2.1. *For $\mathcal{H} \hookrightarrow \overline{\mathcal{H}}$ as above the Euler characteristic of \mathcal{H} can be computed as integral*

$$(21) \quad \chi(\mathcal{H}) = (-1)^d \int_{\overline{\mathcal{H}}} c_d(\Omega_{\overline{\mathcal{H}}}^1(\log \partial\mathcal{H}))$$

over the top Chern class of the logarithmic cotangent bundle.

PROOF. If $\overline{\mathcal{H}}$ is smooth, this is well known, a self-contained proof was given in [CMZ22, Proposition 2.1]. In general we use an allowable modification. By definition this restricts to an isomorphism $W \rightarrow \mathcal{H}$, hence does not change the left hand side. The right hand side also stays the same by push-pull and the pullback formula along an allowable smooth modification. \square

In all our applications, $\overline{\mathcal{H}}$ will be a proper Deligne-Mumford stack with toroidal singularities. We work throughout with orbifold Euler characteristics, and since then both sides of (21) are multiplicative in the degree of a covering, we can apply Proposition 2.1 verbatim.

3. The closure of linear submanifolds

The compactification of a linear submanifold we work with has (currently) no intrinsic definition. Rather we consider the normalization of the closure of a linear submanifold inside the moduli space of multi-scale differentials $\Xi\overline{\mathcal{M}}_{g,n}(\mu)$. We recall from [BDG22] the basic properties of such closures. The goal of this section is to make precise and to explain the following two slogans:

- Near boundary points without horizontal edges, the closure is determined as for the ambient abelian stratum by the combinatorics of the level graph and it is smooth. The *ghost automorphisms*, the stack structure at the boundary that stems from twist groups, agrees with the ghost automorphisms of the ambient

stratum and the intersection pattern is essentially determined by the *profiles* of the level graph, a subset of the profiles of the ambient stratum.

- In the presence of horizontal edges there are toric singularities. Working with the appropriate definition of the logarithmic cotangent sheaf these singularities don't matter. This sheaf decomposes into summands from horizontal nodes, from the level structure, and the deformation of the differentials at the various levels, just as in the ambient stratum.

3.1. Linear submanifolds in generalized strata. Let $\Omega\mathcal{M}_{g,n}(\mu)$ denote the moduli space of Abelian differential of possibly meromorphic signature μ . Despite calling them 'moduli space' or 'strata' we always think of them as quotient stacks or orbifolds and intersection numbers etc. are always understood in that sense. These strata come with a linear structure given by period coordinates (e.g. [Zor06] for an introduction). A *linear submanifold* $\Omega\mathcal{H}$ of $\Omega\mathcal{M}_{g,n}(\mu)$ is an algebraic stack with a map $\Omega\mathcal{H} \rightarrow \Omega\mathcal{M}_{g,n}(\mu)$ which is the normalization of its image and whose image is locally given as a finite union of linear subspaces in period coordinate charts. See [Fil20, Example 4.1.10] for an example that illustrates why we need to pass to the normalization for $\Omega\mathcal{H}$ to be a smooth stack. In the context of holomorphic signatures and $\mathrm{GL}_2(\mathbb{R})$ -orbit closures, the linear manifolds obtained in this way can locally be defined by equations with \mathbb{R} -coefficients ([EM18], [EMM15]). We refer to them as \mathbb{R} -*linear submanifolds*. In this context, the algebraicity follows from being closed by the result of Filip ([Fil16]), but in general algebraicity is an extra hypothesis.

To set up for clutching morphisms and a recursive description of the boundary of compactified linear submanifolds, we now define *generalized strata*, compare [CMZ22, Section 4]. For a tuple $\mathbf{g} = (g_1, \dots, g_k)$ of genera and a tuple $\mathbf{n} = (n_1, \dots, n_k)$ together with a collection of types $\boldsymbol{\mu} = (\mu_1, \dots, \mu_k)$ with $|\mu_i| = n_i$ we first define the disconnected stratum $\Omega\mathcal{M}_{\mathbf{g},\mathbf{n}}(\boldsymbol{\mu}) = \prod_{i=1}^k \Omega\mathcal{M}_{g_i, n_i}(\mu_i)$. Then, for a linear subspace \mathfrak{R} inside the space of the residues at all poles of $\boldsymbol{\mu}$ we define the generalized stratum $\Omega\mathcal{M}_{\mathbf{g},\mathbf{n}}^{\mathfrak{R}}(\boldsymbol{\mu})$ to be the subvariety with residues lying in \mathfrak{R} . Generalized strata obviously come with period coordinates and we thus define a *generalized linear submanifold* $\Omega\mathcal{H}$ to be an algebraic stack together with a map to $\Omega\mathcal{M}_{\mathbf{g},\mathbf{n}}^{\mathfrak{R}}(\boldsymbol{\mu})$ whose image is locally linear in period coordinates and where $\Omega\mathcal{H}$ is the normalization of its image.

Rescaling the differential gives an action of \mathbb{C}^* on strata and the quotient are projectivized strata $\mathbb{P}\Omega\mathcal{M}_{g,n}(\mu)$. The image of a linear submanifold in $\mathbb{P}\Omega\mathcal{M}_{g,n}(\mu)$ is called *projectivized linear manifold* \mathcal{H} , but we usually omit the 'projectivized'.

We refer with an index B to quantities of the ambient projectivized stratum, such as its dimension d_B and the unprojectivized dimension $N_B = d_B + 1$. The same letters without additional index are used for the linear submanifold, e.g. $N = d + 1$, and we write $d_{\mathcal{H}}$ and $N_{\mathcal{H}}$ only if ambiguities may arise.

3.2. Multi-scale differentials: boundary combinatorics. We will work inside the moduli stack of multi-scale differentials, that is the compactification $\overline{B} := \mathbb{P}\overline{\mathcal{M}}_{g,n}(\mu)$ of a stratum $B := \mathbb{P}\Omega\mathcal{M}_{g,n}(\mu)$ defined in [BCGGM19b] and recall some of its properties, see also [CMZ22, Section 3]. Everything carries over with obvious modifications to the compactification $\mathbb{P}\overline{\mathcal{M}}_{\mathbf{g},\mathbf{n}}^{\mathfrak{R}}(\boldsymbol{\mu})$ of generalized strata, see [CMZ22, Proposition 4.1].

Each boundary stratum of $\mathbb{P}\overline{\mathcal{M}}_{g,n}(\mu)$ has its associated level graph Γ , a stable graph of the underlying pointed stable curve together with a weak total order on the vertices, usually given by a level function normalized to have top level zero, and an enhancement $\kappa_e \geq 0$ associated to the edges. Edges are called *horizontal*, if they start and end at the same level, and *vertical otherwise*. Moreover $\kappa_e = 0$ if and only if the edge is horizontal. We denote the closure of the boundary stratum of points with level graph Γ by D_{Γ}^B and denote in general the complement of more degenerate boundary strata by an extra \circ , i.e., here by $D_{\Gamma}^{B,\circ}$. These D_{Γ}^B are in general not connected, and might be empty (e.g. for unsuitably large κ_e).

We let $\mathrm{LG}_L(B)$ be the set of all enhanced $(L+1)$ -level graphs without horizontal edges. The structure of the normal crossing boundary of $\mathbb{P}\overline{\mathcal{M}}_{g,n}(\mu)$ is encoded by *undegenerations*. For any subset $I = \{i_1, \dots, i_n\} \subseteq \{1, \dots, L\}$ there are undegeneration map

$$\delta_{i_1, \dots, i_n} : \mathrm{LG}_L(B) \rightarrow \mathrm{LG}_n(B),$$

that preserves the level passage given as a horizontal line just above level $-i$ and contracts the remaining level passages. We define $\delta_I^c = \delta_{Ic}$.

The boundary strata D_Γ^B for $\Gamma \in \text{LG}_L(B)$ are commensurable to a product of generalized strata $B_\Gamma^{[i]} = \mathbb{P}\Xi\overline{\mathcal{M}}_{\mathbf{g}_i, \mathbf{n}_i}^{\mathfrak{R}_i}(\boldsymbol{\mu}_i)$ defined via the following diagram.

$$(22) \quad \begin{array}{ccccc} & c_\Gamma^{-1}(U_\Gamma) & \xrightarrow{\subset} & D_\Gamma^{B,s} & \\ & \searrow^{p_\Gamma} & & \searrow^{q_\Delta} & \\ & & & & U_\Delta^s \\ & & & & \searrow^{c_\Gamma} \\ & & & & D_\Gamma^B \\ & & & & \swarrow^{c_\Gamma^\Delta} \\ & & & & U_\Gamma \\ & & & & \xrightarrow{\subset} \\ \prod_{i=-L}^0 B_\Gamma^{[i]} =: B_\Gamma & \xleftarrow{\supset} & B_{\Gamma,\Delta} & & \end{array}$$

Here $\mathbf{g}_i, \mathbf{n}_i$ and $\boldsymbol{\mu}_i$ are the tuples of the genera, marked points and signatures of the components at level i of the level graph and \mathfrak{R}_i is the global residue condition induced by the levels above. The covering space $D_\Gamma^{B,s}$ and the moduli stack U_Δ^s of *simple multi-scale differentials compatible with an undegeneration of Δ* were constructed in [CMZ22, Section 4.2].

3.3. Multi-scale differentials: Prong-matchings and stack structure. The notion of a multi-scale differential is based on the following construction. Given a pointed stable curve (X, \mathbf{z}) , a *twisted differential* is a collection of differentials η_v on each component X_v of X , that is *compatible with a level structure* on the dual graph Γ of X , i.e. vanishes as prescribed by μ at the marked points z , satisfies the matching order condition at vertical nodes, the matching residue condition at horizontal nodes and global residue condition of [BCGM18]. A *multi-scale differential of type μ* on a stable curve (X, \mathbf{z}) consists of an enhanced level structure $(\Gamma, \ell, \{\kappa_e\})$ on the dual graph Γ of X , a twisted differential ω of type μ compatible with the enhanced level structure, and a prong-matching for each node of X joining components of non-equal level. Here a *prong-matching σ* is an identification of the (outgoing resp. incoming) real tangent vectors at a zero resp. a pole corresponding to each vertical edge of Γ . Multi-scale differentials are equivalence classes of $(X, \mathbf{z}, \Gamma, \sigma)$ up to the action of the level rotation torus that rescales differentials on lower levels and rotates prong-matchings at the same time.

To an enhanced two-level graph we associate the quantity

$$(23) \quad \ell_\Gamma = \text{lcm}(\kappa_e : e \in E(\Gamma)).$$

which appears in several important place of the construction of $\mathbb{P}\Xi\overline{\mathcal{M}}_{g,n}(\mu)$:

- i) It is the size of the orbit of prong-matchings when rotating the lower level differential. Closely related:
- ii) The local equations of a node are $xy = t_1^{\ell_\Gamma/\kappa_e}$, where t_1 is a local parameter (a *level parameter*) transverse to the boundary. As a consequence a family of differential forms that tends to a generator on top level scales with $t_1^{\ell_\Gamma}$ on the bottom level of Γ .

For graphs with L level passages we define $\ell_i = \ell_{\Gamma,i} = \ell_{\delta_i(\Gamma)}$ to be the lcm of the edges crossing the i -th level passage and $\ell_\Gamma = \prod_{i=1}^L \ell_{\Gamma,i}$.

There are two sources of automorphisms of multi-scale differentials: on the one hand, there are automorphism of pointed stable curves that respect the additional structure (differential, prong-matching). On the other hand, there are *ghost automorphisms*, whose group we denote by $\text{Gh}_\Gamma = \text{Tw}_\Gamma / \text{Tw}_\Gamma^s$, that stem from the toric geometry of the compactification. We emphasize that the twist group Tw_Γ and the simple twist Tw_Γ^s , hence also the ghost group Gh_Γ , depend only on the data of the enhanced level graph and will be inherited by linear submanifolds below. The local isotropy group of $\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ sits in a exact sequence

$$0 \rightarrow \text{Gh}_\Gamma \rightarrow \text{Iso}(X, \omega) \rightarrow \text{Aut}(X, \omega) \rightarrow 0$$

and locally near $(X, \mathbf{z}, \Gamma, \sigma)$ the stack of multi-scale differentials is the quotient stack $[U/\text{Iso}(X, \omega)]$ for some open $U \subset \mathbb{C}^{N_B}$. The same holds for $\mathbb{P}\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ where the automorphism group is potentially larger since ω is only required to be fixed projectively.

3.4. Decomposition of the logarithmic tangent bundle. We now define a Γ -adapted basis, combining [BDG22] and [CMZ22] with the goal of giving a decomposition of the logarithmic tangent bundle that is inherited by a linear submanifold, if the Γ -adapted basis is suitably chosen.

We work on a neighborhood U of a point $p = (X, [\omega], \mathbf{z}) \in D_\Gamma^B$, where Γ is an arbitrary level graph with L levels below zero. We let $\alpha_j^{[i]}$ for $i = 0, \dots, -L$ be the vanishing cycles around the horizontal nodes at level i . Let $\beta_j^{[i]}$ be a dual horizontal-crossing cycle, i.e. i is the top level (in the sense of [BDG22]) of this cycle, $\langle \alpha_j^{[i]}, \beta_j^{[i]} \rangle = 1$ and $\beta_j^{[i]}$ does not cross any other horizontal node at level i . Let $h(i)$ be the number of those horizontal nodes at level i .

We complement the cycles $\beta_j^{[i]}$ by a collection of relative cycles $\gamma_j^{[i]}$ such that for any fixed level i their top level restrictions form a basis of the cohomology at level i relative to the poles and zeros of ω and holes at horizontal nodes quotiented by the subspace of global residue conditions. In particular the span of the $\gamma_j^{[i]}$ contains the $\alpha_j^{[i]}$, and moreover the union

$$\bigcup_{j=-L}^0 \{ \beta_1^{[j]}, \dots, \beta_{h(j)}^{[j]}, \gamma_1^{[j]}, \dots, \gamma_{s(j)}^{[j]} \} \text{ is a basis of } H_1(X \setminus P, Z, \mathbb{C}).$$

Next, we define the ω -periods of these cycles and exponentiate to kill the monodromy around the vanishing cycles. The functions

$$a_j^{[i]} = \int_{\alpha_j^{[i]}} \omega, \quad b_j^{[i]} = \int_{\beta_j^{[i]}} \omega, \quad q_j^{[i]} = \exp(2\pi I b_j^{[i]} / a_j^{[i]}), \quad c_j^{[i]} = \int_{\gamma_j^{[i]}} \omega.$$

are however still not defined on U (only on sectors of the boundary complement) due to monodromy around the vertical nodes.

Coordinates on U are given by *perturbed period coordinates* ([BCGGM19b]), which are related to the periods above as follows. For each level passage there is a *level parameter* t_i that stem from the construction of the moduli space via plumbing. On the bottom level passage L we may take $t_L = c_1^{[-L]}$ as a period. For the higher level passage, the t_i are closely related to the periods of a cycle with top level $-i$, but the latter are in general not monodromy invariant. It will be convenient to write

$$(24) \quad t_{[i]} = \prod_{j=1}^i t_j^{\ell_j}, \quad i \in \mathbb{N}.$$

There are perturbed periods $\tilde{c}_j^{[-i]}$ obtained by integrating $\omega/t_{[i]}$ against a cycle with top level $-i$ over the part of level $-i$ to points nearby the nodes, cutting off the lower level part. By construction, on each sector of the boundary complement we have

$$(25) \quad \tilde{c}_j^{[-i]} - c_j^{[-i]} / t_{[i]} = \sum_{s>i} \frac{t_{[s]}}{t_{[i]}} E_{j,i}^{[-s]}$$

for some linear ('error') forms $E_{j,i}^{[-s]}$ depending on the variables $c_j^{[-s]}$ on the lower level $-s$. Similarly, we can exponentiate the ratio over $a_j^{[-i]}$ of the similarly perturbed $\tilde{b}_j^{[-i]}$ and obtain perturbed exponentiated periods $\tilde{q}_j^{[-i]}$, such that on each sector

$$(26) \quad \log \tilde{q}_j^{[-i]} - \log q_j^{[-i]} = \sum_{s>i} \frac{t_{[s]}}{t_{[i]}} E_{j,i}'^{[-s]}$$

for some linear forms $E_{j,i}'^{[-s]}$. In these coordinates the boundary is given by $\tilde{q}_j^{[-i]} = 0$ and $t_i = 0$. If we let

$$\begin{aligned} \Omega_{i,B}^{\text{hor}}(\log) &= \langle d\tilde{q}_1^{[i]} / \tilde{q}_1^{[i]}, \dots, d\tilde{q}_{h(i)}^{[i]} / \tilde{q}_{h(i)}^{[i]} \rangle, & \Omega_{i,B}^{\text{lev}}(\log) &= \langle dt_{-i} / t_{-i} \rangle \\ \Omega_{i,B}^{\text{rel}} &= \langle d\tilde{c}_2^{[i]} / \tilde{c}_2^{[i]}, \dots, d\tilde{c}_{N(i)-h(i)}^{[i]} / \tilde{c}_{N(i)-h(i)}^{[i]} \rangle, \end{aligned}$$

with $\Omega_{0,B}^{\text{lev}}(\log) = 0$ by convention, we thus obtain a decomposition

$$(27) \quad \Omega_B^1(\log \partial B)|_U = \bigoplus_{i=-L}^0 \left(\Omega_{i,B}^{\text{hor}}(\log) \oplus \Omega_{i,B}^{\text{lev}}(\log) \oplus \Omega_{i,B}^{\text{rel}} \right).$$

3.5. The closure of linear submanifolds. For a linear submanifold \mathcal{H} we denote by $\overline{\mathcal{H}}$ the normalization of the closure of the image of \mathcal{H} as a substack of $\Xi\overline{\mathcal{M}}_{g,n}(\mu)$. We denote by $D_\Gamma = D_\Gamma^{\mathcal{H}}$ the preimage of the boundary divisor D_Γ^B in $\overline{\mathcal{H}}$. Again, a \circ denotes the complement of more degenerate boundary strata, i.e., D_Γ° is the preimage of $D_\Gamma^{B,\circ}$ in $\overline{\mathcal{H}}$.

We will now give several propositions that explain that $\overline{\mathcal{H}}$ is a compactification of \mathcal{H} almost as nice as the compactification $\mathbb{P}\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ of strata. The first statement explains the 'almost'.

PROPOSITION 3.1. *Let Γ be a level graph with only horizontal nodes, i.e., with one level only. Each point in $D_\Gamma^{B,\circ}$ has a neighborhood where the image of $\overline{\mathcal{H}}$ has at worst toric singularities.*

More precisely, the linear submanifold is cut out by linear and binomial equations, see (30) below.

Second, the intersection with non-horizontal boundary components is transversal in the strong sense that each level actually causes dimension drop.

PROPOSITION 3.2. *Let $\Gamma \in \text{LG}_L(B)$ be a level graph without horizontal nodes. Each point in $D_\Gamma^{B,\circ}$ has a neighborhood where each branch of $\overline{\mathcal{H}}$ mapping to that neighborhood is smooth and the boundary $\partial\mathcal{H} = \overline{\mathcal{H}} \setminus \mathcal{H}$ is a normal crossing divisor, the intersection of L different divisors $D_{\delta_i(\Gamma)}^{\mathcal{H}}$.*

In particular the image of $D_\Gamma^{\mathcal{H}}$ has codimension L in D_Γ^B .

The previous proposition allows to show, via the same argument as the proof of [CMZ22, Proposition 5.1], the key result in order to argue inductively.

COROLLARY 3.3. *If $\bigcap_{j=1}^L D_{\Gamma_{i_j}}^{\mathcal{H}}$ is not empty, there is a unique ordering $\sigma \in \text{Sym}_L$ on the set $I = \{i_1, \dots, i_L\}$ of indices such that*

$$D_{\sigma(I)} = \bigcap_{j=1}^L D_{\Gamma_{i_j}}^{\mathcal{H}}.$$

Moreover if $i_k = i_{k'}$ for a pair of indices $k \neq k'$, then $D_{i_1, \dots, i_L} = \emptyset$.

The next statement is crucial to inductively apply the formulas in this paper. Recall that p_Γ and c_Γ are the projection and clutching morphisms of the diagram (22).

PROPOSITION 3.4. *There are generalized linear submanifolds $\Omega\mathcal{H}_\Gamma^{[i]} \rightarrow \Omega\mathcal{M}_{\mathbf{g}_i, \mathbf{n}_i}^{\text{st}}(\boldsymbol{\mu}_i)$ of dimension d_i with projectivization $\mathcal{H}_\Gamma^{[i],\circ}$, such that*

$$\sum_{i=-L}^0 d_i = d_{\mathcal{H}} - L$$

and such that the normalizations $\mathcal{H}_\Gamma^{[i]} \rightarrow B_\Gamma^{[i]}$ of closures of $\mathcal{H}_\Gamma^{[i],\circ}$ together give a product decomposition $\mathcal{H}_\Gamma = \prod_{i=-L}^0 \mathcal{H}_\Gamma^{[i]}$ of the normalization of the p_Γ -image of the c_Γ -preimage of $\text{Im}(D_\Gamma^{\mathcal{H}}) \subset \mathbb{P}\Xi\overline{\mathcal{M}}_{g,n}(\mu)$.

We will call $\mathcal{H}_\Gamma^{[i]} \rightarrow B_\Gamma^{[i]}$ the i -th level linear manifold. Our ultimate goal here is to show the following decomposition. The terminology is explained along with the definition of coordinates.

PROPOSITION 3.5. *Let Γ be an arbitrary level graph with L levels below zero. In a small neighborhood U of a point in $D_\Gamma^{\mathcal{H}}$ there is a direct sum decomposition*

$$(28) \quad \Omega_{\overline{\mathcal{H}}}^1(\log \partial\mathcal{H})|_U = \bigoplus_{i=-L}^0 \left(\Omega_i^{\text{hor}}(\log) \oplus \Omega_i^{\text{lev}}(\log) \oplus \Omega_i^{\text{rel}} \right)$$

for certain subsheaves such that the natural restriction map induces surjections

$$\Omega_{i,B}^{\text{hor}}(\log)|_{\overline{\mathcal{H}}} \rightarrow \Omega_i^{\text{hor}}(\log), \quad \Omega_{i,B}^{\text{lev}}(\log)|_{\overline{\mathcal{H}}} \simeq \Omega_i^{\text{lev}}(\log) \quad \text{and} \quad \Omega_{i,B}^{\text{rel}}|_{\overline{\mathcal{H}}} \rightarrow \Omega_i^{\text{rel}}.$$

Moreover the statements in items i) and ii) of Section 3.3 hold verbatim for the linear submanifold with the same ℓ_Γ .

As a consequence we may use the symbols ℓ_Γ and ℓ_{Γ_i} ambiguously for strata and their linear submanifolds.

We summarize the relevant parts of [BDG22]. Equations of \mathcal{H} are interpreted as homology classes and we say that a *horizontal node is crossed by an equation*, if the corresponding vanishing cycles has non-trivial intersection with the equation. The horizontal nodes are partitioned into \mathcal{H} -cross-equivalence classes by simultaneous appearance in equations for \mathcal{H} . A main observation is that ω -periods of the vanishing cycles in an \mathcal{H} -cross-equivalence class are proportional. Similarly, for each equation and for any level passage the intersection numbers of the equation with the nodes crossing that level add up to zero when weighted appropriately with the residue times ℓ_Γ/κ_e ([BDG22, Proposition 3.11]).

Next, in [BDG22] they sort the equations by level and then write them in reduced row echelon form. One may order the periods so that the distinguished $c_1^{[i]}$ (whose period is close to the level parameter t_{-i}) is among the pivots of the echelon form for each i . The second main observation is that each defining equation of \mathcal{H} can be split into a sum of defining equations, denoted by $F_k^{[i]}$, with the following properties. The upper index i indicate the highest level, whose periods are involved in the equation. Moreover, either $F_k^{[i]}$ has non-trivial intersection with some (vanishing cycles of a) horizontal node at level i and then no intersection with a horizontal node at lower level, or else no intersection with a horizontal node at all.

As a result \mathcal{H} is cut out by two sets of equations, see [BDG22, Equations (4.2), (4.3), (4.4)]. First, there are the equations $G_k^{[i]}$ that are $t_{[-i]}$ -rescalings of linear functions

$$(29) \quad G_k^{[i]} = L_k^{[i]}(\tilde{c}_{2-\delta_{i,0}^{[i]}}, \dots, \tilde{c}_{N(i)-h(i)}^{[i]})$$

in the periods at level i . (To get this form from the version in [BDG22] absorb the terms from lower level periods into the function $c_j^{[i]}$ where $j = j(k, i)$ is the pivot of the equation $F_k^{[i]}$. This does not effect the truth of (25)).

Second, there are multiplicative monomial equations among the exponentiated periods, that can be written as bi-monomial equations with positive exponents

$$(30) \quad H_k^{[i]} = (\tilde{\mathbf{q}}^{[i]})^{J_{1,k}} - (\tilde{\mathbf{q}}^{[i]})^{J_{2,k}}$$

where $\tilde{\mathbf{q}}^{[i]}$ is the tuple of the variables $\tilde{q}_j^{[i]}$ and $J_{1,k}, J_{2,k}$ are tuples of non-negative integers. (In the multiplicative part [BDG22] already incorporated the lower level blurring into the pivot variable.)

PROOF OF PROPOSITION 3.1. This follows directly from the form of the binomial equations (30), see [BDG22, Theorem 1.6]. \square

PROOF OF PROPOSITION 3.2. Smoothness and normal crossing is contained in [BDG22, Corollary 1.8]. The transversality claimed there contains the dimension drop claimed in the proposition. The more precise statement in [BDG22, Theorem 1.5] says that after each intersection of $\overline{\mathcal{H}}$ with a vertical boundary divisor the result is empty or contained in the open boundary divisor $D_\Gamma^{B,\circ}$. \square

PROOF OF PROPOSITION 3.4. This is the main result of [Ben20] or the restatement in [BDG22, Proposition 3.3] and this together with the Proposition 3.2 implies the dimension statement. \square

PROOF OF PROPOSITION 3.5. Immediate from (29) and (30), which are equations among the respective set of generators of the decomposition in (27). The additional claim item ii) follows from the isomorphism of level parameters and transversality. Item i) is a consequence of this. \square

3.6. Push-pull comparison for linear submanifolds. For recursive computations, we will transfer classes from $\mathcal{H}_\Gamma^{[i]}$, which were defined via Proposition 3.4, to $D_\Gamma^{\mathcal{H}}$ essentially via p_Γ -pullback and c_Γ -pushforward. More precisely, taking the normalizations into account, we have to use the maps $c_{\Gamma, \mathcal{H}}$ and $p_{\Gamma, \mathcal{H}}$ defined on the normalization \mathcal{H}_Γ^s of the c_Γ -preimage of the image of $D_\Gamma^{\mathcal{H}}$ in D_Γ^B . To compute degrees we use the analog of the inner triangle in (22) and give a concrete description of \mathcal{H}_Γ^s .

Recall from the introduction that $K_\Gamma^{\mathcal{H}}$ is the number of prong-matchings of Γ that are reachable from within \mathcal{H} .

$$(31) \quad \begin{array}{ccccc} (\Omega\mathcal{H}_\Gamma^o)^{\text{pm}} & \xrightarrow{\quad} & \mathcal{H}_\Gamma^{s,o} & & \\ \downarrow & & \downarrow & \searrow^{c_{\Gamma, \mathcal{H}}} & \\ \Omega\mathcal{H}_\Gamma^o & \xrightarrow{\quad} & \mathcal{H}_\Gamma^o & \xrightarrow{\quad} & B_{\Gamma, \Gamma} \\ & & \swarrow^{p_{\Gamma, \mathcal{H}}} & \swarrow^{p_\Gamma} & \downarrow^{c_\Gamma} \\ & & & & U_\Gamma \longleftarrow D_\Gamma^{\mathcal{H}, o} \end{array}$$

Consider $\Omega\mathcal{H}_\Gamma^o := \prod \Omega\mathcal{H}_\Gamma^{[i]}$ as a moduli space of differentials subject to some (linear) conditions imposed on its periods. Consider now the moduli space $(\Omega\mathcal{H}_\Gamma^o)^{\text{pm}} := (\prod \Omega\mathcal{H}_\Gamma^{[i]})^{\text{pm}}$ where we add the additional datum of one of the $K_\Gamma^{\mathcal{H}}$ prong-matchings reachable from the interior. The torus $(\mathbb{C}^*)^{L+1}$ acts on $\Omega\mathcal{H}_\Gamma^o$ with quotient $\mathcal{H}_\Gamma^o = \prod \mathcal{H}_\Gamma^{[i], o}$. On the other hand, if we take the quotient of $(\Omega\mathcal{H}_\Gamma^o)^{\text{pm}}$ by $(\mathbb{C}^*)^{L+1} = (\mathbb{C}^*) \times (\mathbb{C}^L / \text{Tw}_\Lambda^s)$ we obtain a space $\mathcal{H}_\Gamma^{s,o}$ which is naturally the normalization of a subspace of U_Γ^s , since it covers $D_\Gamma^{\mathcal{H}, o}$ with marked (legs and) edges and whose generic isotropy group does not stem from Gh_Γ (it might be non-trivial, e.g. if a level of Γ consists of a hyperelliptic stratum), while the generic isotropy group of $D_\Gamma^{\mathcal{H}, o}$ is an extension of Gh_Γ by possibly some group of graph automorphisms and possibly isotropy groups of the level strata.

LEMMA 3.6. *The ratio of the degrees the maps in 31 on \mathcal{H}_Γ^s is*

$$\frac{\deg(p_{\Gamma, \mathcal{H}})}{\deg(c_{\Gamma, \mathcal{H}})} = \frac{K_\Gamma^{\mathcal{H}}}{|\text{Aut}_{\mathcal{H}}(\Gamma)|\ell_\Gamma},$$

where $\text{Aut}_{\mathcal{H}}(\Gamma)$ is the subgroup of $\text{Aut}(\Gamma)$ whose induced action on a neighborhood of $D_\Gamma^{\mathcal{H}}$ preserves $\overline{\mathcal{H}}$.

PROOF. We claim that the degree of $p_{\Gamma, \mathcal{H}}$ is the number of prong-matchings equivalence classes, i.e., $\deg(p_{\Gamma, \mathcal{H}}) = K_\Gamma^{\mathcal{H}} / [R_\Gamma : \text{Tw}_\Gamma]$ where $R_\Gamma \cong \mathbb{Z}^L \subset \mathbb{C}^L$ is the level rotation group. In fact this follows since $\text{Tw}_\Gamma^s \subseteq \text{Tw}_\Gamma$ and $\mathcal{H}_\Gamma^{s,o}$ is given by taking the quotient by the action of the level rotation group, which has Tw_Γ as its stabilizer subgroup.

On the other side $c_{\Gamma, \mathcal{H}}$ factors through the quotient by $\text{Gh}_\Gamma = [\text{Tw}_\Gamma : \text{Tw}_\Gamma^s]$ acting by fixing every point. In the remaining quotient map c_Γ^{Γ} of the ambient stratum two points have the same image only if they differ by an automorphism of Γ . However only the subgroup $\text{Aut}_{\mathcal{H}}(\Gamma) \subset \text{Aut}(\Gamma)$ acts on $\text{Im}(H_\Gamma^s)$ and its normalization and contributes to the local isotropy group of the normalization. Thus only this subgroup contributes to the degree of $c_{\Gamma, \mathcal{H}}$. The claimed equality now follows because $[R_\Gamma : \text{Tw}_\Gamma^s] = \ell_\Gamma$. \square

Consider a graph $\Delta \in \text{LG}_1(\mathcal{H}_\Gamma^{[i]})$ defining a divisor in $\mathcal{H}_\Gamma^{[i]}$. We aim to compute its pullback to D_Γ^s and the push forward to D_Γ and to $\overline{\mathcal{H}}$. For this purpose we need extend the commensurability diagram (31) to include degenerations of the boundary strata. This works by copying verbatim the construction that lead in [CMZ22] to the commensurability diagram (22). We will indicate with subscripts \mathcal{H} to the morphisms that we work in this adapted setting. Recall from this construction that in $D_\Gamma^{B,s}$ (and hence in D_Γ^s) the edges of Γ have been labeled once and for all (we write Γ^\dagger for this labeled graph) and that the level strata $\mathcal{H}_\Gamma^{[i]}$ inherit these labels. Consequently, there is unique graph $\tilde{\Delta}^\dagger$ which is a degeneration of Γ^\dagger and such that extracting the levels i and $i-1$ of $\tilde{\Delta}^\dagger$ equals Δ . The resulting unlabeled graph will simply be denoted by $\tilde{\Delta}$. For a fixed labeled graph Γ^\dagger we denote by $J(\Gamma^\dagger, \tilde{\Delta})$ the set of $\Delta \in \text{LG}_1(\mathcal{H}_\Gamma^{[i]})$ such that $\tilde{\Delta}$ is the result of that procedure.

Obviously the graphs in $J(\Gamma^\dagger, \tilde{\Delta})$ differ only by the labeling of their half-edges and the following lemma computes its cardinality.

LEMMA 3.7. *The cardinality of $J(\Gamma^\dagger, \tilde{\Delta})$ is determined by*

$$|J(\Gamma^\dagger, \tilde{\Delta})| \cdot |\text{Aut}_{\mathcal{H}}(\tilde{\Delta})| = |\text{Aut}_{\mathcal{H}_\Gamma^{[i]}(\Delta)}| \cdot |\text{Aut}_{\mathcal{H}}(\Gamma)|.$$

PROOF. The proof is analogous to the one of [CMZ22, Lemma 4.6], where one considers the kernel and cokernel of the map $\varphi : \text{Aut}_{\mathcal{H}}(\tilde{\Delta}) \rightarrow \text{Aut}_{\mathcal{H}}(\Gamma)$ given by undegeneration. \square

We now determine the multiplicities of the push-pull procedure. Recall from Section 3.3 the definition of $\ell_{\Gamma, j} = \ell_{\delta_j(\Gamma)}$ for $j \in \mathbb{Z}_{\geq 1}$.

PROPOSITION 3.8. *For a fixed $\Delta \in \text{LG}_1(\mathcal{H}_\Gamma^{[i]})$, the divisor classes of $D_\Delta^{\mathcal{H}}$ and the clutching of $D_\Delta^{\mathcal{H}}$ are related by*

$$(32) \quad \frac{|\text{Aut}_{\mathcal{H}}(\tilde{\Delta})|}{|\text{Aut}_{\mathcal{H}_\Gamma^{[i]}(\Delta)}| |\text{Aut}_{\mathcal{H}}(\Gamma)|} \cdot c_{\Gamma, \mathcal{H}}^*[D_\Delta^{\mathcal{H}}] = \frac{\ell_\Delta}{\ell_{\tilde{\Delta}, -i+1}} \cdot p_{\Gamma, \mathcal{H}}^{[i], *}[D_\Delta^{\mathcal{H}}].$$

in $\text{CH}^1(D_\Gamma^s)$ and consequently by

$$(33) \quad \frac{|\text{Aut}_{\mathcal{H}}(\tilde{\Delta})|}{|\text{Aut}_{\mathcal{H}}(\Gamma)|} \cdot \ell_{\tilde{\Delta}, -i+1} \cdot [D_\Delta^{\mathcal{H}}] = \frac{|\text{Aut}_{\mathcal{H}_\Gamma^{[i]}(\Delta)}|}{\deg(c_{\Gamma, \mathcal{H}})} \cdot \ell_\Delta \cdot c_{\Gamma, \mathcal{H}, *}(p_{\Gamma, \mathcal{H}}^{[i], *}[D_\Delta^{\mathcal{H}}])$$

in $\text{CH}^1(D_\Gamma)$.

Here (32) is used later for the proofs of the main theorems while (33) is implemented in **diffstrata** for the special case of k -differentials to compute the pull-back of tautological classes from $D_\Delta^{\mathcal{H}}$ to $\tilde{D}_\Delta^{\mathcal{H}}$, see also Section 7.

PROOF. The proof is similar to the one of [CMZ22, Proposition 4.7] and works by comparing the ramification orders of the maps $c_{\Gamma, \mathcal{H}}^{\tilde{\Delta}}$ and $p_{\Gamma, \mathcal{H}}^{\tilde{\Delta}}$. The main difference to the original proof is only that the automorphism factors appearing in the clutching morphisms are the ones fixing \mathcal{H} . \square

The final part of this section is to compare various natural vector bundles under pull-back along the maps $c_{\Gamma, \mathcal{H}}$ and $p_{\Gamma, \mathcal{H}}$. The first of this is \mathcal{E}_Γ^\top , a vector bundle of rank $N_\Gamma^\top - 1$ on $D_\Gamma^{\mathcal{H}}$ that should be thought of as the top level version of the logarithmic cotangent bundle. Formally, let $U \subset D_\Gamma^{\mathcal{H}}$ be an open set centered at a degeneration of the top level of Γ into k level passages. Then we define

$$(34) \quad \mathcal{E}_\Gamma^\top|_U = \bigoplus_{i=-k}^0 \Omega_i^{\text{lev}}(\log)|_U \oplus \Omega_i^{\text{hor}}(\log)|_U \oplus \Omega_i^{\text{rel}}|_U.$$

Let moreover $\xi_{\Gamma, \mathcal{H}}^{[i]}$ be the first Chern class of the line bundle on $D_\Gamma^{\mathcal{H}}$ generated by the multi-scale component at level i and $\mathcal{L}_\Gamma^{[i]}$ be the line bundle whose divisor is given by the degenerations of the i -th level of Γ , as defined more formally in (44) below.

We have the following compatibilities.

LEMMA 3.9. *The first Chern classes of the tautological bundles on the levels of a boundary divisor are related by*

$$(35) \quad c_{\Gamma, \mathcal{H}}^* \xi_{\Gamma, \mathcal{H}}^{[i]} = p_{\Gamma, \mathcal{H}}^{[i], *} \xi_{\mathcal{H}_\Gamma^{[i]}} \quad \text{in } \text{CH}^1(D_\Gamma^s).$$

It is also true that

$$(36) \quad p_{\Gamma, \mathcal{H}}^{[i], *} \mathcal{L}_{\mathcal{H}_\Gamma^{[i]}} = c_{\Gamma, \mathcal{H}}^* \mathcal{L}_\Gamma^{[i]} \quad \text{where } \mathcal{L}_{\mathcal{H}_\Gamma^{[i]}} = \mathcal{O}_{\mathcal{H}_\Gamma^{[i]}} \left(\sum_{\Delta \in \text{LG}_1(\mathcal{H}_\Gamma^{[i]})} \ell_\Delta D_\Delta \right).$$

Similarly for the logarithmic cotangent bundles we have

$$(37) \quad p_{\Gamma, \mathcal{H}}^{[0], *} \Omega_{\mathcal{H}_\Gamma^{[0]}}^1(\log D_{\mathcal{H}_\Gamma^{[0]}}) = c_{\Gamma, \mathcal{H}}^* \mathcal{E}_{\Gamma, \mathcal{H}}^\top.$$

PROOF. The first claim is just the global compatibility of the definitions of the bundles $\mathcal{O}(-1)$ on various spaces, compare [CMZ22, Proposition 4.9].

The second claim is a formal consequence of Lemma 3.7 and Proposition 3.8, just as in [CMZ22, Lemma 7.4].

The last claim follows as in [CMZ22, Lemma 9.6] by considering local generators, which are given in (34) and have for linear submanifolds the same shape as for strata. \square

In the final formulas we will use these compatibilities together with the following restatement of Lemma 3.6.

LEMMA 3.10. *Suppose that $\alpha_\Gamma \in \text{CH}_0(D_\Gamma^{\mathcal{H}})$ is a top degree class and that $c_{\Gamma, \mathcal{H}}^* \alpha_\Gamma = \prod_{i=0}^{-L(\Gamma)} p_{\Gamma, \mathcal{H}}^{[i],*} \alpha_i$ for some α_i . Then*

$$\int_{D_\Gamma^{\mathcal{H}}} \alpha_\Gamma = \frac{K_\Gamma^{\mathcal{H}}}{|\text{Aut}_{\mathcal{H}}(\Gamma)| \ell_\Gamma} \prod_{i=0}^{-L(\Gamma)} \int_{\mathcal{H}_\Gamma^{[i]}} \alpha_i.$$

4. Evaluation of tautological classes

This section serves two purposes. First, we briefly sketch a definition of the tautological ring of linear submanifolds and how the results of the previous section can be used to evaluate expressions in the tautological ring, provided the classes of the linear manifold are known. Second, we provide formulas to compute the first Chern class of the normal bundle $\mathcal{N}_\Gamma^{\mathcal{H}} = \mathcal{N}_{D_\Gamma^{\mathcal{H}}}$ to a boundary divisor $D_\Gamma^{\mathcal{H}}$ of a projectivized linear submanifold $\overline{\mathcal{H}}$. This is needed both for the evaluation algorithm and as an ingredient to prove our main theorems.

4.1. Vertical tautological ring. We denote by $\psi_i \in \text{CH}^1(\overline{\mathcal{H}})$ the pull-backs of the classes $\psi_i \in \text{CH}^1(\overline{\mathcal{M}}_{g,n})$ to a linear submanifold $\overline{\mathcal{H}}$. The *clutching maps* are defined as $\text{cl}_{\Gamma, \mathcal{H}} = \text{i}_{\Gamma, \mathcal{H}} \circ c_{\Gamma, \mathcal{H}}$, where $\text{i}_{\Gamma, \mathcal{H}} : D_\Gamma^{\mathcal{H}} \rightarrow \overline{\mathcal{H}}$ is the inclusion map of the boundary divisor. We define the (*vertical*) *tautological ring* $R_v^\bullet(\overline{\mathcal{H}})$ of $\overline{\mathcal{H}}$ to be the ring with additive generators

$$(38) \quad \text{cl}_{\Gamma, \mathcal{H}, *}\left(\prod_{i=0}^{-L} p_{\Gamma, \mathcal{H}}^{[i],*} \alpha_i\right)$$

where Γ runs over all level graphs without horizontal edges for all boundary strata of \mathcal{H} , including the trivial graph, and where α_i is a monomial in the ψ -classes supported on level i of the graph Γ . That this is indeed a ring follows from the excess intersection formula [CMZ22, Proposition 8.1] that works exactly the same for linear submanifolds, and the normal bundle formula Proposition 4.4 which allows together with Proposition 4.1 to rewrite products in terms of our standard generators. We do not claim that pushforward $R_v^\bullet(\overline{\mathcal{H}}) \rightarrow \text{CH}^\bullet(\overline{\mathcal{M}}_{g,n})$ maps to the tautological ring $R^\bullet(\overline{\mathcal{M}}_{g,n})$, since the fundamental classes of linear submanifolds, e.g. loci of double covers of elliptic curves, may be non-tautological in $\overline{\mathcal{M}}_{g,n}$ (see e.g. [GP03]).

To evaluate a top-degree class of the form $\alpha := \psi_1^{p_1} \cdots \psi_n^{p_n} \cdot [D_\Gamma^{\mathcal{H}}] \in \text{CH}_0(\overline{\mathcal{H}})$ there are (at least) two possible ways to proceed: If one knows the class $[\overline{\mathcal{H}}] \in \text{CH}_{\dim(\mathcal{H})}(\mathbb{P}\Xi\mathcal{M}_{g,n}(\mu))$ and this class happens to be tautological, one may evaluate

$$\int_{\overline{\mathcal{H}}} \alpha = \int_{\mathbb{P}\Xi\mathcal{M}_{g,n}(\mu)} \psi_1^{p_1} \cdots \psi_n^{p_n} \cdot [D_\Gamma] \cdot [\overline{\mathcal{H}}]$$

using the methods described in [CMZ22]. Alternatively one may apply Lemma 3.6 to obtain

$$(39) \quad \int_{\overline{\mathcal{H}}} \alpha = \frac{K_\Gamma^{\mathcal{H}}}{|\text{Aut}_{\mathcal{H}}(\Gamma)| \ell_\Gamma} \prod_{i=0}^{-L} \int_{\mathcal{H}_\Gamma^{[i]}} \prod_{j \in l(i)} \psi_j^{p_j},$$

where $l(i)$ denotes the set of legs on level i of Γ . To evaluate this expression, one needs to determine the fundamental classes of the level linear submanifolds $\mathcal{H}_\Gamma^{[i]}$ in their corresponding generalized strata, which is in general a non-trivial task.

4.2. Evaluation of $\xi_{\mathcal{H}}$. If we want to evaluate a top-degree class in $\text{CH}_0(\overline{\mathcal{H}})$ that is not just a product of ψ -classes and a boundary stratum, but also involves the $\xi_{\mathcal{H}}$ -class, we can reduce to the previous case by applying the following proposition.

PROPOSITION 4.1. *The class $\xi_{\mathcal{H}}$ on the closure of a projectivized linear submanifold $\overline{\mathcal{H}}$ can be expressed as*

$$(40) \quad \xi_{\mathcal{H}} = (m_i + 1)\psi_i - \sum_{\Gamma \in {}_i\text{LG}_1(\mathcal{H})} \ell_{\Gamma}[D_{\Gamma}^{\mathcal{H}}]$$

where ${}_i\text{LG}_1(\mathcal{H})$ are two-level graphs with the leg i on lower level.

PROOF. The formula is obtained by pulling-back the formula in [CMZ22, Proposition 8.1] to $\overline{\mathcal{H}}$ and thereby using the transversality statement from Proposition 3.2. \square

We remark here that in some cases it is possible to directly evaluate the top $\xi_{\mathcal{H}}$ -powers by using that we can represent the powers of the $\xi_{\mathcal{H}}$ -class via an explicit closed current.

Let $\mathbb{P}\Xi\mathcal{M}_{g,n}(\mu)$ be a *holomorphic stratum*, i.e. a stratum of flat surfaces of finite area or equivalently all the entries of μ are non-negative. Then there is a canonical hermitian metric on the tautological bundle $\mathcal{O}_{\mathbb{P}\Xi\mathcal{M}_{g,n}(\mu)}(-1)$ given by the flat area form

$$(41) \quad h(X, \omega, \mathbf{z}) = \text{area}_X(\omega) = \frac{i}{2} \int_X \omega \wedge \bar{\omega}$$

which extends to an hermitian metric of the tautological bundle on $\mathbb{P}\Xi\overline{\mathcal{M}}_{g,n}(\mu)$. If $\overline{\mathcal{H}} \rightarrow \mathbb{P}\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ is the compactification of a linear submanifold of such a holomorphic stratum, then the area metric induces an hermitian metric, which we denote again by h , on the pull-back $\mathcal{O}_{\overline{\mathcal{H}}}(-1)$ of the tautological bundle to $\overline{\mathcal{H}}$. Recall from Proposition 3.1 (combined with the level-wise decomposition in Proposition 3.4) that the singularities of $\overline{\mathcal{H}}$ are toric. Let $\overline{\mathcal{H}}^{\text{tor}} \rightarrow \overline{\mathcal{H}}$ be a resolution of singularities which is locally toric.

PROPOSITION 4.2. *Let $\overline{\mathcal{H}}^{\text{tor}} \rightarrow \mathbb{P}\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ be a resolution of a compactified linear submanifold of a holomorphic stratum. The curvature form $\frac{i}{2\pi}[F_h]$ of the pull metric h to $\overline{\mathcal{H}}^{\text{tor}}$ is a closed current that represents the first Chern class $c_1(\mathcal{O}_{\overline{\mathcal{H}}^{\text{tor}}}(-1))$. More generally, the d -th wedge power of the curvature form represents $c_1(\mathcal{O}_{\overline{\mathcal{H}}^{\text{tor}}}(-1))^d$ for any $d \geq 1$.*

PROOF. In [CMZ19, Proposition 4.3] it was shown that on the neighborhood U of a boundary point of $\mathbb{P}\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ in the interior of the stratum D_{Γ} the metric h has the form

$$(42) \quad h(X, q) = \sum_{i=0}^L |t_{[i]}|^2 \left(h_{(-i)}^{\text{tck}} + h_{(-i)}^{\text{ver}} + h_{(-i)}^{\text{hor}} \right)$$

where $h_{(-i)}^{\text{tck}}$ (coming from the 'thick' part) are smooth positive functions bounded away from zero and

$$(43) \quad h_{(-i)}^{\text{ver}} := - \sum_{p=1}^i R_{(-i),p}^{\text{ver}} \log |t_p|, \quad h_{(-i)}^{\text{hor}} := - \sum_{j=1}^{E_{(-i)}^h} R_{(-i),j}^{\text{hor}} \log |q_j^{[i]}|,$$

where $R_{(-i),p}^{\text{ver}}$ is a smooth non-negative function and $R_{(-i),j}^{\text{hor}}$ is a smooth positive function bounded away from zero, both involving only perturbed period coordinates on levels $-i$ and below.

The statement of the proposition in loc. cit. follows by formal computations from the shape of (42) and the properties of its coefficients, see [CMZ19, Proposition 4.4 and 4.5]. We thus only need to show that in local coordinates of a point in $\overline{\mathcal{H}}^{\text{tor}}$ (mapping to the given stratum D_{Γ}) the metric has the same shape (42). For this purpose, recall that by Proposition 3.4, the level parameters t_i are among the coordinates. On the other hand, a toric resolution of the toric singularities arising from (30) is given by fan subdivision and thus by a collection of variables $y_j^{[i]}$ for each level i , each of which is a product of integral powers of the $q_j^{[i]}$ at that level i . Conversely the map $\overline{\mathcal{H}}^{\text{tor}} \rightarrow \mathbb{P}\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ is given locally by $q_j^{[i]} = \prod_k (y_k^{[i]})^{b_{i,j,k}}$ for some $b_{i,j,k} \in \mathbb{Z}_{\geq 0}$, not all of the $b_{i,j,k} = 0$ for fixed (i, j) . Plugging this into (42) and (43) gives an expression of the same shape and with coefficients

satisfying the same smoothness and positivity properties. Mimicking the proof in loc. cit. thus implies the claim. \square

For a linear submanifold \mathcal{H} consider the vector space given in local period coordinates by the intersection of the tangent space of the unprojectivized linear submanifold with the span of relative periods. We call this space the REL space of \mathcal{H} and we denote by $R_{\mathcal{H}}$ its dimension.

Using Proposition 4.2 we can now generalize the result about vanishing of top ξ -powers on non-minimal strata of differentials to linear submanifolds with non-zero REL (see [Sau18, Proposition 3.3] for the holomorphic abelian strata case).

COROLLARY 4.3. *Let $\overline{\mathcal{H}} \rightarrow \mathbb{P}\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ be a linear submanifold of a holomorphic stratum. Then*

$$\int_{\overline{\mathcal{H}}} \xi_{\overline{\mathcal{H}}}^i \alpha = 0 \quad \text{for } i \geq d_{\mathcal{H}} - R_{\mathcal{H}} + 1,$$

where $d_{\mathcal{H}}$ is the dimension of \mathcal{H} and $R_{\mathcal{H}}$ is the dimension of the REL space and where α is any class of dimension $d_{\mathcal{H}} - i$.

PROOF. Since the area is given by an expression in absolute periods, the pullback of ξ to $\overline{\mathcal{H}}^{\text{tor}}$ is represented by Proposition 4.2 by a $(1, 1)$ -form involving only absolute periods (see [Sau18, Lemma 2.1] for the explicit expression in the case of strata). Taking a wedge power that exceeds the dimension of the space of absolute periods gives zero. \square

4.3. Normal bundles. Finally we state the normal bundle formula, which is necessary to evaluate self-intersections, which is for example needed to evaluate powers of $\xi_{\mathcal{H}}$. More generally, we provide formulas for the normal bundle of an inclusion $j_{\Gamma, \Pi}: D_{\Gamma}^{\mathcal{H}} \hookrightarrow D_{\Pi}^{\mathcal{H}}$ between non-horizontal boundary strata of relative codimension one, say defined by the L -level graph Π and one of its $(L + 1)$ -level graph degenerations Γ . This generalization is needed for recursive evaluations. Such an inclusion is obtained by splitting one of the levels of Π , say the level $i \in \{0, -1, \dots, -L\}$. We define

$$(44) \quad \mathcal{L}_{\Gamma}^{[i]} = \mathcal{O}_{D_{\Gamma}^{\mathcal{H}}} \left(\sum_{\Gamma \xrightarrow{[i]} \tilde{\Delta}} \ell_{\tilde{\Delta}, -i+1} D_{\tilde{\Delta}}^{\mathcal{H}} \right) \quad \text{for any } i \in \{0, -1, \dots, -L\},$$

where the sum is over all graphs $\tilde{\Delta} \in \text{LG}_{L+2}(\mathcal{H})$ that yield divisors in $D_{\Gamma}^{\mathcal{H}}$ by splitting the i -th level, which in terms of undegenerations means $\delta_{-i+1}^{\mathbb{C}}(\tilde{\Delta}) = \Gamma$. The following result contains the formula for the normal bundle as the special case where Π is the trivial graph.

PROPOSITION 4.4. *For $\Pi \xrightarrow{[i]} \Gamma$ (or equivalently for $\delta_{-i+1}^{\mathbb{C}}(\Gamma) = \Pi$) the Chern class of the normal bundle $\mathcal{N}_{\Gamma, \Pi}^{\mathcal{H}} := \mathcal{N}_{D_{\Gamma}^{\mathcal{H}}/D_{\Pi}^{\mathcal{H}}}$ is given by*

$$(45) \quad c_1(\mathcal{N}_{\Gamma, \Pi}^{\mathcal{H}}) = \frac{1}{\ell_{\Gamma, (-i+1)}} (-\xi_{\Gamma, \mathcal{H}}^{[i]} - c_1(\mathcal{L}_{\Gamma, \mathcal{H}}^{[i]}) + \xi_{\Gamma, \mathcal{H}}^{[i-1]}) \quad \text{in } \text{CH}^1(D_{\Gamma}^{\mathcal{H}}).$$

PROOF. We use the transversality statement Proposition 3.2 of \mathcal{H} with a boundary stratum $D_{\Gamma}^{\mathcal{H}}$ in order to have that the transversal parameter is given by t_i . Then the proof is as the same as the one in the case of abelian strata, see [CMZ22, Proposition 7.5]. \square

Since in Section 8 we will need to compute the normal bundle to horizontal divisors for strata of k -differentials, we provide here the general formula for the case of smooth horizontal degenerations of linear submanifolds.

PROPOSITION 4.5. *Let $D_h^{\mathcal{H}} \subset D^{\mathcal{H}}$ be a divisor in a boundary component $D^{\mathcal{H}}$ obtained by horizontal degeneration. Suppose that the linear submanifold is smooth along $D_h^{\mathcal{H}}$ and let e be one of the new horizontal edges in the level graph of $D_h^{\mathcal{H}}$. Then the first Chern class of the normal bundle $\mathcal{N}_{D_h^{\mathcal{H}}}^{\mathcal{H}}$ is given by*

$$c_1(\mathcal{N}_{D_h^{\mathcal{H}}}^{\mathcal{H}}) = -\psi_{e^+} - \psi_{e^-} \in \text{CH}^1(D^{\mathcal{H}})$$

where e^+ and e^- are the half-edges associated to the two ends of e .

PROOF. Similarly to the proof of [CMZ22, Proposition 7.2], consider the divisor D_e in $\overline{\mathcal{M}}_{g,n}$ corresponding to the single edge e and denote by \mathcal{N}_e its normal bundle. The forgetful map $f : D_h \rightarrow D_e$ induces an isomorphism $\mathcal{N}_{D_h}^{\mathcal{H}} \rightarrow f^* \mathcal{N}_{D_e}$ (compare local generators!) and the formula follows from the well-known expression of \mathcal{N}_{D_e} in terms of ψ -classes. \square

We will need the following result about pullbacks of normal bundles to apply the same arguments as in [CMZ22] recursively over inclusions of boundary divisors. The proof is the same as in [CMZ22, Corollary 7.7], since it follows from Proposition 4.4 that we can j-pullback properties of ξ and $\mathcal{L}_\Gamma^{[j]}$ that hold on the whole stratum and hence on linear submanifolds.

LEMMA 4.6. *Let $\Gamma \in \text{LG}_L(\mathcal{H})$ and let $\tilde{\Delta}$ be a codimension one degeneration of the $(-i+1)$ -th level of Γ , i.e., such that $\Gamma = \delta_i^{\mathbb{C}}(\tilde{\Delta})$, for some $i \in \{1, \dots, L+1\}$. Then*

$$j_{\tilde{\Delta}, \Gamma}^* \left(\ell_{\Gamma, j} c_1(\mathcal{N}_{\Gamma/\delta_j^{\mathbb{C}}(\Gamma)}^{\mathcal{H}}) \right) = \begin{cases} \ell_{\tilde{\Delta}, j} c_1 \left(\mathcal{N}_{\tilde{\Delta}/\delta_j^{\mathbb{C}}(\tilde{\Delta})}^{\mathcal{H}} \right), & \text{for } j < i \\ \ell_{\tilde{\Delta}, j+1} c_1 \left(\mathcal{N}_{\tilde{\Delta}/\delta_{(j+1)}^{\mathbb{C}}(\tilde{\Delta})}^{\mathcal{H}} \right) & \text{otherwise.} \end{cases}$$

5. Chern classes of the cotangent bundle via the Euler sequence

The core of the computation of the Chern classes is given by two exact sequences that are the direct counterparts of the corresponding theorems for abelian strata. The proof should be read in parallel with [CMZ22, Section 6 and 9] and we mainly highlight the differences and where the structure theorems of the compactification from Section 3.5 are needed.

THEOREM 5.1. *There is a vector bundle \mathcal{K} on $\overline{\mathcal{H}}$ that fits into an exact sequence*

$$(46) \quad 0 \longrightarrow \mathcal{K} \xrightarrow{\psi} (\overline{\mathcal{H}}_{\text{rel}}^1)^\vee \otimes \mathcal{O}_{\overline{\mathcal{H}}}(-1) \xrightarrow{\text{ev}} \mathcal{O}_{\overline{\mathcal{H}}} \longrightarrow 0,$$

where $\overline{\mathcal{H}}_{\text{rel}}^1$ is the Deligne extension of the local subsystem that defines the tangent space to $\Omega\mathcal{H}$ inside the relative cohomology $\overline{\mathcal{H}}_{\text{rel}, B|_{\overline{\mathcal{H}}}}^1$, such that the restriction of \mathcal{K} to the interior \mathcal{H} is the cotangent bundle $\Omega_{\mathcal{H}}^1$ and for U as in Proposition 3.5 we have

$$\mathcal{K}|_U = \bigoplus_{i=-L}^0 t_{\Gamma[-i]} \cdot \left(\Omega_i^{\text{hor}}(\log) \oplus \Omega_i^{\text{lev}}(\log) \oplus \Omega_i^{\text{rel}} \right).$$

The definition of the evaluation map and the notion of Deligne extension on a stack with toric singularities requires justification given in the proof. For the next result we define the abbreviations

$$(47) \quad \mathcal{E}_{\mathcal{H}} = \Omega_{\overline{\mathcal{H}}}^1(\log \overline{\partial}\mathcal{H}) \quad \text{and} \quad \mathcal{L}_{\mathcal{H}} = \mathcal{O}_{\overline{\mathcal{H}}} \left(\sum_{\Gamma \in \text{LG}_1(B)} \ell_{\Gamma} D_{\Gamma}^{\mathcal{H}} \right)$$

that are consistent with the level-wise definitions in (34) and (44).

THEOREM 5.2. *There is a short exact sequence of quasi-coherent $\mathcal{O}_{\overline{\mathcal{H}}}$ -modules*

$$(48) \quad 0 \longrightarrow \mathcal{E}_{\mathcal{H}} \otimes \mathcal{L}_{\mathcal{H}}^{-1} \rightarrow \mathcal{K} \rightarrow \mathcal{C} \longrightarrow 0$$

where $\mathcal{C} = \bigoplus_{\Gamma \in \text{LG}_1(\mathcal{H})} \mathcal{C}_{\Gamma}$ is a coherent sheaf supported on the non-horizontal boundary divisors, whose precise form is given in Proposition 5.4 below.

PROOF OF THEOREM 5.1. We start with the definition of the maps in the Euler sequence for the ambient stratum, see the middle row in the commutative diagram below. It uses the evaluation map

$$(49) \quad \text{ev}_B : (\overline{\mathcal{H}}_{\text{rel}, B}^1)^\vee \otimes \mathcal{O}_{\overline{B}}(-1) \rightarrow \mathcal{O}_{\overline{B}}, \quad \gamma \otimes \omega \mapsto \int_{\gamma} \omega,$$

restricted to $\overline{\mathcal{H}}$. The first map in the sequence is

$$(50) \quad \text{dc}_i \mapsto \left(\gamma_i - \frac{c_i}{c_k} \alpha_k \right) \otimes \omega, \quad i = 1, \dots, \widehat{k}, \dots, N,$$

as usual in the Euler sequence, on a chart of \mathcal{H} where c_k is non-zero. The exactness of the middle row is the content of [CMZ22, Theorem 6.1].

We next define the sheaf Eq. In the interior, Eq is the local system of equations cutting out $\Omega\mathcal{H}$, and thus the quotient $(\mathcal{H}_{\text{rel}}^1)^\vee = (\mathcal{H}_{\text{rel},B}^1)^\vee/\text{Eq}$ is the relative homology local system, by definition of a linear manifold. The proof in [CMZ22, Section 6.1] concerning the restriction of the sequence to the interior \mathcal{H} uses that \mathcal{H} has a linear structure with tangent space modeled on the local system $\mathbb{H}_{\text{rel}}^1$. In particular it gives the claim about $\mathcal{K}|_{\mathcal{H}}$.

As an interlude, we introduce notation for the Deligne extension of $(\mathcal{H}_{\text{rel},B}^1)^\vee$. For each $\gamma_j^{[i]}$ we let $\widehat{\gamma}_j^{[i]}$ be its extension, the sum of the original cycles and vanishing cycles times logarithms of the coordinates of the boundary divisors to kill monodromies. The functions

$$\widehat{c}_j^{[i]} = \frac{1}{t_{[-i]}} \int_{\widehat{\gamma}_j^{[i]}} \omega$$

are called *log periods* in [BDG22].

We now *define* Eq at the boundary, say locally near a point $p \in D_\Gamma$, to be the subsheaf of $(\overline{\mathcal{H}}_{\text{rel},B}^1)^\vee$ generated by the defining equations $F_k^{[i]}$ constructed in Section 3.5, but with each variable replaced by its Deligne extension. It requires justification that this definition near the boundary agrees with the previous definition in the interior. We can verify this for the distinguished basis consisting of the $F_k^{[i]}$. Equations that do not intersect horizontal nodes agree with their Deligne extension. This cancellation of the compensation terms is [BDG22, Proposition 3.11] (see also the expression for $F_k^{[i]}$ after [BDG22, Proposition 4.1]) which displays the ω -integrals of the terms to be compared. For equations $F_k^{[i]}$ that do intersect horizontal nodes (thus only at level i by construction) the difference $F_k^{[i]}(c_j^{[s]}, \text{all } (j, s)) - F_k^{[i]}(\widehat{c}_j^{[s]}, \text{all } (j, s))$ vanishes thanks to the proportionality of the periods of horizontal nodes in an \mathcal{H} -equivalence class and since on $\overline{\mathcal{H}}$ the equation $H_k^{[i]}$ holds.

By the very definition of defining equation its periods evaluate to zero, explaining the right arrow in the top row of the following diagram and showing that ev is well-defined on the quotient.

$$\begin{array}{ccccccc} 0 & \longrightarrow & \mathcal{K}_{\text{Eq}} & \longrightarrow & \text{Eq} \otimes \mathcal{O}_{\overline{\mathcal{H}}}(-1) & \longrightarrow & 0 \\ & & \downarrow & & \downarrow & & \downarrow \\ 0 & \longrightarrow & \bigoplus_{i=0}^L t_{[-i]} \cdot \Omega_B^{[i]}|_{\overline{\mathcal{H}}} & \xrightarrow{\psi} & (\overline{\mathcal{H}}_{\text{rel},B}^1)^\vee \otimes \mathcal{O}_{\overline{\mathcal{H}}}(-1) & \xrightarrow{\text{ev}_B} & \mathcal{O}_{\overline{\mathcal{H}}} \longrightarrow 0 \\ & & \downarrow q_\Omega & & \downarrow & & \downarrow \\ 0 & \longrightarrow & \bigoplus_{i=0}^L t_{[-i]} \cdot \Omega^{[i]} & \longrightarrow & (\overline{\mathcal{H}}_{\text{rel}}^1)^\vee \otimes \mathcal{O}_{\overline{\mathcal{H}}}(-1) & \xrightarrow{\text{ev}} & \mathcal{O}_{\overline{\mathcal{H}}} \longrightarrow 0. \end{array}$$

Here we used the abbreviations

$$\Omega_B^{[i]} = \Omega_{i,B}^{\text{hor}}(\log) \oplus \Omega_{i,B}^{\text{lev}}(\log) \oplus \Omega_{i,B}^{\text{rel}}, \quad \Omega^{[i]} = \Omega_i^{\text{hor}}(\log) \oplus \Omega_i^{\text{lev}}(\log) \oplus \Omega_i^{\text{rel}}.$$

The surjectivity of q_Ω follows from the definition of the summands in (28). It requires justification that the image is not larger, since the derivatives of the local equations of \mathcal{H} do not respect the direct sum decomposition 27. More precisely we claim that \mathcal{K}_{Eq} is generated by two kinds of equations. Before analyzing them, note that the log periods satisfy by construction an estimate of the form

$$(51) \quad \widehat{c}_j^{[-i]} - c_j^{[-i]} = \sum_{s>i} \frac{t_{[s]}}{t_{[i]}} \widehat{E}_{j,i}^{[-s]}$$

with some error term $\widehat{E}_{j,i}^{[-k]}$ depending on the variables $c_j^{[-s]}$ on the lower level $-s$ as in (25).

For each of the equations (29) the corresponding linear function $L_k^{[i]}$ in the variables $c_j^{[i]}$ is an element in Eq. We use the comparisons (51) and (25) to compute its ψ -preimage in \mathcal{K}_{Eq} via (50). It is $t_{[-i]}$ times the corresponding expression in the $\widehat{c}_j^{[i]}$ plus a linear

combination of the terms $t_{[-s]} \widehat{E}_{j,i}^{[s]}$. The quotient by such a relation does not yield any quotient class beyond those in $\bigoplus_{i=0}^s t_{[-i]} \cdot \Omega^{[i]}$.

We write the other equations (30) as $(\mathbf{q}^{[i]})^{J_{1,k}-J_{2,k}} = 1$ since we are interested in torus-invariant differential forms and can compute on the boundary complement. Consider $d \log$ of this equation. Under the first map ψ of the Euler sequence

$$(52) \quad dq_j^{[i]}/q_j^{[i]} = d \log(q_j^{[i]}) = d \left(2\pi I \frac{b_j^{[i]}}{a_j^{[i]}} \right) \mapsto \frac{2\pi I}{a_j^{[i]}} \left(\beta_j^{[i]} - \frac{b_j^{[i]}}{a_j^{[i]}} \alpha_j^{[i]} \right) \otimes \omega$$

Recall from summary of [BDG22] in Section 3.5 that the functions $a_j^{[i]}$ for all j where $(v_1, \dots, v_{N(i)-h(i)}) := J_{1,k} - J_{2,k}$ is non-zero are rational multiples of each other. Note moreover that $\beta_j^{[i]} - \frac{b_j^{[i]}}{a_j^{[i]}} \alpha_j^{[i]} = \beta_j^{[i]} - \frac{1}{2\pi I} \log(q_j^{[i]}) \alpha_j^{[i]}$ is the Deligne extension of $\beta_j^{[i]}$ across all the boundary divisors that stem from horizontal nodes at level i . For the full Deligne extension $\widehat{\beta}_j^{[i]}$ the correction terms for the lower level nodes have to be added. Together with (26) we deduce that the ψ -image of

$$\sum_{m=1}^{h(i)} v_m a_m^{[i]} \frac{dq_m^{[i]}}{\widetilde{q}_m^{[i]}} = \sum_{m=1}^{h(i)} v_j c_{j(m)}^{[i]} \frac{dq_m^{[i]}}{\widetilde{q}_m^{[i]}}$$

differs from the element in Eq responsible for the equation $H_k^{[i]}$ only by terms from lower level s , which come with a factor $t_{[-s]}$. In this equation used that $a_m^{[i]} = c_{j(m)}^{[i]}$ for an appropriate $j(m)$. Since $c_{j(m)}^{[i]}$ is close to $t_{[-i]} \widetilde{c}_{j(m)}^{[i]}$, compare with (25) this element indeed belongs to the kernel of ψ as claimed in the commutative diagram. The quotient by such a relation does not yield any quotient class beyond those above either. Since the (30) and (29) correspond to a basis (in fact: in reduced row echelon form) of Eq, this completes the proof. \square

PROOF OF THEOREM 5.2. Uses that the summands of $\mathcal{K}|_U$ are, up to t -powers, the decomposition of the logarithmic tangent sheaf by Proposition 3.5. \square

COROLLARY 5.3. *The Chern character and the Chern polynomial of the kernel \mathcal{K} of the Euler sequence are given by*

$$\text{ch}(\mathcal{K}) = N e^{\xi_{\mathcal{H}}} - 1 \quad \text{and} \quad c(\mathcal{K}) = \sum_{i=0}^{N-1} \binom{N}{i} \xi_{\mathcal{H}}^i.$$

PROOF. As a Deligne extension of a local system, $(\overline{\mathcal{H}}_{\text{rel},B}^1)^\vee|_{\overline{\mathcal{H}}}$ has trivial Chern classes except for c_0 . By construction, the pullback of the sheaf Eq to an allowable modification (toric resolution with normal crossing boundary, see the proof of Proposition 2.1) is the Deligne extension of a local system. It follows that all Chern classes but c_0 of this pullback vanish and by push-full this holds for Eq, too. The Chern class vanishing for $(\mathcal{H}_{\text{rel}}^1)^\vee$ and the corollary follows. \square

To start with the computation of \mathcal{C} , we will also need an infinitesimal thickening of the boundary divisor $D_\Gamma^{\mathcal{H}}$, namely we define $D_{\Gamma,\bullet}^{\mathcal{H}}$ to be its ℓ_Γ -th thickening, the non-reduced substack of $\overline{\mathcal{H}}$ defined by the ideal $\mathcal{I}_{D_\Gamma^{\mathcal{H}}}^{\ell_\Gamma}$. We will factor the above inclusion using the notation

$$\mathfrak{i}_\Gamma = \mathfrak{i}_{\Gamma,\bullet} \circ j_{\Gamma,\bullet}: D_\Gamma^{\mathcal{H}} \xrightarrow{j_{\Gamma,\bullet}} D_{\Gamma,\bullet}^{\mathcal{H}} \xrightarrow{\mathfrak{i}_{\Gamma,\bullet}} \overline{\mathcal{H}}.$$

We will denote by $\mathcal{L}_{\Gamma,\bullet}^\top = (j_{\Gamma,\bullet})_*(\mathcal{L}_\Gamma^\top)$ and $\mathcal{E}_{\Gamma,\bullet}^\top = (j_{\Gamma,\bullet})_*(\mathcal{E}_\Gamma^\top)$ the push-forward to the thickening of the vector bundles defined in (44) and (34).

PROPOSITION 5.4. *The cokernel of (48) is given by*

$$(53) \quad \mathcal{C} = \bigoplus_{\Gamma \in \text{LG}_1(B)} \mathcal{C}_\Gamma \quad \text{where} \quad \mathcal{C}_\Gamma = (\mathfrak{i}_{\Gamma,\bullet})_*(\mathcal{E}_{\Gamma,\bullet}^\top \otimes (\mathcal{L}_{\Gamma,\bullet}^\top)^{-1}).$$

Moreover, there is an equality of Chern characters

$$\mathrm{ch}\left((i_{\Gamma,\bullet})_*(\mathcal{E}_{\Gamma,\bullet}^\top \otimes (\mathcal{L}_{\Gamma,\bullet}^\top)^{-1})\right) = \mathrm{ch}\left((i_{\Gamma})_*\left(\bigoplus_{j=0}^{\ell_{\Gamma}-1} \mathcal{N}_{\Gamma}^{\otimes -j} \otimes \mathcal{E}_{\Gamma}^\top \otimes (\mathcal{L}_{\Gamma}^\top)^{-1}\right)\right).$$

PROOF. The second part of the statement is justified by the original argument in [CMZ19, Lemma 9.3].

The first part of the statement follows since, from Theorem 5.1 we know that

$$\mathcal{K}|_U = \bigoplus_{i=-L}^0 \prod_{j=1}^{-i} t_j^{\ell_j} \cdot \left(\Omega_i^{\mathrm{hor}}(\log) \oplus \Omega_i^{\mathrm{lev}}(\log) \oplus \Omega_i^{\mathrm{rel}}\right)$$

and from Proposition 3.5 we also know that

$$(54) \quad (\mathcal{E}_{\mathcal{H}} \otimes \mathcal{L}_{\mathcal{H}}^{-1})|_U = \bigoplus_{i=-L}^0 \prod_{j=1}^L t_j^{\ell_j} \cdot \left(\Omega_i^{\mathrm{hor}}(\log) \oplus \Omega_i^{\mathrm{lev}}(\log) \oplus \Omega_i^{\mathrm{rel}}\right)$$

where Γ is an arbitrary level graph with L levels below zero and U is a small neighborhood of a point in $D_{\Gamma}^{\mathcal{H},\circ}$. \square

We can finally compute

PROPOSITION 5.5. *The Chern character of the twisted logarithmic cotangent bundle $\mathcal{E}_{\mathcal{H}} \otimes \mathcal{L}_{\mathcal{H}}^{-1}$ can be expressed in terms of the twisted logarithmic cotangent bundles of the top levels of non-horizontal divisors as*

$$\mathrm{ch}(\mathcal{E}_{\mathcal{H}} \otimes \mathcal{L}_{\mathcal{H}}^{-1}) = Ne^{\xi} - 1 - \sum_{\Gamma \in \mathrm{LG}_1(\mathbb{B})} i_{\Gamma*} \left(\mathrm{ch}(\mathcal{E}_{\Gamma}^\top) \cdot \mathrm{ch}(\mathcal{L}_{\Gamma}^\top)^{-1} \cdot \frac{(1 - e^{-\ell_{\Gamma} c_1(\mathcal{N}_{\Gamma})})}{c_1(\mathcal{N}_{\Gamma})} \right).$$

PROOF. The proof [CMZ19, Prop. 9.5] works in the same way, since the only tool that was used is the Grothendieck-Riemann-Roch Theorem applied to the map $f = i_{\Gamma}$, which is still a regular embedding. \square

PROOF OF THEOREM 1.1 AND THEOREM 1.2. The final formulas of the full twisted Chern character, Chern polynomials and Euler characteristic follow from the arguments used for Abelian strata in [CMZ19, Section 9], since they were purely formal starting from the previous proposition. The relevant inputs needed are the compatibility statement of Lemma 3.9, the formula for pulling back normal bundles given in Lemma 4.6 and Corollary 3.3. \square

PROOF OF THEOREM 1.3. A formal consequence of Theorem 1.2 and the rewriting in [CMZ22, Theorem 9.10] (with the reference to [CMZ22, Proposition 4.9] replaced by Lemma 3.9) is

$$(55) \quad \chi(\mathcal{H}) = (-1)^d \sum_{L=0}^d \sum_{\Gamma \in \mathrm{LG}_L(\mathcal{H})} N_{\Gamma}^\top \cdot \ell_{\Gamma} \cdot \int_{D_{\Gamma}^{\mathcal{H}}} \prod_{i=-L}^0 (\xi_{\Gamma,\mathcal{H}}^{[i]})^{d_{\Gamma}^{[i]}}$$

We now use Lemma 3.10 to convert integrals on a boundary component into the product of integrals of its the level strata. \square

6. Example: Euler characteristic of the eigenform locus

For a non-square $D \in \mathbb{N}$ with $D \equiv 0$ or $1 \pmod{4}$ let

$$\Omega E_D(1,1) \subseteq \Omega \mathcal{M}_{2,2}(1,1) \quad \text{and} \quad \Omega W_D \subseteq \Omega \mathcal{M}_{2,1}(2)$$

be the eigenform loci for real multiplication by \mathcal{O}_D in the given stratum, see [McM03], [Cal04], [McM07a] for the first proofs that these loci are linear submanifolds and some background. We define $E_D := \mathbb{P}\Omega E_D(1,1)$ as the projectivized eigenform locus. Associating with the curve its Jacobian, the projectivized eigenform locus maps to the *Hilbert modular surface*

$$X_D = \mathbb{H} \times \mathbb{H} / \mathrm{SL}(\mathcal{O}_D \oplus \mathcal{O}_D^{\vee}).$$

Inside X_D let $P_D \subseteq X_D$ denote the *product locus*, i.e. the curve consisting of those surfaces which are polarized products of elliptic curves. The *Weierstrass curve* W_D is defined to be the image of ΩW_D . It is contained in the complement $X_D \setminus P_D$.

The goal of this section is to provide references and details for the proof of Theorem 1.4 and in particular (19). The numerical input is

$$\chi(X_D) = 2\zeta(-1) \quad \text{and} \quad \chi(P_D) = -\frac{5}{2}\chi(X_D) = -5\zeta(-1),$$

where $\zeta = \zeta_{\mathbb{Q}(\sqrt{D})}$ is the Dedekind zeta function. The first formula is due to Siegel [Sie36], see also [Gee88, Theorem IV.1.1], the second is given in [Bai07, Theorem 2.22]




FIGURE 1. The boundary divisors of the eigenform locus E .

To apply Theorem 1.3 to the linear manifold E_D we need to list the boundary strata without horizontal curves. This list consists of two divisorial strata only, given in Figure 1, namely the product locus and the Weierstrass locus. To justify the coefficients in (19) we need:

LEMMA 6.1. *The top-powers of ξ on the respective level strata evaluate to*

$$\int_E \xi^2 = 0, \quad \int_{D_{\Gamma_P}^\perp} 1 = 1, \quad \text{and} \quad \int_{D_{\Gamma_W}^\perp} 1 = 1.$$

PROOF. The first integral is an application of Corollary 4.3. For the second note that there is unique differential up to scale of type $(1, 1, -2, -2)$ on a \mathbb{P}^1 with vanishing residues, the third is obvious. \square

The proof is completed by noticing that that automorphism groups in Theorem 1.3 are trivial and that all three prong-matchings for Γ_W are reachable since they belong to one orbit of the prong rotation group.

7. Strata of k -differentials

Our goal here is to prove Corollary 1.5 that gives a formula for the Euler characteristic of strata $\mathbb{P}\Omega^k \mathcal{M}_{g,n}(\mu)$ of k -differentials. Those strata can be viewed as linear submanifolds of strata of Abelian differentials $\mathbb{P}\Omega \mathcal{M}_{\hat{g}, \hat{n}}(\hat{\mu})$ via the canonical covering construction and thus Theorem 1.3 applies. This is however of little practical use as we do not know the classes of k -differential strata in $\mathbb{P}\Omega \mathcal{M}_{\hat{g}, \hat{n}}(\hat{\mu})$. However, we do know their classes in $\overline{\mathcal{M}}_{g,n}$ via Pixton's formulas for the DR-cycle ([HS21], [BHPSS20]). As a consequence the formula in Corollary 1.5 can be implemented, and the `diffstrata` package does provide such an implementation. In this section we thus recall the basic definitions of the compactification and collect all the statements to perform evaluation of expressions in the tautological rings on strata of k -differentials.

7.1. Compactification of strata of k -differentials. We want to work on the multi-scale compactification $\overline{\mathcal{Q}} := \overline{\mathcal{Q}}_k := \mathbb{P}\Xi^k \overline{\mathcal{M}}_{g,n}(\mu)$ of the space of k -differentials. As topological space this compactification was given in [CMZ19], reviewing the plumbing construction from [BCGGM19b], but without giving the stack structure. Here we consider a priori the compactification of Section 3. We give some details, describing auxiliary stacks usually by giving \mathbb{C} -valued points and morphisms, from which the reader can easily deduce the notion of families following the procedure in [BCGGM19b]. From this description it should become clear that the two compactifications, the one of Section 3 and [CMZ19], agree up to explicit isotropy groups (see Lemma 7.2). In particular the compactification $\overline{\mathcal{Q}}_k$ is smooth. This follows also directly from the definition of Section 3, since the only potential singularities are at the horizontal nodes. There however the local equations (30) simply compare monomials (with exponent one), the various q -parameters of the k preimages of a horizontal node.

We start by recalling notation for the canonical k -cover *in the primitive case*. Let X be a Riemann surface of genus g and let q be a *primitive* meromorphic k -differential of type $\mu = (m_1, \dots, m_n)$, i.e. not the d -th power of a k/d -differential for any $d > 1$. This datum defines (see e.g. [BCGGM19a, Section 2.1]) a connected k -fold cover $\pi: \widehat{X} \rightarrow X$ such that $\pi^*q = \omega^k$ is the k -power of an abelian differential. This differential ω is of type

$$\widehat{\mu} := \left(\underbrace{\widehat{m}_1, \dots, \widehat{m}_1}_{g_1 := \gcd(k, m_1)}, \underbrace{\widehat{m}_2, \dots, \widehat{m}_2}_{g_2 := \gcd(k, m_2)}, \dots, \underbrace{\widehat{m}_n, \dots, \widehat{m}_n}_{g_n := \gcd(k, m_n)} \right),$$

where $\widehat{m}_i := \frac{k+m_i}{\gcd(k, m_i)} - 1$. (Here and throughout marked points of order zero may occur.)

We let $\widehat{g} = g(\widehat{X})$ and $\widehat{n} = \sum_i \gcd(k, m_i)$. The type of the covering determines a natural subgroup $S_{\widehat{\mu}} \subset S_{\widehat{n}}$ of the symmetric group that allows only the permutations of each the $\gcd(k, m_i)$ points corresponding to a preimage of the i -th point. In the group $S_{\widehat{\mu}}$ we fix the element

$$(56) \quad \tau_0 = \left(12 \cdots g_1 \right) \left(g_1 + 1 \ g_1 + 2 \cdots g_1 + g_2 \right) \cdots \left(1 + \sum_{i=1}^{n-1} g_i \cdots \sum_{i=1}^n g_n \right),$$

i.e. the product of cycles shifting the g_i points in the π -preimage of each point in \mathbf{z} . We fix a primitive k -th root of unity ζ_k throughout.

We consider the stack $\Omega\mathcal{H}_k := \Omega\mathcal{H}_k(\widehat{\mu})$ whose points are

$$(57) \quad \{(\widehat{X}, \widehat{\mathbf{z}}, \omega, \tau) : \tau \in \text{Aut}(\widehat{X}), \text{ord}(\tau) = k, \tau^*\omega = \zeta_k\omega, \tau|_{\widehat{\mathbf{z}}} = \tau_0\}.$$

Families are defined in the obvious way. Morphisms are morphisms of the underlying pointed curves that commute with τ . Since the marked points determine the differential up to scale, the differentials are identified by the pullback of morphisms up to scale. Commuting with τ guarantees that morphisms descend to the quotient curves by $\langle \tau \rangle$ (for a morphism f to descend, a priori $f\tau f^{-1} = \tau^a$ for some a would be sufficient, but the action on ω implies that in fact $a = 1$). It will be convenient to label the tuple of points $\widehat{\mathbf{z}}$ by tuples (i, j) with $i = 1, \dots, n$ and $j = 1, \dots, \gcd(k, m_i)$. There is a natural forgetful map $\Omega\mathcal{H}_k \rightarrow \Omega\mathcal{M}_{\widehat{g}, \widehat{n}}$ and period coordinates (say, after providing both sides locally with a Teichmüller marking) show that this map is the normalization of its image and the image is cut out by linear equations, i.e. that $\Omega\mathcal{H}_k$ is a linear submanifold as defined in Section 3.1.

The subgroup

$$(58) \quad G = \left\langle \left(12 \cdots g_1 \right), \left(g_1 + 1 \ g_1 + 2 \cdots g_1 + g_2 \right), \dots, \left(1 + \sum_{i=1}^{n-1} g_i \cdots \sum_{i=1}^n g_n \right) \right\rangle \subset S_{\widehat{\mu}}$$

generated by the cycles that τ_0 is made from acts on $\Omega\mathcal{H}_k$ and on the projectivization \mathcal{H}_k . We denote the quotient of the latter by $\mathcal{H}_k^{\text{mp}} := \mathcal{H}_k/G$, where the upper index is an abbreviation of *marked (only) partially*.

Since τ has ω as eigendifferential, its k -th power naturally descends to (projectivized) k -differential $[q]$ on the quotient $X = \widehat{X}/\langle \tau \rangle$, which is decorated by the marked points \mathbf{z} , the images of $\widehat{\mathbf{z}}$.

We denote by \mathcal{Q} the stack with the same underlying set as $\mathcal{H}_k^{\text{mp}}$, but where morphisms are given by the morphisms of $(X/\langle \tau \rangle, \mathbf{z}, [q])$ in $\mathbb{P}\Omega^k\mathcal{M}_{g,n}(\mu)$. Written out on curves, a morphism in \mathcal{Q} is a map $f: \widehat{X}/\langle \tau \rangle \rightarrow \widehat{X}'/\langle \tau' \rangle$, such that there exists a commutative diagram

$$(59) \quad \begin{array}{ccc} \widehat{X} & \xrightarrow{\quad g \quad} & \widehat{X}' \\ \downarrow & & \downarrow \\ X = \widehat{X}/\langle \tau \rangle & \xrightarrow{\quad f \quad} & X' = \widehat{X}'/\langle \tau' \rangle. \end{array}$$

If two such maps g exist, they differ by pre- or postcomposition with an automorphism of \widehat{X} resp. \widehat{X}' . Via the canonical cover construction, the stack \mathcal{Q} is isomorphic to $\mathbb{P}\Omega^k\mathcal{M}_{g,n}(\mu)$. The non-uniqueness of g exhibits $\mathcal{H}_k^{\text{mp}} = \mathcal{Q}/\langle \tau \rangle$ as the quotient stack by a group of order k , acting trivially.

As in Section 3, we denote by $\overline{\Omega\mathcal{H}_k} := \overline{\Omega\mathcal{H}_k}(\mu)$ the normalization of the closure of $\Omega\mathcal{H}_k$ in $\Xi\overline{\mathcal{M}}_{\widehat{g},\widehat{n}}(\mu)$ and let $\overline{\mathcal{H}_k} := \overline{\mathcal{H}_k}(\mu)$ be the corresponding projectivizations. We next describe the boundary strata of $\overline{\mathcal{H}_k}$. These are indexed by enhanced level graphs $\widehat{\Gamma}$ together with an $\langle\tau\rangle$ -action on them. We will leave the group action implicit in our notation. The following lemma describes the objects parametrized by the boundary components $D_{\widehat{\Gamma}}^{\mathcal{H}_k}$ (using the notation from Section 3) of the compactification $\overline{\mathcal{H}_k}$.

LEMMA 7.1. *A point in the interior of the boundary stratum $D_{\widehat{\Gamma}}^{\mathcal{H}_k}$ is given by a tuple*

$$\{(\widehat{X}, \widehat{\Gamma}, \widehat{\mathbf{z}}, [\boldsymbol{\omega}], \boldsymbol{\sigma}, \tau) : \tau \in \text{Aut}(\widehat{X}), \quad \text{ord}(\tau) = k, \quad \tau^*\boldsymbol{\omega} = \zeta_k\boldsymbol{\omega}, \quad \tau|_{\widehat{\mathbf{z}}} = \tau_0\}$$

where $(\widehat{X}, \widehat{\Gamma}, \widehat{\mathbf{z}}, [\boldsymbol{\omega}], \boldsymbol{\sigma}) \in \mathbb{P}\Xi\overline{\mathcal{M}}_{\widehat{g},\widehat{n}}(\widehat{\mu})$ is a multi-scale differential and where moreover the prong-matching $\boldsymbol{\sigma}$ is equivariant with respect to the action of $\langle\tau\rangle$.

The equivariance of prong-matching requires an explanation: Suppose x_i and y_i are standard coordinates near the node corresponding to an edge e of Γ , so that the prong-matching at e is given by $\sigma_e = \frac{\partial}{\partial x_i} \otimes -\frac{\partial}{\partial y_i}$ (compare [BCGGM19b, Section 5] for the relevant definitions). Then τ^*x_i and τ^*y_i are standard coordinates near $\tau(e)$. We say that a global prong-matching $\boldsymbol{\sigma} = \{\sigma_e\}_{e \in E(\widehat{\Gamma})}$ is *equivariant* if $\sigma_{\tau(e)} = \frac{\partial}{\partial \tau^*x_i} \otimes -\frac{\partial}{\partial \tau^*y_i}$ for each edge e .

PROOF. The necessity of the conditions on the boundary points is obvious from the definition in (57), except for the prong-matching equivariance. This follows from the construction of the induced prong-matching in a degenerating family in [BCGGM19b, Proposition 8.4] and applying τ to it.

Conversely, given $(\widehat{X}, \widehat{\Gamma}, \widehat{\mathbf{z}}, [\boldsymbol{\omega}], \boldsymbol{\sigma}, \langle\tau\rangle)$ as above with equivariant prong-matchings, we need to show that it is in the boundary of \mathcal{H}_k . This is achieved precisely by the equivariant plumbing construction given in [BCGGM19a]. \square

The group G still acts on the compactification $\overline{\Omega\mathcal{H}_k}$ and on its projectivization $\overline{\mathcal{H}_k}$. As above we denote the quotient by $\overline{\mathcal{H}_k}^{\text{mp}} = \overline{\mathcal{H}_k}/G$ to indicate that the points $\widehat{\mathbf{z}}$ are now marked only partially. By Lemma 7.1 we may construct $\overline{\mathcal{Q}}$ just as in the uncompactified case.

The map $\overline{\mathcal{H}_k}^{\text{mp}} \rightarrow \overline{\mathcal{Q}}$ is in general non-representable due to the existence of additional automorphisms of objects in $\overline{\mathcal{H}_k}^{\text{mp}}$. This resembles the situation common for Hurwitz spaces, where the target map is in general non-representable, too. We denote by $d : \overline{\mathcal{H}_k}^{\text{mp}} \rightarrow \overline{\mathcal{Q}}$ the composition of the maps.

7.2. Generalized strata of k -differentials. Our notion of generalized strata is designed for recursion purposes so that the extraction of levels of a boundary stratum of $\overline{\mathcal{Q}}$ is an instance of a generalized stratum (of k -differentials). This involves incorporating disconnected strata, differentials that are non-primitive on some components, and residue conditions. Moreover, we aim for a definition of a space of k -fold covers on which the group G acts, to match with the previous setup. The key is to record which of the marked points is adjacent to which component, an information that is obviously trivial in the case of primitive k -differentials.

A map $\mathcal{A} : \widehat{\mathbf{z}} \rightarrow \pi_0(\widehat{X})$ that records which marked point is adjacent to which component of \widehat{X} is called an *adjacency datum*. (Such an adjacency datum is equivalent to specifying a one-level graph of a generalized stratum, which is indeed the information we get when we extract level strata.) The subgroup G from (58) acts on the triples $(\widehat{X}, \widehat{\mathbf{z}}, \mathcal{A})$ of pointed stable curves with adjacency map by acting simultaneously on $\widehat{\mathbf{z}}$ and on \mathcal{A} by precomposition. For a fixed adjacency datum \mathcal{A} we consider the stack $\Omega\widetilde{\mathcal{H}_k}(\widehat{\mu}, \mathcal{A})$ whose points are

$$\{(\widehat{X}, \widehat{\mathbf{z}}, \boldsymbol{\omega}, \tau) : (\widehat{X}, \widehat{\mathbf{z}}) \text{ have adjacency } \mathcal{A}, \quad \tau \in \text{Aut}(\widehat{X}), \\ \text{ord}(\tau) = k, \quad \tau^*\boldsymbol{\omega} = \zeta_k\boldsymbol{\omega}, \quad \tau|_{\widehat{\mathbf{z}}} = \tau_0, \}$$

We denote by $\Omega\mathcal{H}_k(\widehat{\mu}, [\mathcal{A}]) := G \cdot \Omega\widetilde{\mathcal{H}_k}(\widehat{\mu}, \mathcal{A})$ the G -orbit of this space.

A *residue condition* is given by a τ -invariant partition $\lambda_{\mathfrak{R}}$ of a subset of the set $H_p \subseteq \{1, \dots, \widehat{n}\}$ of marked points such that $\widehat{m}_i < -1$. We often also call the associated linear subspace

$$\mathfrak{R} := \left\{ (r_i)_{i \in H_p} \in \mathbb{C}^{H_p} : \sum_{i \in \lambda} r_i = 0 \text{ for all } \lambda \in \lambda_{\mathfrak{R}} \right\}.$$

the residue condition. This space will typically not be G -invariant. We denote by $\Omega\mathcal{H}_k^{\mathfrak{R}}(\widehat{\mu}, \mathcal{A}) \subseteq \Omega\mathcal{H}_k(\widehat{\mu}, \mathcal{A})$ the subset where for each $R \in \mathfrak{R}$ the residues of $\widehat{\omega}$ at all the points $z_i \in R$ add up to zero. If $(\widehat{X}, \widehat{\mathbf{z}}, \omega, \tau)$ is contained in $\Omega\mathcal{H}_k^{\mathfrak{R}}(\widehat{\mu}, \mathcal{A})$, then $g \cdot (\widehat{X}, \widehat{\mathbf{z}}, \omega, \tau)$ is contained in $\Omega\mathcal{H}_k^{g \cdot \mathfrak{R}}(\widehat{\mu}, g \cdot \mathcal{A})$ for any $g \in G$. That is, the G -action simultaneously changes the residue condition and the adjacency datum. We denote by $[\mathfrak{R}, \mathcal{A}]$ the G -orbit of this pair and use the abbreviation

$$(60) \quad \Omega\mathcal{H}_k^{[\mathfrak{R}, \mathcal{A}]} := G \cdot \Omega\mathcal{H}_k^{\mathfrak{R}}(\widehat{\mu}, \mathcal{A})$$

for the G -orbit of the spaces, $\widehat{\mu}$ being tacitly fixed throughout.

As above, we denote by $\mathcal{H}_k^{[\mathfrak{R}, \mathcal{A}]}$ the projectivization of $\Omega\mathcal{H}_k^{[\mathfrak{R}, \mathcal{A}]}$ and by $\mathcal{H}_k^{\mathfrak{R}, \text{mp}} := \mathcal{H}_k^{[\mathfrak{R}, \mathcal{A}]} / G$ the G -quotient, dropping the information about adjacency and the connected components to ease notation. Finally, we denote by $\mathcal{Q}^{\mathfrak{R}}$ the stack with the same underlying set as $\mathcal{H}_k^{\mathfrak{R}, \text{mp}}$ and with morphisms defined in the same way as above for \mathcal{Q} . Recall that the curves in $\mathcal{Q}^{\mathfrak{R}}$ may be disconnected. We call such a stratum with possibly disconnected curves and residue conditions a *generalized stratum of k -differentials*. Since $\mathcal{H}_k^{[\mathfrak{R}, \mathcal{A}]}$ is a linear submanifold, we can still compactify them as before and a version of Lemma 7.1 with adjacency data still holds.

We will now compute the degree of the map d from the linear submanifolds to the strata of k -differential. Our definition of generalized strata of k -differentials makes the degree of this map the same in the usual and in the generalized case.

LEMMA 7.2. *The map $d : \overline{\mathcal{H}}_k^{[\mathfrak{R}, \mathcal{A}]} \rightarrow \overline{\mathcal{Q}}^{\mathfrak{R}}$ is proper, quasi-finite, unramified and of degree*

$$\deg(d) = \frac{1}{k} \prod_{m_i \in \mu} \gcd(m_i, k).$$

PROOF. The degree is a consequence of being composed of a quotient by a group of order $|G| = \prod_{m_i \in \mu} \gcd(m_i, k)$ and the non-representable inverse of a quotient by a group of order k .

The map is unramified as both quotient maps are unramified. \square

7.3. Decomposing boundary strata. Having constructed strata of k -differentials, we now want to decompose their boundary strata again as a product of generalized strata of k -differentials and argue recursively. In fact, the initial stratum should be a generalized stratum $\overline{\mathcal{Q}}^{\mathfrak{R}}$, thus coming with its own residue condition, but we suppress this in our notation, focusing on the new residue condition that arise when decomposing boundary strata. Here 'decomposition' of the boundary strata should be read as a construction of a space finitely covering both of them, as given by the following diagram,

$$(61) \quad \begin{array}{ccc} & D_{\pi}^{\circ, \mathcal{H}_k, s} & \\ & \swarrow p_{\pi} \quad \searrow c_{\pi} & \\ \mathcal{H}_k(\pi) := \prod_{i=0}^{-L} \mathcal{H}_k(\pi_{[i]}) \hookrightarrow \text{Im}(p_{\pi}) & & D_{\pi}^{\circ, \mathcal{H}_k} \\ \mathbf{d}_{\pi} \downarrow & & \downarrow d_{\pi} \\ \mathcal{Q}(\pi) := \prod_{i=0}^{-L} \mathcal{Q}(\pi_{[i]}) & & D_{\pi}^{\circ, \mathcal{Q}}, \end{array}$$

whose notation we now start to explain. Note that the diagram is for the open boundary strata throughout, since we mainly need the degree all these maps as in Lemma 3.6 (the existence of a similar diagram over the completions follows as at the beginning of Section 3.2).

We denote by $\widehat{\Gamma}$ the level graphs indexing the boundary strata of $\mathbb{P}\Xi\overline{\mathcal{M}}_{\widehat{g},\widehat{n}}(\widehat{\mu})$ and thus of $\overline{\mathcal{H}}_k$. Following our general convention for strata their legs are labeled, but not the edges. In $\overline{\mathcal{H}}_k^{\text{mp}}$ the leg-marking is only well-defined up to the action of G . A graph with such a marking is said to be *marked (only) partially* and denoted by $\widehat{\Gamma}_{\text{mp}}$. Even though curves in $\overline{\mathcal{H}}_k$ are marked (and not only marked up to the action of G), the boundary strata of $\overline{\mathcal{H}}_k$ are naturally indexed by partially marked graphs as well: If $\widehat{\Gamma}$ is the dual graph of one stable curve in the boundary of $\overline{\mathcal{H}}_k$, then for all $g \in G$ the graph $g \cdot \widehat{\Gamma}$ is the dual graph of another stable curve in the boundary of $\overline{\mathcal{H}}_k$. The existence of τ implies that level graphs $\widehat{\Gamma}$ at the boundary of $\overline{\mathcal{H}}_k$ come with the quotient map by this action. To each boundary stratum of $\overline{\mathcal{Q}}$ we may thus associate a k -cyclic covering of graphs $\pi : \widehat{\Gamma}_{\text{mp}} \rightarrow \Gamma$ (see [CMZ19, Section 2] for the definitions of such covers). We denote the corresponding (open) boundary strata by $D_\pi^{\circ,\mathcal{Q}} \subset \overline{\mathcal{Q}}$ and the (open) boundary strata corresponding to such a G -orbit of graphs by $D_\pi^{\circ,\mathcal{H}_k} \subset \overline{\mathcal{H}}_k$. The map $d_\pi : D_\pi^{\circ,\mathcal{H}_k} \rightarrow D_\pi^{\circ,\mathcal{Q}}$ is the restriction of the map $d : \overline{\mathcal{H}}_k \rightarrow \overline{\mathcal{Q}}$.

Next we construct the commensurability roof just as in (31), though for each $\widehat{\Gamma}$ in the G -orbit separately, so that $D_\pi^{\circ,\mathcal{H}_k,s}$ is the disjoint union of a G -orbit of the roofs in (31).

Next we define the spaces $\mathcal{H}_k(\pi_{[i]})$. Consider the linear submanifolds of generalized strata of k -differentials with signature and adjacency datum given by the i -th level of one marked representative $\widehat{\Gamma}$ of $\widehat{\Gamma}_{\text{mp}}$ (the resulting strata are independent of the choice of a representative). Their product defines the image $\text{Im}(p_\pi)$. For every level i , consider the orbit under $G(\mathcal{H}_k(\pi_{[i]}))$, where $G(\mathcal{H}_k(\pi_{[i]}))$ is the group as in (58) for the i -th level, of the linear submanifolds we extracted from the levels. We define $\mathcal{H}_k(\pi_{[i]})$ to be these orbits, which in particular are then linear submanifolds associated to generalized strata of k -differentials as we defined them above. We can hence consider, for every level, the morphism given by the quotient by $G(\mathcal{H}_k(\pi_{[i]}))$ composed with the non-representable map that kills the $\langle \tau \rangle$ -isotropy groups at each level and denote by $\mathcal{Q}(\pi_{[i]})$ its image, which is called the generalized stratum of k -differentials at level i . The map \mathbf{d}_π in diagram 61 is just a product of maps like the map d above, thus Lemma 7.2 immediately implies:

LEMMA 7.3. *The degree of the map \mathbf{d}_π in the above diagram (61) is*

$$\deg(\mathbf{d}_\pi) = \frac{1}{k^{L+1}} \prod_{i=1}^n \gcd(m_i, k) \prod_{e \in E(\Gamma)} \gcd(\kappa_e, k)^2$$

where κ_e is the k -enhancement of the edge e .

We recall Lemma 3.6 and compute explicitly the coefficients appearing in our setting here. Note that the factor $|\text{Aut}_{\mathcal{H}}(\Gamma)|$ there should be called $|\text{Aut}_{\mathcal{H}_k}(\widehat{\Gamma})|$ in the notation used in this section.

LEMMA 7.4. *The ratio of the degrees of the topmost maps in (61) is*

$$\frac{\deg(p_\pi)}{\deg(c_\pi)} = \frac{K_{\widehat{\Gamma}}^{\mathcal{H}_k}}{|\text{Aut}_{\mathcal{H}_k}(\widehat{\Gamma})| \cdot \ell_{\widehat{\Gamma}}}$$

where the number of reachable prong-matchings is given by

$$K_{\widehat{\Gamma}}^{\mathcal{H}_k} = \prod_{e \in E(\Gamma)} \frac{\kappa_e}{\gcd(\kappa_e, k)}$$

and $\text{Aut}_{\mathcal{H}_k}(\widehat{\Gamma})$ is the subgroup of automorphisms of $\widehat{\Gamma}$ commuting with τ .

We remark that the quantity $\ell_{\widehat{\Gamma}}$ is intrinsic to Γ , for a two-level graph it is given by $\ell_{\widehat{\Gamma}} = \text{lcm}\left(\frac{\kappa_e}{\gcd(\kappa_e, k)} \text{ for } e \in E(\Gamma)\right)$.

PROOF. The first statement is exactly the one of Lemma 3.6 since the topmost maps in (61) are given by a disjoint union of the topmost maps in (31).

For the second statement, consider an edge $e \in E(\Gamma)$. The edge e has $\gcd(\kappa_e, k)$ preimages, each with an enhancement $\frac{\kappa_e}{\gcd(\kappa_e, k)}$. The prong-matching at one of the preimages determines the prong-matching at the other preimages by Lemma 7.1, as they are related by the action of the automorphism.

For the third statement, we need to prove that the subgroup of $\text{Aut}(\widehat{\Gamma})$ fixing setwise the linear subvariety $\overline{\mathcal{H}}_k$ is precisely the subgroup commuting with τ . If $\rho \in \text{Aut}(\widehat{\Gamma})$ commutes with τ , then it descends to a graph automorphism of Γ and gives an automorphism of families of admissible covers of stable curves, thus preserving $\overline{\mathcal{H}}_k$. Conversely, if ρ fixes $\overline{\mathcal{H}}_k$, it induces an automorphism of families of admissible covers of stable curves, thus of coverings of graphs. A priori this implies only that ρ normalizes the subgroup generated by τ . Note however that on $\overline{\mathcal{H}}_k$ the automorphism τ acts by a fixed root of unity ζ_k . If $\rho\tau\rho^{-1}$ is a non-trivial power of τ , this leads to another (though isomorphic) linear subvariety. We conclude that ρ indeed commutes with τ . \square

The aim of the following paragraphs is to rewrite the evaluation Lemma 3.10 in our context in order to find the shape of the formula in Corollary 1.5. We elaborate on basic definitions to distinguish notions of isomorphisms and automorphisms. The underlying graph of an enhanced (k-)level graph can be written as a tuple $\Gamma = (V, H, L, a : H \cup L \rightarrow V, i : H \rightarrow H)$, where V , H and L are the sets of vertices, half-edges and legs, a is the attachment map and i is the fixpoint free involution that specifies the edges. An isomorphism of graphs $\sigma : \Gamma \rightarrow \Gamma'$ is a pair of bijections $\sigma = (\sigma_V : V \rightarrow V', \sigma_H : H \rightarrow H')$ that preserve the attachment of the half-edges and legs and the the identification of the half-edges to edges, i.e. the diagrams

$$(62) \quad \begin{array}{ccc} H \cup L & \xrightarrow{a} & V \\ \downarrow \sigma_H \cup \text{id}_L & & \downarrow \sigma_V \\ H' \cup L & \xrightarrow{a'} & V' \end{array} \quad \begin{array}{ccc} H & \xrightarrow{i} & H \\ \downarrow \sigma_H & & \downarrow \sigma_H \\ H' & \xrightarrow{i'} & H' \end{array}$$

commute. If the graph is an enhanced level graph, we additionally ask that σ preserves the enhancements and level structure. In the presence of a deck transformation τ , we moreover ask that σ commutes with τ .

In the sequel we will encounter isomorphisms of graphs with the same underlying sets of vertices and half-edges. We emphasize that in this case an isomorphism σ is an *automorphism* if and only if it preserves the maps a and i , i.e. if

$$(63) \quad \sigma_V^{-1} \circ a \circ (\sigma_H \cup \text{id}_L) = a \quad \text{and} \quad \sigma_H^{-1} \circ i \circ \sigma_H = i.$$

We now define the group of level-wise half-edge permutations compatible with the cycles of τ , i.e., we let

$$\mathbf{G} := \mathbf{G}_\pi = \prod_{i=0}^{-L} G(\mathcal{H}_k(\pi_{[i]})),$$

where $G(\mathcal{H}_k(\pi_{[i]}))$ is the group G from (58) applied to the i -th level stratum. An element of the group \mathbf{G} is a permutation $g : H \cup L \rightarrow H \cup L$ and acts on a graph $\widehat{\Gamma}$ via $g \cdot \widehat{\Gamma} = (V, H, L, a \circ g, i)$.

There is a natural action of the group \mathbf{G} on the set of all (possibly disconnected) graphs with the same set of underlying vertices as $\widehat{\Gamma}_{\text{mp}}$. We denote by

$$(64) \quad \text{Stab}_{\mathbf{G}}(\widehat{\Gamma}) := \{g \in \mathbf{G} : g\widehat{\Gamma} \cong \widehat{\Gamma}\}$$

the stabilizer. Note that this is in general not a group, as it is not the stabilizer of an element but of an isomorphism class. We also denote by $\text{Stab}_{\mathbf{G}}(\mathcal{H}(\pi))$ the set of elements of \mathbf{G} which fix the adjacency data (or equivalently the 1-level graphs) of the level-wise linear manifolds $\mathcal{H}(\pi_{[i]})$, i.e., elements which permute vertices with the same signature and permute legs of the same order on the same vertex.

LEMMA 7.5. *We have*

$$|\text{Aut}_{\mathcal{H}_k}(\widehat{\Gamma})| \cdot |\text{Stab}_{\mathbf{G}}(\widehat{\Gamma})| = |\text{Aut}(\Gamma)| \prod_{e \in E(\Gamma)} \gcd(\kappa_e, k) \cdot |\text{Stab}_{\mathbf{G}}(\mathcal{H}(\pi))|$$

PROOF. Fix a cover $\widehat{\Gamma} \rightarrow \Gamma$. We may assume that the vertices of Γ are $\{1, \dots, v_\Gamma\}$, the legs are $\{1, \dots, n\}$ and the half-edges are $\{1^\pm, \dots, h_\Gamma^\pm\}$ with the convention that $i(h^\pm) = h^\mp$. For $\widehat{\Gamma}$, we may assume that the preimages of vertex v are $(v, 1), \dots, (v, p_v)$ such that $\tau((v, q)) = (v, q + 1)$, where equality in the second entry is to be read mod p_v . Similarly, we index the legs of $\widehat{\Gamma}$ by tuples $(m, 1), \dots, (m, p_m)$ for $m = 1, \dots, n$, and the half-edges by

tuples $(h^\pm, 1), \dots, (h^\pm, p_{h^\pm})$ for $h^\pm = 1, \dots, h_\Gamma^\pm$, again such that (h^+, q) and (h^-, q) form an edge.

We consider the group \mathcal{P} of pairs of permutations $\sigma = (\sigma_V, \sigma_H)$ of the vertices and half-edges of $\widehat{\Gamma}$ that are of the following form: There exists a $\gamma = (\gamma_V, \gamma_H) \in \text{Aut}(\Gamma)$, integers $\lambda_v \in \mathbb{Z}/p_v\mathbb{Z}$ for any $v \in V(\Gamma)$ and integers $\mu_{h^\pm} \in \mathbb{Z}/p_{h^\pm}\mathbb{Z}$ for any $h^\pm \in E(\Gamma)$ such that

$$\sigma_V = \{(v, q) \mapsto (\gamma_V(v), q + \lambda_v)\} \quad \text{and} \quad \sigma_H = \{(h^\pm, q) \mapsto (\gamma_H(h^\pm), q + \mu_{h^\pm})\}.$$

We let this group act on $\widehat{\Gamma}$ via $\sigma \cdot \widehat{\Gamma} = (V, H, L, \sigma_V^{-1} \circ a \circ (\sigma_H \cup \text{id}_L), i)$. An element $\sigma \in \mathcal{P}$ acts always as an isomorphism since the diagrams (62) commute. If we denote by e the edge given by h^\pm , we have $p_{h^\pm} = \text{gcd}(\kappa_e, k)$. Hence the group \mathcal{P} has cardinality

$$|\mathcal{P}| = |\text{Aut}(\Gamma)| \cdot \prod_{e \in E(\Gamma)} \text{gcd}(\kappa_e, k) \cdot \prod_{v \in V(\Gamma)} p_v.$$

Recall that the group \mathbf{G} is a product cyclic groups and thus abelian. The stabilizer $\text{Stab}_{\mathbf{G}}(\mathcal{H}_k(\pi))$ has a subgroup Stab^f where only half-edges and legs attached to the same vertex are permuted (the superscript f is for *fixed*), i.e. the elements $g \in \text{Stab}^f$ are exactly those for which $a \circ g = a$. The quotient $\text{Stab}^p := \text{Stab}_{\mathbf{G}}(\mathcal{H}_k(\pi)) / \text{Stab}^f$ can be identified with those elements of \mathbf{G} that permute legs and half-edges in such a way that whenever a leg or half-edge attached to a vertex v_1 is moved to another vertex v_2 , then all the legs and half-edges attached to v_1 are moved to v_2 . So we may alternatively identify Stab^p with τ -invariant permutations of the vertices of $\widehat{\Gamma}$ (hence the superscript p for *permutation*). This yields $|\text{Stab}^p| = \prod_{v \in V(\Gamma)} p_v$.

The group \mathcal{P} comes with a commutative triangle

$$\begin{array}{ccc} \text{Aut}_{\mathcal{H}}(\widehat{\Gamma}) & \hookrightarrow & \mathcal{P} \\ & \searrow & \downarrow \\ & & \text{Aut}(\Gamma) \end{array}$$

where the vertical map is the forgetful map, the diagonal map is the quotient by G -map and the horizontal map is natural injection. Since we computed above $|\mathcal{P}|$, we know that the kernel of the surjective map $\mathcal{P} \rightarrow \text{Aut}(\Gamma)$ has cardinality $\prod_{e \in E(\Gamma)} \text{gcd}(\kappa_e, k) \cdot \prod_{v \in V(\Gamma)} p_v$.

Note now that the group Stab^f acts on the set $\text{Stab}_{\mathbf{G}}(\widehat{\Gamma})$ and we denote by $\text{Stab}_{\mathbf{G}}(\widehat{\Gamma}) / \text{Stab}^f$ the space of orbits. We are done if we can identify elements of $\text{Stab}_{\mathbf{G}}(\widehat{\Gamma}) / \text{Stab}^f$ with elements of the cosets in $\mathcal{P} / \text{Aut}_{\mathcal{H}}(\widehat{\Gamma})$.

For this identification, first consider $g \in \text{Stab}_{\mathbf{G}}(\widehat{\Gamma})$. By definition, there exists an isomorphism $\sigma(g) : g \cdot \widehat{\Gamma} \rightarrow \widehat{\Gamma}$ such that $g \cdot \widehat{\Gamma} = \sigma(g)(\widehat{\Gamma})$. This induces a map $\sigma : \text{Stab}_{\mathbf{G}}(\widehat{\Gamma}) \rightarrow \mathcal{P}$. Note that Stab^f is a subgroup of $\text{Aut}_{\mathcal{H}}(\widehat{\Gamma})$. If we had chosen a different representative g' in the orbit $g \cdot \text{Stab}^f$, the resulting element $\sigma(g') \in \mathcal{P}$ would differ by an element of $\text{Aut}_{\mathcal{H}}(\widehat{\Gamma})$. Hence σ induces a well-defined map $\text{Stab}_{\mathbf{G}}(\widehat{\Gamma}) / \text{Stab}^f \rightarrow \mathcal{P} / \text{Aut}_{\mathcal{H}}(\widehat{\Gamma})$. We now construct an inverse map for σ . For any $\rho \in \mathcal{P}$, we need to find an element $g \in \mathbf{G}$ such that $\sigma(g) = \rho$, i.e. such that $g \cdot \widehat{\Gamma} = \rho(\widehat{\Gamma})$. This implies that g must satisfy the equation

$$a \circ g = \rho_V^{-1} \circ a \circ (\rho_H \cup \text{id}_L),$$

which determines the element g up to the action of Stab^f . The resulting g does not depend on the choice of a representative of the coset $\rho / \text{Aut}_{\mathcal{H}}(\widehat{\Gamma})$ because of (63). \square

We let now

$$(65) \quad S(\pi) = \frac{|G|}{|\mathbf{G}|} \cdot \frac{|\text{Stab}_{\mathbf{G}}(\widehat{\Gamma})|}{|\text{Stab}_G(\widehat{\Gamma})|} = \frac{|\text{Stab}_{\mathbf{G}/G}(\widehat{\Gamma})|}{\prod_e \text{gcd}(\kappa_e, k)^2}$$

where the stabilizers are defined in a way analogous to (64).

REMARK 7.6. *The ratio $S(\pi) = 1$ for many coverings of graphs $\pi : \widehat{\Gamma} \rightarrow \Gamma$, e.g. when all vertices of Γ have exactly one preimage in $\widehat{\Gamma}$. In this case \mathbf{G}/G only permutes half-edges adjacent to one vertex, and this always stabilizes the graph. Thus $S(\pi) = 1$, as $|\mathbf{G}/G| = \prod_e \text{gcd}(\kappa_e, k)^2$. More generally $S(\pi) = 1$ if each edge of Γ is adjacent to at least*




FIGURE 2. A covering of graphs $\pi : \widehat{\Gamma} \rightarrow \Gamma$ in $\Xi^2 \overline{\mathcal{M}}_{3,1}(8)$ with non-trivial $S(\pi)$.

one vertex which has exactly one preimage in $\widehat{\Gamma}$. In this case it is straightforward to verify that the obvious generators of \mathbf{G}/G are stabilizing the graph.

If there are vertices of Γ with more than one pre-image in $\widehat{\Gamma}$, then $S(\pi)$ is in general non-trivial. Consider for example the covering of graphs π depicted in Figure 2, for which $S(\pi) = \frac{1}{2}$.

As a consequence of the degree computation in Lemma 7.4 and Lemma 7.5, we can write an evaluation lemma for k -differentials analogous to Lemma 3.10. We give two versions, for \mathcal{H}_k and \mathcal{Q} respectively.

LEMMA 7.7. *Let $(\pi : \widehat{\Gamma}_{\text{mp}} \rightarrow \Gamma) \in \text{LG}_L(\mathcal{H}_k^{\text{mp}})$ and $\widehat{\Gamma}$ a marked version of $\widehat{\Gamma}_{\text{mp}}$. Suppose that $\alpha_\pi \in \text{CH}_0(D_\pi^{\mathcal{H}_k})$ and $\beta_\pi \in \text{CH}_0(D_\pi^{\mathcal{Q}})$ are top degree classes and that*

$$c_\pi^* \alpha_\pi = p_\pi^* \prod_{i=0}^{-L} \alpha_i \quad \text{and} \quad c_\pi^* d_\pi^* \beta_\pi = p_\pi^* \mathbf{d}_\pi^* \prod_{i=0}^{-L} \beta_i$$

for some α_i and β_i . Then

$$\int_{D_\pi^{\mathcal{H}_k}} \alpha_\pi = S(\pi) \cdot \frac{\prod_{e \in E(\Gamma)} \kappa_e}{|\text{Aut}(\Gamma)| \cdot \prod_{e \in E(\Gamma)} \gcd(\kappa_e, k)^2 \cdot \ell_{\widehat{\Gamma}}} \cdot \prod_{i=0}^{-L} \int_{\mathcal{H}_k(\pi_{[i]})} \alpha_i$$

and

$$\int_{D_\pi^{\mathcal{Q}}} \beta_\pi = S(\pi) \cdot \frac{\prod_{e \in E(\Gamma)} \kappa_e}{k^L \cdot |\text{Aut}(\Gamma)| \cdot \ell_{\widehat{\Gamma}}} \cdot \prod_{i=0}^{-L} \int_{\mathcal{Q}(\pi_{[i]})} \beta_i.$$

PROOF. In order to show the first statement, we first apply Lemma 7.4 and note that the map p_π is not surjective in general. It is now enough to check that the number of adjacency data appearing in $\mathcal{H}_k(\pi)$ is $|\mathbf{G}|/|\text{Stab}_{\mathbf{G}}(\mathcal{H}_k(\pi))|$, while the one appearing in the image of p_π is $|G|/|\text{Stab}_G \widehat{\Gamma}|$. We finally use Lemma 7.5 to rewrite the prefactor. For the second statement, we additionally apply Lemma 7.2 and Lemma 7.3. \square

We are finally ready to prove Corollary 1.5.

PROOF OF COROLLARY 1.5. The orbifold Euler characteristics of $\mathcal{Q} = \mathbb{P}\Omega^k \mathcal{M}_{g,n}(\mu)$ and \mathcal{H}_k are related by

$$\chi(\mathbb{P}\Omega^k \mathcal{M}_{g,n}(\mu)) = \frac{1}{\deg(d)} \cdot \chi(\mathcal{H}_k).$$

We apply the general Euler characteristic formula in the form (55) to \mathcal{H}_k and group the level graphs $\widehat{\Gamma} \in \text{LG}_L(\mathcal{H}_k)$ by those with the same graph $\widehat{\Gamma}_{\text{mp}}$ that is marked partially. Since the integrals do not depend on the marking, we obtain

$$\chi(\mathcal{Q}) = \frac{k}{|G|} (-1)^d \sum_{L=0}^d \sum_{(\pi: \widehat{\Gamma}_{\text{mp}} \rightarrow \Gamma) \in \text{LG}_L(\mathcal{H}_k^{\text{mp}})} N_\pi^\top \cdot \ell_{\widehat{\Gamma}} \cdot \int_{D_\pi^{\mathcal{H}_k}} \prod_{i=-L}^0 (\xi_{\widehat{\Gamma}, \mathcal{H}_k}^{[i]})^{d_\Gamma^{[i]}}$$

where we used the notation that $\widehat{\Gamma}$ is a fully marked representative of $\widehat{\Gamma}_{\text{mp}}$. Thanks to Lemma 3.9 we can apply Lemma 7.7 and convert the integral over $D_\pi^{\mathcal{H}_k}$ into a ξ -integral

over the product of $\mathcal{H}_k(\pi_{[i]})$. We hence obtain

$$\begin{aligned} & \chi(\mathbb{P}\Omega^k \mathcal{M}_{g,n}(\mu)) \\ &= \frac{k}{|G|} \cdot (-1)^d \sum_{L=0}^d \sum_{(\pi: \widehat{\Gamma}_{\text{mp}} \rightarrow \Gamma) \in \text{LG}_L(\mathcal{H}_k^{\text{mp}})} S(\pi) \frac{\prod_{e \in E(\Gamma)} \kappa_e \cdot N_\pi^\top}{|\text{Aut}(\Gamma)| \cdot \prod_e \gcd(\kappa_e, k)^2} \cdot \prod_{i=0}^{-L} \int_{\mathcal{H}_k(\pi_{[i]})} \xi^{d_{\pi}^{[i]}} \\ &= \left(\frac{-1}{k}\right)^d \sum_{L=0}^d \sum_{(\pi: \widehat{\Gamma}_{\text{mp}} \rightarrow \Gamma) \in \text{LG}_L(\mathcal{Q})} S(\pi) \cdot \frac{\prod_{e \in E(\Gamma)} \kappa_e \cdot N_\pi^\top}{|\text{Aut}(\Gamma)|} \cdot \prod_{i=0}^{-L} \int_{\mathcal{Q}(\pi_{[i]})} \zeta^{d_{\pi}^{[i]}}. \end{aligned}$$

For the second equality, we used that

$$(66) \quad d^* \zeta = k \xi, \quad \text{and hence} \quad d_* \xi = \frac{\deg(d)}{k} \zeta$$

for any level stratum, together with the dimension statement of Proposition 3.4. The final result is what we claimed in Corollary 1.5. \square

7.4. Evaluating tautological classes. In this section we explain how to evaluate any top degree class of the form

$$(67) \quad \beta := \zeta^{p_0} \psi_1^{p_1} \cdots \psi_n^{p_n} \cdots [D_{\pi_1}^{\mathcal{Q}}] \cdots [D_{\pi_w}^{\mathcal{Q}}] \in \text{CH}_0(\overline{\mathcal{Q}})$$

for any generalized stratum $\overline{\mathcal{Q}}$ of k -differentials. First, we show how to transform the previous class into the form

$$\beta = \sum_i \psi_1^{q_{i,1}} \cdots \psi_1^{q_{i,n}} [D_{\sigma_i}^{\mathcal{Q}}].$$

Then by Lemma 7.7, we can write every summand of β as a product of ψ -classes evaluated on generalized strata of k -differentials. We finally will explain how to evaluate such classes.

Let us start with the first task. The relations in the Chow ring of a general linear submanifold we obtained in Section 4 immediately apply to the covering $\overline{\mathcal{H}}_k$ and we want to restate them in the Chow ring of the generalized stratum $\overline{\mathcal{Q}}$ of k -differentials. Let i be the index of a marked point in $\overline{\mathcal{Q}}$ and (i, j) be the index of a preimage of this point in $\overline{\mathcal{H}}_k$. Moreover, let m_i denote the order of the k -differential at the i -th marked point, and let $\widehat{m}_{i,j}$ denote the order of the abelian covering at the (i, j) -th marked point. Then the relation

$$(68) \quad \psi_{i,j} = \frac{\gcd(m_i, k)}{k} \cdot d^* \psi_i$$

holds, see for example [SZ20, Lemma 3.9]. Using the relation

$$\widehat{m}_{i,j} + 1 = (m_i + k) / \gcd(m_i, k)$$

and applying push-pull we obtain

$$(69) \quad (\widehat{m}_{i,j} + 1) d_* \psi_{i,j} = \frac{\deg(d)}{k} (m_i + k) \psi_i.$$

We can now write the analogue of Proposition 4.1 for the first Chern class $\zeta \in \text{CH}^1(\overline{\mathcal{Q}})$ of the tautological line bundle on the stratum of k -differentials.

COROLLARY 7.8. *The class ζ can be expressed as*

$$\begin{aligned} \zeta &= (m_i + k) \psi_i - \sum_{(\pi: \widehat{\Gamma}_{\text{mp}} \rightarrow \Gamma) \in {}_i\text{LG}_1(\overline{\mathcal{Q}})} k \ell_{\widehat{\Gamma}_{\text{mp}}} [D_\pi^{\mathcal{Q}}] \\ &= (m_i + k) \psi_i - \sum_{(\pi: \widehat{\Gamma}_{\text{mp}} \rightarrow \Gamma) \in {}_i\text{LG}_1(\overline{\mathcal{Q}})} S(\pi) \frac{\prod_{e \in E(\Gamma)} \kappa_e}{|\text{Aut}(\Gamma)|} \text{cl}_{\pi,*} p_\pi^* \mathbf{d}_\pi^* [\mathcal{Q}(\pi)] \end{aligned}$$

where ${}_i\text{LG}_1(\overline{\mathcal{Q}})$ are covers of two-level graphs with the leg i on lower level and $\text{cl}_\pi = i_\pi \circ d_\pi \circ c_\pi$ is the clutching morphism analogous to (38).

PROOF. The first equation is obtained by pushing forward the equation in Proposition 4.1 along d and using the relations (66) and (69). The second equation is obtained from the first by Lemma 7.7. \square

REMARK 7.9. *The expression given by the second line of Corollary 7.8 reproves the formula of [Sau21, Theorem 3.12] and computes explicitly the coefficients appearing in loc.cit., which were computed only for special two-level graphs.*

To state the formula for the normal bundle, let

$$\mathcal{L}_\pi^\top = \mathcal{O}_{D_\pi^\mathcal{Q}} \left(\sum_{\substack{(\sigma: \widehat{\Delta}_{\text{mp}} \rightarrow \Delta) \in \text{LG}_2(\overline{\mathcal{Q}}) \\ \delta_2(\sigma) = \pi}} \ell_{\widehat{\Delta}, 1} D_\sigma^\mathcal{H} \right)$$

denote the top level correction bundle.

COROLLARY 7.10. *Suppose that D_π is a divisor in $\overline{\mathcal{Q}}$ corresponding to a covering of graphs $(\pi: \widehat{\Gamma}_{\text{mp}} \rightarrow \Gamma) \in \text{LG}_1(\overline{\mathcal{Q}})$. Then the first Chern class of the normal bundle is given by*

$$c_1(\mathcal{N}_\pi) = \frac{1}{\ell_{\widehat{\Gamma}}} \left(-\frac{1}{k} \zeta_\pi^\top - c_1(\mathcal{L}_\pi^\top) + \frac{1}{k} \zeta_\pi^\perp \right) \in \text{CH}^1(D_\pi^\mathcal{Q}),$$

where ζ_π^\top , resp. ζ_π^\perp , is the first Chern class of the line bundle generated by the top, resp. bottom, level multi-scale component.

PROOF. We can pull-back the right and left hand sides of the relation via d . Using the expression (66), we see that the pulled-back relation holds since it agrees with the one of Proposition 4.4. Since d is a quasi-finite proper unramified map, we are done. The same argument, together with Proposition 4.5, works for the second statement about horizontal divisors. \square

Using the same arguments as [CMZ22, Proposition 8.1], it is possible to show an excess intersection formula in this context of k -differentials. We will not explicitly do this here since the methods and the result are exactly parallel to the original ones for Abelian differentials. Using the previous ingredients we can then reduce the computation of the class β in (67) to the computation of a top-degree product of ψ -classes

$$\alpha := \psi_1^{p_1} \cdots \psi_n^{p_n} \in \text{CH}_0(\overline{\mathcal{Q}})$$

on a generalized stratum. If we can describe the class of a generalized stratum in its corresponding moduli space of pointed curves, then we are done since it is possible to compute top-degree tautological classes on the moduli space of curves, e.g. with the sage package *admcycles*, see [DSZ21].

One of the advantages in comparison to the situation with general linear submanifolds (as explained in Section 4) is that the fundamental classes of strata of primitive k -differentials $\mathbb{P}\Xi^k \overline{\mathcal{M}}_{g,n}(\mu)$ are known in $\overline{\mathcal{M}}_{g,n}$, see [BHPSS20].

More generally, if \mathcal{Q} parameterizes k -differentials, on a curve with connected τ -quotient, which are d -th powers of primitive $k' := k/d$ -differentials, we can compare ψ -classes on $\overline{\mathcal{Q}}$ to ψ -classes on the stratum of primitive k' differentials $\mathbb{P}\Xi^{k'} \overline{\mathcal{M}}_{g,n}(\mu/d)$ via the diagram

$$\begin{array}{ccc} \mathcal{H}_k^{\text{mp}}(\mu) & \xrightarrow{\phi} & \mathcal{H}_{k'}^{\text{mp}}(\mu/d) \\ \downarrow d_1 & & \downarrow d_2 \\ \mathcal{Q} & & \mathbb{P}\Xi^{k'} \overline{\mathcal{M}}_{g,n}(\mu/d) \end{array}$$

where the map ϕ sends the disconnected curve $(\bigcup_{i=1}^d \widehat{X}_i, \bigcup_{i=1}^d \widehat{\mathbf{z}}_i, \bigcup_{i=1}^d \omega_i, \tau)$ to $(\widehat{X}_1, \mathbf{z}_1, \omega_1, \tau^d|_{\widehat{X}_1})$. The map ϕ has degree $\deg(\phi) = d^{n-1}$, since up to the action of τ there are such many ways to distribute the marked points $\widehat{\mathbf{z}}$ onto the connected components of \widehat{X} . Using $\deg(d_1) = \frac{1}{k}$ and $\deg(d_2) = \frac{1}{k'}$ we can evaluate α as

$$\int_{\mathcal{Q}} \alpha = d^n \int_{\mathbb{P}\Xi^{k'} \overline{\mathcal{M}}_{g,n}(\mu/d)} \psi_1^{p_1} \cdots \psi_n^{p_n}.$$

If \mathcal{Q} parameterizes primitive differentials on disconnected curves, then $\int_{\mathcal{Q}} \alpha = 0$ since we go down in dimension by looking at the image of the projection to the moduli spaces of curves.

It remains to explain how to evaluate intersection numbers in the presence of residue conditions. In addition to the space \mathfrak{R} defined starting from a τ -invariant partition $\lambda_{\mathfrak{R}}$ we consider the linear subspace

$$R := \left\{ (r_i)_{i \in H_p} \in \mathbb{C}^{H_p} : \begin{array}{l} \sum_{i \in \mathcal{A}^{-1}(\widehat{X}')} r_i = 0 \text{ for all } \widehat{X}' \in \pi_0(\widehat{X}) \\ r_i = \zeta_k^{-1} r_{\tau(i)} \text{ for all } i \in H_p \end{array} \right\}$$

cut out by the residue theorem on each component and the deck transformation. Recall that $\lambda_{\mathfrak{R}}$ is τ -invariant. Let $\lambda_{\mathfrak{R}_0}$ denote a subset of $\lambda_{\mathfrak{R}}$ obtained by removing one element, and let \mathfrak{R}_0 denote the new set of residue conditions. For ease of notation let for now $H_k^{\mathfrak{R}} := \mathbb{P}\Omega\mathcal{H}_k^{[\mathfrak{R}, \mathcal{A}]}$ and $H_k^{\mathfrak{R}_0} := \mathbb{P}\Omega\mathcal{H}_k^{[\mathfrak{R}_0, \mathcal{A}]}$. If $R \cap \mathfrak{R} = R \cap \mathfrak{R}_0$ then $\mathcal{H}_k^{\mathfrak{R}} = \mathcal{H}_k^{\mathfrak{R}_0}$. So assume that $R \cap \mathfrak{R} \neq R \cap \mathfrak{R}_0$, in which case $\mathcal{H}_k^{\mathfrak{R}} \subsetneq \mathcal{H}_k^{\mathfrak{R}_0}$ is a divisor since removing one element from $\lambda_{\mathfrak{R}}$ forces to remove its τ -orbit. For a divisor $D_{\pi}^{\mathcal{H}_k^{\mathfrak{R}}} \subseteq \overline{\mathcal{H}_k^{\mathfrak{R}}}$, we denote by \mathfrak{R}^{\top} the residue conditions induced by \mathfrak{R} on the top-level stratum $\mathcal{H}_k(\pi_{[0]})$. It can be simply computed by discarding from the parts of $\lambda_{\mathfrak{R}}$ all indices of legs that go to lower level in $D_{\pi}^{\mathcal{H}_k^{\mathfrak{R}}}$. Moreover, we denote by R^{\top} the linear subspace belonging to the top-level stratum of π that is cut out by the residue theorem and the deck transformation.

PROPOSITION 7.11. *The class of $\overline{\mathcal{H}_k^{\mathfrak{R}}}$ compares inside the Chow ring of $\overline{\mathcal{H}_k^{\mathfrak{R}_0}}$ to the class ξ by the formula*

$$[\overline{\mathcal{H}_k^{\mathfrak{R}}}] = -\xi - \sum_{(\pi: \widehat{\Gamma}_{\text{mp}} \rightarrow \Gamma) \in \text{LG}_1^{\mathfrak{R}}(\overline{\mathcal{H}_k^{\mathfrak{R}_0}})} \ell_{\widehat{\Gamma}}[D_{\pi}^{\mathcal{H}_k^{\mathfrak{R}_0}}] - \sum_{(\pi: \widehat{\Gamma}_{\text{mp}} \rightarrow \Gamma) \in \text{LG}_{1, \mathfrak{R}}(\overline{\mathcal{H}_k^{\mathfrak{R}_0}})} \ell_{\widehat{\Gamma}}[D_{\pi}^{\mathcal{H}_k^{\mathfrak{R}_0}}],$$

where $\text{LG}_1^{\mathfrak{R}}(\overline{\mathcal{H}_k^{\mathfrak{R}_0}})$ are the two-level graphs with $R^{\top} \cap \mathfrak{R}^{\top} = R^{\top} \cap \mathfrak{R}_0^{\top}$, i.e., where the GRC on top level induced by \mathfrak{R} does no longer introduce an extra condition, and where $\text{LG}_{1, \mathfrak{R}}(\overline{\mathcal{H}_k^{\mathfrak{R}_0}})$ are the two-level graphs where all the legs involved in the condition forming $\mathfrak{R} \setminus \mathfrak{R}_0$ go to lower level.

PROOF. The formula is obtained by intersecting the formula in [CMZ22, Proposition 8.3] with $\overline{\mathcal{H}_k^{\mathfrak{R}_0}}$ and thereby using the transversality statement from Proposition 3.2. \square

By pushing down this relation along d and applying relation (66) we obtain a similar relation for a generalized stratum of k -differentials $\mathcal{Q}^{\mathfrak{R}}$ with residue conditions \mathfrak{R} .

COROLLARY 7.12. *The class of $\overline{\mathcal{Q}^{\mathfrak{R}}}$ compares inside the Chow ring of $\overline{\mathcal{Q}^{\mathfrak{R}_0}}$ to the class ζ by the formula*

$$[\overline{\mathcal{Q}^{\mathfrak{R}}}] = -\frac{1}{k}\zeta - \sum_{(\pi: \widehat{\Gamma}_{\text{mp}} \rightarrow \Gamma) \in \text{LG}_1^{\mathfrak{R}}(\overline{\mathcal{Q}^{\mathfrak{R}_0}})} \ell_{\widehat{\Gamma}}[D_{\pi}^{\mathcal{Q}^{\mathfrak{R}_0}}] - \sum_{(\pi: \widehat{\Gamma}_{\text{mp}} \rightarrow \Gamma) \in \text{LG}_{1, \mathfrak{R}}(\overline{\mathcal{Q}^{\mathfrak{R}_0}})} \ell_{\widehat{\Gamma}}[D_{\pi}^{\mathcal{Q}^{\mathfrak{R}_0}}],$$

where $\text{LG}_1^{\mathfrak{R}}(\overline{\mathcal{Q}^{\mathfrak{R}_0}})$ are the two-level graphs with $R^{\top} \cap \mathfrak{R}^{\top} = R^{\top} \cap \mathfrak{R}_0^{\top}$, i.e. where the GRC on top level induced by \mathfrak{R} does no longer introduce an extra condition and where $\text{LG}_{1, \mathfrak{R}}(\overline{\mathcal{Q}^{\mathfrak{R}_0}})$ are the two-level graphs where all the legs involved in the condition forming $\mathfrak{R} \setminus \mathfrak{R}_0$ go to lower level.

The last expression allows us, in the presence of residue conditions, to reduce to the previous situations without residue conditions when we want to evaluate α .

7.5. Values and cross-checks. In this section we provide in Table 2 and Table 3 some Euler characteristics for strata of k -differentials. We abbreviate $\chi_k(\mu) := \chi(\mathbb{P}\Omega^k \mathcal{M}_{g,n}(\mu))$. Moreover we provide several cross-checks for our values.

The second power of the projectivized Hodge bundle over \mathcal{M}_2 is the union of the strata of quadratic differentials of type (4), (2, 2), (2, 1²) and (1⁴), if all of them are taken with unmarked zeros. (Note that there are no quadratic differentials of type (3, 1).) All quadratic differentials of type (4) are second powers of abelian differentials of type (2). The stratum (2, 2) contains both primitive quadratic differentials and second powers of abelian differentials of type (1, 1). From Table 2 and [CMZ22, Table 1] we read off that

$$\chi_1(2) + \frac{1}{2}\chi_2(2, 2) + \frac{1}{2}\chi_1(1, 1) + \frac{1}{2}\chi_2(2, 1^2) + \frac{1}{4!}\chi_2(1^4) = -\frac{1}{80} = \chi(\mathbb{P}^2)\chi(\mathcal{M}_2).$$

μ	(2, 2)	(2, 1 ²)	(1 ⁴)	(5, -1)	(4, 1, -1)
$\chi_2(\mu)$	$-\frac{1}{8}$	$\frac{1}{5}$	-1	$-\frac{7}{15}$	$\frac{6}{5}$
μ	(3, 2, -1)	(3, 1 ² , -1)	(2 ² , 1, -1)	(2, 1 ³ , -1)	(1 ⁵ , -1)
$\chi_2(\mu)$	$\frac{5}{3}$	-5	-6	26	-147

TABLE 2. Euler characteristics of the strata of quadratic differentials in genus 2 with at most one simple pole

Similarly, one checks for the third power of the projectivized Hodge bundle over \mathcal{M}_2 that the numbers in provided in Table 3 add up to $-\frac{1}{48} = \chi(\mathbb{P}^4)\chi(\mathcal{M}_2)$.

μ	(6)	(5, 1)	(4, 2)	(3, 3)	(4, 1 ²)	(3, 2, 1)
$\chi_3(\mu)$	$\frac{1}{3}$	$-\frac{4}{5}$	$-\frac{9}{8}$	$-\frac{4}{3}$	$\frac{16}{5}$	4
μ	(2 ³)	(3, 1 ³)	(2 ² , 1 ²)	(2, 1 ⁴)	(1 ⁶)	
$\chi_3(\mu)$	$\frac{41}{10}$	-16	$-\frac{52}{3}$	90	-567	

TABLE 3. Euler characteristics of the strata of holomorphic 3-differentials in genus 2

Now consider the second power of the projectivized Hodge bundle twisted by the universal section over $\mathcal{M}_{2,1}$. It decomposes into the unordered strata (4), (5, -1), (4, 1, -1), (3, 2, -1), (2, 1²), (3, 1², -1), (2², 1, -1), (2, 1³, -1), (1⁵, -1), (4, 0), (2², 0), (2, 1², 0), (1⁴, 0), the ordered stratum (2²), (2, 1²) (since the zero at the unique marked point is distinguished) and the partially ordered stratum (1⁴). The stratum (2, 1²) appears two times in the list: the first time the unique marked point is the zero of order 2, the second time it is one of the simple zeros. On the stratum (1⁴) one of the simple zeros is distinguished, while the others may be interchanged. Note that $\chi_k(m_1, \dots, m_n, 0) = (2 - 2g - n)\chi_k(m_1, \dots, m_n)$. The contributions in Table 2 and [CMZ22, Table 1] add up to $\frac{1}{30} = \chi(\mathbb{P}^3)\chi(\mathcal{M}_{2,1})$.

8. Ball quotients

The goal of this section is to prove Theorem 1.7, which gives an independent proof of the Deligne-Mostow-Thurston construction ([DM86], [Thu98]) of ball quotients via cyclic coverings. For this proof of concept we consider the special case of surfaces, i.e. lattices in $\text{PU}(1, 2)$.

We first prove a criterion for showing that a two dimensional smooth Deligne-Mumford stack is a ball quotient via the Bogomolov-Miyaoka-Yau equality. Even though such a criterion exists in many contexts, typically pairs of a variety and a \mathbb{Q} -divisor with various hypothesis on the singularities a priori allowed, see for example [GKPT19]; [GT22b], we found no criterion for stacks in the literature. Only the inequality was proven in [CT20] and only in the compact case.

We then investigate the special two dimensional strata of k -differentials of genus zero considered in Deligne-Mostow-Thurston, compute all the relevant intersection numbers and construct, via a contraction of some specific divisor, the smooth surface stack which we finally show to be a ball quotient.

8.1. Ball quotient criterion. We provide a version of the Bogomolov-Miyaoka-Yau inequality for stacks in the surface case, based on [KNS89]. Singularity terminology and basics about the minimal model program can be found e.g. in [KM98].

PROPOSITION 8.1. *Suppose that $\overline{\mathfrak{B}}$ is a smooth Deligne-Mumford stack of dimension 2 with trivial isotropy group at the generic point and let \mathcal{D}_1 be a normal crossing divisor. Moreover, suppose that $K_{\overline{\mathfrak{B}}}(\log \mathcal{D}_1)^2 > 0$ and that $K_{\overline{\mathfrak{B}}}(\log \mathcal{D}_1)$ intersects positively any curve not contained in \mathcal{D}_1 . Then the Miyaoka-Yau inequality*

$$(70) \quad c_1^2(K_{\overline{\mathfrak{B}}}(\log \mathcal{D}_1)) \leq 3c_2(K_{\overline{\mathfrak{B}}}(\log \mathcal{D}_1))$$

holds, with equality if and only if $\mathfrak{B} = \overline{\mathfrak{B}} \setminus \mathcal{D}_1$ is a ball quotient, i.e. there is a cofinite lattice $\Gamma \in \mathrm{PU}(1, n)$ such that $\mathfrak{B} = [\mathbb{B}^2/\Gamma]$ as quotient stack, where $\mathbb{B}^2 = \{(z_1, z_2) \in \mathbb{C}^2 : |z_1|^2 + |z_2|^2 < 1\}$ is the 2-ball.

PROOF. Let \mathcal{D} be the divisor defined as \mathcal{D}_1 together with the sum \mathcal{D}_2 of the divisors \mathcal{D}_2^i with non-trivial isotropy groups of order b_i . Let $\pi : \overline{\mathfrak{B}} \rightarrow \overline{B}$ be the map to the coarse space and let $D_1 = \pi(\mathcal{D}_1)$, $D_2 = \sum(1 - 1/b_i)\pi(\mathcal{D}_2^i)$ and $D = D_1 + D_2$.

We start by assuming that the pair (\overline{B}, D) is log-canonical and the pair (\overline{B}, D_2) is log-terminal. We will show that this assumptions holds in our situation at the end of the proof.

Let \overline{B}' be a log-minimal model given by contracting all the log-exceptional curves in D_1 , i.e., contracting all irreducible curves $C \subseteq D_1$ with the properties $C^2 < 0$ and $(c_1(K_{\overline{B}}) + [D_1] + [D_2]) \cdot C \leq 0$, and let D'_i be the image of D_i , for $i = 1, 2$. Then

$$K_{\overline{B}}(\log D_1) + D_2 = \pi^*(K_{\overline{B}'}(\log D'_1) + D'_2).$$

Moreover the log-canonical bundle satisfies

$$(71) \quad K_{\overline{\mathfrak{B}}}(\log \mathcal{D}_1) = \pi^*(K_{\overline{B}}(\log D_1) + D_2).$$

The fact that the support of the log-exceptional curves is in \mathcal{D}_1 , together with (71), implies that $K_{\overline{B}'} + D'_1 + D'_2$ is numerically ample. By the assumption above on the singularities we know that (\overline{B}, D) is log-canonical. Hence we are in the situation of applying [KNS89, Theorem 12].

As a consequence of (71) we know that $c_1^2(K_{\overline{\mathfrak{B}}}(\log \mathcal{D}_1))$ coincides with the left hand side of the Miyaoka-Yau inequality of [KNS89, Theorem 12] applied to \overline{B}' with boundary divisor $D'_1 + D'_2$.


Moreover, by the Gauss-Bonnet theorem for DM-stacks (see e.g. [CMZ22, Proposition 2.1]) we can also identify $c_2(K_{\overline{\mathfrak{B}}}(\log \mathcal{D}_1))$ with the right hand side of the inequality of [KNS89, Theorem 12] applied to \overline{B}' with boundary divisor $D'_1 + D'_2$, up to non-log-terminal singularities (similarly as it was done in [CT20, Section 3.2]). By the assumption above, the pair (\overline{B}, D_2) is log-terminal and so the previous identification of the right hand side of [KNS89, Theorem 12] with $c_2(K_{\overline{\mathfrak{B}}}(\log \mathcal{D}))$ is true without corrections.

This shows inequality (70) and that in the case of equality $\overline{B}' \setminus D'_1 \cong \overline{B} \setminus D_1$ is a ball quotient, i.e. $\overline{B} \setminus D_1 \cong \mathbb{B}^2/\Gamma$. Moreover, in this case, the divisors D_2^i are the branch loci of π with branch indices b_i .

Since $\overline{B} \setminus D_1$ is the coarse space associated both to $\overline{\mathfrak{B}} \setminus \mathcal{D}_1$ and to $[\mathbb{B}^2/\Gamma]$, this implies that these two DM stacks have to differ by a composition of root constructions along divisors (see e.g. [CT20, Section 3.1]). But since the branch indices of D_2^i can be identified with the isotropy groups of the corresponding divisors in $[\mathbb{B}^2/\Gamma]$, and since they coincide with the isotropy groups of the corresponding divisor $\overline{B} \setminus D_1$, we can identify $\overline{B} \setminus D_1$ with $[\mathbb{B}^2/\Gamma]$, as non-trivial root constructions would have changed the size of such isotropy groups.

We are finally left to show the assumption on the singularities. First, there exists a resolution $\widetilde{\mathfrak{B}}$ of $\overline{\mathfrak{B}}$ where the proper transform $\widetilde{\mathcal{D}}$ of \mathcal{D} is a normal crossing divisor and the exceptional divisors \mathcal{E}_i are log-exceptional, i.e. $\mathcal{E}_i^2 < 0$ and $(c_1(K_{\widetilde{\mathfrak{B}}}) + [\widetilde{\mathcal{D}}_1]) \cdot \mathcal{E}_i \leq 0$. Indeed such a resolution can be obtained by blowing-up smooth points of the DM stack, where the numerical conditions can be checked on an étale chart just as for the usual blow-up of a smooth point of a variety.

In this situation the corresponding exceptional divisors E_i for the coarse space resolution \widetilde{B} of \overline{B} are also log-exceptional, i.e., $(c_1(K_{\widetilde{B}}) + [\widetilde{D}_1] + [\widetilde{D}_2]) \cdot E_i \leq 0$ and $E_i^2 \leq 0$. Since contracting log-exceptional divisors does not change the singularity type, this implies that to show that $(\overline{B}, D_1 + D_2)$ is log-canonical and (\overline{B}, D_2) is log-terminal, it is enough to show that $(\widetilde{B}, \widetilde{D}_1 + \widetilde{D}_2)$ is log-canonical and $(\widetilde{B}, \widetilde{D}_2)$ is log-terminal.

FIGURE 3. Level graphs of boundary divisors for strata $\Omega\mathcal{M}_{0,5}(a_1, \dots, a_5)$

In order to do this, we observe that in general since $(\mathfrak{B}, \tilde{D})$ is a smooth DM stack with normal crossing divisor, then $(\tilde{B}, \tilde{D}_1 + \sum_i \tilde{D}_2^i)$ is log-canonical. Details are given in [CCM22, Theorem 5.1], using [HH09, Proposition A.13]. Then we can use that \tilde{B} has at worst klt singularities (since it is a surface with quotient singularities and by [KM98, Prop. 4.18]). It is easy to show that this implies that $(\tilde{B}, \tilde{D}_1 + \sum_i t_i \tilde{D}_2^i)$ has log-canonical singularities and $(\tilde{B}, \sum_i t_i \tilde{D}_2^i)$ has log-terminal singularities, for any $0 \leq t_i < 1$. The desired statement follows then by setting $t_i = 1 - 1/b_i$. \square

8.2. Strata of genus zero satisfying (INT). Let (a_1, \dots, a_5) be positive integers such that $\gcd(a_1, \dots, a_5, k) = 1$ with

$$\sum_{i=1}^5 a_i = 2k, \quad \text{and for all } i \neq j \quad \left(1 - \frac{a_i}{k} - \frac{a_j}{k}\right)^{-1} \in \mathbb{Z} \quad \text{if } a_i + a_j < k.$$

The first condition states that $\mu = (-a_1, \dots, -a_5)$ is a type of a stratum of k -differentials on 5-pointed rational lines and that the intersection form on eigenspace giving period coordinates has the desired signature $(1, 2)$. Imposing the gcd-condition lets us work without loss of generality with primitive k -differentials. The last condition is the condition (INT) of [DM86]. For Deligne-Mostow this condition is key to ensure that the period map extends as an étale map over all boundary divisors. Thurston [Thu98] uses this condition to show that his cone manifolds are indeed orbifolds. Mostow completed in [Mos88] the $g = 0$ picture by showing that up to the variant Σ INT from [Mos86] these are the only ball quotient surfaces uniformized by the VHS of a cyclic cover of 5-punctured projective line. We recall from [DM86, Section 14] that there are exactly 27 five-tuples satisfying INT, and all of them satisfy in fact the integrality condition INT for all $i \neq j$ with $a_i + a_j \neq k$.

For us the condition INT has the most important consequence that the enhancements $\hat{\kappa}_e$ of the abelian covers of the level graphs are all one. This implies that ghost groups of all strata in this section are trivial. However the condition INT also enters at other places of the following computations of automorphism groups and intersection numbers.

In the sequel we will use the notation $\mathcal{Q} = \Omega^k \mathcal{M}_{0,5}(a_1, \dots, a_5)$. We now list the boundary divisors without horizontal edges. A short case inspection shows that the only possibilities are the level graphs $\Gamma = \Gamma_{ij}$, see Figure 3 left, and $L = L_{ij}$, see Figure 3 middle, that yield the 'dumbbell' divisors with two or three legs on bottom level under the condition that the a_i 's on lower level add up to less than k , and the level graphs $\Lambda = {}_{i,j}\Lambda_{p,q}$ that yield 'cherry' divisors, see Figure 3 right (V -shaped graphs are ruled out by $\sum a_i = 2k$). We define $\kappa_{i,j} := k - (a_i + a_j)$, which is both the k -enhancement of the single edge of $\Gamma_{i,j}$ and the negative of the k -enhancement of the single edge of $L_{i,j}$.

LEMMA 8.2. *Each of the graphs $\Gamma_{i,j}$, $L_{i,j}$ and ${}_{i,j}\Lambda_{p,q}$ is the codomain of an unique covering of graphs $\pi \in \text{LG}_1(\mathcal{Q})$ and for each such covering $S(\pi) = 1$.*

PROOF. We will give the argument for $\Gamma_{1,2}$, the argument for the other graphs is similar. The number of preimages of the vertices of $\Gamma_{1,2}$ is $\gcd(k, a_1, a_2)$ for the bottom level and $\gcd(k, a_3, a_4, a_5)$ for the top level, while the edge has $\kappa_{1,2}$ preimages.

We claim that for any cover of graphs $\pi : \hat{\Gamma}_{\text{mp}} \rightarrow \Gamma_{1,2}$ the domain is connected. In fact, suppose there are k' components. This subdivides the top level and the bottom level into subset of equal size. This implies $k' \mid \gcd(k, a_1, a_2)$ and $k' \mid \gcd(k, a_3, a_4, a_5)$, and hence $k' = 1$ because of $\gcd(k, a_1, \dots, a_5) = 1$.

To construct such a cover of graphs it suffices to prescribe one edge of $\hat{\Gamma}_{\text{mp}}$, the other edges are then forced, since τ -acts transitively on edges. Since the vertices on top and

bottom level are indistinguishable (forming each one orbit τ -orbit) the resulting graph is independent of the choice of the first edge. In particular $\widehat{\Gamma}_{\text{mp}}$ is unique and $S(\pi) = 1$. \square

Next we compute (self)-intersection numbers of boundary divisors.

LEMMA 8.3. *The self-intersection numbers of the boundary divisors of $\overline{\mathcal{Q}}$ are*

$$[D_{\Gamma}^{\mathcal{Q}}]^2 = -\frac{\kappa_{i,j}^2}{k^2} - \sum_{\substack{p < q, a_p + a_q < k \\ p, q \notin \{i, j\}}} \frac{\kappa_{i,j} \kappa_{p,q}}{k^2},$$

$$[D_L^{\mathcal{Q}}]^2 = -\frac{\kappa_{i,j}^2}{k^2} \quad \text{and} \quad [D_{\Lambda}^{\mathcal{Q}}]^2 = -\frac{\kappa_{i,j} \kappa_{p,q}}{k^2}.$$

The mutual intersection numbers are

$$[D_{\Gamma}^{\mathcal{Q}}] \cdot [D_L^{\mathcal{Q}}] = \begin{cases} \frac{|\kappa_{i,j} \kappa_{p,q}|}{k^2} & \text{if } \Gamma \cap L \neq \emptyset \\ 0 & \text{otherwise} \end{cases}$$

$$[D_{\Gamma}^{\mathcal{Q}}] \cdot [D_{\Lambda}^{\mathcal{Q}}] = \begin{cases} \frac{\kappa_{i,j} \kappa_{p,q}}{k^2} & \text{if } \Gamma \cap \Lambda \neq \emptyset \\ 0 & \text{otherwise.} \end{cases}$$

PROOF. For the self-intersection numbers consider the formula in Corollary 7.10. As remarked above, the condition (INT) implies that all enhancements of the abelian coverings are 1 and hence the same is true for the $\widehat{\ell}$ -factor in the corollary. Let $\Delta_{i,j}^{p,q}$ denote the slanted cherry with points i, j on bottom level and points p, q on middle level. Together with Corollary 7.8 and Corollary 7.10 we obtain

$$[D_{\Gamma_{i,j}}^{\mathcal{Q}}]^2 = \frac{-1}{k} \zeta^{\top} - c_1(\mathcal{L}^{\top}) = -\frac{\kappa_{i,j}^2}{k^2} \int_{\mathcal{M}_{0,4}} \psi_1 - \sum_{\substack{p < q, a_p + a_q < k \\ p, q \notin \{i, j\}}} [D_{\Delta_{i,j}^{p,q}}^{\mathcal{Q}}].$$

The degree of the slanted cherry is

$$(72) \quad \int_{\overline{\mathcal{Q}}} [D_{\Delta_{i,j}^{p,q}}^{\mathcal{Q}}] = \frac{\kappa_{i,j} \kappa_{p,q}}{k^2}$$

by applying the second formula in Lemma 7.7 and Lemma 8.2. The other numbers are obtained similarly. \square

8.3. The contracted spaces. We want to construct the compactified ball quotient candidate $\overline{\mathfrak{B}}$ from $\overline{\mathcal{Q}}$ by contracting the all the divisors $D_L^{\mathcal{Q}}$ and $D_{\Lambda}^{\mathcal{Q}}$. This is in fact possible:

LEMMA 8.4. *The divisors $D_L^{\mathcal{Q}}$ and $D_{\Lambda}^{\mathcal{Q}}$ of $\overline{\mathcal{Q}}$ are contractible. The DM-stack $\overline{\mathfrak{B}}$ obtained from $\overline{\mathcal{Q}}$ by contracting those divisors is smooth. If $D_L^{\mathfrak{B}}$ and $D_{\Lambda}^{\mathfrak{B}}$ denote the points in $\overline{\mathfrak{B}}$ obtained by contracting the corresponding divisors in $\overline{\mathcal{Q}}$ then*

$$\int_{\overline{\mathfrak{B}}} [D_L^{\mathfrak{B}}] = \frac{\kappa_{i,j}^2}{k^2} \quad \text{and} \quad \int_{\overline{\mathfrak{B}}} [D_{\Lambda}^{\mathfrak{B}}] = \frac{\kappa_{i,j} \kappa_{p,q}}{k^2}.$$

PROOF. For each of the two types of boundary divisors $D_L^{\mathcal{Q}}$ and $D_{\Lambda}^{\mathcal{Q}}$, we will write a neighborhood U as quotient stack $[\widetilde{U}/G]$ with \widetilde{U} smooth, and show that the preimage of the boundary divisor in \widetilde{U} is a \mathbb{P}^1 with self-intersection number -1 . Castelnuovo's criterion then implies that this curve is smoothly contractible. The order of G will be $\frac{k^2}{\kappa_{i,j}^2}$ for $D_L^{\mathcal{Q}}$ and $\frac{k^2}{\kappa_{i,j} \kappa_{p,q}}$ for $D_{\Lambda}^{\mathcal{Q}}$. After contracting the covering \mathbb{P}^1 , the quotient is a point with isotropy group G and the claim on the degrees follows.

We first consider a cherry divisor $D_{\Lambda}^{\mathcal{Q}}$. Let $D_{\Lambda}^{\mathcal{H}_k^{\text{mp}}}$ denote its preimage in $\mathcal{H}_k^{\text{mp}}$. As all the abelian enhancements of the cover of $_{i,j}\Lambda_{p,q}$ are one, the divisor $D_{\Lambda}^{\mathcal{H}_k^{\text{mp}}}$ is irreducible, in fact isomorphic to \mathbb{P}^1 with coordinates the scales of the differential forms on the cherries.

We compute the order of the automorphism group of any point $(\widehat{X}, \widehat{\omega})$ in $D_{\Lambda}^{\mathcal{H}_k^{\text{mp}}}$. Suppose first that $(\widehat{X}, \widehat{\omega})$ is generic. The irreducible components of \widehat{X} group into three τ -orbits: The components \widehat{X}^{\top} corresponding to the top-level vertex of $_{i,j}\Lambda_{p,q}$, the components $\widehat{X}_{i,j}^{\perp}$

corresponding to the vertex with marked points i, j , and the components $\widehat{X}_{p,q}^\perp$ corresponding to the vertex with marked points p, q . Observe that there are $\kappa_{i,j}$ edges between \widehat{X}^\top and $\widehat{X}_{i,j}^\perp$ and $\kappa_{p,q}$ edges between \widehat{X}^\top and $\widehat{X}_{p,q}^\perp$. The restriction of τ to each of the three (not necessarily connected) curves $\widehat{X}^\top, \widehat{X}_{i,j}^\perp, \widehat{X}_{p,q}^\perp$ has order k . Given an automorphism of the complete curve \widehat{X} its restrictions to \widehat{X}^\top and $\widehat{X}_{i,j}^\perp$ need to agree on the $\kappa_{i,j}$ nodes, and the analogue argument applies to $\widehat{X}_{p,q}^\perp$. Hence after fixing the automorphism on the top-level curve \widehat{X}^\top , there are $\frac{k^2}{\kappa_{i,j}\kappa_{p,q}}$ possible choices for the automorphism on the two bottom-level curves left. Together with the k choices for the top-level automorphism, we obtain

$$|\text{Aut}(\widehat{X}, \widehat{\omega})| = \frac{k^3}{\kappa_{i,j}\kappa_{p,q}}.$$

As the non-representable map $\mathcal{H}_k^{\text{mp}} \rightarrow \mathcal{Q}$ has degree $\frac{1}{k}$, this yields that the generic point of $D_\Lambda^\mathcal{Q}$ has an isotropy group of size $r := \frac{k^2}{\kappa_{i,j}\kappa_{p,q}}$. Exactly the same argument also applies to the two boundary points of $D_\Lambda^\mathcal{Q}$ corresponding to the slanted cherries.

The automorphism group is thus generated by multiplying the transversal t -parameter (compare Section 3.4) by an r -th root of unity in local charts covering all of $i,j\Lambda_{p,q}$. We may thus take for U any tubular neighborhood of $D_\Lambda^\mathcal{Q}$ and take a global cover \widetilde{U} of degree $\frac{k^2}{\kappa_{i,j}\kappa_{p,q}}$. Comparing with the degree of the normal bundle in Lemma 8.3 shows that preimage of $D_\Lambda^\mathcal{Q}$ in \widetilde{U} is a (-1) -curve.

We now consider a dumbbell divisor $D_L^\mathcal{Q}$. As above one checks that the isotropy group at the generic point of $D_L^\mathcal{Q}$ is of order $\frac{k}{|\kappa_{i,j}|}$ and that the isotropy groups of the boundary points of the divisor have a quotient group of that order. Consider a tubular neighborhood of $D_L^\mathcal{Q}$ and a degree $\frac{k}{|\kappa_{i,j}|}$ cover that trivializes the isotropy group at the generic point. Let $\widetilde{D}_L^\mathcal{Q}$ be the preimage of the boundary divisor in this cover.

Let p, q, r denote the three marked points on the bottom level of a point in $L_{i,j}$. By applying the above line of arguments again, the three boundary points of $\widetilde{D}_L^\mathcal{Q}$ have cyclic isotropy groups of sizes $\frac{k}{\kappa_{p,q}}, \frac{k}{\kappa_{p,r}}$ and $\frac{k}{\kappa_{q,r}}$ respectively. The triangle group $T = T(\frac{k}{\kappa_{p,q}}, \frac{k}{\kappa_{p,r}}, \frac{k}{\kappa_{q,r}})$ is always spherical, because $a_i + a_j > k$ implies $a_p + a_q + a_r < k$ and hence

$$2 - (1 - \frac{\kappa_{p,q}}{k}) - (1 - \frac{\kappa_{p,r}}{k}) - (1 - \frac{\kappa_{q,r}}{k}) = 2 - 2\frac{a_p + a_q + a_r}{k} > 0.$$

This implies that the T -cover of $\widetilde{D}_L^\mathcal{Q}$ ramified to order $k/\kappa_{p,q}$ along the divisor where $\{p, q\}$ have come together etc, trivializes the isotropy groups on the boundary divisor $\widetilde{D}_L^\mathcal{Q}$ and the preimage of $\widetilde{D}_L^\mathcal{Q}$ is a \mathbb{P}^1 . More precisely, the isotropy groups of order $k/\kappa_{p,q}$ do not fix isolated points on the boundary divisor but have one-dimensional stabilizer, the boundary divisors intersecting $\widetilde{D}_L^\mathcal{Q}$. This implies that the above T -cover actually provides a chart of a full tubular neighborhood.

It remains to show that $|T| = k/|\kappa_{i,j}|$ in order to conclude with the normal bundle degree from Lemma 8.3 that this \mathbb{P}^1 is a (-1) -curve. To show this, recall that as T is spherical, there are only the cases $(\frac{k}{\kappa_{p,q}}, \frac{k}{\kappa_{p,r}}, \frac{k}{\kappa_{q,r}}) = (2, 2, n)$ for $n \in \mathbb{N}_{\geq 2}$ and $(\frac{k}{\kappa_{p,q}}, \frac{k}{\kappa_{p,r}}, \frac{k}{\kappa_{q,r}}) = (2, 3, n)$ for $n \in \{3, 4, 5\}$ to consider. In the first case the order of $T(2, 2, n)$ is $2n$, and assuming that $\frac{k}{\kappa_{p,q}} = \frac{k}{\kappa_{p,r}} = 2$, one easily checks that $2\frac{k}{\kappa_{q,r}} = \frac{k}{|\kappa_{i,j}|}$ by using $\sum_i a_i = 2k$. In the second case the order of $T(2, 3, n)$ is $2\text{lcm}(6, n)$, and the claimed equality follows with a similar argument. \square

We will now compute the Chern classes of $\overline{\mathfrak{B}}$. Let $c : \overline{\mathcal{Q}} \rightarrow \overline{\mathfrak{B}}$ denote the contraction map. Let

$$\mathbf{I} := \{(i, j) : i < j, a_i + a_j < k\} \quad \text{and} \quad \mathbf{L} := \{(i, j) : i < j, a_i + a_j > k\}$$

be the pairs of integers appearing as indices of the $\Gamma_{i,j}$ and $L_{i,j}$. Let $\mathbf{I} = \mathbf{I}_{ij}^{pq}$ denote the common degeneration of Γ_{ij} and L_{pq} , i.e. the three-level graph with points p, q on bottom

level, i, j on top level and the remaining point on the middle level. Accordingly, we write

$$\mathbf{\Lambda} := \{(i, j, p, q) : i < j, i < p < q, j \notin \{p, q\}, a_i + a_j < k, a_p + a_q < k\} \quad \text{and}$$

$$\mathbf{I} := \{(i, j, p, q) : i < j, i < p < q, j \notin \{p, q\}, a_i + a_j > k, a_p + a_q < k\}$$

for the quadruples of possible indices. Recall that D_{hor} is the union of all boundary divisors $D_{H_{ij}}$ whose level graph has a horizontal edge, i.e. corresponding to pairs (i, j) with $a_i + a_j = k$. We write

$$\mathbf{H} := \{(i, j) : i < j, a_i + a_j = k\}.$$

We summarize the intersections of the boundary divisors: The cherry $D_{i,j\Lambda_{p,q}}^{\mathcal{Q}}$ intersects precisely $D_{\Gamma_{ij}}^{\mathcal{Q}}$ and $\Gamma_{pq}^{\mathcal{Q}}$. The divisor $D_{L_{ij}}$ intersects precisely the three divisors $D_{\Gamma_{ab}}^{\mathcal{Q}}$ for any pair (a, b) disjoint from $\{i, j\}$. For the divisor $D_{\Gamma_{ij}}^{\mathcal{Q}}$ consider any pair (p, q) of the three remaining points as $\{p, q, r\}$. This gives an intersection with a cherry if $a_p + a_q < k$, with a horizontal divisor if $a_p + a_q = k$ and with an L -divisor if $a_p + a_q > k$. Consequently, the divisor $D_{H_{ij}}^{\mathcal{Q}}$ intersects precisely the three divisors $D_{\Gamma_{ab}}^{\mathcal{Q}}$ for any pair (a, b) disjoint from $\{i, j\}$.

LEMMA 8.5. *The self-intersection numbers of the boundary divisors of $\overline{\mathfrak{B}}$ are*

$$[D_{\Gamma_{i,j}}^{\mathfrak{B}}]^2 = -\frac{\kappa_{i,j}^2}{k^2} + \sum_{\substack{p < q, a_p + a_q > k \\ p, q \notin \{i, j\}}} \frac{\kappa_{i,j}^2}{k^2} \quad \text{and} \quad [D_{H_{i,j}}^{\mathfrak{B}}]^2 = -1.$$

The mutual intersection numbers are for $\{i, j\} \cap \{p, q\} = \emptyset$ given by

$$[D_{\Gamma_{i,j}}^{\mathfrak{B}}] \cdot [D_{\Gamma_{p,q}}^{\mathfrak{B}}] = \frac{\kappa_{i,j} \kappa_{p,q}}{k^2} \quad \text{and} \quad [D_{\Gamma_{i,j}}^{\mathfrak{B}}] \cdot [D_{H_{p,q}}^{\mathfrak{B}}] = \frac{\kappa_{i,j}}{k}$$

and for $|\{i, j, p\}| = 3$ by

$$[D_{\Gamma_{i,j}}^{\mathfrak{B}}] \cdot [D_{\Gamma_{i,p}}^{\mathfrak{B}}] = \begin{cases} \frac{\kappa_{i,j} \kappa_{i,p}}{k^2} & \text{if } a_i + a_j + a_p < k \\ 0 & \text{otherwise.} \end{cases}$$

PROOF. We claim that the pull back of $[D_{\Gamma_{i,j}}^{\mathfrak{B}}]$ is given by

$$c^*[D_{\Gamma_{i,j}}^{\mathfrak{B}}] = [D_{\Gamma_{i,j}}^{\mathcal{Q}}] + \sum_{\substack{p < q, a_p + a_q > k \\ p, q \notin \{i, j\}}} \frac{\kappa_{i,j}}{|\kappa_{p,q}|} [D_{L_{p,q}}^{\mathcal{Q}}] + \sum_{\substack{p < q, a_p + a_q < k \\ p, q \notin \{i, j\}}} [D_{i,j\Lambda_{p,q}}^{\mathcal{Q}}].$$

To determine the coefficients in the above expression, one may intersect the equation $c^*[D_{\Gamma_{i,j}}^{\mathfrak{B}}] = [D_{\Gamma_{i,j}}^{\mathcal{Q}}] + \sum_{p,q} l_{p,q} [D_{L_{p,q}}^{\mathcal{Q}}] + \sum_{p,q} \lambda_{p,q} [D_{i,j\Lambda_{p,q}}^{\mathcal{Q}}]$ with unknown coefficients with each of the divisors $[D_{L_{p,q}}^{\mathcal{Q}}]$ and $[D_{i,j\Lambda_{p,q}}^{\mathcal{Q}}]$ in turn. The left hand side vanishes by push-pull, and the intersection numbers on the right hand side are given by Lemma 8.3. The claimed intersection numbers involving only Γ -divisors follow again by Lemma 8.3.

The pull back of the horizontal divisor is given by $c^*[D_{H_{i,j}}^{\mathfrak{B}}] = [D_{H_{i,j}}^{\mathcal{Q}}]$. The intersection number $[D_{\Gamma_{i,j}}^{\mathfrak{B}}] \cdot [D_{H_{p,q}}^{\mathfrak{B}}] = [D_{\Gamma_{i,j}}^{\mathcal{Q}}] \cdot [D_{H_{p,q}}^{\mathcal{Q}}]$ follows from Lemma 7.7 and Lemma 8.2. Finally by Proposition 4.5 and (68), the normal bundle of $[D_{H_{i,j}}^{\mathcal{Q}}]$ is given by $-\psi_e$ in $\text{CH}(D_{H_{i,j}}^{\mathcal{Q}})$, where ψ_e is the ψ -class supported on the half edge of $H_{i,j}$ that is adjacent to the vertex with three adjacent marked points. \square

PROPOSITION 8.6. *The log canonical bundle on $\overline{\mathfrak{B}}$ has first Chern class*

$$(73) \quad c_1(\Omega_{\overline{\mathfrak{B}}}^1(\log D_{\text{hor}})) = \sum_{i,j \in \Gamma} \left(\frac{k}{2\kappa_{i,j}} - 1 \right) [D_{\Gamma_{i,j}}^{\mathfrak{B}}] + \frac{1}{2} [D_{\text{hor}}^{\mathfrak{B}}] \quad \text{in } \text{CH}_1(\mathfrak{B})$$

Its square and the second Chern class are given by

$$(74) \quad c_1(\Omega_{\overline{\mathfrak{B}}}^1(\log D_{\text{hor}}))^2 = 6 - 3 \sum_{i,j \in \Gamma} \frac{\kappa_{i,j}}{k} + 3 \sum_{i,j \in \mathbf{L}} \frac{\kappa_{i,j}^2}{k^2} + 3 \sum_{i,j,p,q \in \mathbf{\Lambda}} \frac{\kappa_{i,j} \kappa_{p,q}}{k^2}$$

and

$$(75) \quad c_2(\Omega_{\mathfrak{B}}^1(\log D_{\text{hor}})) = 2 - \sum_{i,j \in \Gamma} \frac{\kappa_{i,j}}{k} + \sum_{i,j \in \mathbf{L}} \frac{\kappa_{i,j}^2}{k^2} + \sum_{i,j,p,q \in \Lambda} \frac{\kappa_{i,j} \kappa_{p,q}}{k^2}.$$

respectively.

PROOF. To derive (73) from Theorem 1.1 we insert into

$$c_1(\Omega_{\mathcal{Q}}^1(\log D_{\text{hor}})) = \frac{3}{k} \cdot \zeta + \sum_{\mathbf{L}} [D_{\mathbf{L}}^{\mathcal{Q}}] + \sum_{\Lambda} [D_{\Lambda}^{\mathcal{Q}}]$$

that $5\xi - \sum(m_i + k)\psi_i$ is a sum of boundary terms by the relation (7.8). Consider Keel's relation

$$\psi_i = \frac{1}{6} \sum_{\substack{c < d \\ i \notin \{c,d\}}} \Delta_{cd} + \frac{1}{2} \sum_{a \neq i} \Delta_{ia},$$

where Δ_{ij} is the boundary divisor in $\overline{\mathcal{M}}_{0,5}$ where the points (i, j) have come together. We pull back this relation via the forgetful map $\pi : \mathbb{P}\Xi^k \overline{\mathcal{M}}_{0,5}(\mu) \rightarrow \overline{\mathcal{M}}_{0,5}$. Since this map is a root-stack construction and the isotropy groups of the divisors were computed in the proof of Lemma 8.4, we obtain

$$\pi^* \Delta_{ab} = \begin{cases} \frac{1}{|\kappa_{ab}|} [D_{\mathbf{L}_{ab}}^{\mathcal{Q}}] & \text{if } a + b < -k \\ [D_{\mathbf{H}_{ab}}] & \text{if } a + b = -k \\ \frac{1}{\kappa_{ab}} [D_{\Gamma_{ab}}^{\mathcal{Q}}] + \sum_{\substack{i < j, a_i + a_j < k \\ i,j \notin \{a,b\}}} \frac{1}{\kappa_{ab}} [D_{i,j\Lambda_{a,b}}^{\mathcal{Q}}] & \text{if } a + b > -k. \end{cases}$$

Putting everything together we find in $\text{CH}_1(\mathcal{Q})$ that

$$(76) \quad c_1(\Omega_{\mathcal{Q}}^1(\log D_{\text{hor}})) = \sum_{i,j \in \Gamma} \left(\frac{k}{2\kappa_{i,j}} - 1 \right) [D_{\Gamma_{i,j}}^{\mathcal{Q}}] + \sum_{i,j \in \mathbf{L}} \left(\frac{k}{2|\kappa_{i,j}|} - 1 \right) [D_{\mathbf{L}_{i,j}}^{\mathcal{Q}}] \\ + \sum_{i,j,p,q \in \Lambda} \left(\frac{k}{2\kappa_{i,j}} + \frac{k}{2\kappa_{p,q}} - 1 \right) [D_{i,j\Lambda_{p,q}}^{\mathcal{Q}}] + \frac{1}{2} [D_{\text{hor}}^{\mathcal{Q}}]$$

and since the divisors $D_{\mathbf{L}_{i,j}}^{\mathcal{Q}}$ and $D_{i,j\Lambda_{p,q}}^{\mathcal{Q}}$ are smoothly contractible we deduce (73).

To derive (74) we first note that $-\frac{1}{4}|\Gamma| + \frac{1}{2}|\Lambda| + \frac{5}{4}|\mathbf{H}| + \frac{5}{4}|\mathbf{L}| = 5$ and that for $(i, j) \in \mathbf{L}$ the relation

$$1 + \sum_{\substack{p \in \{1, \dots, 5\} \setminus \{i,j\} \\ \{q,r\} = \{1, \dots, 5\} \setminus \{i,j,p\}}} \left(-\frac{\kappa_{p,q} + \kappa_{p,r}}{k} + 2\frac{\kappa_{p,q}\kappa_{p,r}}{k^2} + \frac{\kappa_{q,r}^2}{k^2} \right) = 4\frac{\kappa_{i,j}^2}{k^2}$$

holds because of $\sum_i a_i = 2k$. Using those relations and the intersection numbers in Lemma 8.5 squaring (73) yields

$$c_1(\Omega_{\mathfrak{B}}^1(\log D_{\text{hor}}))^2 = 5 - \sum_{i,j \in \Gamma} \left(2\frac{\kappa_{i,j}}{k} + \frac{\kappa_{i,j}^2}{k^2} \right) + 2 \sum_{i,j,p,q \in \Lambda} \frac{\kappa_{i,j} \kappa_{p,q}}{k^2} + 4 \sum_{i,j \in \mathbf{L}} \frac{\kappa_{i,j}^2}{k^2}$$

and (74) follows because $\sum_i a_i = 2k$ implies

$$(77) \quad 1 + \sum_{i,j \in \Gamma} \left(-\frac{\kappa_{i,j}}{k} + \frac{\kappa_{i,j}^2}{k^2} \right) + \sum_{i,j,p,q \in \Lambda} \frac{\kappa_{i,j} \kappa_{p,q}}{k^2} - \sum_{i,j \in \mathbf{L}} \frac{\kappa_{i,j}^2}{k^2} = 0.$$

The second Chern class can be computed as

$$c_2(\Omega_{\mathfrak{B}}^1(\log D_{\text{hor}})) = \chi(\mathcal{M}_{0,5}) + \sum_{i,j \in \Gamma} \chi(D_{\Gamma_{i,j}}^{\mathfrak{B}, \circ}) + \sum_{i,j \in \mathbf{L}} \chi(D_{\mathbf{L}_{i,j}}^{\mathfrak{B}}) + \sum_{i,j,p,q \in \Lambda} \chi(D_{i,j\tilde{\Lambda}_{p,q}}^{\mathfrak{B}}),$$

where $\chi(D_{\Gamma_{i,j}}^{\mathfrak{B}, \circ}) = \chi(D_{\Gamma_{i,j}}^{\mathcal{Q}, \circ}) = \frac{\kappa_{i,j}}{k}$ be Lemma 7.7 and Lemma 8.2 and the Euler characteristics of the points are given in Lemma 8.4. \square

8.4. The ball quotient certificate. We can finally put together the previous intersection numbers and use our ball quotient criterion to show that the contracted spaces are ball quotients.

PROOF OF THEOREM 1.7. We apply Proposition 8.1 and check that first that the only log-exceptional curves for $c_1(\Omega_{\overline{\mathfrak{B}}}^1(\log D_{\text{hor}}))$ are the components of D_{hor} . In fact since the expression (73) is an effective divisor and since $\overline{\mathfrak{B}} \setminus \mathcal{D} \cong \mathcal{M}_{0,5}$ is affine, we only have to check positivity of c_1^2 and the intersection with $D_{H_{ab}}$ and $D_{\Gamma_{i,j}}^{\mathfrak{B}}$. For the $D_{\Gamma_{i,j}}^{\mathfrak{B}}$ -intersections this follows from the intersection numbers in Lemma 8.5. In fact, the self-intersection number of $D_{\Gamma_{i,j}}^{\mathfrak{B}}$ is negative only if $a_p + a_q \leq k$ for any pair $\{p, q\}$ disjoint from $\{i, j\}$. Using Lemma 8.3 we compute in this case that

$$[D_{\Gamma_{i,j}}^{\mathfrak{B}}] \cdot c_1(\Omega_{\overline{\mathfrak{B}}}^1(\log D_{\text{hor}})) = \frac{\kappa_{ij}}{k} \left(\frac{2a_p + 2a_q + 2a_r - a_i - a_j}{k} - 1 \right),$$

where $\{a_1, a_2, a_3, a_4, a_5\} = \{a_i, a_j, a_p, a_q, a_r\}$. Since $a_i + a_j < k$, this expression is positive. Moreover, one directly computes

$$[D_{H_{a,b}}] \cdot c_1(\Omega_{\overline{\mathfrak{B}}}^1(\log D_{\text{hor}})) = 0.$$

That $c_1(\Omega_{\overline{\mathfrak{B}}}^1(\log D_{\text{hor}}))^2 > 0$ is a consequence of the above, as $c_1(\Omega_{\overline{\mathfrak{B}}}^1(\log D_{\text{hor}}))$ is by Equation (73) a linear combination of the divisors $D_{\Gamma_{i,j}}^{\mathfrak{B}}$ and $D_{\text{hor}}^{\mathfrak{B}}$ with positive coefficients. \square

The multi-scale boundary of the gothic locus

1. Introduction

Let $\mu = (m_1, \dots, m_n) \in \mathbb{N}^n$ such that $\sum_i m_i = 2g - 2$. We denote by $\mathbb{P}\Omega\mathcal{M}_{g,n}(\mu)$ the projectivized stratum of abelian differentials (X, ω) , where X is a Riemann surface of genus g and ω is an one-form with zeros as prescribed by μ . Affine invariant subvarieties of $\Omega\mathcal{M}_{g,n}(\mu)$, or equivalently $\mathrm{GL}_2(\mathbb{R})^+$ -orbit closures, are locally given by \mathbb{R} -linear equations in period coordinates. One orbit closures of special interest is the gothic locus $\Omega G \subseteq \Omega\mathcal{M}_{4,6}(0^3, 2^3)$, first described in [MMW17]. The gothic locus ΩG was the first known example of a primitive rank two \mathbb{R} -linear submanifold and counterexamples to an earlier conjecture of Mirzakhani. The interest in this locus stems moreover from the fact that it contains a dense set of formerly unknown primitive Teichmüller curves. The name gothic locus origins from [MMW17], as a translation surface in the locus resembles the outline of a gothic cathedral.

Even after projectivization, the moduli spaces $\mathbb{P}\Omega\mathcal{M}_{g,n}(\mu)$, and in particular orbit closures, are in general non-compact. A well-behaved compactification of $\mathbb{P}\Omega\mathcal{M}_{g,n}(\mu)$, the so called *moduli space of multi-scale differentials* $\mathbb{P}\Xi\overline{\mathcal{M}}_{g,n}(\mu)$, has been constructed by Bainbridge-Chen-Gendron-Grushevsky-Möller [BCGGM19b]. By a recent result of Benirschke [Ben20], the boundary of a $\mathrm{GL}_2(\mathbb{R})^+$ -orbit closure in the moduli space of multi-scale differentials is locally given by \mathbb{R} -linear equations, but few non-trivial examples of such boundaries appear in the literature. In this article, which is still work in progress, we will highlight some aspects of the boundary of the closure of the gothic locus $\mathbb{P}\Xi\overline{G} := \overline{\mathbb{P}\Omega G} \subseteq \mathbb{P}\Xi\overline{\mathcal{M}}_{4,6}(0^3, 2^3)$. The long term goal of this project is to compute the fundamental class of the gothic locus $\mathbb{P}\Xi\overline{G}$, the Euler characteristic of $\mathbb{P}\Omega G$ by using the results in Chapter II, a complete description of the boundary of $\mathbb{P}\Xi\overline{G}$, and the number of ends of ΩG .

Our first statement is about the intersection of the gothic locus with non-horizontal boundary divisors.

PROPOSITION 1.1 (Proposition 5.1). *The gothic locus $\mathbb{P}\Xi\overline{G}$ intersects the non-horizontal boundary divisors $D_{\Gamma_4}, \dots, D_{\Gamma_9}$ whose dual graphs are listed in Figure 1.*




FIGURE 1. Vertical divisors intersected by the gothic locus

REMARK 1.2. *The gothic locus $\mathbb{P}\Xi\overline{G}$ might intersect additional non-horizontal boundary divisor. Those divisors are listed in Section 4.*

The intersection of the gothic locus with a non-horizontal boundary divisor consists of a \mathbb{R} -linear submanifold on every level. For the non-horizontal strata $D_{\Gamma_4}, \dots, D_{\Gamma_9}$ we will see in Section 5 that the top-levels of the intersection with the gothic locus are well-known \mathbb{R} -linear submanifolds.

Recall that a *Teichmüller curve* is an immersed algebraic curve $C \rightarrow \mathcal{M}_g$ which is the image under the forgetful map of a 2-dimensional variety $M \rightarrow \Omega\mathcal{M}_{g,n}(\mu)$ which is locally cut out by \mathbb{R} -linear equations in period coordinates. In strata of holomorphic differentials several equivalent characterisations of Teichmüller curves exist. However, in strata of meromorphic differentials those characterizations do no longer agree, and we are using the above as our definition. If \mathcal{T} is a Teichmüller curve and $(X, \omega) \in \mathcal{T}$ is a differential, then in the abelian case the $\mathrm{GL}_2(\mathbb{R})^+$ -orbit of (X, ω) will agree with \mathcal{T} . In the meromorphic case this is no longer true: in fact the $\mathrm{GL}_2(\mathbb{R})^+$ -orbit is never equal to \mathcal{T} . Instead it sweeps out only one of the chambers of \mathcal{T} bounded by loci of parallel saddle connections; see [MM23] for details. Following [MM23], we say that the differential $(X, \omega) \in \mathcal{T}$ *generates* \mathcal{T} if its $\mathrm{GL}_2(\mathbb{R})^+$ -orbit is equal to \mathcal{T} on an open subset of \mathcal{T} .

A Teichmüller curve in a stratum of meromorphic differentials is called *obvious* if it is the intersection of a Hurwitz space above another stratum of abelian differentials and a locus prescribed by residue conditions. By analyzing the bottom levels of the intersection of the gothic locus with D_{Γ_4} in Section 6 we obtain an example of a non-obvious Teichmüller curve.

THEOREM 1.3. *Let $(X, \omega) \subseteq \Omega\mathcal{M}_{1,6}(-3^2, 2^3)$ be the canonical cover of the 6-differential of type $(-10, -5, 3)$. The differential (X, ω) generates a non-obvious Teichmüller curve. In the chart in Figure 2 this Teichmüller curve is given by the equations*

$$w_i = -w_{i+3} \quad \text{for } i = 1, 2, 3 \quad \text{and} \quad w_1 + w_3 + w_5 = 0.$$




FIGURE 2. A surface of infinite area generating a non-obvious Teichmüller curve in the stratum $\Omega\mathcal{M}_{1,6}(-3^2, 2^3)$

The fact that the generating differential in Theorem 1.3 is the canonical cover of a 6-differential is a shadow of the fact that the gothic locus is the $\mathrm{GL}_2(\mathbb{R})^+$ -orbit closure of the locus of unfoldings of quadrilaterals with angles $\frac{\pi}{6}, \frac{\pi}{6}, \frac{\pi}{6}, \frac{3\pi}{2}$, or equivalently the $\mathrm{GL}_2(\mathbb{R})^+$ -orbit closure of the canonical covers of the stratum of 6-differentials $\Omega^6\mathcal{M}_{0,4}(-5^3, 3)$. The 6-differential of type $(-10, -5, 3)$ corresponds to the bottom level of the (up to permutation of the marked points) unique boundary divisor of $\mathbb{P}\Xi^6\mathcal{M}_{0,4}(-5^3, 3)$. The other rank 2 orbit closures constructed in [EMMW20] arise similarly to the gothic locus as the orbit closures of canonical covers of strata of k -differentials in genus zero. We expect that an analysis of the boundary divisors of those strata will yield more non-obvious Teichmüller curves which are generated by canonical covers of k -differentials.

For the horizontal boundary divisors, we can provide a list of all boundary strata that might possibly intersect the gothic locus.

PROPOSITION 1.4 (Proposition 7.1). *The gothic locus $\mathbb{P}\Xi\overline{G}$ only intersects the horizontal strata listed in Figure 3.*




FIGURE 3. The purely horizontal boundary strata in the gothic locus

In some sense the intersections with the horizontal boundary divisors is a bit more subtle. For example it is not straight forward to calculate the dimension of those intersections. In Section 7 we will prove partial results that back the following expectation.

EXPECTATION 1.5 (Expectation 7.2). *We expect the following:*

- *The gothic locus intersects the three strata D_{Γ_2} , D_{Γ_3} and $D_{\Gamma_{15}}$ (depicted in the top row of Figure 3) in a divisor.*
- *The gothic locus intersects the three strata D_{Γ_1} , $D_{\Gamma_{19}}$ and $D_{\Gamma_{18}}$ (depicted in the middle row of Figure 3) in codimension 2.*
- *The gothic locus intersects the two strata $D_{\Gamma_{20}}$ and $D_{\Gamma_{22}}$ (depicted in the bottom row of Figure 3) in codimension 3.*
- *The gothic locus does not intersect the two strata $D_{\Gamma_{17}}$ and $D_{\Gamma_{21}}$.*

Recall from above that the gothic locus contains a dense set of Teichmüller curves. Those curves are non-compact, and the points in the boundary are called cups.

PROPOSITION 1.6 (Proposition 7.6). *The interior of each of the four horizontal strata $D_{\Gamma_1}^G$, $D_{\Gamma_2}^G$, $D_{\Gamma_3}^G$ and $D_{\Gamma_{20}}^G$ contains cusps of a primitive Teichmüller curve contained in the gothic locus ΩG . The interior of the stratum $D_{\Gamma_{19}}^G$ contains cusps of a non-primitive Teichmüller curve.*

To obtain the information about the boundary we proceed as follows. As we will recall in Section 2, the gothic locus can be defined, following [EMMW20], via a certain Hurwitz space \mathcal{H} of dihedral covers. There is a subspace $\mathcal{D} \subseteq \mathcal{H}$ of codimension 1 corresponding to the gothic locus. Let $\overline{\mathcal{H}}$ denote the admissible covers compactification and $\overline{\mathcal{D}} \subseteq \overline{\mathcal{H}}$ the closure of \mathcal{D} therein. Then there are two natural forgetful maps

$$\begin{array}{ccc} \mathbb{P}\Xi\overline{G} & & \overline{\mathcal{D}} \\ & \searrow & \swarrow \\ & \overline{\mathcal{M}}_{4,6} & \end{array}$$

whose images agree. In Section 3 we will analyze which of the boundary divisors of $\overline{\mathcal{H}}$ are intersected by $\overline{\mathcal{D}}$. In Section 4 we will relate this to the boundary of the gothic locus $\mathbb{P}\Xi\overline{G}$ and obtain a list of all strata of the ambient moduli space that can possibly contain

a boundary divisor of the gothic locus. The results in Section 4 rely heavily on a computer program written in Sage that will be made available separately.

In the last Section 8, we outline a possible approach to compute the fundamental class of the gothic locus. This would allow for the computation of its Euler characteristic via Theorem II.1.3.

2. The gothic locus

We will recall the different definitions of the gothic locus ΩG from [MMW17] and [EMMW20], and collect some facts that will be useful in the sequel.

2.1. Quadrilaterals and period coordinates. We begin by giving a geometric definition of the gothic locus and recalling its equation in local period coordinates. We denote by ΩZ_G the family of curves in $\Omega M_{4,6}(0^3, 2^3)$ that is obtained by unfolding the quadrilaterals with angles $\frac{\pi}{6}, \frac{\pi}{6}, \frac{\pi}{6}, \frac{3\pi}{2}$. We will refer to ΩZ_G as the *cyclic locus*. Via cutting and gluing of the unfolded polygon one will eventually arrive at the polygons depicted in Figure 4.




FIGURE 4. A curve in the locus of cyclic forms ΩZ_G

DEFINITION 2.1 (Gothic locus). *We define the gothic locus ΩG as the orbit closure*

$$\Omega G := \overline{\mathrm{GL}_2(\mathbb{R})^+ \cdot \Omega Z_G} \subseteq \Omega \mathcal{M}_{4,6}(0^3, 2^3).$$

REMARK 2.2. *In contrast to [MMW17] and [EMMW20], we explicitly mark the three fixed points of the cyclic deck transformation.*

Let J denote the involution of the surface depicted in Figure 4 that is given by rotating the polygons by π . Observe that this involution must exist on all curves in the gothic locus ΩG . Let $\Sigma := Z(\omega) \cup \mathrm{Fix}(J)$. We regard the vectors $v_i, w_i, \alpha, \beta, \gamma$ in Figure 4 as period coordinates in $H^1(X; \Sigma)$. We choose the orientation of v_i, w_i as counterclockwise in the hexagons and the orientation of α, β and γ from left to right. This 15 vectors span $H^1(X; \Sigma) \cong \mathbb{C}^{13}$ with the two relations

$$\sum v_i = \sum w_i = 0.$$

In these coordinates the gothic locus ΩG is locally defined by the equations

$$(78) \quad \begin{aligned} v_i &= -v_{i+3}, & w_i &= -w_{i+3} & \text{for } i = 1, 2, 3, \\ v_1 + v_3 + v_5 &= w_1 + w_3 + w_5 = 0, \\ \alpha + v_4 &= \beta + w_4 = \alpha + \beta - \gamma = 0. \end{aligned}$$

The first two lines of equations can be found in [MMW17, Equation 9.2]. For the third line notice that the action of J implies $\alpha' = \alpha, \beta' = \beta$ and $\gamma' = \gamma$. We can use this to simplify the obvious relations $\alpha' + \alpha + v_3 + v_4 + v_5 = 0, \beta' + \beta + w_3 + w_4 + w_5 = 0$ and $\gamma' + \gamma - w_5 - v_5 - w_4 - v_4 - w_3 - v_3 = 0$.

The cyclic forms $\Omega Z_G \subseteq \Omega G$ are cut out by the additional non-linear equations

$$(79) \quad \begin{aligned} |v_i| &= |v_j| & \text{for } 1 \leq i, j \leq 6 \\ |w_i| &= |w_j| & \text{for } 1 \leq i, j \leq 6 \\ v_1 &= \lambda \cdot w_1 & \text{for some } \lambda \in \mathbb{R}_{>0}. \end{aligned}$$

2.2. Dihedral triples. We will now recall from [EMMW20] the relation of the gothic locus to a certain Hurwitz space of dihedral covers. Let

$$D_{12} = \langle r, f \mid r^6 = f^2 = (rf)^2 = e \rangle$$

denote the dihedral group of order 12. Let $Y \rightarrow \mathbb{P}^1$ be a normal D_{12} -cover with monodromy datum

$$(f, f, f, rf, rf, rf, r^3),$$

let $X := Y/f$ denote the quotient, and let $\pi: X \rightarrow \mathbb{P}^1$ be the induced (non-Galois) cover.

We denote by \mathcal{H} the Hurwitz space of those covers $\pi: X \rightarrow \mathbb{P}^1$. Moreover, we denote the fiber above the branch point with ramification associated to r^3 by $X^* \subseteq X$ and refer to it as the *special fiber* of π . One checks that $|X^*| = 3$ and $g(X) = 4$.

We denote by $\mathcal{D} \subseteq \mathcal{H}$ the locus of covers $(\pi: X \rightarrow \mathbb{P}^1) \in \mathcal{H}$ such that X admits an one-form $\omega \in \Omega(X)$ subject to

$$(80) \quad (\omega) = 2 \cdot X^* \quad \text{and} \quad (r + r^{-1})^*(\omega) = (\zeta_6 + \zeta_6^{-1}) \cdot \omega,$$

where $\zeta_6 = e^{2\pi i/6}$. We call the triples (X, ω, π) satisfying (80) *dihedral triples*, and we call \mathcal{D} the *dihedral locus*. As $fr^3 = r^3f$, the action of r^3 on Y commutes with the action of f . Thus the action of r^3 descends to an involution on $X = Y/f$. We will denote this involution by J . The involution J has 3 fixed points, one in each fiber above the ramification points corresponding to rf .

We denote by $S_{3+3} := S_{\{1,2,3\}} \times S_{\{4,5,6\}} \subseteq S_6$ the indicated subgroup of the symmetric group. The Hurwitz space \mathcal{H} comes with a map

$$(81) \quad \phi: \mathcal{H} \rightarrow \mathcal{M}_{4,6}/S_{3+3}$$

where the three fixed points of J are mapped to the points 1, 2, 3 and the three points in the special fiber are mapped to the points 4, 5, 6. If we denote by $\Omega\tilde{D}$ the set of all dihedral triples, we have a map

$$\Omega\tilde{D} \rightarrow \Omega\mathcal{M}_{4,6}(0^3, 2^3)/S_{3+3}$$

and we denote by

$$\Omega D \subseteq \Omega\mathcal{M}_{4,6}(0^3, 2^3)$$

the preimage of $\Omega\tilde{D}$ under the quotient map by S_{3+3} .

Recall that we defined the gothic locus ΩG in Definition 2.1 as the $\mathrm{GL}_2(\mathbb{R})^+$ -orbit closure of certain cyclic forms. The locus ΩD gives us another description.

THEOREM 2.3 ([EMMW20, Theorem 5.3]). *The gothic locus is the closure*

$$\Omega G = \overline{\Omega D} \subseteq \Omega\mathcal{M}_{4,6}(0^3, 2^3)$$

of the dihedral locus.

Note that the cyclic forms ΩZ_G are not contained in ΩD (the corresponding covers $\pi: X \rightarrow \mathbb{P}^1$ would need to be Galois), but are in fact contained in the closure $\Omega G = \overline{\Omega D}$ by [EMMW20, Theorem 4.2]. Using the fact that $Y \rightarrow \mathbb{P}^1$ is a D_{12} -cover, we can form the commutative diagram

$$(82) \quad \begin{array}{ccc} & Y & \\ & \downarrow & \\ & X = Y/f & \\ \swarrow & & \searrow \\ A := X/J = Y/\langle f, r^3 \rangle & & B := Y/\langle f, r^2 \rangle \\ \swarrow & & \searrow \\ & \mathbb{P}^1 \cong Y/D_{12} & \end{array}$$

Observe that the map j is of degree 2 and thus Galois. The genera of the quotient curves is $g(A) = g(B) = 1$.

The Hurwitz space \mathcal{H} comes with the target morphisms $\delta: \mathcal{H} \rightarrow \mathcal{M}_{0,7}$. The image of the dihedral locus \mathcal{D} under this morphism will be the main object of our interest in the next section.

DEFINITION 2.4 (Moduli space of critical values). *Let*

$$V_G := \delta(\mathcal{D}) \subseteq \mathcal{M}_{0,7}$$

denote the image of the dihedral locus. We call V_G the moduli space of critical values.

The dimension of the moduli space V_G has been computed in [EMMW20].

THEOREM 2.5 ([EMMW20, Theorem 4.4]). *The moduli space of critical values is of dimension $\dim(V_G) = 3$. In particular is V_G a divisor in $\mathcal{M}_{0,7}$ and $\dim(\Omega G) = 4$.*

2.3. Original definition. We obtained the maps in Diagram (82) as a consequence of the construction of the gothic locus ΩG via the Hurwitz space \mathcal{H} . To complete the picture, we recall that the maps in the diagram can in fact be used to define the gothic locus in the first place. In fact, this was the first published description.

To this end, let $\Omega\mathcal{M}_{4,6}(2^3, 0^3)^-$ denote the subvariety of $\Omega\mathcal{M}_{4,6}(2^3, 0^3)$ where on the curves (X, ω) there exists an involution $J: X \rightarrow X$ that fixes all the marked points and such that ω is J -antiinvariant. For such a curve X we say that a holomorphic map $p: X \rightarrow B$ is *odd*, if there exists an involution $j: B \rightarrow B$ such that $p \circ J = j \circ p$. In [MMW17], the gothic locus was defined as

$$\Omega G = \left\{ (X, \omega) \in \Omega\mathcal{M}_{4,6}(2^3, 0^3)^- \left| \begin{array}{l} \exists \text{ a curve } B \in \mathcal{M}_1 \text{ and an odd,} \\ \text{degree three rational map } p: X \rightarrow B \\ \text{such that } |p(Z(\omega))| = 1. \end{array} \right. \right\}.$$

This definition agrees with our previous definitions as was shown in [EMMW20].

2.4. The quadratic quotients. For an abelian differential $(X, \omega) \in \Omega G$ the differential ω is J -antiinvariant by (80). Therefore, the quotient $(X, \omega^2)/J$ gives a well-defined quadratic differential. We denote the quotient map by

$$j: \Omega G \rightarrow \mathcal{Q}_{1,6}(-1^3, 1^3).$$

and refer to the image $\mathcal{Q}G := j(\Omega G)$ as *the quadratic gothic locus*

The stratum $\mathcal{Q}_{1,6}(-1^3, 1^3)$ has a natural forgetful map to $\mathcal{M}_{1,3}$ by forgetting the markings at the zeros and only remembering the markings at the simple poles. The image of $\mathcal{Q}G$ under this map is the so-called flex locus F . It is of dimension 2 and was the first known example of a totally geodesic surface, see [MMW17] for details. Its fundamental class has been computed by Chen.

THEOREM 2.6 ([Che22, Theorem 1.1]). *The fundamental class of the flex locus is*

$$[F] = \frac{4}{3}\delta_{\text{irr}} + 4(\delta_{0;\{1,2\}} + \delta_{0;\{1,3\}} + \delta_{0;\{2,3\}}) + 4\delta_{0;\{1,2,3\}} \in R^1(\overline{\mathcal{M}}_{1,3}).$$

3. The moduli space of critical values

To understand the compactification of the gothic locus ΩG , it will be useful to analyze the closure of the moduli space of critical values V_G inside the Deligne-Mumford compactification $\overline{\mathcal{M}}_{0,7}$. This will help us later on to understand the admissible covers compactification $\overline{\mathcal{D}} \subseteq \overline{\mathcal{H}}$. The main goal of this section is to prove Corollary 3.3.

REMARK 3.1. *In theory it should be possible to push the techniques in this section further to directly determine the complete list of boundary strata of the gothic locus $\mathbb{P}\Xi\overline{G}$. However this is a very tedious and error-prone task. Instead we hope to determine the complete list of boundary strata a posteriori from the fundamental class, see Remark 8.4.*

By a result of [EMMW20], the dihedral triples $\Omega\tilde{\mathcal{D}}$ can be parametrized by certain tuples of polynomials $P_G \subseteq (\mathbb{C}[x])^2$. We will reformulate this result in terms of triples of polynomials $B_G \subseteq (\mathbb{C}[x])^3/\mathbb{C}^\times$ that parametrize the projectivization $\mathbb{P}\Omega\tilde{\mathcal{D}}$. Our parametrization has the advantage that the polynomials do depend more directly on the branch points of the cover $\pi: X \rightarrow \mathbb{P}^1$ and allows us to show:

PROPOSITION 3.2. *The point $(b_1, \dots, b_6, \infty) \in \mathcal{M}_{0,7}$ belongs to the subvariety V_G if and only if there exist $a, b \in \mathbb{C}$ and $c_1 \in \mathbb{C}^\times$, such that*

$$(x - b_4)(x - b_5)(x - b_6) - c_1(x - b_1)(x - b_2)(x - b_3) = (ax + b)^3.$$

Observe that the expression is symmetric in b_1, b_2, b_3 and b_4, b_5, b_6 as expected. To describe which boundary divisors of $\overline{\mathcal{M}}_{0,7}$ are intersected by the gothic locus ΩG it is therefore natural to work on the quotient $\overline{\mathcal{M}}_{0,7}/S_{3+3}$, where $S_{3+3} := S_{\{1,2,3\}} \times S_{\{4,5,6\}} \subseteq S_7$. We denote by $\Delta_{i,j} \subset \overline{\mathcal{M}}_{0,7}/S_{3+3}$ the reducible boundary divisor where i of the marked points 1, 2, 3 and j of the marked points 4, 5, 6 lie on one irreducible component.

COROLLARY 3.3. *The variety \overline{V}_G/S_{3+3} does not intersect the boundary divisors $\Delta_{1,2}$ and $\Delta_{2,1}$.*

PROOF. We will discuss the case $\Delta_{2,1}$, the case $\Delta_{1,2}$ follow by essentially the same argument. In Proposition 3.10 we will see that we may assume that $b_1 = b_2 = b_4 = 0$, which implies $b = 0$. After dividing by x we are left with the equation

$$(x - b_5)(x - b_6) - c_1 x(x - b_3) = a^3 x^2.$$

Hence $b_5 = 0$ or $b_6 = 0$, a contradiction. \square

REMARK 3.4. *By considering the equation in Proposition 3.2, one can check that the variety \overline{V}_G/S_{3+3} intersects all other boundary divisors, that is $\Delta_{0,2}$, $\Delta_{0,3}$, $\Delta_{1,1}$, $\Delta_{1,3}$, $\Delta_{2,0}$, $\Delta_{2,2}$, $\Delta_{2,3}$, $\Delta_{3,0}$, $\Delta_{3,1}$ and $\Delta_{3,2}$.*

We begin to prepare for the proof of Proposition 3.2 by recalling some facts and notation from [EMMW20, §4].

Let $\mathbb{C}[x]_k$ denote the space of polynomials of degree k or less. For $(p, q) \in \mathbb{C}[x]_3 \times \mathbb{C}[x]_1$ we have a factorization

$$(83) \quad p^2 - 4q^6 = (p - 2q^3)(p + 2q^3) = (d_1 s_1^2)(d_2 s_2^2),$$

where the $d_i \in \mathbb{C}[x]$ are square free polynomials. We define the subvariety

$$(84) \quad P_G := \{(p, q) \mid d = d_1 d_2 \text{ is separable of degree 6 and } \deg(d_1) = 3\} \subseteq \mathbb{C}[x]_3 \times \mathbb{C}[x]_1.$$

Let T_t and U_t denote the Chebyshev polynomials of the first and second kind, respectively. The space of pairs P_G can be used to parametrize the dihedral locus:

THEOREM 3.5 ([EMMW20, Theorem 4.3]). *There is a surjective algebraic map*

$$P_G \rightarrow \Omega \tilde{D}, \quad (p, q) \mapsto (X, \omega, \pi).$$

The D_{12} -cover $Y \rightarrow \mathbb{P}^1$ is defined in $\mathbb{C}[x, y]$ by the equation

$$(85) \quad y^{12} - p(x)y^6 + q(x)^6 = 0,$$

the covering is given by $(x, y) \mapsto y$ and the action of D_{12} is given by

$$r \cdot (x, y) = (x, \zeta_6 y) \quad \text{and} \quad f \cdot (x, y) = \left(x, \frac{q(x)}{y} \right).$$

Let $u := (y + q/y)/2$. The curve $X = Y/f$ is defined in $\mathbb{C}[x, u]$ by the equation

$$2q(x)^3 T_6(q(x)^{-1/2} u) = p(x)$$

and the map $\pi: X \rightarrow \mathbb{P}^1$ is given by $\pi(x, u) = x$. The one-form ω is given by

$$\omega = \frac{dx}{q(x)^{5/2} U_5(q(x)^{-1/2} u)} \in \Omega(X).$$

Moreover, ω is the pushforward of the one-form

$$(86) \quad \nu := y \cdot \frac{dx}{2y^6 - p(x)} \in \Omega(Y)$$

by the quotient map $Y \rightarrow X$.

The special fiber of the dihedral map $\pi: X \rightarrow \mathbb{P}^1$ is the fiber above ∞ . The other branch points of the dihedral map π are precisely the 6 zeros of the separable polynomial $d_1 d_2$ in (84). More precisely, the 3 zeros of d_1 are the ramification points with monodromy associated to $[f]$, while the 3 zeros of d_2 are the ramification points with monodromy associated to $[rf]$. In order to understand the variety V_G we need to understand which polynomials d_1, d_2 can appear in (84).

We first observe that the set P_G is in some sense “to large”. The group of units \mathbb{C}^\times acts on P_G by

$$z \cdot (p, q) = (z^3 p, zq).$$

Note that this in particular implies a group action of the subgroup of the 3th roots of unity $\mu_3 \subseteq \mathbb{C}^\times$ by

$$\zeta \cdot (p, q) = (p, \zeta q).$$

COROLLARY 3.6. *There is a commutative diagram*

$$\begin{array}{ccc} P_G & \longrightarrow & \mathbb{P}\Omega\tilde{D} \\ & \searrow & \nearrow \\ & P_G/\mathbb{C}^\times & \end{array}$$

where all maps are surjective.

PROOF. Let $(p, q) \in P_G$ and let

$$\tau: P_G \rightarrow \mathbb{P}\Omega\tilde{D}$$

denote the surjective map from the previous Theorem. We need to show that this map factors through P_G/\mathbb{C}^\times , i.e. that for all $z \in \mathbb{C}^\times$ it is $\tau((p, q)) = \tau((z^3 p, zq))$. It suffices to prove this for Y and ν , as X , ω and π are determined by those two. The equation (85) for the curve Y_1 corresponding to $\tau((p, q))$ is

$$y^{12} - p(x)y^6 + q(x)^6 = 0,$$

the equation for the curve Y_z corresponding to $\tau((z^3 p, zq))$ is

$$y^{12} - z^3 p(x)y^6 + z^6 q(x)^6 = 0$$

and an isomorphism between both curves is given by

$$\begin{aligned} \sigma: Y_1 &\rightarrow Y_2 \\ (x, y) &\mapsto (x, z^{1/2}y). \end{aligned}$$

The equations (86) for the corresponding one-forms are

$$\nu_1 := y \cdot \frac{dx}{2y^6 - p(x)} \quad \text{and} \quad \nu_z := y \cdot \frac{dx}{2y^6 - z^3 p(x)}$$

and the pull-back is

$$\sigma^* \nu_z = z^{1/2} y \cdot \frac{dx}{2z^3 y^6 - z^3 p(x)} = z^{1/6} \nu_1.$$

Thus $\tau((p, q)) = \tau((z^3 p, zq))$ agree in the projectivization $\mathbb{P}\Omega\tilde{D}$. \square

We gather some more or less obvious facts about all the polynomials floating around.

LEMMA 3.7. *For $(p, q) \in P_G$ and d_1, d_2, s_1, s_2 as in (83), we have*

- (i) $p = \frac{1}{2}(d_2 s_2^2 + d_1 s_1)$, $q^3 = \frac{1}{4}(d_2 s_2^2 - d_1 s_1)$,
- (ii) $\deg(q) \leq 1$,
- (iii) $\deg(d_2) = 3$,
- (iv) $\deg(s_1) = 0$ and
- (v) $\deg(s_2) = 0$.

In particular, we can choose $s_1 = s_2 = 1$ and hence

$$(87) \quad p = \frac{1}{2}(d_2 + d_1), \quad q^3 = \frac{1}{4}(d_2 - d_1).$$

PROOF. Recall that $(p, q) \in \mathbb{C}[x]_3 \times \mathbb{C}[x]_1$; in particular $\deg(q^3) \leq 3$ (implying (ii)). Thus all the inequalities in

$$3 = \deg(d_1) \leq \deg(d_1 s_1^2) = \deg(p - 2q^3) \leq 3$$

need to be in fact equalities, implying (iv). Claim (i) is seen by solving (83) for p and q . Claim (iii) follows from $\deg(d_1 d_2) = 6$ and $\deg(d_1) = 3$, and we finally see claim (v) by

$$3 = \deg(d_2) \leq \deg(d_2 s_2^2) = \deg(p + 2q^3) \leq 3.$$

\square

The relations obtained in the previous lemma motivate the definition of the subvariety

$$(88) \quad B_G := \left\{ (d_1, d_2) \in (\mathbb{C}[x])^2 \left| \begin{array}{l} \deg(d_1) = \deg(d_2) = 3 \\ d_1 d_2 \text{ separable} \\ d_2 - d_1 \in (\mathbb{C}[x]_1)^3 \end{array} \right. \right\} / \mathbb{C}^\times,$$

where \mathbb{C}^\times acts on B_G by rescaling the polynomials. The locus B_G provides an alternative parametrization of the dihedral triples:

PROPOSITION 3.8. *There is a surjective map $B_G \rightarrow V_G$. If b_1, \dots, b_3 are the roots of d_1 and b_4, \dots, b_6 are the roots of d_2 the map is given by*

$$(d_1, d_2) \mapsto (b_1, \dots, b_6, \infty).$$

PROOF. By Corollary 3.6 we have a surjective map $P_G/\mathbb{C}^\times \rightarrow \mathbb{P}\Omega\tilde{D} \rightarrow V_G$. Thus it suffices to give a surjective map $B_G \rightarrow \mathbb{P}\Omega\tilde{D}$. Consider the map

$$B_G \rightarrow P_G/\mathbb{C}^\times \\ (d_1, d_2) \mapsto \left(\frac{1}{2}(d_2 + d_1), \left(\frac{1}{4}(d_2 - d_1) \right)^{1/3} \right).$$

This map is well-defined and surjective by Lemma 3.7. \square

PROOF OF PROPOSITION 3.2. This is a direct consequence of Proposition 3.8. \square

We let $\mathrm{PGL}_2(\mathbb{C})$ act on the polynomials $\mathbb{C}[x]$ by Möbius transformation of the roots: for a polynomial $f := c(x - a_1) \cdots (x - a_n) \in \mathbb{C}[x]$ and a matrix $A \in \mathrm{PGL}_2(\mathbb{C})$ we define

$$A \cdot f := c(x - A \cdot a_1) \cdots (x - A \cdot a_n).$$

The group action of the stabilizer subgroup $\mathrm{PGL}_2(\mathbb{C})_\infty$ is partially compatible with the ring structure on $\mathbb{C}[x]$.

LEMMA 3.9. *For all $A \in \mathrm{PGL}_2(\mathbb{C})$ and all $f, g \in \mathbb{C}[x]$ the formal distributivity*

$$(A \cdot f) \cdot (A \cdot g) = A \cdot (fg).$$

holds. If moreover $\deg(f) = \deg(g)$ and $A \in \mathrm{PGL}_2(\mathbb{C})_\infty$, then there exists a $\lambda \in \mathbb{C}^\times$ such that

$$(A \cdot f) + (A \cdot g) = A \cdot (\lambda(f + g)).$$

PROOF. The distributivity with multiplication is immediate and the second equation can be checked on the generators

$$T_z := \begin{pmatrix} 1 & z \\ 0 & 1 \end{pmatrix} \text{ for } z \in \mathbb{C} \quad \text{and} \quad D_z := \begin{pmatrix} z & 0 \\ 0 & 1 \end{pmatrix} \text{ for } z \in \mathbb{C}^\times,$$

by checking that $(A \cdot f) + (A \cdot g)$ vanishes on all the zeros of $A \cdot (f + g)$. \square

PROPOSITION 3.10. *The group action of the stabilizer subgroup $\mathrm{PGL}_2(\mathbb{C})_\infty$ on B_G defined by*

$$A \cdot (d_1, d_2) = (A \cdot d_1, A \cdot d_2)$$

is well-defined and there is a commutative diagram

$$\begin{array}{ccc} B_G & \xrightarrow{\quad} & \mathbb{P}\Omega\tilde{D} \\ & \searrow & \nearrow \\ & B_G/\mathrm{PGL}_2(\mathbb{C})_\infty & \end{array}$$

where all maps are surjective.

PROOF. If the group action is well-defined, the existence of the diagram is immediate, as the action of $\mathrm{PGL}_2(\mathbb{C})_\infty$ on B_G is obviously compatible with the group action on \mathbb{P}^1 . We need to prove that the action is indeed well-defined. Let $(d_1, d_2) \in B_G$. The degree and separability conditions in the definition of B_G (see (88)) are invariant under the group action. Assume that $d_2 - d_1 = f^3$ for some $f \in \mathbb{C}[x]_1$. By the previous lemma there exists an $\lambda \in \mathbb{C}^\times$ such that

$$(A \cdot d_2) - (A \cdot d_1) = A \cdot (\lambda(d_2 - d_1)) = \lambda A \cdot f^3 = (\lambda^{1/3} A \cdot f)^3.$$

Thus $A \cdot (d_1, d_2) \in B_G$ as claimed. \square

This completes the proof of Corollary 3.3.

4. The boundary superset

We denote by $\mathbb{P}\Xi^2\overline{G} := \overline{\mathbb{P}QG} \subset \mathbb{P}\Xi^2\overline{\mathcal{M}}_{1,6}(-1^3, 1^3)$ the closure of the quadratic gothic locus inside the moduli space of multi-scale differentials. The goal of this section is to determine a superset of the boundary divisors of $\mathbb{P}\Xi^2\overline{\mathcal{M}}_{1,6}(-1^3, 1^3)$ which are intersected by the quadratic gothic locus $\mathbb{P}\Xi^2\overline{G}$. For this, it is more convenient to work with the unordered quadratic gothic locus

$$\mathbb{P}\Xi^2\overline{G}/S_{3+3} \subseteq \mathbb{P}\Xi^2\overline{\mathcal{M}}_{1,6}(-1^3, 1^3)/S_{3+3},$$

as this allows for a more concise listing of the relevant enhanced level graphs. For an enhanced quadratic level graph Γ^Q we denote by $D_{\Gamma^Q}^Q \subseteq \mathbb{P}\Xi^2\overline{\mathcal{M}}_{1,6}(-1^3, 1^3)$ the boundary stratum parametrized by Γ^1 . Our goal in this section is to obtain a list of all the boundary strata $D_{\Gamma^Q}^Q$ whose intersection with the gothic locus might be boundary divisor of the gothic locus. By work of Benirschke-Dozier-Grushevsky (see Theorem 4.4 for the precise statement) such strata might either be vertical divisors, or horizontal strata with an arbitrary number of horizontal edges. Relying on the results of a computer program, in this section we will prove:

PROPOSITION 4.1. *The boundary of the quadratic gothic locus $\partial\mathbb{P}\Xi^2\overline{G}/S_{3+3}$ is contained in the subspace*

$$\partial\mathbb{P}\Xi^2\overline{G}/S_{3+3} \subseteq \bigcup_i D_{\Gamma_i^Q}^Q$$

where the union is taken over all graphs in Figure 5. More precisely, each generic point of the boundary $\partial\mathbb{P}\Xi^2\overline{G}$ is contained in the interior of one of the listed boundary strata $D_{\Gamma_i^Q}^Q$.

For each covering $(X \rightarrow \mathbb{P}^1) \in \mathcal{H}$ we may consider the intermediate covering $X/J = A \rightarrow \mathbb{P}^1$. Let \mathcal{H}_Q denote the Hurwitz space parametrizing those coverings and let $\overline{\mathcal{H}}_Q$ denote the compactification with admissible covers. As for \mathcal{H} (see (81)) there is a forgetful map $\phi : \overline{\mathcal{H}}_Q \rightarrow \overline{\mathcal{M}}_{1,6}/S_{3+3}$, where the marked points are the three points in the special fiber and the images of the three fixed points of J .

Recall that there is the dihedral locus $\mathcal{D} \subset \mathcal{H}$. Denote by $\mathcal{D}_Q \subset \mathcal{H}_Q$ its image and by $\overline{\mathcal{D}}_Q \subset \overline{\mathcal{H}}_Q$ the closure. It follows from the discussion in Section 2 that

$$\phi(\overline{\mathcal{D}}_Q) = \pi(\mathbb{P}\Xi^2\overline{G})/S_{3+3},$$

where $\pi : \mathbb{P}\Xi\overline{\mathcal{M}}_{1,6}(-1^3, 1^3) \rightarrow \overline{\mathcal{M}}_{1,6}$ is the natural forgetful map. We will make use of this equality to prove Proposition 4.1.

For a boundary stratum $D_{\Gamma^Q}^Q$ we denote by $D_{\Gamma^Q}^{QG} := D_{\Gamma^Q}^Q \cap \mathbb{P}\Xi^2\overline{G}$ its intersection with the quadratic gothic locus. Our intermediate goal is to prove:

PROPOSITION 4.2. *Assume that $D_{\Gamma^Q}^{QG} \subseteq \mathbb{P}\Xi^2\overline{G}$ is a divisor. Then there exists a boundary stratum $\Delta^{\overline{\mathcal{H}}_Q} \subseteq \overline{\mathcal{H}}_Q$ such that $\pi(D_{\Gamma^Q}^{QG})/S_{3+3} \subseteq \phi(\Delta^{\overline{\mathcal{H}}_Q})$ and $\dim \Delta^{\overline{\mathcal{H}}_Q} \geq 1$.*

The forgetful map $\pi : \mathbb{P}\Xi^2\overline{\mathcal{M}}_{1,6}(-1^3, 1^3) \rightarrow \overline{\mathcal{M}}_{1,6}$ may have positive fiber dimension when restricted to a boundary divisor. An upper bound for this fiber dimension is given as follows.

LEMMA 4.3. *If $D_{\Gamma^Q}^Q \subseteq \mathbb{P}\Xi^2\overline{\mathcal{M}}_{1,6}(-1^3, 1^3)$ is a divisor, then the fiber dimension of $\pi|_{D_{\Gamma^Q}^Q}$ above a generic point of $\pi(D_{\Gamma^Q}^Q)$ is at most 1.*

PROOF. The fiber dimension can only be greater than one if there exists a boundary divisor with either at least three vertices on one level or two vertices on top- and bottom level. By listing all the boundary strata (for example with `diffstrata`) one checks that there is only one two-level graph with more than three vertices on one level, the one

¹In Chapter I and II we indexed the boundary strata of quadratic strata $\mathbb{P}\Xi^2\overline{\mathcal{M}}_{g,n}(\mu)$ by coverings of enhanced level graphs. Here we only index them by “the lower half” of the covering, i.e. by the quadratic level graph. A priori this gives a coarser indexing, but as we will see it doesn’t matter here as all the coverings are uniquely determined by the quadratic level graphs.

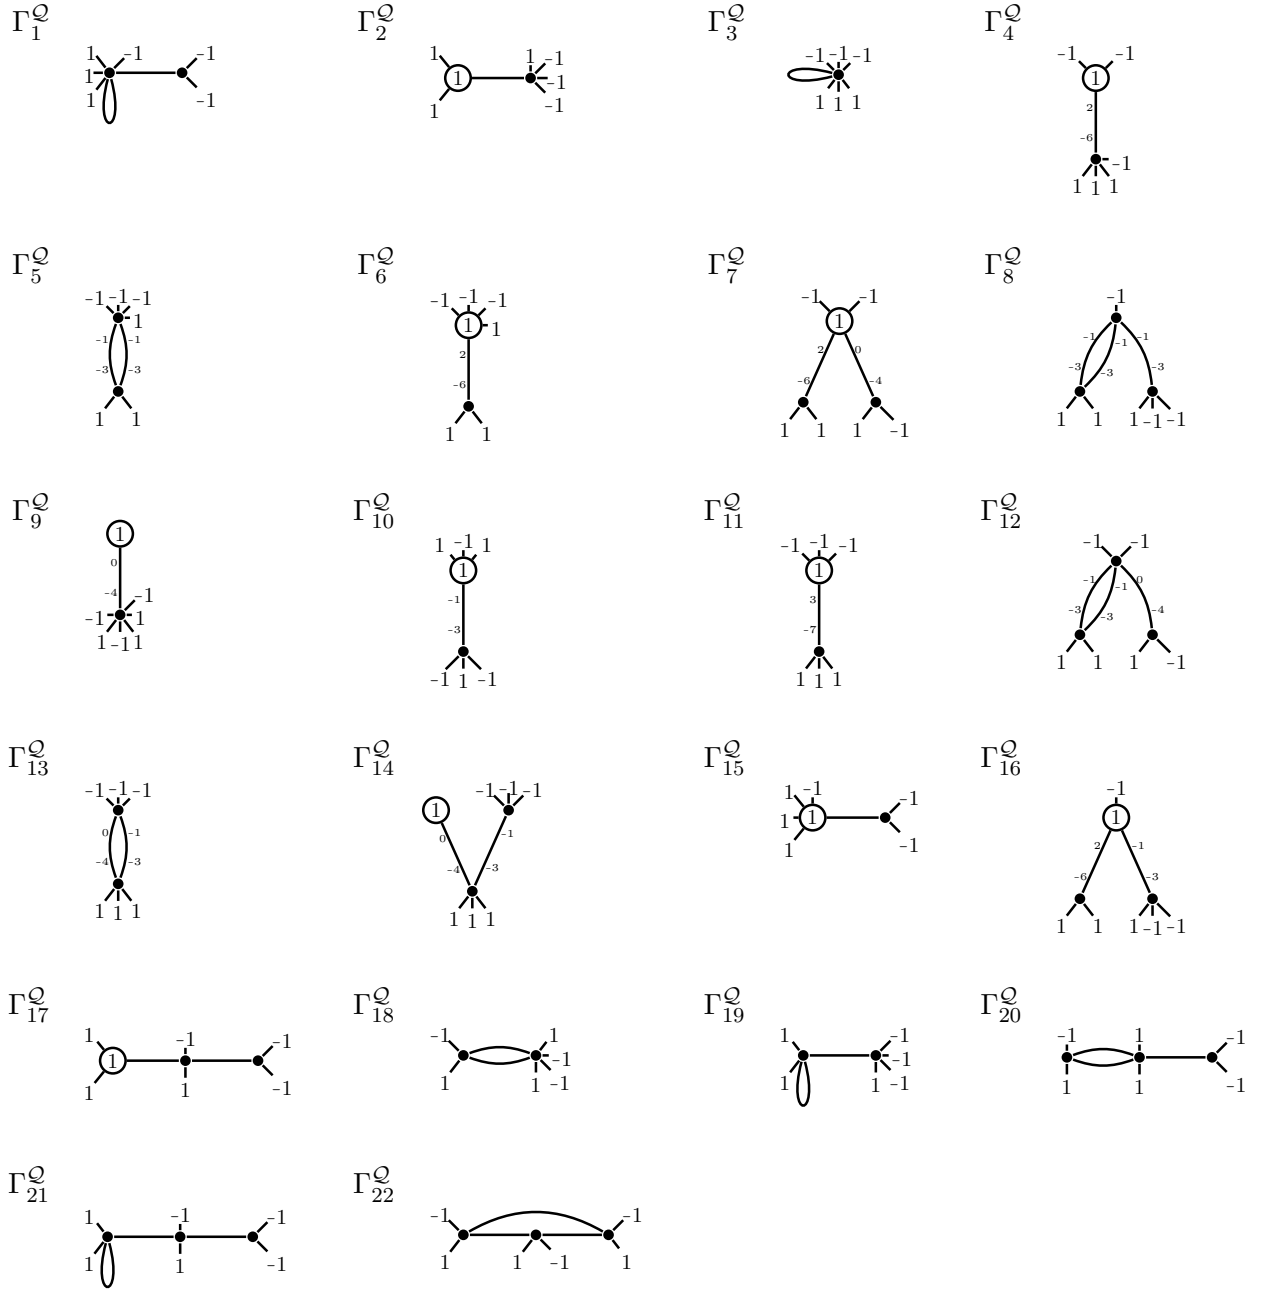


FIGURE 5. The boundary strata of $\mathbb{P}\Xi^2\overline{\mathcal{M}}_{1,6}(-1^3, 1^3)/S_{3+3}$ that might be intersected by the quadratic gothic locus $\mathbb{P}\Xi^2\overline{\mathcal{G}}/S_{3+3}$

depicted in Figure 6. Because of the GRC, the vertices on bottom level can not be scaled independently, hence the fiber dimension of π restricted to this boundary divisor is 1. \square




FIGURE 6. A graph with three vertices on bottom level

PROOF OF PROPOSITION 4.2. Let $D_{\Gamma^Q}^{QG} \subseteq \mathbb{P}\Xi^2\overline{\mathcal{G}}$ be a divisor. By Lemma 4.3 the dimension of its image is $1 \leq \dim \pi(D_{\Gamma^Q}^{QG}) \leq 2$. Hence there must be a boundary stratum $\Delta^{\overline{\mathcal{H}}_Q} \subseteq \overline{\mathcal{H}}_Q$ such that $\phi(\Delta^{\overline{\mathcal{H}}_Q} \cap \overline{\mathcal{D}}_Q)$ contains $\pi(D_{\Gamma^Q}^{QG})/S_{3+3}$ and is of dimension at least one. \square

We want to proceed by listing all the boundary divisors of $\overline{\mathcal{H}}_{\mathcal{Q}}$ and check whether or not their image under ϕ admits a level graph structure. As we are, at least for now, only interested in divisors of the gothic locus, we do not need to consider all level graph structures by the following theorem.

THEOREM 4.4 ([BDG22, Theorem 1.5]). *Let $D_{\Gamma_{\mathcal{Q}}}^{\mathcal{Q}G} \subseteq \mathbb{P}\Xi^2\overline{\mathcal{G}}$ be a divisor. Then $\Gamma^{\mathcal{Q}}$ has either*

- *two level and no horizontal edge or*
- *one level and at least one horizontal edge.*

We prove Proposition 4.1 by relying on the output of a computer program written in Sage that does the following:

- List all boundary strata $\Delta^{\overline{\mathcal{H}}_{\mathcal{Q}}} \subseteq \overline{\mathcal{H}}_{\mathcal{Q}}/S_{3+3}$ of dimension at least 1.
- Throw away all strata that intersect the divisors $\Delta_{1,2}$ and $\Delta_{2,1}$.
- For each remaining stratum, consider its image under the forgetful map to $\overline{\mathcal{M}}_{1,6}/S_{3+3}$.
- For each stratum in the image check if it admits the structure of an enhanced level graph as in Theorem 4.4, otherwise discard it.

The enhanced level graphs obtained in this way are exactly those listed in Figure 5.

REMARK 4.5. *Note that Figure 5 includes all horizontal level graphs that appear as images of boundary strata of $\overline{\mathcal{H}}_{\mathcal{Q}}$ up to codimension 3. By also listing the boundary strata of codimension 4 of the Hurwitz space $\overline{\mathcal{H}}_{\mathcal{Q}}$ one can check that there are no additional horizontal level graphs.*

5. Non-horizontal divisors

In Proposition 4.1 we have determined a superset for the boundary divisors of the quadratic gothic locus $\mathbb{P}\Xi^2\overline{\mathcal{G}}$. Note that for all the quadratic enhanced level graphs $\Gamma_i^{\mathcal{Q}}$ listed there the covering abelian enhanced level graph, which we denote by Γ_i , is uniquely determined, see Section I.3 for details on the construction of the coverings. Hence this also determines a superset for the boundary divisors of the gothic locus $\mathbb{P}\Xi\overline{\mathcal{G}}$. In this section, we will provide explicit families of flat surfaces converging to some of the non-horizontal divisors, and thus prove that those divisors are in fact intersected by the gothic locus. For those divisors we list the dual graphs of the double covers in Figure 7. The intersection $D_{\Gamma_i}^G := D_{\Gamma_i} \cap \mathbb{P}\Xi G$ on each level is a \mathbb{R} -linear submanifold. On the left of each level we indicate either the dimension of the respective linear submanifold, or the linear submanifold itself if it is a locus of double covers: For a given stratum of quadratic differentials $\mathbb{P}\Xi^2\overline{\mathcal{M}}_{g,n}(\mu)$ one can define a linear submanifold in a stratum of abelian differentials via the double cover construction. We will denote this linear submanifold by $\mathbb{P}\Xi^2\overline{\mathcal{M}}_{g,n}(\mu)$. For details of the definition see Section II.7.²

In this section we will mostly be concerned with the top-level of each stratum. We remark that the dimension of the top-level determines the dimension of the bottom-level. In the next Section 6 we will have a closer look at some of the bottom-levels. We will treat the horizontal divisors in Section 7.

PROPOSITION 5.1. *The gothic locus $\mathbb{P}\Xi\overline{\mathcal{G}}$ intersects the boundary divisors $D_{\Gamma_4}, \dots, D_{\Gamma_9}$.*

REMARK 5.2. *We are currently unable to decide whether or not the gothic locus intersects the vertical divisors that are listed in Proposition 4.1 but not in Proposition 5.1.*

To prove Proposition 5.1 we will handle each of the boundary strata listed there separately. For each stratum, we will show that it is intersected by the gothic locus $\mathbb{P}\Xi\overline{\mathcal{G}}$ by giving an explicit degeneration of the surface in Figure 4 along the Equations (78).

REMARK 5.3. *To see that the limiting surface is actually contained in the claimed boundary stratum, one should work with the conformal topology on the moduli space: A sequence (X_t, ω_t) converges to a twisted differential (X, ω) if there exists an exhaustion K_t of $X \setminus \{\text{nodes of } X\}$ and a sequence of conformal maps $g_t : K_t \rightarrow X_t$ such that $g_t^*\omega_t$ converges to ω uniformly on compact sets, see [BCGGM19b, Section 3.3] for details. This*

²The linear submanifolds are denote by $\overline{\mathcal{H}}_k$ in loc. cit.




FIGURE 7. Divisors intersected by the gothic locus

convergence can be verified in the flat pictures by choosing an appropriate exhaustion. In the following we will suppress this technicality.

5.1. The stratum D_{Γ_4} . The stratum $D_{\Gamma_4}^G$ contains a point of the closure of the cyclic forms $\mathbb{P}\Xi\bar{Z}_G \subseteq \mathbb{P}\Xi\bar{G}$ (and is in fact the only stratum that does so). By letting $w_i \rightarrow 0$ with equal speed, one gets to the polygons depicted in Figure 8. The boundary of the gothic locus is cut out by the equations

$$\sum_i v_i = 0, \quad v_i = -v_{i+3}, \quad v_1 + v_3 + v_5 = 0, \quad \beta + v_2 = 0 \quad \text{and} \quad \alpha - \beta = 0,$$

thus $\dim D_{\Gamma_4}^{G,\top} = 1$. This implies $\dim D_{\Gamma_4}^{G,\perp} = 1$. We remark that the torsion equations $\beta' = \beta = u_5$ still exists, but no longer relates the two zeros of order 3. Instead it relates the two nodes at top-level.

The two hexagons in Figure 8 can be identified by translation. (We emphasize that this is not the action of the involution J , which rotates the polygons). This identification exhibits the top-level as a two-fold covering of the Teichmüller curve in $\Omega\mathcal{M}_1(0^2)$ with marked points P, Q , which is given by the condition that $[P - Q]$ is 3-torsion.




FIGURE 8. The top-level of the cyclic curve in $D_{\Gamma_4}^G$

5.2. The stratum D_{Γ_5} . The top-level component of the stratum $D_{\Gamma_5}^G$ is again a Teichmüller curve. The flat picture can be obtained by pinching the two small hexagons in Figure 4 in different directions. The result of pinching the left hexagon vertically and the right hexagon orthogonally to w_2 is depicted on the left of Figure 9. This polygon can be re-glued to the L-shaped surface on the right of the same figure. The top-level $D_{\Gamma_5}^G$ is cut out by the equations

$$\alpha_1 = \alpha_2 \quad \text{and} \quad \alpha_3 = \alpha_4$$

and is thus a threefold cover of an elliptic curve where the three 2-torsion points are marked.




FIGURE 9. The top-level component of $D_{\Gamma_5}^G$ and a re-gluing

5.3. The stratum D_{Γ_6} . A flat picture of the top-level component of the stratum $D_{\Gamma_6}^G$ can be obtained by pinching only one of the small hexagons in Figure 4. The resulting flat picture is given in Figure 10. In the figure, we highlight in gray the part of the surface that is coming from the large hexagon. This component has the property that for $v_2 = v_4 = w_1/2 = w_2 = w_5 = \alpha' = \alpha$ (note that this is only one additional condition) the depicted polygon is an origami and thus gives rise to a Teichmüller curve.




FIGURE 10. The top-level component of $D_{\Gamma_6}^G$

We want to give a more precise description of the top-level $D_{\Gamma_6}^{G,\top}$. Let $\mathcal{T} \subseteq \Omega\mathcal{M}_1(0^3)$ be the two-dimensional subspace, where the marked points P, Q, Z are subject to the relation $P + Q = Z$. If we choose Z as zero, this is equivalent to $P + Q$ being 1-torsion. We claim that the top-level of the stratum $D_{\Gamma_6}^G$ is a sixfold covering of \mathcal{T} as follows. Consider the covering that is ramified exactly above P, Q, Z . More precisely, let the ramification profile above Z be $(3, 1^3)$ and the ramification profiles above P and Q by $(2, 1^5)$. The double zero (marked with a black square) is the ramification point of order three above Z , the two simple zeros (marked with black and white circles) are the two ramification points of order two above P and Q , respectively. The three unramified points in the fiber above Z are the three marked regular points (marked with a cross). In Figure 11 we depict the tessellation of a generic curve in the stratum with the preimages of the torus.




FIGURE 11. A generic curve in the top-level component of $D_{\Gamma_6}^G$ and the covered torus

5.4. The stratum D_{Γ_7} . The flat picture of the stratum $D_{\Gamma_7}^G$ can be obtained by letting $w_1, w_4 \rightarrow 0$ in Figure 4, the resulting polygon is depicted in Figure 12.




FIGURE 12. The top-level component of $D_{\Gamma_7}^G$

We claim that the top-level component $D_{\Gamma_7}^{G,\top}$ is a covering of a Hilbert modular surface for a square discriminant $\mathcal{Q}(\sqrt{d^2})$ in \mathcal{M}_2 . We can cut and re-glue the polygons to the polygon in Figure 13, and by forgetting the marked regular points we obtain the image under the forgetful map $\pi: \Omega\mathcal{M}_2(0^4, 1^2) \rightarrow \Omega\mathcal{M}_2(1^2)$. The image of $D_{\Gamma_7}^{G,\top}$ is locally given by the equations

$$\alpha_1 = \alpha_5 \quad \text{and} \quad \alpha_2 = \alpha_4.$$

The forgetful map π is locally injective on $D_{\Gamma_7}^{G,\top}$: The periods v_i, w_2 and w_3 in Figure 12 are completely determined by the α_i . For example it is $\alpha_1 = v_1 + v_2, \alpha_3 = \beta' + \beta = -2v_1$, and similar equations hold for the other periods. By construction, the component $D_{\Gamma_7}^{G,\top}$ is $\text{GL}_2(\mathbb{R})^+$ -invariant and so is its image under the forgetful map π . The $\text{GL}_2(\mathbb{R})^+$ -orbit closures in genus two have been classified by McMullen [McM07b]. As the dimension of the image is two, its projection to \mathcal{M}_2 needs to be a Hilbert modular surface by [McM07b, Theorem 1.2].




FIGURE 13. Image of the top-level component of $D_{\Gamma_7}^G$ in $\Omega\mathcal{M}_2(1^2)$

5.5. The stratum D_{Γ_8} . A surface in this stratum can be obtained by letting not only w_1 and w_4 converge to zero as for $D_{\Gamma_7}^G$, but letting also v_3 and v_6 converge to zero, all with equal speed. The resulting locus is 1-dimensional and coincides with the double cover of the stratum $\mathbb{P}\Xi^2\mathcal{M}_{0,4}(-1^4)$ for dimensional reasons.

5.6. The stratum D_{Γ_9} . After rotating the right hexagon in Figure 4 by 30 degrees we can find two pairs of simply crossing geodesics, see Figure 14. After cutting along these geodesics and re-gluing the surface as depicted in Figure 15 the surface converges to the top-level of $D_{\Gamma_9}^G$ by letting all the labeled periods converge to zero.

6. Non-obvious Teichmüller curves

In this section we will prove Theorem 1.3 by exhibiting $D_{\Gamma_4}^{G,\perp}$ as a non-obvious Teichmüller curve. As we have seen in Section 5 the top-level of $D_{\Gamma_4}^G$ has dimension 1. Hence the bottom level has dimension 1, too. In particular is $D_{\Gamma_4}^{G,\perp}$ a Teichmüller curve. We can regard this either as a curve in $\Omega\mathcal{M}_{1,6}(-3^2, 0, 2^3)$, or, after forgetting the marked regular point, in the stratum $\Omega\mathcal{M}_{1,5}(-3^2, 2^3)$. Before we prove the theorem, we recall some general fact about coverings of flat surfaces.




FIGURE 14. Two pairs of simply crossing geodesics




FIGURE 15. The same curve as in Figure 14 after re-gluing

LEMMA 6.1. *Let $f : (X, \omega) \rightarrow (Y, \eta)$ be a covering of flat surfaces and $q \in Y$. Then η has a poles at q if and only if ω has a pole at every point in the fiber above q .*

PROOF. Let $p \in X$ be a preimage of q . Assume that $\text{ord}_q(\eta) = a$ for some $p \in X$ and that f is ramified to order k at p . Then $\text{ord}_p(\omega) = (a + 1)k - 1$, and thus

$$\begin{cases} \text{ord}_p(\omega) \geq a & \text{if } a \geq -1 \\ \text{ord}_p(\omega) \leq a & \text{if } a \leq -1. \end{cases}$$

In particular we see that η has a pole at q if and only if ω has a pole at every point in the fiber above q . \square

PROOF OF THEOREM 1.3. We begin by showing that $D_{\Gamma_4}^{G, \perp}$ is not obvious. First note that the condition on the residues is trivial, i.e. the only condition is the one imposed by the residue theorem. Hence we need to check that $D_{\Gamma_4}^{G, \perp}$ does not coincide with a Hurwitz space above a 1-dimensional stratum of abelian differentials. By the Riemann-Hurwitz formula, any such stratum would necessarily parametrize curves in genus 0. We want to rule out the existence of such a map $D_{\Gamma_4}^{G, \perp} \rightarrow \Omega\mathcal{M}_{0,4}(m_1, \dots, m_4)$.

For a contradiction assume that a map $D_{\Gamma_4}^{G, \perp} \rightarrow \Omega\mathcal{M}_{0,4}(m_1, \dots, m_4)$ exists and let $f : (X, \omega) \rightarrow (Y, \eta)$ be a corresponding covering. Let $p_1, p_2 \in X$ be the two points where ω vanishes to order -3 . In the notation of Lemma 6.1 observe that

$$-3 = \text{ord}_{p_i}(\omega) = (a + 1)k - 1,$$

hence $k \in \{1, 2\}$. By Lemma 6.1 there are two possibilities for the points p_i :

- (1) They might lie in different fibers, and then necessarily are total ramification points of f . In particular is $\deg(f) = 2$ (because of $k \in \{1, 2\}$). In this case f must be totally ramified above all four marked points of Y . On the other hand, two of the points where ω vanishes to order 2 must lie in the same fiber (as all three of those points must lie in the fibers above the two remaining marked points), a contradiction.
- (2) They might lie in the same fiber. Then $\deg(f) \in \{2, 4\}$ (because of $k \in \{1, 2\}$). For $\deg(f) = 2$ we arrive at the same contradiction as in Case (1): in this case f would need to be totally ramified above all four marked points again. For $\deg(f) = 4$ one checks with the Riemann-Hurwitz formula that it is not possible for two of the marked double zeros to lie in one fiber. Hence the ramification profiles in the other three fibers must be $(3, 1)$, and η of type $(-2, 0^3)$. But then $D_{\Gamma_4}^{G, \perp}$ would have REL while the gothic locus ΩG has REL zero, a contradiction.

The cyclic locus $\Omega Z_G \subseteq \Omega G$ agrees with the locus of unfoldings of quadrilaterals of type $(1, 1, 1, 9)$, or equivalently with the locus of canonical covers of the 6-differentials of type $(-5^3, 3)$. Hence the canonical cover of the (up to permutation of the marked points) unique boundary point of the stratum $\mathbb{P}\Xi^6\mathcal{M}_{0,4}(-5^3, 3)$ must be contained in $\mathbb{P}\Xi G$. One easily checks that this canonical cover is actually contained in D_{Γ_4} . The bottom level of the canonical cover is precisely the 6-differential of type $(-10, -5, 3)$, hence the $\mathrm{GL}_2(\mathbb{R})^+$ -orbit closure of this differential is contained in $D_{\Gamma_4}^{G,\perp}$ and generates a non-obvious Teichmüller curve.³

To obtain the flat picture in Figure 2 and the equations one traces the degeneration discussed in Section 5.1 on the bottom level. \square

We end this section with some remarks on the definition of *obvious*. The definition we recalled in the introduction is the one introduced in [MM23]. Recall that strata of meromorphic abelian differentials are in general not connected. The connected components have been classified in [Boi15]. One of the components is the hyperelliptic component. With their definition of obvious Möller-Mullane proved:

THEOREM 6.2 ([MM23, Theorem 1.1]). *The only Teichmüller curves in the hyperelliptic connected component of a stratum of meromorphic differentials are obvious Teichmüller curves.*

As we have seen in Theorem 1.3, this theorem does not hold in the other components of strata. The hope is that there aren't "too many" non-obvious Teichmüller curves to allow for a possible classification. However there is a large class of easily constructed non-obvious Teichmüller curves: given a (projectively) 1-dimensional stratum of quadratic differentials, for example $\Omega^2\mathcal{M}_{0,4}(m_1, m_2, m_3, m_4)$, the canonical cover $\widehat{\Omega}^2\mathcal{M}_{0,4}(m_1, m_2, m_3, m_4)$ will be a Teichmüller curve in a stratum of meromorphic differentials, and in general this Teichmüller curve will not be obvious. For example it follows from our discussion in Section 5 that

$$D_{\Gamma_5}^{QG,\perp} \cong \mathbb{P}\Xi^2\overline{\mathcal{M}}_{0,4}(-3^2, 1^2),$$

and, using similar arguments to the ones in the proof of Theorem 1.3, one can check that $D_{\Gamma_5}^{G,\perp}$ is in fact non-obvious.

7. Horizontal strata, cylinders, and cusps of Teichmüller curves

We are now turning our attention towards the horizontal boundary strata. Our first observation is that Proposition 4.1 and Remark 4.5 imply that we know all horizontal strata that might possibly be intersected by the gothic locus.

PROPOSITION 7.1. *The gothic locus $\mathbb{P}\Xi\overline{G}$ only intersects the horizontal strata listed in Figure 3.*

In contrast to the non-horizontal strata, it is not obvious what the dimension of the intersection of the gothic locus with a horizontal stratum is (even if we know that the intersection is non-empty) as a consequence of Theorem 4.4. We arranged the graphs in Figure 3 to match our expectation about the horizontal boundary.

EXPECTATION 7.2. *We expect the following:*

- *The gothic locus intersects the three strata D_{Γ_2} , D_{Γ_3} and $D_{\Gamma_{15}}$ (depicted in the top row of Figure 3) in a divisor.*
- *The gothic locus intersects the three strata D_{Γ_1} , $D_{\Gamma_{19}}$ and $D_{\Gamma_{18}}$ (depicted in the middle row of Figure 3) in codimension 2.*
- *The gothic locus intersects the two strata $D_{\Gamma_{20}}$ and $D_{\Gamma_{22}}$ (depicted in the bottom row of Figure 3) in codimension 3.*
- *The gothic locus does not intersect the two strata $D_{\Gamma_{17}}$ and $D_{\Gamma_{21}}$.*

We are currently unable to prove all of the above expectations, but we will provide partial results and evidence in the following.

³The Teichmüller curve generated by the canonical cover of the 6-differential might not agree with $D_{\Gamma_4}^{G,\perp}$, but might be an irreducible component.

PROPOSITION 7.3. *The intersection of the gothic locus $\mathbb{P}\Xi\overline{G}$ with the two horizontal strata D_{Γ_2} and D_{Γ_3} is a divisor.*

Before proving Proposition 7.3 we need to recall a technical tool we will use throughout this section. Given a flat surface (X, ω) , a cylinder decomposition \mathcal{C} is the collection of all cylinders in X in a given direction. We can go from (X, ω) to a nearby surface in the moduli space by stretching all the cylinders in the decomposition \mathcal{C} by the same speed. If (X, ω) was contained in a $\mathrm{GL}_2(\mathbb{R})^+$ -orbit closure $\Omega\mathcal{H}$, then all those nearby surfaces will also be contained in $\Omega\mathcal{H}$ as a consequence of the cylinder deformation theorem [Wri15, Theorem 1.1]. Hence if we start with a surface in the gothic locus, choose a cylinder deformation and stretch all the cylinders at the same speed in the direction orthogonal to their core curve to infinity, we will obtain a surface in the boundary of the gothic locus $\mathbb{P}\Xi\overline{G}$. As noted in Remark 5.3, this should be verified in the conformal topology. The necessary exhaustion can be obtained by cutting smaller and smaller pieces from the infinite cylinders of the limiting surface.

PROOF OF PROPOSITION 7.3. By stretching all cylinders in the direction orthogonal to their core curves, the horizontal cylinder decomposition in Figure 16 gives rise to a generic curve in the stratum D_{Γ_3} . Similarly, the vertical cylinder decomposition in Figure 17 gives rise to a generic curve in the stratum D_{Γ_2} .

For D_{Γ_2} and D_{Γ_3} the intersection with the gothic locus must be a divisor, as $D_{\Gamma_2}^{\mathcal{Q}}$ and $D_{\Gamma_3}^{\mathcal{Q}}$ are divisors. \square




FIGURE 16. Four vertical and two horizontal geodesics in a curve in the gothic locus




FIGURE 17. Two geodesics in a curve in the gothic locus

PROPOSITION 7.4. *The gothic locus $\mathbb{P}\Xi\overline{G}$ intersects the two horizontal strata D_{Γ_1} and $D_{\Gamma_{19}}$. Moreover, the intersection with each of those strata has an irreducible component which is of codimension 2 in the gothic locus.*

PROOF. To prove that the stratum $D_{\Gamma_{19}}$ is intersected consider the curve in Figure 18a, which is the same curve as in Figure 17 after some re-gluing. In the coordinates in Figure 18a the gothic locus is cut out by the equations

$$(89) \quad \begin{aligned} x_i &= x_{i+3} \quad \text{for } i = 1, 2, 3, \\ y_i &= \frac{1}{3}(x_1 + x_2 + x_3) \quad \text{for } i = 1, \dots, 6. \end{aligned}$$

Hence the curve in Figure 18b is contained in $D_{\Gamma_2}^G$. Stretching the cylinders of the cylinder decomposition given by the two finite geodesics depicted in the figure gives rise to a curve in $D_{\Gamma_{19}}^G$. Moreover, around this curve the intersection is of codimension 2.

To obtain a curve in D_{Γ_1} one can stretch the vertical cylinder decomposition in Figure 16. The claim about the codimension follows as for $D_{\Gamma_{19}}$. \square




FIGURE 18. Degenerations toward $D_{\Gamma_{19}}^G$

REMARK 7.5. *Our guess is that the strata $D_{\Gamma_{17}}$ and $D_{\Gamma_{21}}$ are not intersected by the gothic locus. To see why, consider the surface in Figure 18b. Note that there can not be a horizontal degeneration of the irreducible component that contains the three marked points \times : As all y_i agree by (89) there can not be a finite geodesic that is disjoint from the already infinite cylinder. Hence the connected component of $D_{\Gamma_2}^G$ to which this surface belongs can not intersect $D_{\Gamma_{17}}$ or $D_{\Gamma_{21}}$.*

Note that this argument alone is not sufficient to prove that the gothic locus does not intersect $D_{\Gamma_{17}}$ and $D_{\Gamma_{21}}$, as the intersection $D_{\Gamma_2}^G = D_{\Gamma_2} \cap \mathbb{P}\Xi\bar{G}$ might have multiple irreducible components.

Recall that the gothic locus ΩG contains a dense set of primitive Teichmüller curves, see [EMMW20, Theorem 1.4]. Let $\mathcal{T} \subseteq \Omega G$ be such a Teichmüller curve. Teichmüller curves are never compact, and hence we may consider the closure $\mathbb{P}\bar{\mathcal{T}} \subseteq \mathbb{P}\Xi\bar{G}$. The boundary $\partial\mathbb{P}\bar{\mathcal{T}}$ consists of a finite number of points, the *cusps* of the curve.

PROPOSITION 7.6. *The interior of each of the four horizontal strata $D_{\Gamma_1}^G$, $D_{\Gamma_2}^G$, $D_{\Gamma_3}^G$ and $D_{\Gamma_{20}}^G$ contains cusps of a primitive Teichmüller curve contained in the gothic locus ΩG . The interior of the stratum $D_{\Gamma_{19}}^G$ contains cusps of a non-primitive Teichmüller curve.*

PROOF. The surface depicted in Figure 16 generates, for the “correct“ choice of side length, a primitive Teichmüller curve, as it is the unfolding of the quadrilateral in [EMMW20,

Theorem 1.4]. As we have seen in the proof of Proposition 7.3, the horizontal and vertical cylinder decompositions give rise to curves in the strata $D_{\Gamma_1}^{G,\circ}$ and $D_{\Gamma_3}^{G,\circ}$.

We claim that the surface depicted in Figure 17 generates, again for the “correct” choice of side length, a primitive Teichmüller curve. To see this, note that the moduli of the vertical and horizontal cylinder decompositions agree with the moduli of the cylinders in Figure 16, and hence are commensurable. Moreover, the curve needs to be defined over some number field $\mathbb{Q}(\sqrt{d})$ and hence is primitive. As we have seen in the proof of Proposition 7.3, the vertical cylinder decomposition gives rise to curves in the stratum $D_{\Gamma_2}^{G,\circ}$.

The horizontal cylinder decomposition of the surface depicted in Figure 14 gives rise to a point in $D_{\Gamma_{20}}^G$. This surface is in fact the so-called duck-shaped surface depicted in [MT20, Figure 3]. By [MT20, Proposition 2.3] this surface generates, again for the “correct” choice of side length, a Teichmüller curve. The cylinder decomposition of the cathedral-shaped surface depicted in [MT20, Figure 2] gives rise to a cusp in $D_{\Gamma_{20}}^G$, too.⁴

The curve in Figure 18a can be chosen such that it is square tiled and has a vertical cylinder decomposition that gives rise to a point in the stratum $D_{\Gamma_{19}}^G$: In complex coordinates choose $z_1 = x_1 = -1$, $x_2 = i$ and $x_3 = 1$ (recall that those suffices to specify the surface because of the Equations (89)). \square

COROLLARY 7.7. *The gothic locus intersects the stratum $D_{\Gamma_{20}}$.*

8. Towards the fundamental class

In this section we give an outline of a possible approach to compute the fundamental class of the image of the quadratic gothic locus $\mathbb{P}\Xi^2\overline{G}$ in $\overline{\mathcal{M}}_{1,6}$. We denote this class by

$$[\mathbb{P}\Xi^2\overline{G}] \in H^6(\overline{\mathcal{M}}_{1,6}).$$

We work with $\mathbb{P}\Xi^2\overline{G}$ instead of $\mathbb{P}\Xi\overline{G}$ for two main reasons:

- By [Pet14] all cohomology $H^\bullet(\overline{\mathcal{M}}_{1,6})$ is tautological, in particular is $[\mathbb{P}\Xi^2\overline{G}] \in R^3(\overline{\mathcal{M}}_{1,6})$ a tautological class.
- We want to work with the Sage package `admcycles` [DSZ21]. `Admcycles` is barely able to compute a generating set for the tautological ring $R^3(\overline{\mathcal{M}}_{1,6})$ on a typical computer. The tautological ring $R_3(\overline{\mathcal{M}}_{4,6})$ is currently out of reach (at least without special hardware).

8.1. Why this class is of interest. Let $\mathcal{H} \subseteq \mathbb{P}\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ be an arbitrary linear submanifold. If one knows the fundamental class of \mathcal{H} in the Chow ring of $\mathbb{P}\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ one can compute the Euler characteristic of this submanifold by using Equation (55). Actually it suffices to know the fundamental class of the image of \mathcal{H} in $\overline{\mathcal{M}}_{g,n}$: instead of intersecting the ξ -classes in (55) with the fundamental class in $\mathbb{P}\Xi\overline{\mathcal{M}}_{g,n}(\mu)$, one can as well push those ξ -classes to $R^\bullet(\overline{\mathcal{M}}_{g,n})$ and intersect with the class of the image of \mathcal{H} there.

Hence knowing the fundamental class of the gothic locus $[\mathbb{P}\Xi\overline{G}]$ would allow to compute its Euler characteristic. However, if we can determine $[\mathbb{P}\Xi^2\overline{G}]$, the work done in Section II.7 allows to pull this class back from $\mathbb{P}\Xi^2\overline{\mathcal{M}}_{1,6}(-1^3, 1^3)$ to $\mathbb{P}\Xi\overline{\mathcal{M}}_{4,6}(0^3, 2^3)$ to obtain $[\mathbb{P}\Xi\overline{G}]$.

As we recalled in Section 2.4, the class of the flex locus $[F] \in R^1(\mathcal{M}_{1,3})$ has been computed in [Che22]. Even though the flex locus is the image of the quadratic gothic locus under the forgetful map $\pi : \mathcal{M}_{1,6} \rightarrow \mathcal{M}_{1,3}$ there is unfortunately no way to directly compare those classes. This is because the fiber dimension of π above a generic point is 3, while the fiber dimension of the restriction $\pi|_{\mathcal{Q}G}$ is 1.

8.2. Our setup. By applying the Faber-Zagier-Pixton relations [PP21], one checks (using `admcycles`) that the vector space $R^3(\overline{\mathcal{M}}_{1,6})$ is generated by 756 elements. Let $\alpha_1, \dots, \alpha_{756} \in R^3(\overline{\mathcal{M}}_{1,6})$ be a generating set. Then there exist rational numbers λ_i such that

$$(90) \quad \sum_{i=1}^{756} \lambda_i \alpha_i = [\mathbb{P}\Xi^2\overline{G}].$$

⁴The reader should be warned that in the picture in loc. cit. there are two cylinders with the same color (black).

In the rest of the section we will outline possible approaches to obtain linear equations for the λ_i .

8.3. Homogeneous equations. In a first step, we want to use our results from Section 4 to obtain a list of homogeneous equations for the λ_i . In the previous sections we have been slightly sloppy with notation and we need to be more precise now. The underlying stable graphs of the graphs $\Gamma_i^{\mathcal{Q}}$ listed in Figure 5 index a priori boundary strata of the quotient $\overline{\mathcal{M}}_{1,6}/S_{3+3}$. For such a graph $\Gamma_i^{\mathcal{Q}}$ we denote by $\mathbf{\Gamma}_i^{\mathcal{Q}}$ the set of all possible marked versions of $\Gamma_i^{\mathcal{Q}}$. Our first observation is the following corollary.

COROLLARY 8.1. *Let $\mathcal{M}_{\Delta} \subseteq \overline{\mathcal{M}}_{1,6}$ be a boundary stratum of codimension $d < 3$. If \mathcal{M}_{Δ} has empty intersection with all the strata $\mathcal{M}_{\Gamma^{\mathcal{Q}}}$ for $\Gamma^{\mathcal{Q}} \in \bigcup_i \mathbf{\Gamma}_i^{\mathcal{Q}}$ for $\Gamma_i^{\mathcal{Q}}$ in Proposition 4.1, then*

$$\sum_{i=1}^{756} \lambda_i(\alpha_i \cdot \gamma \cdot [\mathcal{M}_{\Delta}]) = \gamma \cdot [\mathcal{M}_{\Delta}] \cdot [\mathbb{P}\Xi^2\overline{\mathcal{G}}] = 0$$

for all tautological classes $\gamma \in R^{3-d}(\overline{\mathcal{M}}_{1,6})$.

PROOF. This is a direct consequence of Proposition 4.1. \square

To check whether or not \mathcal{M}_{Δ} and $\mathcal{M}_{\Gamma^{\mathcal{Q}}}$ have empty intersection is a combinatorial question about the existence of a common degeneration of the two graphs Δ and $\Gamma^{\mathcal{Q}}$. Unfortunately, using `admcycles` one checks that Corollary 8.1 is completely useless: In the notation of the corollary there is not a single boundary stratum \mathcal{M}_{Δ} that has empty intersection with all the strata $\mathcal{M}_{\Gamma^{\mathcal{Q}}}$.

We can drastically improve the situation by using the fact that the quadratic gothic locus $\mathbb{P}\Xi^2\overline{\mathcal{G}}$ is an irreducible subvariety of the moduli space $\mathbb{P}\Xi^2\overline{\mathcal{M}}_{1,6}(-1^3, 1^3)$. We state the following proposition in a more general context as it might be useful elsewhere.

PROPOSITION 8.2. *Let $\overline{\mathcal{H}} \subseteq \mathbb{P}\Xi^k\overline{\mathcal{M}}_{g,n}(\mu)$ be a closed irreducible subvariety, and let $\Theta_1, \dots, \Theta_l$ be the enhanced level graphs corresponding to the zero dimensional boundary strata $D_{\Theta_i}^{\mathcal{H}}$ of $\overline{\mathcal{H}}$. Let $\mathcal{M}_{\Delta} \subseteq \overline{\mathcal{M}}_{g,n}$ be a boundary stratum. Then $\mathcal{M}_{\Delta} \cap \pi(\overline{\mathcal{H}})$ is non-empty if and only if there is an $i \in \{1, \dots, l\}$ such that Θ_i is a degeneration of Δ .*

PROOF. Assume that Θ_i is a degeneration of Δ , then $\pi(D_{\Theta_i}^{\mathcal{H}}) \subseteq \mathcal{M}_{\Delta} \cap \pi(\overline{\mathcal{H}}) \neq \emptyset$.

For the converse implication, assume that $\mathcal{M}_{\Delta} \cap \pi(\overline{\mathcal{H}}) \neq \emptyset$. Then there is a degeneration $\tilde{\Theta}$ of Δ that admits the structure of an enhanced level graph such that $D_{\tilde{\Theta}}^{\mathcal{H}}$ is non-empty. If $\tilde{\Theta} = \Theta_i$ for some i we are done. So let us assume that $\tilde{\Theta} \neq \Theta_i$ for all i . Then $\dim D_{\tilde{\Theta}}^{\mathcal{H}} \geq 1$ by assumption. We claim that one of the Θ_i must be a degeneration of $\tilde{\Theta}$. This follows from the fact that a stratum of k -differentials never contains a complete curve and hence every boundary stratum must degenerate further if it is of positive dimension, compare the proof of Proposition I.5.17. In particular is this Θ_i a degeneration of Δ . \square

We do not know the precise list of zero dimensional boundary strata of the gothic locus at the moment, but we can obtain a list that certainly contains all the zero dimensional boundary strata by choosing an approach similar to the one in Proposition 4.1:

- List all boundary strata $\Delta^{\overline{\mathcal{H}}_{\mathcal{Q}}} \subseteq \overline{\mathcal{H}}_{\mathcal{Q}}/S_{3+3}$ (in all codimensions)
- Throw away all strata that intersect the divisors $\Delta_{1,2}$ and $\Delta_{2,1}$.
- For each remaining stratum, consider its image under the forgetful map to $\overline{\mathcal{M}}_{1,6}/S_{3+3}$.
- For each stratum in the image check if it admits the structure of an enhanced level graph such that
 - the number of levels plus the number of levels with horizontal edges is between 1 and 3,
 - if the above number is smaller than 3 then there is a level with multiple horizontal edges,
 - after contracting all horizontal edges the graph is an intersection of a number of vertical divisors listed in Proposition 4.1.

The conditions on the number of levels and the number of horizontal edges stems from Theorem 4.4 and the fact that every stratum of codimension 3 must be the intersection of precisely three divisors. With this list of graphs at hand we can apply:

COROLLARY 8.3. *Let $\overline{\mathcal{H}} \subseteq \mathbb{P}\Xi^k \overline{\mathcal{M}}_{g,n}(\mu)$ be a closed irreducible subvariety, and let $\Theta_1, \dots, \Theta_L$ be a set of enhanced level graphs that contains all enhanced level graphs corresponding to the zero dimensional boundary strata $D_{\Theta_i}^{\mathcal{H}}$ of $\overline{\mathcal{H}}$. Let $\mathcal{M}_{\Delta} \subseteq \overline{\mathcal{M}}_{g,n}$ be a boundary stratum. Then $\mathcal{M}_{\Delta} \cap \pi(\overline{\mathcal{H}})$ is empty if non of the Θ_i is a degeneration of Δ .*

If $\mathbf{\Delta}$ is the list of all stable graph such that for all $\Delta \in \mathbf{\Delta}$ the stratum $\mathcal{M}_{\Delta} \subseteq \overline{\mathcal{M}}_{1,6}$ has codimension $d_{\Delta} < 3$ and \mathcal{M}_{Δ} fulfills the condition of the above corollary for our list of possible zero-dimensional strata, then the same idea as in Corollary 8.1 gives us linear equations

$$\sum_{i=1}^{756} \lambda_i (\alpha_i \cdot \gamma \cdot [\mathcal{M}_{\Delta}]) = 0 \quad \text{for all } \Delta \in \mathbf{\Delta} \text{ for all } \gamma \in R^{3-d_{\Delta}}(\overline{\mathcal{M}}_{1,6}).$$

Let A_{hom} be the rational matrix corresponding to this linear system. Using sage we check that

$$\dim(\ker(A_{\text{hom}})) = 79.$$

REMARK 8.4. *Knowing the fundamental class $[\mathbb{P}\Xi^2 \overline{G}]$ would allow us to a posteriori determine which of the possible zero dimensional strata are actually intersected by the gothic locus: A zero dimensional stratum $D_{\Theta}^{\mathcal{Q}}$ is intersected if and only if $[\pi(D_{\Theta}^{\mathcal{Q}})] \cdot [\mathbb{P}\Xi^2 \overline{G}] \neq 0$.*

8.4. Towards inhomogeneous equations from $\overline{\mathcal{M}}_{1,6}$ -divisors. We would like to use the information we obtained in Section 5 about the non-horizontal divisors to obtain more relations. However, as we will see in the following, there are several obstructions on our way.

Let us first consider the divisor $D_{\Gamma_6^{\mathcal{Q}}}^{\mathcal{Q}G}$. For the ambient divisor we have the maps

$$\begin{array}{ccc} D_{\Gamma_6^{\mathcal{Q}}}^{\mathcal{Q}} & \xrightarrow{\pi} & \mathcal{M}_{\Gamma_6^{\mathcal{Q}}} \\ & & \begin{array}{l} \nearrow \tau^{\top} \\ \searrow \tau^{\perp} \end{array} \\ & & \begin{array}{l} \overline{\mathcal{M}}_{1,5} \\ \overline{\mathcal{M}}_{0,3} \end{array} \end{array}$$

where the two rightmost arrows are the projections to the irreducible components at the top and bottom level. Up to a rational factor c that can be worked out explicitly it is

$$(91) \quad c \cdot [\pi(D_{\Gamma_6^{\mathcal{Q}}}^{\mathcal{Q}})] = [\mathcal{M}_{\Gamma_6^{\mathcal{Q}}}] \cdot [\mathbb{P}\Xi^2 \overline{G}]$$

because $\mathcal{M}_{\Gamma_6^{\mathcal{Q}}}$ is a divisor. Assuming that we know the class $[\pi(D_{\Gamma_6^{\mathcal{Q}}}^{\mathcal{Q}})]$ we could evaluate

$$\sum_{i=1}^{756} \lambda_i (\alpha_i \cdot \gamma \cdot [\mathcal{M}_{\Gamma_6^{\mathcal{Q}}})) = \gamma \cdot [\mathcal{M}_{\Gamma_6^{\mathcal{Q}}}] \cdot [\mathbb{P}\Xi^2 \overline{G}] = c \cdot \gamma \cdot [\pi(D_{\Gamma_6^{\mathcal{Q}}}^{\mathcal{Q}})]$$

for any class $\gamma \in R^2(\overline{\mathcal{M}}_{1,6})$ and obtain new relations for the λ_i . The class of $\pi(D_{\Gamma_6^{\mathcal{Q}}}^{\mathcal{Q}})$ is determined by the two classes $\tau^{\bullet}(\pi(D_{\Gamma_6^{\mathcal{Q}}}^{\mathcal{Q}}))$. The image of $\tau^{\perp}(\pi(D_{\Gamma_6^{\mathcal{Q}}}^{\mathcal{Q}}))$ is conveniently given by a point. From the picture in Figure 11 it is not hard to see that the component of $D_{\Gamma_6^{\mathcal{Q}}}^{\mathcal{Q}G, \top}$ to which the depicted surface belongs is a locus of covers of \mathbb{P}^1 branched above 5 points: one simply divides both depicted curves by the involution J (which acts by rotation by π). The class of this locus in $R^2(\overline{\mathcal{M}}_{1,6})$ can be computed with the methods described in [Lia21]. However, we are currently unable to prove that this is all of $D_{\Gamma_6^{\mathcal{Q}}}^{\mathcal{Q}G, \top}$.

PROBLEM 8.5. *Is $D_{\Gamma_6^{\mathcal{Q}}}^{\mathcal{Q}G, \top}$ irreducible?*

Let us given an estimate of how useful this approach can be. There are five non-horizontal graphs with one edge listed in Figure 5, namely $\Gamma_j^{\mathcal{Q}}$ for $j \in J := \{4, 6, 9, 10, 11\}$. We denote by $\Gamma_{\text{vd}}^{\mathcal{Q}} := \bigcup_{j \in J} \Gamma_j^{\mathcal{Q}}$ the set of all marked versions of those vertical divisors. Let B be a basis of $R^2(\overline{\mathcal{M}}_{1,6})$ and consider the matrix

$$A_{\text{vd}} := (\alpha_i \cdot \gamma \cdot [\mathcal{M}_{\Gamma_j^{\mathcal{Q}}}))_{\substack{(\gamma, \Gamma_j^{\mathcal{Q}}) \in B \times \Gamma_{\text{vd}}^{\mathcal{Q}} \\ i=1, \dots, 756}} \in \mathbb{Q}^{|B| \cdot |\Gamma_{\text{vd}}^{\mathcal{Q}}| \times 756}.$$

Then there is a vector $b \in \mathbb{Q}^{|B| \cdot |\Gamma_{\text{vd}}^{\mathcal{Q}}|}$ given by the intersection products on the right hand side of (91) such that

$$A_{\text{vd}} \cdot (\lambda_i)_i = b.$$

While we are currently unable to compute b we can easily compute A with `admcycles` and check that

$$\dim(\ker(A_{\text{vd}})) = 148.$$

In our opinion this is surprisingly small given that our generating set α_i has 756 elements and we only considered 5 of the divisors listed in Figure 5.

More importantly we may consider the stacked matrix

$$A := \begin{pmatrix} A_{\text{hom}} \\ A_{\text{vd}} \end{pmatrix}$$

where A_{hom} is the matrix from the previous Section 8.3 and check that

$$\dim(\ker(A)) = 0.$$

Hence determining the vector b would yield the class of the quadratic gothic locus $[\mathbb{P}\Xi^2\overline{G}]$.

8.5. Towards inhomogeneous equations from other strata. Let us now consider $D_{\Gamma_4^{\mathcal{Q}}}^{\mathcal{Q}G}$, and let π and τ^\bullet be the analogue of the above maps. From our discussion in Section 5.1 it follows that $\tau^\top(\pi(D_{\Gamma_4^{\mathcal{Q}}}^{\mathcal{Q}G}))$ is a locus of pillowcase covers. The class of this locus in $R^2(\overline{\mathcal{M}}_{1,3})$ can again be computed with the methods described in [Lia21]. For the bottom-level we are not aware of any such description. Hence the following problem is currently unsolved.

PROBLEM 8.6. *Compute the class of $\tau^\perp(\pi(D_{\Gamma_4^{\mathcal{Q}}}^{\mathcal{Q}G}))$.*

The situation is slightly different for $D_{\Gamma_5^{\mathcal{Q}}}^{\mathcal{Q}G}$. Again following our discussion in Section 5, one can check that $D_{\Gamma_5^{\mathcal{Q}}}^{\mathcal{Q}G,\top}$ is a locus of pillowcase covers. Hence the class of its images under the analog of $\tau^\top \circ \pi$ can again be computed with the methods described in [Lia21]. The bottom level $D_{\Gamma_5^{\mathcal{Q}}}^{\mathcal{Q}G,\perp}$ is a stratum of quadratic differentials, and the fundamental classes of such strata are known by [BHPSS20] and can be computed using `admcycles`. However, $\mathcal{M}_{\Gamma_5^{\mathcal{Q}}}$ is not a divisor, hence the analogue of Equation (91) does not hold. There is a possible way forward. One could pull back Equation (90) to the tautological Ring of $\mathbb{P}\Xi^2\overline{\mathcal{M}}_{1,6}(-1^3, 1^3)$ in the hope to intersect with the class of the divisor $D_{\Gamma_5^{\mathcal{Q}}}^{\mathcal{Q}}$. There are the following two problems. In theory this pullback can be computed with `diffstrata`. But because `diffstrata` does currently not support graphs with horizontal edges it is not possible to pull back all classes α_i . The other problem is that $\pi^*[\mathbb{P}\Xi^2\overline{G}]$ is a class of codimension 3, while the gothic locus itself is of codimension 2 in $\mathbb{P}\Xi^2\overline{\mathcal{M}}_{1,6}(-1^3, 1^3)$. Hence $\pi^*[\mathbb{P}\Xi^2\overline{G}]$ is actually the intersection of the gothic locus with a class of codimension 1 in the Chow ring of $\mathbb{P}\Xi^2\overline{\mathcal{M}}_{1,6}(-1^3, 1^3)$.

PROBLEM 8.7. *Determine this codimension 1 class and extend `diffstrata` to support horizontal edges.*

All the other non-horizontal divisors have at least one of the above problems. One can hope to obtain the fundamental class of the gothic locus if one solves those problems for a sufficient number of divisors.

Visible Lagrangians for Hitchin systems and pillowcase covers

1. Introduction

Mirror symmetry questions about Higgs bundle moduli spaces have been intensively studied in recent years. The work of Hausel and Thaddeus [HT03] initiated the direction of research by observing the SYZ mirror symmetry of the $\mathrm{SL}(n, \mathbb{C})$ - and $\mathrm{PGL}(n, \mathbb{C})$ -Hitchin system and proving topological mirror symmetry for $n = 2, 3$. Later Donagi and Pantev [DP12] established the duality between Hitchin systems associated to a complex reductive Lie group G and its Langlands dual group G^L . To a complex reductive group G , we associate a moduli space of G -Higgs bundles \mathcal{M}_G with a Hitchin map $\mathrm{Hit} : \mathcal{M}_G \rightarrow \mathcal{B}_G$ to a half-dimensional vector space. Then Donagi and Pantev showed that there is an isomorphism $\mathcal{B}_G \cong \mathcal{B}_{G^L}$, such that the generic fibers over corresponding points under this isomorphism are torsors over dual abelian varieties. Furthermore, the Fourier-Mukai transform yields an equivalence of derived categories of the regular loci of the G and G^L -Hitchin system.

About the same time the work of Kapustin and Witten [KW07] raised the question about mirror symmetry of special subvarieties referred to as branes in the physical literature. A brane is a pair (\mathcal{N}, F) of a subvariety \mathcal{N} and a sheaf F supported on \mathcal{N} with special geometric properties. This initiated plenty of mathematical research to find examples of branes or their supports for G -Hitchin system [BS16]; [Hit16]; [Hit17]; [HS18]; [BS19]; [FJ22]; [HH22]; [FP23]. However, [KW07] also propose a correspondence between branes under mirror symmetry. This seems to be less considered in the mathematical literature (see [Hit16]; [FJ22]; [HH22]; [FP23] for exceptions).

In this paper, we consider so-called (B, A, A) -branes, that is pairs (\mathcal{N}, F) , where \mathcal{N} is a complex Lagrangian subvariety and F is a flat bundle on \mathcal{N} . We also describe the subvarieties and sheaves related to these (B, A, A) -branes by the Fourier-Mukai transform. The work of Kapustin and Witten [KW07] suggests that these are (B, B, B) -branes, i.e. hyperholomorphic subvarieties with a hyperholomorphic sheaf. We give indication for this conjecture for our main example.

More specifically, we are interested in complex Lagrangians \mathcal{L} such that the restriction of the Hitchin map factors through a proper subvariety $\mathcal{B}' = \mathrm{Hit}(\mathcal{L}) \subsetneq \mathcal{B}$. Such Lagrangians are called *visible* in the symplectic geometry literature [Eva23]. This is complementary to the recent work of Hausel and Hitchin [HH22], who studied the upward flow to certain points in the nilpotent cone.

We first abstractly consider visible Lagrangians in the $G = \mathrm{GL}(n, \mathbb{C})$ and $G = \mathrm{SL}(n, \mathbb{C})$ -Hitchin system and describe their proposed mirror dual by computing the Fourier-Mukai transform of flat sheaves on them. The first main result is

THEOREM 1.1 (Theorem 3.2). *Let $\mathcal{L} \subset \mathcal{M}_G$ be a visible Lagrangian, such that $\mathcal{B}' = \mathrm{Hit}(\mathcal{L}) \subsetneq \mathcal{B}_G$ and $\mathcal{B}' \cap \mathcal{B}_G^{\mathrm{reg}} \neq \emptyset$. Let $s : \mathcal{B}' \rightarrow \mathcal{L}$ be a section of $\mathrm{Hit}|_{\mathcal{L}}$. The fiber-wise Fourier-Mukai transform of the structure sheaf $\mathcal{O}_{\mathcal{L}}$ is supported on a holomorphic symplectic subvariety $\mathcal{I}_{\mathcal{L}, s} \subset \mathcal{M}_{G^L}$, such that $\mathrm{Hit}|_{\mathcal{I}} : \mathcal{I}_{\mathcal{L}, s} \rightarrow \mathcal{B}'$ is an algebraically completely integrable system.*

Note that every hyperholomorphic subvariety is holomorphic symplectic. Hence, this observation fits well with the mirror symmetry proposal in [KW07].

On the question of existence, we mention three situations under which we expect visible Lagrangians, examples of which appeared in the literature. The first is Lie-theoretic. An inclusion of a semisimple Lie group G_1 into a reductive Lie group G_2 defines a morphism of Higgs bundle moduli spaces $\mathcal{M}_{G_1} \rightarrow \mathcal{M}_{G_2}$. Its image is a hyperholomorphic subvariety. We expect the mirror dual to this hyperholomorphic subvarieties to be visible Lagrangians.

We consider the example of $\mathrm{SL}(n, \mathbb{C}) \subset \mathrm{GL}(n, \mathbb{C})$ in Section 3.3. Another example where $G_1 = \mathrm{Sp}(2n, \mathbb{C})$ and $G_2 = \mathrm{GL}(2n, \mathbb{C})$ was considered by Hitchin in [Hit16]. Other than that, we hope to return to this type of visible Lagrangians in a subsequent work and will mainly focus on $G = \mathrm{SL}(2, \mathbb{C})$ in the remainder of the paper. The second type of visible Lagrangians appeared in the work of [FGOP21]; [FP23]; [FHHO23]. They are completely contained in the singular locus of the Hitchin map and substantially use the geometry of the singular Hitchin fibers. In particular, they do not fall within the scope of this work.

The focus of this work is a third type of visible Lagrangians related to the symmetries of the underlying Riemann surface. The general fiber of the Hitchin system is a torsor over an abelian variety. A necessary condition for the existence of a visible Lagrangian $\mathcal{L} \rightarrow \mathcal{B}'$ is that the Hitchin fibers over \mathcal{B}' correspond to reducible (i.e. non-simple) abelian varieties. Comparing the dimension of the $\mathrm{SL}(2, \mathbb{C})$ -Hitchin base and the reducible locus in the corresponding moduli space of abelian varieties suggests that there are finitely many directions in the Hitchin base, where the Hitchin fiber is isomorphic to a reducible abelian variety. Hence, it is natural to look for visible Lagrangians $\mathcal{L} \rightarrow \mathcal{B}'$ over lines $\mathbb{C}a \subset \mathcal{B}_G$ in the Hitchin base. We have the following second main theorem, of which a $\mathrm{SL}(n, \mathbb{C})$ -version is proven in Theorem 6.1.

THEOREM 1.2 (Corollary 6.2). *Let $q \in H^0(X, K_X^2)$ be a quadratic differential with simple zeros only. Then there exists a visible Lagrangian*

$$\mathcal{L} \rightarrow \mathcal{B}' = \{tq \mid t \in \mathbb{C}\} \subset \mathcal{B}_{\mathrm{SL}(2, \mathbb{C})}(X)$$

if and only if (X, q) is a pillowcase cover.

The notion of pillowcase cover stems from the theory of flat surfaces. It means that there is a covering $X \rightarrow \mathbb{P}^1$, such that the quadratic differential q is the pullback of a quadratic differential on \mathbb{P}^1 with four simple poles. The later should be figured as a pillowcase, see Figure 2. We give a short introduction to the idea of flat surfaces in Section 4.

Motivated by the above considerations on the moduli space of abelian varieties, in Section 5, we study examples of Riemann surfaces, where there exist several lines in the $\mathrm{SL}(2, \mathbb{C})$ -Hitchin base associated to visible Lagrangians. We prove the following result, which might be of independent interest from the point of view of flat surfaces.

THEOREM 1.3 (Proposition 5.1). *There exist infinitely many genera g , such that there exists a Riemann surface X of genus g with two quadratic differentials q_1, q_2 with simple zeros only, such that (X, q_i) are pillowcase covers and q_1, q_2 are not related by pullback along an automorphism of X .*

Finally, we consider the subintegrable system $\mathcal{I} \subset \mathcal{M}_{\mathrm{PGL}(2, \mathbb{C})}$ of Theorem 1.1 associated to the visible Lagrangian \mathcal{L} of Theorem 1.2. We observe that \mathcal{I} is birational to Hausel's toy model [Hau98]. Under the natural extra condition on the pillowcase cover to be *uniform*, we prove the following theorem that confirms the Kapustin-Witten picture for visible Lagrangians of this kind. All the pillowcase covers of Theorem 1.3 are uniform.

THEOREM 1.4 (Corollary 6.5). *Let (X, q) be a uniform pillowcase cover with simple zeros only and $\mathcal{L} \rightarrow \mathbb{C}q$ the visible Lagrangian of Theorem 1.2. Then the associated subintegrable system $\mathcal{I} \subset \mathcal{M}_{\mathrm{PGL}(2, \mathbb{C})}$ of Theorem 1.1 is a hyperholomorphic subvariety.*

2. Symplectic geometry

2.1. Completely integrable systems.

DEFINITION 2.1. A completely integrable system is a holomorphic symplectic manifold (M, Ω) together with a proper flat morphism $H : M \rightarrow B$ to a complex manifold B , such that on the complement $B \setminus S$ of some proper closed subvariety S the fibers of H are complex Lagrangian tori. It is called algebraically completely integrable, if the Lagrangian tori are endowed with continuously varying polarizations $\rho_b \in H^{(1,1)}(M_b) \cap H^2(M_b, \mathbb{Z})$, i.e. they are abelian varieties.

We will refer to $B^{\mathrm{reg}} = B \setminus S$ as the regular locus and to S as the singular locus of a completely integrable system.

DEFINITION 2.2. An integral affine structure on a smooth manifold B is a torsion-free flat connection on the tangent bundle TB .

For completely integrable systems there is a natural identification of the cotangent bundle to the base with the torus-invariant vector fields along to the fibers of H given as follows. Let $\alpha \in T_b^\vee B$ be a (holomorphic) one-form then there exists an invariant vector field X on M_b , such that $\Omega(X, \cdot) = H^*\alpha$. Denoting by VM the bundle of invariant vector fields along to the fibers over B^{reg} we obtain an identification $T^\vee B^{\text{reg}} \cong VM$. Locally over $U \subset B^{\text{reg}}$ we can choose a section of $H : M \rightarrow B$ and identify $H^{-1}(U) \cong V_U M / \Lambda$ for a family of lattices $\Lambda \subset V_U M$. This yields a family of lattices $\Lambda \subset T^\vee B^{\text{reg}}$. The dual family of lattices $\Lambda^\vee \subset TB^{\text{reg}}$ defines a torsion-free flat connection on TB^{reg} , where a section is flat if and only if it is constant with respect to lattice coordinates (see [Fre99, §3] for more details).

2.2. Visible Lagrangians. This idea goes back to lecture notes of Jonathan David Evans [Eva23, Chapt. 5] in the context of real completely integrable systems.

DEFINITION 2.3 (Visible Lagrangians). Let $H : M \rightarrow B$ be a completely integrable system. A Lagrangian subvariety $\mathcal{L} \subset M$ is called *visible*, if $H|_{\mathcal{L}} : \mathcal{L} \rightarrow B$ factors as $H|_{\mathcal{L}} = f \circ g$, where $f : B' \rightarrow B$ is an embedding of a proper subvariety B' , such that $g|_{B' \setminus S'} : \mathcal{L}|_{B' \setminus S'} \rightarrow B' \setminus S'$ is a smooth fiber bundle on the complement of some proper subvariety $S' \subsetneq B'$.

The simplest example of a visible Lagrangian is a complex torus fiber M_b with $B' = \{b\}$. On the other hand, a Lagrangian section $s : B \rightarrow M$ is an example of a Lagrangian that is not visible. We denote $B'^{\text{reg}} = B' \setminus S'$. For the visible Lagrangians considered in the present work we will mostly have $B'^{\text{reg}} = B' \cap B^{\text{reg}}$.

THEOREM 2.4. Let $H : M \rightarrow B$ be a completely integrable system and $\mathcal{L} \rightarrow B' \subsetneq B$ a visible Lagrangian with $B'^{\text{reg}} \subset B^{\text{reg}}$. Then at each smooth point $b \in B'^{\text{reg}}$ the base locus B' is rational with respect to the integral affine structure on B^{reg} and for $b \in B'^{\text{reg}}$ the fiber \mathcal{L}_b is a union of complex tori generated by the invariant vector fields $T_b B'^\perp \subset V_b M$.

PROOF. Locally at a smooth point $b \in B'^{\text{reg}}$ the subvariety $B' \subset B$ is cut out by $k = \text{codim} B'$ many functions $f_1, \dots, f_k \in \mathcal{O}_B$. We can associate invariant vector fields $X_i \in VM|_{B'}$ along the torus fibers so that $\Omega(X_i, \cdot) = H^*df_i$ for $i = 1, \dots, k$. Let $m \in \mathcal{L}_b$ and $Y \in T_m M$, such that $DH(Y) \in T_b B'$, then

$$(92) \quad \Omega(X_i, Y) = H^*df_i Y = df_i(DH(Y)) = 0.$$

Therefore, the connected components of the fiber of \mathcal{L} over $b \in B'^{\text{reg}}$ are integral submanifolds of the distribution $V\mathcal{L} = \text{span}(X_1, \dots, X_k)$. Hence the connected components of the fibers of $\mathcal{L}_b \rightarrow B'^{\text{reg}}$ are complex subtori. In particular, the subspace $V_b \mathcal{L} \subset V_b M$ is rational with respect to the lattice Λ_b . By definition this is equivalent to $H_* T\mathcal{L} = TB'$ being rational with respect to the integral affine structure on B . Finally, by (92) we have $V\mathcal{L} \subset (TB'^{\text{reg}})^\perp \subset T^\vee B^{\text{reg}} \cong VM$. The first inclusion is an equality both being of rank k . \square

3. Hitchin systems

In this section, we will briefly review Hitchin systems - the algebraically completely integrable systems of interest in this work. Then we will give an example of a visible Lagrangian of $\mathcal{M}_{\text{GL}(n, \mathbb{C})}$ that stems from the embedding of $\text{SL}(n, \mathbb{C}) \subset \text{GL}(n, \mathbb{C})$. In the remainder of the paper we will focus on visible Lagrangians that do not come from Lie theory.

3.1. Preliminaries about the $\text{GL}(n, \mathbb{C})$ and $\text{SL}(n, \mathbb{C})$ -Hitchin systems. Let G be a complex reductive Lie group and X a Riemann surface, then there is a moduli space of stable G -Higgs bundles. More precisely, we denote by $\mathcal{M}_G(X)$ the neutral component of the moduli space of stable G -Higgs bundles. It is a hyperkähler manifold, in particular, holomorphic symplectic. There is the Hitchin map $\text{Hit} : \mathcal{M}_G(X) \rightarrow \mathcal{B}_G(X)$ to a half-dimensional vector space $\mathcal{B}_G(X)$, which is an algebraically completely integrable system in the sense of Definition 2.1. So it is sensible to ask for the existence of visible Lagrangians.

Let us be more concrete about the cases $G = \mathrm{GL}(n, \mathbb{C})$ and $G = \mathrm{SL}(n, \mathbb{C})$. A $\mathrm{GL}(n, \mathbb{C})$ -Higgs bundles is a pair (E, Φ) of a holomorphic vector bundle E together with a section $\Phi \in H^0(X, \mathrm{End}(E) \otimes K_X)$, where K_X is the canonical bundle of X . Considering the neutral component of the $\mathrm{GL}(n, \mathbb{C})$ -moduli space means to fix the degree of E to be 0. For (E, Φ) to be a $\mathrm{SL}(n, \mathbb{C})$ -Higgs bundle we add the condition of the determinant bundle of E and trace of Φ to be trivial. In the $\mathrm{GL}(n, \mathbb{C})$ -case the Hitchin map is given by

$$\mathrm{Hit} : \mathcal{M}_{\mathrm{GL}(n, \mathbb{C})} \rightarrow \mathcal{B}_{\mathrm{GL}(n, \mathbb{C})} = \bigoplus_{i=1}^n H^0(X, K^i), \quad (E, \Phi) \mapsto \underline{a}(\Phi) = (a_1(\Phi), \dots, a_n(\Phi)),$$

where $a_i \in \mathbb{C}[\mathfrak{gl}(n, \mathbb{C})]^{\mathrm{GL}(n, \mathbb{C})}$ is the i -th coefficient of the characteristic polynomial. In the $\mathrm{SL}(n, \mathbb{C})$ -case it is given by

$$\mathrm{Hit} : \mathcal{M}_{\mathrm{SL}(n, \mathbb{C})} \rightarrow \mathcal{B}_{\mathrm{SL}(n, \mathbb{C})} = \bigoplus_{i=2}^n H^0(X, K^i), \quad (E, \Phi) \mapsto \underline{a}(\Phi) = (a_2(\Phi), \dots, a_n(\Phi)).$$

Fixing a point $\underline{a} \in \mathcal{B}_G$ in the Hitchin base the eigenvalues of Φ define a branched n -sheeted cover $\pi : \Sigma_{\underline{a}} \rightarrow X$ - the so-called spectral curve. The discriminant of the characteristic polynomial defines a map $\mathrm{disc}_G : \mathcal{B}_G \rightarrow H^0(X, K^{r(r-1)})$. The discriminant locus $\Delta_G \subset \mathcal{B}_G$ is the preimage of the sections of $H^0(X, K^{r(r-1)})$ with higher order zeros under disc_G . Its complement $\mathcal{B}_G^{\mathrm{reg}} = \mathcal{B}_G \setminus \Delta_G$ is referred to as the regular locus. In particular, for $\underline{a} \in \mathcal{B}_G^{\mathrm{reg}}$ the spectral curve $\Sigma_{\underline{a}}$ is smooth. For $G = \mathrm{SL}(2, \mathbb{C})$ the regular locus $\mathcal{B}_{\mathrm{SL}(2, \mathbb{C})}^{\mathrm{reg}}$ is the locus of quadratic differentials with simple zeros only.

The fibers of the $\mathrm{GL}(n, \mathbb{C})$ -Hitchin map over $\mathcal{B}_{\mathrm{GL}(n, \mathbb{C})}^{\mathrm{reg}}$ - the so-called regular fibers - are torsors over $\mathrm{Jac}(\Sigma_{\underline{a}})$ via the spectral correspondence. The branched cover $\pi : \Sigma_{\underline{a}} \rightarrow X$ defines a norm map $\mathrm{Nm}_{\Sigma/X} : \mathrm{Jac}(\Sigma_{\underline{a}}) \rightarrow \mathrm{Jac}(X)$. The kernel of this morphism defines the Prym variety $\mathrm{Prym}(\Sigma_{\underline{a}})$. A regular fiber of the $\mathrm{SL}(n, \mathbb{C})$ -Hitchin system is a torsor over $\mathrm{Prym}(\Sigma_{\underline{a}})$. The neutral component of the Higgs bundles moduli space allows for the existence of a section $s_H : \mathcal{B}_G \rightarrow \mathcal{M}_G$ of the Hitchin map - the so-called Hitchin section. For $\mathrm{GL}(n, \mathbb{C})$ it is given by

$$s_H(\underline{a}) = \left(K^{\frac{n-1}{2}} \oplus \dots \oplus K^{-\frac{n-1}{2}}, \begin{pmatrix} a_1 & a_2 & \dots & a_n \\ 1 & a_1 & \ddots & \vdots \\ & \ddots & \ddots & a_2 \\ & & & 1 & a_1 \end{pmatrix} \right).$$

The tangent space to $\mathrm{Jac}(\Sigma_{\underline{a}})$ at the identity is $H^1(\Sigma_{\underline{a}}, \mathcal{O}_{\Sigma_{\underline{a}}})$ by the exponential sequence. The inclusion of the Hitchin fiber into $\mathcal{M}_{\mathrm{GL}(n, \mathbb{C})}$ yields an exact sequence of tangent spaces

$$0 \rightarrow H^1(\Sigma_{\underline{a}}, \mathcal{O}_{\Sigma_{\underline{a}}}) \rightarrow T_{(E, \Phi)} \mathcal{M}_{\mathrm{GL}(n, \mathbb{C})} \rightarrow H^0(\Sigma_{\underline{a}}, K_{\Sigma_{\underline{a}}}) \rightarrow 0$$

(see [Mar94, Prop. 8.2]). The holomorphic symplectic form identifies the vertical tangent vectors with the dual of the tangent space to the Hitchin base. Combining this with Serre duality yields the following identification of the tangent space to base with differentials on $\Sigma_{\underline{a}}$.

PROPOSITION 3.1 ([Bar16, Proposition 3.4]). *Let $\underline{a} \in \mathcal{B}_{\mathrm{GL}(n, \mathbb{C})}^{\mathrm{reg}}$. The identification of the tangent space $T_{\underline{a}} \mathcal{B}_{\mathrm{GL}(n, \mathbb{C})}$ with the dual of the tangent space to the fiber $T_L^\vee \mathrm{Jac}(\Sigma_{\underline{a}}) \cong H^0(\Sigma_{\underline{a}}, K_\Sigma)$ is given by*

$$t : \bigoplus_{i=1}^n H^0(X, K_X^i) \rightarrow H^0(\Sigma_{\underline{a}}, K_{\Sigma_{\underline{a}}}) \\ \sum_{i=1}^n \alpha_i X_i \mapsto \frac{1}{s_B} \sum_{i=1}^n \alpha_i \pi^* X_i (\lambda^{n-i} + \pi^* a_2 \lambda^{n-2-i} + \dots + \pi^* a_{n-i}),$$

where $s_B = d\pi \in H^0(X, \pi^* K_X^{n-1})$ is supported at the branch divisor.

PROOF. The original statement gives the isomorphism with values in $H^0(\Sigma_{\underline{a}}, \pi^* K_X^n)$. We have the isomorphism of sheaves $K_{\Sigma_{\underline{a}}} \rightarrow \pi^* K_X^n$ given by $\phi \mapsto \phi s_B|_U$ for $\phi \in K_{\Sigma_{\underline{a}}}|_U$.

This is well-defined by the following: Let (U, w) a coordinate disc centered at $y \in \Sigma$, such that $\pi : U \rightarrow \pi(U)$, $w \mapsto z = w^b$ with $b \geq 1$. Then $s_B|_U = w^{b-1}(\pi^*dz)^{n-1}$. Hence, $\phi \cdot s_B(U) = fdw \cdot w^{b-1}(\pi^*dz)^{n-1} = f(\pi^*dz)^n$, where we used that $\pi^*dz = w^{b-1}dw$. Composing Baraglia's isomorphism with the inverse of the above defines the asserted isomorphism t . \square

In the following proposition, we take an algebro-geometric point of view and will consider the family of smooth curves $\Sigma \rightarrow \mathcal{B}_G^{\text{reg}}$. The $\text{GL}(n, \mathbb{C})$ -Hitchin system is a torsor over the abelian scheme defined by the relative Jacobian $\text{Jac}(\Sigma/\mathcal{B}_G^{\text{reg}})$.

PROPOSITION 3.2. *A subvariety $\mathcal{L} \subset \mathcal{M}_{\text{GL}(n, \mathbb{C})}$ over a proper subvariety $\mathcal{B}' \subsetneq \mathcal{B}_{\text{GL}(n, \mathbb{C})}$ with $\mathcal{B}'^{\text{reg}} \subset \mathcal{B}^{\text{reg}}$ and connected fibers is a visible Lagrangian if and only if*

- i) *There exists an abelian subscheme $A \subset \text{Jac}(\Sigma/\mathcal{B}_G^{\text{reg}})|_{\mathcal{B}'^{\text{reg}}}$ over $\mathcal{B}'^{\text{reg}}$, such that \mathcal{L} is an A -torsor.*
- ii) *The relative tangent bundle $T_{A/\mathcal{B}'^{\text{reg}}} \subset T_{\text{Jac}(\Sigma)/\mathcal{B}'^{\text{reg}}}$ is the kernel of the map*

$$T_{\text{Jac}(\Sigma)/\mathcal{B}'^{\text{reg}}} \cong R^1\pi_*\mathcal{O}_\Sigma \rightarrow \mathbb{C}^{\dim \mathcal{B}'^{\text{reg}}}$$

defined by evaluating on the image of

$$T\mathcal{B}'^{\text{reg}} \rightarrow R^0\pi_*K_\Sigma, \quad X \mapsto t(X)$$

using the Serre pairing on Σ . Here t was defined in Proposition 3.1.

PROOF. By Theorem 2.4 fibers of $\mathcal{L} \rightarrow \mathcal{B}'^{\text{reg}}$ are complex subtori in the Hitchin fiber. Hence, using a local section $s : U \rightarrow \mathcal{L}$ on an open $U \subset \mathcal{B}'^{\text{reg}}$, we can identify the fibers with an abelian subscheme of $A \rightarrow \text{Jac}(\Sigma/\mathcal{B}'^{\text{reg}})$. Condition ii) is the family version of Theorem 2.4 reformulated by using the observations about the holomorphic symplectic form on $\mathcal{M}_{\text{GL}(n, \mathbb{C})}$ in the previous paragraph. Conversely, condition ii) is equivalent to the restriction of the symplectic form of $\mathcal{M}_{\text{GL}(n, \mathbb{C})}$ to be zero on the tangent bundle to \mathcal{L} . Hence, \mathcal{L} is a Lagrangian. \square

In general, a visible Lagrangian might not have connected fibers. We will give an example in Theorem 3.5.

To obtain an analogous statement in the $\text{SL}(n, \mathbb{C})$ -case we have to identify the image of $T\mathcal{B}_{\text{SL}(n, \mathbb{C})}^{\text{reg}} \subset T\mathcal{B}_{\text{GL}(n, \mathbb{C})}^{\text{reg}}$ through the isomorphism t of Proposition 3.1. The pullback $\pi^* : H^0(X, K_X) \rightarrow H^0(\Sigma_{\underline{a}}, K_{\Sigma_{\underline{a}}})$ defines an inclusion of the differentials on X into the differentials on $\Sigma_{\underline{a}}$. We define the linear map

$$\text{pr}_X : H^0(\Sigma_{\underline{a}}, K_{\Sigma_{\underline{a}}}) \rightarrow H^0(X, K_X), \quad \lambda \mapsto \eta,$$

where η is define as follows: Let $U \subset X$ be a trivially covered open, i.e. $\pi^{-1}(U) = \bigsqcup_{i=1}^n U_i$. Define $\eta(U) = \frac{1}{n} \sum_{i=1}^n \lambda(U_i)$. Then η extends to an abelian differential on X by the Riemann extension theorem. Clearly $\text{pr}_X \circ \pi^* = \text{id}$. Denote by $H^0(\Sigma_{\underline{a}}, K_{\Sigma_{\underline{a}}})^-$ the kernel of pr_X . Then this induces a splitting $H^0(\Sigma_{\underline{a}}, K_{\Sigma_{\underline{a}}}) = H^0(X, K_X) \oplus H^0(\Sigma_{\underline{a}}, K_{\Sigma_{\underline{a}}})^-$.

The Prym variety is the kernel of the Norm map $\text{Nm}_{\Sigma/X} : \text{Jac}(\Sigma_{\underline{a}}) \rightarrow \text{Jac}(X)$ induced by the Norm map on structure sheaves $\text{Nm}_{\Sigma/X} : \pi_*\mathcal{O}_{\Sigma_{\underline{a}}} \rightarrow \mathcal{O}_X$, $\pi_*f \rightarrow \det(\pi_*f)$. Tangentially, we have a splitting of $\pi_*\mathcal{O}_{\Sigma_{\underline{a}}} = \mathcal{O}_X \oplus \pi_*\mathcal{O}_{\Sigma_{\underline{a}}}^-$, where $\pi_*\mathcal{O}_{\Sigma_{\underline{a}}}^-$ is the kernel of $\text{nm}_{\Sigma/X} : \pi_*\mathcal{O}_{\Sigma_{\underline{a}}} \rightarrow \mathcal{O}_X$, $\pi_*f \mapsto \text{tr}(\pi_*f)$. This yields a splitting of the tangent space to the $\text{GL}(n, \mathbb{C})$ -Hitchin fiber over \underline{a}

$$\begin{aligned} T\text{Jac}(\Sigma_{\underline{a}}) &= H^1(\Sigma_{\underline{a}}, \mathcal{O}_{\Sigma_{\underline{a}}}) = H^1(X, \mathcal{O}_X) \oplus H^1(X, \pi_*\mathcal{O}_{\Sigma_{\underline{a}}}^-) \\ &= T\pi^*\text{Jac}(X) \oplus T\text{Prym}(\Sigma_{\underline{a}}). \end{aligned}$$

Let $\alpha \in H^0(X, K_X)$ and $\beta \in H^1(\Sigma_{\underline{a}}, \mathcal{O}_{\Sigma_{\underline{a}}}) \cong H^{(0,1)}(\Sigma_{\underline{a}})$, then

$$\int_{\Sigma_{\underline{a}}} \pi^*\alpha \wedge \beta = \int_X \alpha \wedge \text{nm}_{\Sigma/X}(\beta).$$




FIGURE 1. Langland's duality between $\mathrm{SL}(n, \mathbb{C})$ -Hitchin system and $\mathrm{PGL}(n, \mathbb{C})$ -Hitchin system

Therefore, the Serre pairing on $\Sigma_{\underline{a}}$ restricts to a non-degenerate pairing between $H^0(\Sigma_{\underline{a}}, K_{\Sigma_{\underline{a}}})^-$ and $H^1(\Sigma_{\underline{a}}, \mathcal{O}_{\Sigma_{\underline{a}}})^-$. Consequently, the isomorphism t of Proposition 3.1 restricts to

$$t : \bigoplus_{i=2}^n H^0(X, K_X^i) \rightarrow H^0(\Sigma_{\underline{a}}, K_{\Sigma_{\underline{a}}})^-$$

by the same formula.

Now Proposition 3.2 readily translates to a characterization of visible Lagrangians in $\mathcal{M}_{\mathrm{SL}(n, \mathbb{C})}$.

3.2. Langland's duality for Hitchin systems. In this section, we will review the Langland's duality of Hitchin system as considered in [DP12]. We will fix $G = \mathrm{SL}(n, \mathbb{C})$ and $G^L = \mathrm{PGL}(n, \mathbb{C})$ (cf. Remark 3.3). Recall that we consider the neutral components of the moduli spaces. The situation is visualized in Figure 1. First there is an isomorphism of the Hitchin bases for G and G^L , which maps the G -discriminant locus to the G^L -discriminant locus. In the cases under consideration, this isomorphism is the identity. The $\mathrm{SL}(n, \mathbb{C})$ -Hitchin system restricted to $\mathcal{B}_{\mathrm{SL}(n, \mathbb{C})}^{\mathrm{reg}}$ is a torsor over the abelian scheme $\mathrm{Prym}(\Sigma/\mathcal{B}_{\mathrm{SL}(n, \mathbb{C})}^{\mathrm{reg}})$. On the other hand, the (neutral component of the) moduli space of $\mathrm{PSL}(n, \mathbb{C})$ -Higgs bundles is the quotient

$$\mathcal{M}_{\mathrm{PSL}(n, \mathbb{C})} = \mathcal{M}_{\mathrm{SL}(n, \mathbb{C})} / \mathrm{Jac}(X)[n]$$

under the action of $\mathrm{Jac}(X)[n]$ on the moduli spaces of $\mathrm{SL}(n, \mathbb{C})$ -Higgs bundles by tensor product. The $\mathrm{PGL}(n, \mathbb{C})$ -Hitchin system restricted to $\mathcal{B}_{\mathrm{PSL}(n, \mathbb{C})}^{\mathrm{reg}}$ is a torsor for the dual abelian scheme $\mathrm{Prym}(\Sigma/\mathcal{B}_{\mathrm{PSL}(n, \mathbb{C})}^{\mathrm{reg}})^{\vee} = \mathrm{Prym}(\Sigma/\mathcal{B}_{\mathrm{SL}(n, \mathbb{C})}^{\mathrm{reg}})/\pi^* \mathrm{Jac}(X)[n]$. In the case of $\mathrm{SL}(n, \mathbb{C})$ and $\mathrm{PGL}(n, \mathbb{C})$ the quotient map $\delta : \mathcal{M}_G \rightarrow \mathcal{M}_{G^L} = \mathcal{M}_G / \mathrm{Jac}(X)[n]$ extends the polarization to a finite morphism between the moduli spaces. This morphism is holomorphic symplectic, i.e. it induces an symplectic isomorphism of tangent spaces at the points with non-trivial stabilizer.

THEOREM 3.1. *Let $\mathcal{L} \rightarrow \mathcal{B}' \subset \mathcal{B}_G$ be a visible Lagrangian in \mathcal{M}_G with $\mathcal{B}'^{\mathrm{reg}} \subset \mathcal{B}^{\mathrm{reg}}$, connected fibers and a section $s : \mathcal{B}'^{\mathrm{reg}} \rightarrow \mathcal{L}|_{\mathcal{B}'^{\mathrm{reg}}}$. Then there exists a holomorphic symplectic subvariety $\mathcal{I} \subset \mathcal{M}_{G^L}$, such that $\mathcal{I} \rightarrow \mathcal{B}' \subset \mathcal{B}_{G^L}$ is an algebraically completely integrable system and $s' = \delta \circ s$ defines a section $s' : \mathcal{B}'^{\mathrm{reg}} \rightarrow \mathcal{I}|_{\mathcal{B}'^{\mathrm{reg}}}$.*

PROOF. Using the section s we can identify the Hitchin system over $\mathcal{B}'^{\mathrm{reg}}$ with the abelian scheme $\mathrm{Prym}(\Sigma/\mathcal{B}'^{\mathrm{reg}}) \rightarrow \mathcal{B}'^{\mathrm{reg}}$. By Theorem 2.4 the fiber \mathcal{L}_b for $b \in \mathcal{B}'^{\mathrm{reg}}$ is complex subtorus that by assumption contains $s(b)$. Hence, \mathcal{L} defines an abelian subscheme $A \subset \mathrm{Prym}(\Sigma/\mathcal{B}'^{\mathrm{reg}})$. We obtain an exact sequence of abelian schemes over $\mathcal{B}'^{\mathrm{reg}}$

$$(93) \quad 0 \rightarrow A \rightarrow \mathrm{Prym}(\Sigma/\mathcal{B}'^{\mathrm{reg}}) \rightarrow Q \rightarrow 0,$$

where the quotient Q is again an abelian scheme over $\mathcal{B}'^{\mathrm{reg}}$. Dually, we obtain an exact sequence of abelian schemes

$$(94) \quad 0 \rightarrow Q^{\vee} \rightarrow \mathrm{Prym}(\Sigma/\mathcal{B}'^{\mathrm{reg}})^{\vee} \rightarrow A^{\vee} \rightarrow 0.$$

We can define a section $s' = \delta \circ s : \mathcal{B}'^{\text{reg}} \rightarrow \text{Hit}_{G^L}^{-1}(\mathcal{B}'^{\text{reg}}) \subset \mathcal{M}_{G^L}$. Acting by Q^\vee on $s' : \mathcal{B}'^{\text{reg}} \rightarrow \mathcal{M}_{G^L}$ defines a submanifold $I \subset \text{Hit}^{-1}(\mathcal{B}'^{\text{reg}}) \subset \mathcal{M}_{G^L}$. We define $\mathcal{I} = \bar{I}$. We want to show that this is subintegrable system with regular locus I , i.e. the holomorphic symplectic form restricts to a non-degenerate form on I and $\text{Hit}|_{\mathcal{I}} : \mathcal{I} \rightarrow \mathcal{B}'^{\text{reg}}$ is an integrable system. First note that for all $b \in \mathcal{B}'^{\text{reg}}$

$$\dim \mathcal{I}_b = \dim(\text{Prym}(\Sigma_b)) - \dim \mathcal{L}_b = \frac{1}{2} \dim \mathcal{M}(X, G) - \dim \mathcal{L}_b = \dim B'^{\text{reg}}.$$

Hence, the restricted Hitchin map will lead the correct number of commuting Hamiltonian functions and the fibers are complex tori by definition. To show that it forms a subintegrable system it suffices to show that the tangent space of \mathcal{I} at the section s' is a symplectic vector space with the restriction of the symplectic form. The argument at a general point follows by a translation along the fibers of \mathcal{I} . We identify the tangent spaces $T_{s(b)}\mathcal{M}_G \cong T_{s'(b)}\mathcal{M}_{G^L}$ using δ . Then the vertical tangent spaces to \mathcal{L} and \mathcal{I} are complementary by definition. The tangent vectors in $Ds'(T\mathcal{B}'^{\text{reg}})$ pair to zero with the tangent vectors to \mathcal{L} being Lagrangian. One the other hand, the symplectic form on $T_{s'}\mathcal{M}_{G^L}$ is non-degenerate. Hence, it restricts to a non-degenerate symplectic form on $T_{s'}\mathcal{I}$. \square

THEOREM 3.2. *Let $\mathcal{L} \rightarrow \mathcal{B}' \subset \mathcal{B}_G$ be a visible Lagrangian in \mathcal{M}_G with $\mathcal{B}'^{\text{reg}} \subset \mathcal{B}^{\text{reg}}$, connected fibers and a section $s : \mathcal{B}'^{\text{reg}} \rightarrow \mathcal{L}$. We identify the G - (respectively G^L -) Hitchin system over $\mathcal{B}'^{\text{reg}}$ with the abelian schemes using the section s (respectively $s' = \delta \circ s$). Then the fiber-wise Fourier-Mukai transform of the structure sheaf $\mathcal{O}_{\mathcal{L}}$ over $\mathcal{B}'^{\text{reg}}$ is the structure sheaf of the holomorphic symplectic subvariety $\mathcal{I} \subset \mathcal{M}_{G^L}$ defined in Theorem 3.1.*

PROOF. Let $\underline{a} \in \mathcal{B}'^{\text{reg}}$ and denote by $P = \text{Prym}(\Sigma_{\underline{a}})$. As in the previous proof we use the section s to obtain the exact sequences of abelian varieties

$$0 \rightarrow A \rightarrow P \rightarrow Q \rightarrow 0 \quad (93) \quad \text{and} \quad 0 \rightarrow Q^\vee \rightarrow P^\vee \rightarrow A^\vee \rightarrow 0 \quad (94).$$

We use the symmetric Fourier-Mukai transform introduced in [Sch22]. More precisely, we are going to use Proposition 1.6 therein. We denote by $\mathcal{P} \rightarrow P \times P^\vee$ the Poincaré bundle. Then the symmetric Fourier-Mukai transform is defined by

$$\text{SFM}_P : D^b(P) \rightarrow D^b(P^\vee), \quad \text{FM}_{\mathcal{P}} \circ \Delta_P,$$

where $\Delta_P = \text{Hom}(\cdot, \omega_P[\dim P])$ is the Serre duality functor. First we have to show that \mathcal{O}_A is a GV-sheaf. Let $\xi \in P^\vee$. We have

$$H^i(P, \iota_* \mathcal{O}_A \otimes \mathcal{P}_\xi^{-1}) = H^i(A, \mathcal{P}_\xi^{-1}|_A).$$

It is zero for $i > \dim A$. For $i \leq \dim A$ it is non-zero if and only if $\mathcal{P}_\xi^{-1}|_A$ is trivial. Hence, if and only if $\xi \in Q^\vee$. Hence, the support loci

$$\{\xi \in P^\vee \mid H^i(P, \iota_* \mathcal{O}_A \otimes \mathcal{P}_\xi^{-1}) \neq 0\}$$

have codimension $\geq i$ for all i . Therefore, $\iota_* \mathcal{O}_A$ is GV on P . In particular, it is WIT - that is $\text{SFM}_P(\mathcal{O}_A)$ is a sheaf - by [PP11, Proposition 3.2].

The symmetric Fourier-Mukai transform on A has the property

$$(95) \quad \text{SFM}_A(\mathcal{O}_A) = \mathbb{C}_0 \quad \text{SFM}_{A^\vee}(\mathbb{C}_0) = \mathcal{O}_A,$$

where \mathbb{C}_0 is the skyscraper sheaf of length 1 at $0 \in P^\vee$. Now [Sch22, Proposition 1.6] says that if a sheaf \mathcal{F} is GV on A and $\iota_* \mathcal{F}$ is GV on P then

$$\text{SFM}_P(\iota_* \mathcal{F}) = (\iota^\vee)^* \text{SFM}_A(\mathcal{F}).$$

By (95) this yields $\text{SFM}_P(\mathcal{O}_A) = \mathcal{O}_{Q^\vee}$. \square

REMARK 3.3. The above arguments similarly work for the case of $G = G^L = \text{GL}(n, \mathbb{C})$. Here the Hitchin system restricted to $\mathcal{B}_{\text{GL}(n, \mathbb{C})}^{\text{reg}}$ is a torsor over the abelian scheme $\text{Jac}(\Sigma) \rightarrow \mathcal{B}_{\text{GL}(n, \mathbb{C})}^{\text{reg}}$, which is self-dual due to the principal polarization of the Jacobians. Hence, again given a visible Lagrangian with connected fibers together with a section we obtain a holomorphic symplectic submanifold \mathcal{I} by the arguments of the proof of Theorem 3.1. Furthermore, the arguments in the proof of Theorem 3.2 work for the structure sheaf of any abelian subvariety of an abelian variety.

3.3. Visible Lagrangians in Hitchin systems. In this section, we will give the first examples of visible Lagrangians in Hitchin systems. These examples are independent of the choice of the Riemann surface. The first one is associated to the embedding of the complex Lie groups $\mathrm{SL}(n, \mathbb{C}) \subset \mathrm{GL}(n, \mathbb{C})$. The second example appeared in the work of [FGOP21] and is a subvariety of the singular fibers of the Hitchin map. The remainder of the paper will deal with visible Lagrangians defined on special Riemann surfaces.

Visible Lagrangian associated to the subgroup $\mathrm{SL}(n, \mathbb{C}) \subset \mathrm{GL}(n, \mathbb{C})$:

Consider the embedding $\mathcal{B}_{\mathrm{SL}(n, \mathbb{C})} \subset \mathcal{B}_{\mathrm{GL}(n, \mathbb{C})}$. In this section, we will define a visible Lagrangian over $\mathcal{B}' = \mathcal{B}_{\mathrm{SL}(n, \mathbb{C})}$.

THEOREM 3.4. *Let $\mathcal{B}' = \mathcal{B}_{\mathrm{SL}(n, \mathbb{C})}(X) \subset \mathcal{B}_{\mathrm{GL}(n, \mathbb{C})}$ with $\mathcal{B}'^{\mathrm{reg}} = \mathcal{B}_{\mathrm{SL}(n, \mathbb{C})}^{\mathrm{reg}}$. Then we can act by the trivial torsor $\mathrm{Jac}(X) \times \mathcal{B}'$ on the Hitchin section by tensoring the underlying bundle. The orbit defines a visible Lagrangian $\mathcal{L} \rightarrow \mathcal{B}'$. The fiber-wise Fourier-Mukai transform of $\mathcal{O}_{\mathcal{L}}$ over $\mathcal{B}'^{\mathrm{reg}}$ is supported on the moduli space of $\mathrm{SL}(n, \mathbb{C})$ -Higgs bundles $\mathcal{M}_{\mathrm{SL}(n, \mathbb{C})} \subset \mathcal{M}_{\mathrm{GL}(n, \mathbb{C})}$.*

PROOF. Tensoring with a line bundle preserves stability. Hence \mathcal{L} is well-defined. The Hitchin fiber over $\underline{a} \in \mathcal{B}'^{\mathrm{reg}}$ reflects the splitting of the differentials on Σ by the isomorphism

$$\mathrm{Jac}(\Sigma_{\underline{a}}) = \pi^* \mathrm{Jac}(X) \times \mathrm{Prym}(\Sigma_{\underline{a}}) / \pi^* \mathrm{Jac}(X)[n],$$

where $\pi^* \mathrm{Jac}(X)[n]$ acts diagonally (see Mumford [Mum74]). In particular, we have an exact sequence of abelian schemes over $\mathcal{B}'^{\mathrm{reg}}$

$$(96) \quad 0 \rightarrow \mathrm{Jac}(X) \times \mathcal{B}'^{\mathrm{reg}} \xrightarrow{L} \mathrm{Jac}(\Sigma / \mathcal{B}'^{\mathrm{reg}}) \rightarrow \mathrm{Prym}(\Sigma / \mathcal{B}'^{\mathrm{reg}}) / \pi^* \mathrm{Jac}(X)[n] \rightarrow 0.$$

As explained in Section 3.1 the holomorphic symplectic form on $\mathrm{Hit}^{-1}(\mathcal{B}_{\mathrm{GL}(n, \mathbb{C})}^{\mathrm{reg}}) \subset \mathcal{M}_{\mathrm{GL}(n, \mathbb{C})}$ restricts to the Serre pairing between the tangent space of the base identified with $H^0(\Sigma_{\underline{a}}, K_{\Sigma_{\underline{a}}})$ by Proposition 3.1 and the tangent space to the fibers $H^1(\Sigma_{\underline{a}}, \mathcal{O}_{\Sigma_{\underline{a}}})$. With respect to this pairing we have

$$H^1(X, \mathcal{O}_X)^\perp = H^0(\Sigma_{\underline{a}}, K_{\Sigma_{\underline{a}}})^-$$

The subtorsor $\mathcal{L} \subset \mathrm{Hit}^{-1}(\mathcal{B}'^{\mathrm{reg}})$ that is defined by the action of $\mathrm{Jac}(X) \times \mathcal{B}'^{\mathrm{reg}}$ on the Hitchin section has vertical tangent bundle $H^1(X, \mathcal{O}_X)$ and the tangent space to the $\mathrm{SL}(n, \mathbb{C})$ Hitchin base are identified with $H^0(\Sigma_{\underline{a}}, K_{\Sigma_{\underline{a}}})^-$. Therefore, \mathcal{L} is a visible Lagrangian.

By Theorem 3.2 the fiber-wise Fourier-Mukai transform of the structure sheaf of \mathcal{L} over $\mathcal{B}'^{\mathrm{reg}}$ is supported on torsor over the abelian scheme that is dual to quotient in the exact sequence (96). That is the abelian scheme $\mathrm{Prym}(\Sigma / \mathcal{B}_{\mathrm{SL}(n, \mathbb{C})}^{\mathrm{reg}})$. The closure of this locus is the moduli space of $\mathrm{SL}(n, \mathbb{C})$ -Higgs bundles. \square

We expect that every embedding of complex reductive Lie groups $G_1 \subset G_2$, such that center of G_1 is mapped to the center of G_2 defines a visible Lagrangian $\mathcal{L} \rightarrow \mathcal{B}_{G_1} \subset \mathcal{B}_{G_2}$ Fourier-Mukai dual to image of the induced morphism of moduli spaces $\mathcal{M}_{G_1} \rightarrow \mathcal{M}_{G_2}$.

By definition, the visible Lagrangian of Theorem 3.4 has connected fibers. We provide an example with disconnected fibers by acting on a multi-section of the $\mathrm{SL}(2, \mathbb{C})$ -Hitchin map instead of a section. Consider the moduli space of $\mathrm{SL}(2, \mathbb{R})$ -Higgs bundles $\mathcal{M}_{\mathrm{SL}(2, \mathbb{R})}$. The Hitchin map restricts to a 2^{6g-6} -covering

$$\mathrm{Hit} | : \mathcal{M}_{\mathrm{SL}(2, \mathbb{R})} \cap \mathrm{Hit}^{-1} \mathcal{B}_{\mathrm{SL}(2, \mathbb{C})}^{\mathrm{reg}} \rightarrow \mathcal{B}_{\mathrm{SL}(2, \mathbb{C})}^{\mathrm{reg}}.$$

THEOREM 3.5. *Let $\mathcal{B}' = \mathcal{B}_{\mathrm{SL}(2, \mathbb{C})} \subset \mathcal{B}_{\mathrm{GL}(2, \mathbb{C})}$ with $\mathcal{B}'^{\mathrm{reg}} = \mathcal{B}_{\mathrm{SL}(2, \mathbb{C})}^{\mathrm{reg}}$. We can act by the trivial torsor $\mathrm{Jac}(X) \times \mathcal{B}'$ on $\mathcal{M}_{\mathrm{SL}(2, \mathbb{R})} \rightarrow \mathcal{B}'$ by tensoring. The orbit defines a visible Lagrangian $\mathcal{L} \rightarrow \mathcal{B}'$. For $\underline{a} \in \mathcal{B}'^{\mathrm{reg}}$ the fiber-wise Fourier-Mukai transform of $\mathcal{O}_{\mathcal{L}_{\underline{a}}}$ is supported on $\mathcal{M}_{\mathrm{SL}(2, \mathbb{C})} \subset \mathcal{M}_{\mathrm{GL}(2, \mathbb{C})}$ and given by a flat vector bundle of rank $6g - 6$ on $\mathrm{Prym}(\Sigma_{\underline{a}})$.*

PROOF. Let $U \subset \mathcal{B}_{\mathrm{SL}(2, \mathbb{C})}^{\mathrm{reg}}$ open, such that there exist sections

$$s_1 = s_H, s_2, \dots, s_{6g-6} : U \rightarrow \mathcal{M}_{\mathrm{SL}(2, \mathbb{C})},$$

such that $\mathrm{Hit}^{-1}(U) \cap \mathcal{M}_{\mathrm{SL}(2, \mathbb{R})} = \bigsqcup_{i=1}^{6g-6} s_i(U)$. Then the orbit of each $s_i(U)$ under tensoring with line bundles in $\mathrm{Jac}(X)$ is Lagrangian by the previous proof. Hence, \mathcal{L} defines a visible Lagrangians.

We identify the Hitchin fibers with the abelian scheme $\text{Jac}(\Sigma)$ using the Hitchin section s_1 . Fix $\underline{a} \in U$ and define $l_i = (s_1 - s_i)(\underline{a}) \in \text{Jac}(\Sigma_{\underline{a}})$. Recall the exact sequence of abelian schemes (96). We have

$$\mathcal{O}_{\mathcal{L}_{\underline{a}}} = \bigoplus_{i=1}^{6g-6} t_{l_i}^* \iota_* \mathcal{O}_{\text{Jac}(X)},$$

where t_{l_i} denotes the translation by l_i on $\text{Jac}(\Sigma_{\underline{a}})$. Let $J = \text{Jac}(\Sigma_{\underline{a}})$ and $\mathcal{P} \rightarrow J \times J$ the Poincaré bundle. Using the relation between tensor product and translations under Fourier-Mukai transform (see [Sch22]) we obtain

$$\text{SFM}_J(t_{l_i}^* \iota_* \mathcal{O}_{\text{Jac}(X)}) = \mathcal{P}_{l_i} \otimes \text{SFM}_J(\iota_* \mathcal{O}_{\text{Jac}(X)}) = \mathcal{P}_{l_i} \otimes (\iota^\vee)^* \mathcal{C}_0 = \mathcal{P}_{l_i} \otimes \mathcal{O}_{\text{Prym}(\Sigma_{\underline{a}})}.$$

The right hand side defines a flat line bundle on $\text{Prym}(\Sigma_{\underline{a}})$. The Fourier-Mukai of the structure sheaf $\mathcal{O}_{\mathcal{L}_{\underline{a}}}$ is the direct sum of all these flat line bundles. \square

Visible Lagrangians over the singular locus of the Hitchin base:

Other examples of visible Lagrangians were considered in [FGOP21]. Here $\mathcal{B}^{\text{reg}} \subset \mathcal{B}_{\text{GL}(n, \mathbb{C})}$ is the locus of spectral curves with the maximal number of $n(n-1)(g-1)$ nodes as their only singularities and $\mathcal{B}' = \overline{\mathcal{B}^{\text{reg}}}$. In particular, these Lagrangians are completely contained in the singular locus of the Hitchin map and hence Theorem 2.4 and 3.2 do not apply. The compactified Jacobians over \mathcal{B}^{reg} contain subvarieties isomorphic to $(\mathbb{P}^1)^{n(n-1)(g-1)}$, which can be interpreted as parameters for Hecke modifications of the Higgs bundles at the node by [Hor22]. Applying these Hecke modifications to the Hitchin section yields a visible Lagrangian over \mathcal{B}' . Interestingly, in this work the authors considered Arinkin’s Fourier-Mukai transform for compactified Jacobians and found that the support of the fiber-wise Fourier-Mukai transform is supported on a hyperholomorphic subvariety - the so-called Narasimhan-Ramanan BBB-brane. In the subsequent work [FHHO23], together with the first author this construction will be generalized to visible Lagrangians $\mathcal{L} \rightarrow \mathcal{B}'$ over the closure of the locus of spectral curves with any number of nodes as their only singularities.

Lagrangians that are not visible: The upward flow of a very stable Higgs bundle considered in [HH22]; [PP19] defines a complex Lagrangian that is supported over the whole Hitchin base and hence is not visible.

4. Parallelogram-tiled surfaces and pillowcase covers

In this section we will briefly review the interpretation of abelian and quadratic differentials in terms of flat geometry. Then we will discuss certain types of these flat geometries on a Riemann surface that will play a special role in the following section.

An abelian differential $\lambda \in H^0(\Sigma, K_\Sigma)$ on a Riemann surface Σ determines a singular flat metric, such that all transition functions are translation. A coordinate of the flat metric at $y \in \Sigma \setminus Z(\lambda)$ is a holomorphic coordinate z at y , such that $\lambda = dz$. In this way, one obtains a flat metric on $\Sigma \setminus Z(\lambda)$, such that coordinate transitions are translation. It extends to a singular flat metric on Σ by cone points of cone angle $(k+1)\pi$ at a zero of λ of order k . This is a so-called translation surface.

Similarly, we can associate a singular flat metric to a quadratic differential (X, q) , where X is a Riemann surfaces and $q \in H^0(X, K_X^2)$. A flat coordinate at $x \in X \setminus Z(q)$ is a holomorphic coordinate z at x , such that $q = dz^{\otimes 2}$. In this case, the coordinate functions are compositions of translations and reflections. It extends to a singular flat metric on X by cone points of cone angle $(k+2)\pi$ at a zero of q of order k . This is a so-called half-translation surface. When q has simple zeros only, the spectral curve defined in section 3 is referred to as the canonical cover of (X, q) from this point of view. It is the universal cover X , such that the pullback of q has a square-root. (Here we consider λ as a section of K_Σ instead of $\pi^* K_X$ as in section 3. We have $\pi^* K_X = K_\Sigma(-R)$. Hence, if q has simple zeros, then the abelian differential λ has double zeros at all branch points.)

We say that a quadratic differential (X, q) is of type $\mu(q) = (m_1, \dots, m_n)$ if the orders of the zeros of the differential are m_1, \dots, m_n . We will use exponential notation if multiple m_i agree, i.e. we write $(1^{4g(X)-4})$ for $(1, \dots, 1)$.

In the following, particularly symmetric (half-)translation surfaces play a special role: Parallelogram-tiled surfaces and pillowcase covers. We obtain coordinates on the moduli




FIGURE 2. Pillowcase with canonical cover: Opposite sides identified, when not indicated otherwise. The involution on the cover acts as central symmetry in the two-torsion points.

space of abelian differentials by recording periods. The periods of (Σ, λ) in $H^1(\Sigma, Z(\lambda), \mathbb{C})$ are given by

$$H_1(\Sigma, Z(\lambda), \mathbb{C}) \rightarrow \mathbb{C}, \quad c \mapsto \int_c \lambda$$

and local coordinates, the so called *period coordinates*, are given by the image of a basis of relative homology under this map. The coordinate changes of period coordinates are induced by diffeomorphisms of $(\Sigma, Z(\lambda))$ and hence preserve the lattice

$$H^1(\Sigma, Z(\lambda), \mathbb{Z} \oplus i\mathbb{Z}) \subset H^1(\Sigma, Z(\lambda), \mathbb{C}).$$

These integral points correspond to square-tiled surfaces: One obtains a cover of an elliptic curve by

$$p : \Sigma \rightarrow \mathbb{C}/\mathbb{Z} \oplus i\mathbb{Z}, \quad y \mapsto \int_{y_0}^y \lambda$$

for a choice of base point $y_0 \in \Sigma$. This cover is branched over one point $0 \in E$ with ramification points $z_i \in Z(\lambda)$ and ramification profile $(\text{ord}_{z_1}(\lambda), \dots, \text{ord}_{z_n}(\lambda))$. In particular, this is a cover of flat surfaces, i.e. $\lambda = p^*\omega$ for some abelian differential ω on E . By scaling the area of the flat torus (Σ, λ) we obtain a dense subset of translation surfaces, that are square-tiled in the above sense. We use a slight generalization.

DEFINITION 4.1. A translation surface (Σ, λ) is parallelogram-tiled if and only if there exists a branched cover $p : \Sigma \rightarrow E$ branched over one point, such that $\lambda = p^*\omega$ for some abelian differential ω on E .

Given a parallelogram-tiled surface we can act by $\text{GL}^+(2, \mathbb{R})$ on the representing polygon and obtain a family of parallelogram-tiled surfaces over the j -line of elliptic curves. For $j = 1728$ we recover a square-tiled surface. The analogue of parallelogram-tiled surfaces for quadratic differentials are pillowcase covers. A pillowcase is a half-translation surface (\mathbb{P}^1, η) , where η has four simple poles $0, 1, \infty, x$, see Figure 2.

DEFINITION 4.2. A half-translation surface (X, q) is called pillowcase cover if there exists a cover $\check{p} : X \rightarrow \mathbb{P}^1$ branched over $(0, 1, \infty, x)$, such that $q = \check{p}^*\eta$ for a quadratic differential η on \mathbb{P}^1 with simple poles at $(0, 1, \infty, x)$.

The canonical cover of (\mathbb{P}^1, η) is the elliptic differential (E, ω) , where $E = \mathbb{C}/(\mathbb{Z} + \tau\mathbb{Z})$ with $\lambda(\tau) = x$ and the involution is given by multiplication with -1 . Here λ is the modular lambda function. For a given x , the number τ can be computed explicitly. Let K denote the complete elliptic integral of the first kind

$$(97) \quad K(k) = \int_0^{\frac{\pi}{2}} \frac{d\theta}{\sqrt{1 - k^2 \sin^2 \theta}}.$$

Then $t(x) = i \frac{K(\sqrt{1-x})}{K(\sqrt{x})}$ is a section of λ . There is the following well-known relation between parallelogram-tiled surfaces and pillowcase covers.

LEMMA 4.3. *Let (X, q) be a half-translation surface. The canonical cover (Σ, λ) is a parallelogram-tiled surface if and only if the quadratic differential is a pillowcase cover.*

PROOF. Starting from a pillowcase cover as above we have the following diagram

$$\begin{array}{ccc}
 & (\Sigma, \lambda) & \\
 \swarrow \pi_{\text{mod } J} & & \searrow p \\
 (X, q) & & (E, \omega) \\
 \searrow \check{p} & & \swarrow 2:1 \\
 & (\mathbb{P}^1, \eta) &
 \end{array}$$

and want to show that there exist the dashed arrow $p : \Sigma \rightarrow E$, such that the diagram commutes. Away from the singularities of the flat structure associated with (X, q) a point on Σ corresponds to a choice of a (local) square of q . This corresponds to a choice of a local square root of η under the covering map to \mathbb{P}^1 . Hence there is a induced map $p : \Sigma \setminus Z(\lambda) \rightarrow E \setminus Z(\omega)$ such that $p^*\omega = \lambda$. The map p uniquely extends to a map of Riemann surfaces $\Sigma \rightarrow E$. A zero $x \in X$ of q of order $k \geq 0$ corresponds to a $(k + 2) : 1$ -branch point of \check{p} . Hence, the map can be extended to $p : \Sigma \rightarrow E$ by gluing in a $k : 1$ -ramification point, if k is odd and a $\frac{k}{2} : 1$ -ramification point, if k is even, at $y \in \pi^{-1}(x)$. By construction the relation of the differentials persists under this extension. Taking the quotient of E by two-torsion points we obtain a map $(\Sigma, \lambda) \rightarrow E/E[2]$, branched over one point such that λ is the pullback of a differential on E . Hence, it is a parallelogram-tiled surface.

For the converse, we start with the configuration

$$\begin{array}{ccc}
 & (\Sigma, \lambda) & \\
 \swarrow \pi_{\text{mod } J} & & \searrow p \\
 (X, q) & & (E, \omega)
 \end{array}$$

The differential λ is anti-symmetric with respect to an involution J . In particular, J sends singularities of the flat structure to singularities of the same type, saddle connections to saddle connections and hence squares to squares. Hence it descends to the elliptic curve (E, ω) to the reflection in the two torsion points of E . The quotient is the pillowcase surface as illustrated in Figure 2. In particular, there is an induced map $\check{p} : \Sigma/J = X \rightarrow \mathbb{P}^1$, such that $q = \check{p}^*\eta$. \square

REMARK 4.4. The proof of the previous lemma shows that one can always assume the pillowcase cover to be non-simply branched over only one point on \mathbb{P}^1 . This can be achieved by taking the quotient by the order four automorphism of \mathbb{P}^1 that permutes the four marked points preserving the cross-ratio.

In the following, a special kind of pillowcase cover will be important.

DEFINITION 4.5. We call a pillowcase cover (X, q) *uniform* if every fiber of $\check{p} : X \rightarrow \mathbb{P}^1$ over $y \in D$ consists of ramification points of the same ramification index i .

5. Multifold pillowcase covers

Motivated by the connection to visible Lagrangians in Theorem 6.1, we are interested in Riemann surfaces X which admit multiple quadratic differentials $q_1, \dots, q_n \in H^0(X, K^2)$, such that

- the vanishing loci $Z(q_1), \dots, Z(q_n)$ are pairwise different,
- the half-translation surfaces (X, q_i) are pillowcase covers.

DEFINITION 5.1. We call a Riemann surface X with quadratic differentials q_1, \dots, q_n as above a n -fold *pillowcase cover*. We call a n -fold pillowcase cover a *multifold pillowcase cover* if $n \geq 2$.

Apart from being uniform, we want the pillowcase covers to have simple zeros only.

Two quadratic differentials q_1, q_2 on X are called isomorphic if there exists an automorphism φ of X such that $\varphi^*q_1 = q_2$. We remark that from the point of view of flat geometry isomorphic differentials are usually not distinguished. However two isomorphic differentials might still have different vanishing loci and therefore correspond to different

points in the $\mathrm{SL}(2, \mathbb{C})$ -Hitchin base. Hence, we will treat these differentials as different from each other in the following. In this section we will prove

THEOREM 5.2. *For infinitely many genera g there exist multifold uniform pillowcase covers with simple zeros only.*

PROOF. Assume that we know a single multifold uniform pillowcase cover X with simple zeros only of some genus $g \geq 2$. Let q_1, \dots, q_n be the corresponding quadratic differentials. Then we can obtain examples in infinitely many genera by taking unramified coverings $f : \widehat{X} \rightarrow X$ in different degrees and the differentials $\widehat{q}_i = f^* q_i$.

We will provide explicit examples in the following. \square

The following examples have been found by a computer search using [GAP22]. There are many more examples to be found, but we will restrict our discussion to three examples in low genus. While the claimed properties of the examples can be checked by hand, it is much more convenient to use a computer algebra system.

EXAMPLE 5.3 (Genus 2). Consider the group $\mathrm{GL}(2, F_3)$ of order 48 and let $f : X \rightarrow \mathbb{P}^1$ be the $\mathrm{GL}(2, F_3)$ -cover branched above three points with monodromy datum

$$\begin{pmatrix} 0 & 1 \\ 1 & 2 \end{pmatrix}, \begin{pmatrix} 2 & 0 \\ 0 & 1 \end{pmatrix}, \begin{pmatrix} 2 & 2 \\ 1 & 0 \end{pmatrix}.$$

The orders of the matrices are 8, 2 and 6, and the genus of X is 2. The group $\mathrm{GL}(2, F_3)$ has 12 subgroups of order 6 which we denote by H_1, \dots, H_{12} . For each such subgroup the quotient X/H_i is of genus 0, and the quotient map $g_i : X \rightarrow X/H_i$ is branched above four points with ramification orders 2 and 3. Let η_i denote the quadratic differential on X/H_i of type (-1^4) whose simple poles are supported at the branch points of g_i . The pullback $q_i := g_i^* \eta_i$ is a uniform pillowcase cover with simple zeros only on X . We claim that the vanishing loci of the differentials q_i are pairwise different. This is a very explicit but lengthy computation which is left to the reader. In particular, X is a 12-fold uniform pillowcase cover with simple zeros only.

EXAMPLE 5.4 (Genus 3). Consider the group $\mathrm{SL}(3, F_2)$ of order 168 and let $f : X \rightarrow \mathbb{P}^1$ be the $\mathrm{SL}(3, F_2)$ -cover branched above three points with monodromy datum

$$\begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 1 & 0 & 1 \end{pmatrix}, \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{pmatrix}, \begin{pmatrix} 1 & 0 & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{pmatrix}.$$

The orders of the matrices are 2, 3 and 7. The Riemann surface X is the Klein quartic and has genus 3. The group $\mathrm{SL}(3, F_2)$ has 14 subgroups of order 24, and as in the previous example each of those subgroups gives rise to a uniform pillowcase covers with simple zeros only. In particular, X is a 14-fold uniform pillowcase cover with simple zeros only.

EXAMPLE 5.5 (Non-isomorphic differentials). Let $G := A_4 \times \mathbb{Z}/3\mathbb{Z}$. We choose generators $\langle a := (123), b := (12)(34) \rangle = A_4$ and $\langle c \rangle = \mathbb{Z}/3\mathbb{Z}$. Consider the subgroups $H_1 := \langle a, c \rangle \cong (\mathbb{Z}/3\mathbb{Z})^2$ and $H_2 := \langle a, b \rangle \cong A_4$. Let $f : X \rightarrow \mathbb{P}^1$ be the G -cover branched above three points with monodromy datum (bc, a^2c^2, ab) . The genus of X is 4.


Consider the two intermediate covers $A_1 := X/H_1$ and $A_2 := X/H_2$, both of genus 0. We define the two differentials q_1 and q_2 on X as in the previous examples.

In this example it is relatively easy to see that the differentials q_1 and q_2 are non-isomorphic. For this it is convenient to consider the respective canonical covers (Σ_i, λ_i) of (X, q_i) .

PROPOSITION 5.1. *The canonical covers Σ_1 and Σ_2 are non-isomorphic. In particular, $(X, q_1) \not\cong (X, q_2)$, i.e. there doesn't exist an automorphism $\varphi : X \rightarrow X$, such that $\varphi^* q_1 = q_2$.*

PROOF. The curves Σ_1 and Σ_2 are covers of \mathbb{P}^1 via $\Sigma_i \rightarrow X \xrightarrow{f} \mathbb{P}^1$, branched over three points. Both covers $\Sigma_i \rightarrow \mathbb{P}^1$ have a monodromy representation with elements in the symmetric group $S_{2|G|}$. The covers Σ_1 and Σ_2 are isomorphic if and only if the elements of the monodromy representation are conjugated in $S_{2|G|}$, which is easily checked not to be the case. \square

In the rest of this section we will produce explicit flat pictures of the two quadratic differentials (X, q_1) and (X, q_2) of Example 5.5. Recall that Σ_i is a square tiled surface by Lemma 4.3. We thus have the diagram



First we determine the tori E_i , which also determine the pillowcases A_i . We start with the torus E_2 . The map $A_2 \rightarrow \mathbb{P}^1$ is cyclic of degree 3 and totally ramified over two points, and unramified otherwise. One of the branch points of $E_2 \rightarrow A_2$ agrees with a ramification point of $A_2 \rightarrow \mathbb{P}^1$, while the three other branch points of $E_2 \rightarrow A_2$ lie in one fiber of $A_2 \rightarrow \mathbb{P}^1$. Hence we may assume that the branch points of $E_2 \rightarrow A_2$ are $0, 1, \zeta_3$ and ζ_3^2 . Those four points have the cross-ratio $D(0, 1; \zeta_3, \zeta_3^2) = \zeta_6^5$, and $\lambda(\zeta_3) = \zeta_6^5$, where λ is the modular lambda function. Hence up to isomorphism

$$E_2 \cong \mathbb{C}/(\mathbb{Z} \oplus \zeta_3\mathbb{Z}) \cong \mathbb{C}/(\mathbb{Z} \oplus \zeta_6\mathbb{Z}).$$

The cross-ratio of the four branch points $0, 1, x, \infty \in \mathbb{P}^1$ of $E_1 \rightarrow A_1$ is again uniquely determined by the ramification profile of the maps $E_1 \rightarrow A_1 \rightarrow \mathbb{P}^1$ and is given by

$$D(0, \infty; 1, x) = 15\sqrt{3} - 26.$$

One numerically computes $\tau := t(15\sqrt{3} - 26) \approx 1 + 2.143182698915i$, where t is the function from (97), and for this τ we have

$$E_1 \cong \mathbb{C}/(\mathbb{Z} \oplus \tau\mathbb{Z}).$$

To obtain pictures of the pillowcase covers $X \rightarrow A_i$, we need to determine how to glue the copies of the pillowcase A_i . We describe how to obtain this information for a general pillowcase cover which is a G -cover. The idea is to compare two G -actions on the $|G|$ -many copies of A_i .

Given a G -cover $X \rightarrow \mathbb{P}^1$ of degree $d = |G|$ with monodromy datum (g_1, \dots, g_d) , the bijective map

$$m_g : G \rightarrow G, \quad h \mapsto gh$$

induces a map $\rho : G \rightarrow S_d$ when we identify $\sigma : G \rightarrow \{1, \dots, d\}$. If $p \in \mathbb{P}^1$ is a point (which is not a branch point), and we identify the fiber of $X \rightarrow \mathbb{P}^1$ above p with $\{1, \dots, d\}$ via σ , then the lift of a simple loop with basepoint p around the i -th ramification point starting in $q \subseteq \{1, \dots, d\}$ will end in a point $\rho(g_i)(q)$.




FIGURE 3. A pillowcase

On the other hand X consists of d copies of the pillowcase \mathbb{P}^1 as depicted in Figure 3. We can again label those copies with $\{1, \dots, d\}$. After choosing an orientation on the vertical and horizontal cylinder, X is uniquely determined by the permutations $h_1, h_2, v \in S_d$, which indicate which copy of the pillowcase we reach when we leave a given copy in the direction indicated in the figure.

We label the four branch points by a, b, c, d as in Figure 3. Lifting a small clockwise cycle around point a will act on the copies as $h_2 \circ v$, around b as $h_1 \circ v^{-1}$, around c as h_2^{-1} and around d as h_1^{-1} .

Hence h_1, h_2, v are given (up to the choice of σ and hence ρ) by

$$h_1 = \rho(g_0)^{-1}, \quad h_2 = \rho(g_1)^{-1}, \quad v = h_2^{-1} \circ \rho(g_3) = \rho(g_4)^{-1} \circ h_1.$$

Now we come back to the cases we are interested in. For $X \rightarrow A_1$ we can choose σ such that

$$\begin{aligned} h_1 = h_2 &= (1\ 2\ 3)(4\ 5\ 6)(7\ 8\ 9) \\ v &= (1\ 5\ 9)(2\ 6\ 7)(3\ 4\ 8), \end{aligned}$$

for $X \rightarrow A_2$ we can choose σ such that

$$\begin{aligned} h_1 &= (1\ 4)(2\ 3)(5\ 7)(6\ 8)(9\ 10)(11\ 12) \\ h_2 = v &= (1\ 7\ 12)(2\ 8\ 11)(3\ 5\ 10)(4\ 6\ 9). \end{aligned}$$

Combined with the information about the tori E_i this gives rise to the pictures in Figure 4, where the horizontal edges are glued by half-translation as indicated by the labeling.

6. Visible Lagrangians over lines in the Hitchin base

In this section, we study visible Lagrangians over a line $\mathcal{B}' = \mathbb{C}\underline{a} \subset \mathcal{B}_{\mathrm{SL}(n, \mathbb{C})}$. In the first part, we give an existence criterion using Proposition 3.2. In second part, for $G = \mathrm{SL}(2, \mathbb{C})$ and $\underline{a} = q$ a quadratic differential, we study the holomorphic symplectic subvariety \mathcal{I}_q , which is the proposed mirror dual by Theorem 3.2. If (X, q) is a uniform pillowcase cover we will show that \mathcal{I}_q is a hyperholomorphic subvariety birational to Hausel's toy model.

THEOREM 6.1. *Let $\underline{a} = (0, \dots, 0, a_n) \in \mathcal{B}_{\mathrm{SL}(n, \mathbb{C})}^{\mathrm{reg}}$. Then there exists a visible Lagrangian over $\mathbb{C}\underline{a}$, if and only if the spectral curve with its abelian differential (Σ, λ) is parallelogram-tiled.*

PROOF. First assume there exists a visible Lagrangian over $\mathbb{C}\underline{a}$. Let (Σ, λ) be the spectral curve to \underline{a} . Then by Proposition 3.2 there is an exact sequence of abelian varieties

$$0 \rightarrow A \rightarrow \mathrm{Prym}(\Sigma_{\underline{a}}) \xrightarrow{\psi} E \rightarrow 0.$$

Here A is of codimension 1 and E is an elliptic curve. Under the assumption on \underline{a} the spectral curve is \mathbb{Z}_n -Galois and $s_B = \lambda^{n-1}$ (up to a constant). Hence, the map t of Proposition 3.1 becomes

$$t : \bigoplus_{i=2}^n H^0(X, K_X^i) \rightarrow H^0(\Sigma, K_\Sigma), \quad \sum_{i=2}^n \alpha_i X_i \mapsto \sum_{i=2}^n \alpha_i \frac{\pi^* X_i}{\lambda^{i-1}}.$$


In particular, the tangent vector $X_n = a_n$ is mapped to $\frac{\pi^* a_n}{\lambda^i} = \lambda$. Hence, the differential of the map ψ can explicitly be written as

$$(98) \quad d\psi : H^1(\Sigma, \mathcal{O}_\Sigma)^- \rightarrow T_0 E, \quad \alpha \mapsto c \int_\Sigma \alpha \wedge \lambda.$$


Therefore, ψ is given by $D = \sum a_i y_i \mapsto \sum \int_{y_0}^{y_i} \lambda$ up to the choice of a point $y_0 \in \Sigma$. Denote by σ a generator of the \mathbb{Z}_n action on Σ . We want show that the composition $\psi \circ \mathrm{AP}$ with the Abel-Prym map

$$\mathrm{AP} : \Sigma \rightarrow \mathrm{Prym}(\Sigma), \quad y \mapsto \mathcal{O}(y - \sigma y)$$

is a branched covering $p : \Sigma \rightarrow E$. Clearly, AP identifies all ramification points of $\pi : \Sigma \rightarrow X$. If Σ is not hyperelliptic it is easy to see, that these are the only points on Σ that are identified. The spectral curve Σ is never hyperelliptic by [Bar16, Lemma 4.1]. By definition $\sigma^* \lambda = \xi_n \lambda$ for some primitive n -th root of unity ξ_n . This implies that the differential of the composition $\psi \circ \mathrm{AP}$ is $(1 - \xi)\lambda$. Hence the differential of the composition is injective away from the ramification points of $\Sigma \rightarrow X$. In particular, $p = \psi \circ \mathrm{AP} : \Sigma \rightarrow E$ is a proper holomorphic map ramified at $Z(\lambda)$ over $0 \in E$. Furthermore, λ considered as abelian differential has an order n zero at each branch point and hence the points in $Z(\lambda)$



(A) The pillowcase cover $X \rightarrow A_1$



(B) The pillowcase cover $X \rightarrow A_2$

FIGURE 4. Two non-isomorphic pillowcase covers on the same curve X

are $n + 1 : 1$ -ramifications. In particular, the pullback of an abelian differential ω on E has the same divisor as λ and we can find a specific ω with $p^*\omega = \lambda$. Therefore, (Σ, λ) is parallelogram-tiled.

For the converse, let $\underline{a} = (0, \dots, 0, a_n) \in \mathcal{B}_{\mathrm{SL}(n, \mathbb{C})}^{\mathrm{reg}}$, such that (Σ, λ) is parallelogram-tiled, i.e. there is a covering $p : \Sigma \rightarrow E$, such that $\lambda = p^*\omega$ for some abelian differential ω on E . The covering p induces a Norm map

$$\mathrm{Nm}_E : \mathrm{Jac}(\Sigma) \rightarrow E, \quad D = \sum a_i y_i \mapsto \sum a_i \int_{y_0}^{y_i} \lambda.$$

The restriction to $\mathrm{Prym}(\Sigma)$ defines the desired map

$$\psi = \mathrm{Nm}_E|_{\mathrm{Prym}} : \mathrm{Prym}(\Sigma) \rightarrow E.$$

Let $D = (n - 1)y - \sigma y - \dots - \sigma^{n-1}y$. Then $D \in \mathrm{Prym}(\Sigma)$ and $\mathrm{Nm}_E(D) = n\mathrm{Nm}_E(y)$. Hence, ψ surjects onto E being a divisible group. Its differential is the map (98). As Σ does not change, when multiplying q with a scalar, we can define an abelian subscheme $\ker(\psi) \subset \mathrm{Prym}(\Sigma/\mathbb{C}^\times q)$ that satisfies the criterion of Proposition 3.2. To obtain a visible Lagrangian over $\mathbb{C}q$ we can act with the abelian scheme $\ker(\psi)$ on any section of $\mathrm{Hit}|_{\mathbb{C}q}$. To obtain a concrete example we may choose the Hitchin section. \square

COROLLARY 6.2. *Let $q \in \mathcal{B}_{\mathrm{SL}(2, \mathbb{C})}^{\mathrm{reg}}$ be a quadratic differential with simple zeros only. Then there exists a visible Lagrangian over $\mathbb{C}q$ if and only if (X, q) is a pillowcase cover.*

PROOF. This is immediate from Theorem 6.1 and Lemma 4.3. \square

6.1. The Fourier-Mukai dual. Now, we consider the Fourier-Mukai dual of the visible Lagrangian defined above in the case of $G = \mathrm{SL}(2, \mathbb{C})$. Let (X, q) be a pillowcase cover and $\mathcal{L}_q \rightarrow \mathcal{B}' = \mathbb{C}q$ the visible Lagrangian defined as the closure of the orbit of the abelian subscheme $A \subset \mathrm{Prym}(\Sigma/\mathbb{C}^\times q)$ on the Hitchin section. In particular, $\mathcal{B}'^{\mathrm{reg}} = \mathbb{C}^\times q$. The fibers of $\mathcal{L} \rightarrow \mathcal{B}'^{\mathrm{reg}}$ are of codimension 1 in the $\mathrm{SL}(2, \mathbb{C})$ -Hitchin fibers over $\mathcal{B}'^{\mathrm{reg}}$. Hence, by Theorem 3.2 the fiber-wise Fourier-Mukai transformation of the structure sheaf of \mathcal{L} is supported on an elliptic surface $I_q \rightarrow \mathbb{C}^\times q$ obtained by acting with the abelian scheme $E \subset \mathrm{Prym}(\Sigma/\mathbb{C}^\times q)^\vee$ on the Hitchin section of $\mathcal{M}_{\mathrm{PGL}(2, \mathbb{C})}$. Its closure $\mathcal{I}_q = \overline{I}_q$ is the proposed mirror dual. We have the following proposition.

PROPOSITION 6.1. *The subvariety $\mathcal{I}_q \subset \mathcal{M}_{\mathrm{PSL}(n, \mathbb{C})}$ is birational to Hausel's toy model $\mathcal{M}_{\mathrm{toy}}$ [Hau98].*

PROOF. Hausel's toy model is constructed as an elliptic surface as follows. Take $(\mathbb{P}^1, p_1, p_2, p_3, p_4)$ and consider the elliptic curve $E \rightarrow \mathbb{P}^1$, which is the canonical cover of the pillowcase with involution τ . Let further $M = \tau'$ be the involution $-1 : \mathbb{C} \rightarrow \mathbb{C}$. Consider $E \times \mathbb{C}/(\tau \times \tau')$. This orbifold has 4 \mathbb{Z}_2 -points $\hat{p}_i \times 0$. The projection to the second factor induces the map $M \rightarrow \mathbb{C}, (x, y) \mapsto y^2$ with generic fiber E . Blowing up the 4 orbifold points we obtain a smooth surface - the toy model $\mathcal{M}_{\mathrm{toy}}$. The subintegrable system \mathcal{I}_q associated to $\mathcal{B}' = \mathbb{C}q$ for a pillowcase cover (X, q) is an elliptic fibration over $\mathbb{C}^\times q \subset \mathcal{B}'$. The modulus of the elliptic curve is constant and determined by the four points p_1, \dots, p_4 on the pillowcase. As for Hausel's toy model the monodromy around $0 \in \mathbb{C}q$ is given by -1 . In fact if we scale our quadratic differential with $e^{i\phi}$ the abelian differential is multiplied by $e^{i\frac{1}{2}\phi}$. \square

For uniform pillowcase covers (X, q) we can indeed prove that $\mathcal{I}_q \subset \mathcal{M}_{\mathrm{SL}(2, \mathbb{C})}(X)$ is a hyperholomorphic subvariety. This will be achieved by defining a morphism Θ from a moduli space of semi-stable parabolic $\mathrm{SL}(2, \mathbb{C})$ -Higgs bundles to $\mathcal{M}_{\mathrm{SL}(2, \mathbb{C})}(X)$, such that the image of Θ is \mathcal{I}_q . Denote by $\mathcal{M}_{\underline{\alpha}} = \mathcal{M}_{\underline{\alpha}}(\mathbb{P}^1, D)$ the moduli space of semi-stable strongly parabolic $\mathrm{SL}(2, \mathbb{C})$ -Higgs bundle (E, Φ) on \mathbb{P}^1 with $D = 0 + 1 + \infty + y$ with parabolic weights $\underline{\alpha} = ((\alpha_{y,1}, \alpha_{y,2})_{y \in D})$. Here E is a rank 2 bundle of determinant $\det(E) = \mathcal{O}(-4)$ together with complete flag $\{0\} \subsetneq E_{y,1} \subsetneq E_{y,2} = E_y$ at each $y \in D$. By our convention flags are ascending and weights descending, i.e. $\alpha_{y,1} > \alpha_{y,2}$. The Higgs field $\Phi \in H^0(\mathbb{P}^1, \mathrm{End}_0(E) \otimes K(D))$ must preserve these flags, in the sense that $\mathrm{Res}(\Phi)(E_y) \subset E_{y,1}$. The Higgs bundle (E, Φ) is called semi-stable if and only if for each sub-Higgs bundle $(L, \psi) \subset (E, \Phi)$ we have $\mathrm{pardeg}(L, \psi) \leq \mathrm{pardeg}(E, \Phi)$. On (\mathbb{P}^1, D) this condition is automatically satisfied

for all Higgs bundles that are not nilpotent. $\mathcal{M}_\alpha(\mathbb{P}^1, D)$ carries a hyperkähler metric defined by interpreting it as the moduli space of flat logarithmic connections with certain fix monodromy at D via non-abelian Hodge theory. See [FMSW22] for more details. For generic weights, the moduli space \mathcal{M}_α is isomorphic to the smooth surface \mathcal{M}_{toy} constructed above.

Recall from Definition 4.5 that for a uniform pillowcase cover (X, q) every fiber of $\check{p} : X \rightarrow \mathbb{P}^1$ over one of the four marked points on \mathbb{P}^1 has a well-defined ramification index.

THEOREM 6.3. *Let (X, q) be a uniform pillowcase cover with simple zeros only. We define so-called compatible parabolic weights at $y \in D$ of ramification index i by $(\alpha_0 = \frac{i+1}{i+2}, \alpha_1 = \frac{1}{i+2})$. Let $\mathcal{M}_\alpha(\mathbb{P}^1, D)$ be the moduli space of strongly parabolic Higgs bundles on (\mathbb{P}^1, D) with compatible parabolic weights. Then there exists a holomorphic map*

$$\Theta : \mathcal{M}_\alpha(\mathbb{P}^1, D) \rightarrow \mathcal{M}_{\text{SL}(2, \mathbb{C})}(X),$$

such that

- i) it maps the Hitchin section of \mathcal{M}_α to the Hitchin section of $\mathcal{M}_{\text{SL}(2, \mathbb{C})}(X)$ restricted to $\mathbb{C}q$.
- ii) for all $c \in \mathbb{C}^\times$, it makes the following diagram commute

$$\begin{array}{ccc} E & \xrightarrow{p^*} & \text{Prym}(\Sigma_q) \\ \downarrow \cong & & \downarrow \cong \\ \text{Hit}^{-1}(c\eta) & \xrightarrow{\Theta} & \text{Hit}^{-1}(cq) \end{array}$$

where we used the Hitchin section to identify the corresponding Hitchin fibers with abelian varieties.

- iii) Θ can be promoted to a morphism of hermitian Higgs bundles, such that solutions to the Hitchin equation on (\mathbb{P}^1, D) are mapped to solutions to the Hitchin equation on X .

PROOF. Recall the square of coverings of Figure 5. The morphism Θ is given by a Hecke modified pullback along \check{p} . Let $(E, \Phi) \in \mathcal{M}_\alpha$. First define

$$(E', \Phi') = \left(\check{p}^*(E \otimes \mathcal{O}(3)) \otimes K_X^{-\frac{1}{2}}, \check{p}^*\Phi \right).$$

This is a meromorphic Higgs bundle on X with $\text{tr}(\Phi') = 0$. By assumption of q having simple zeros only all points in $\check{p}^{-1}(D)$ are ramification points of \check{p} , which are $2 : 1$ or $3 : 1$. Let $R_i \in \text{Div}^+(X)$ be the divisor that has weight 1 at the branch points that are $(i+1) : 1$ and $R = R_1 + R_2$. The pullback of the quasi-parabolic structure defines a quasi-parabolic structure on E' at R given by $E'_{x,1} := \check{p}^*E_{\check{p}x,1} \subset E'_x$. Now we define a Hecke modification

$$0 \rightarrow (\widehat{E}, \widehat{\Phi}) \rightarrow (E', \Phi') \rightarrow \bigoplus_{x \in \text{supp } R} E'_x/E'_{x,1} \rightarrow 0.$$

Then the map Θ is defined as

$$\Theta : (E, \Phi) \mapsto (\widehat{E}, \widehat{\Phi}).$$




FIGURE 5. Square of coverings associated to a pillowcase.

Well-definedness: First, we have to show that Θ is well-defined. We have

$$\check{p}^*\mathcal{O}_{\mathbb{P}^1}(-2) = \check{p}^*K_{\mathbb{P}^1} = K_X(-R_1 - 2R_2)$$

and $\text{div}(q) = R_2$. Hence, the determinant of E' computes to

$$\det(E') = \check{p}^* \mathcal{O}(2) \otimes K_X^{-1} = K_X^{-2}(R_1 + 2R_2) = \mathcal{O}(R_1 + R_2).$$

Therefore, the determinant of \widehat{E} is $\det(\widehat{E}) = \det(E')(-R_1 - R_2) = \mathcal{O}_X$.

To show that the Higgs field $\widehat{\Phi}$ is holomorphic we do a local computation. Locally at each $y \in D$ we can find a frame s_1, s_2 of E adapted to the parabolic structure such that

$$(99) \quad \Phi = \begin{pmatrix} \phi_0(z)dz & \phi_1(z)\frac{dz}{z} \\ \phi_2(z)dz & -\phi_0(z)dz \end{pmatrix}.$$

and $E_{y,1} = \langle s_1|_y \rangle$. Hence, at all ramification points $x \in \check{p}^{-1}D$ the pullback has the form

$$(100) \quad \check{p}^* \Phi = \begin{pmatrix} \phi_0(w^k)w^{k-1}dw & \phi_1(w^k)\frac{dw}{w} \\ \phi_2(w^k)w^{k-1}dw & -\phi_0(w^k)w^{k-1}dw \end{pmatrix},$$

where w is a coordinate centered at x , such that $\check{p} : w \mapsto z = w^k$. By assumption $k = 2, 3$. The pullback quasi-parabolic structure is given by $E'_{x,1} = \langle \check{p}^* s_1|_y \rangle$. Now, it is an easy computation to see that the Hecke modification modifies the Higgs field to

$$\widehat{\Phi} = \begin{pmatrix} \phi_0(w^k)w^{k-1}dw & \phi_1(w^k)dw \\ \phi_2(w^k)w^{k-2}dw & -\phi_0(w^k)w^{k-1}dw \end{pmatrix}.$$

Hence, indeed $(\widehat{E}, \widehat{\Phi})$ defines a $\text{SL}(2, \mathbb{C})$ -Higgs bundle. Its poly-stability will follow from the existence of a solution to Hitchin's equation at the end of the proof.

Hitchin sections: Now, we want to apply this morphism to a point in the Hitchin section. We identify the Hitchin base of $\mathcal{M}_{\underline{\alpha}}$ as $\{c\eta \mid c \in \mathbb{C}\}$ for a fixed quadratic differential η with simple poles at D . Denote by $q = \check{p}^* \eta$ its pullback. Then a point (E, Φ) in the Hitchin section of $\mathcal{M}_{\underline{\alpha}}$ is given by

$$(E, \Phi) = \left(\mathcal{O}(-1) \oplus \mathcal{O}(-3), \begin{pmatrix} 0 & c \\ \eta & 0 \end{pmatrix} \right) \in \mathcal{M}_{\underline{\alpha}}.$$

After pullback and tensoring we obtain

$$(E', \Phi') = \left(K^{\frac{1}{2}}(R_1 + R_2) \oplus K^{-\frac{1}{2}}, \begin{pmatrix} 0 & c\check{p}^*1 \\ \check{p}^*\eta & 0 \end{pmatrix} \right).$$

Here $\check{p}^*1 \in \check{p}^* \mathcal{O}(2) \otimes \check{p}^* K_{\mathbb{P}^1}$ has a zero of order 1 at each 2 : 1 ramification point and a zero of order 2 at each 3 : 1 ramification point. The pullback quasi-parabolic structure at $\check{p}^{-1}D$ is given by the second coordinate with respect to the splitting. Hence, the Hecke modification yields

$$\Theta(E, \Phi) = \left(K^{\frac{1}{2}} \oplus K^{-\frac{1}{2}}, \begin{pmatrix} 0 & cq \\ 1 & 0 \end{pmatrix} \right) = s_H(cq)$$

as asserted.

Compatibility with pullback by p : We show that this map extends the spectral correspondence on the regular locus. First we showed above that the Hitchin section of $\mathcal{M}_{\underline{\alpha}}$ is mapped on the Hitchin section of $\mathcal{M}_{\text{SL}(2, \mathbb{C})}(X)$. It is easy to see that the eigen line bundle on E of the Hitchin section of $\mathcal{M}_{\underline{\alpha}}$ is $\check{\pi}^* \mathcal{O}(-3)$. Similarly, the eigen line bundle on Σ of the Hitchin section of $\mathcal{M}_{\text{SL}(2, \mathbb{C})}(X)$ is $\pi^* K_X^{-\frac{1}{2}}$. We use these two Hitchin sections to identify the fibers with the corresponding abelian varieties. Then $(E, \Phi) \in \mathcal{M}_{\underline{\alpha}}$ corresponds to an element $L_1 \in \text{Jac}(E)$. Similarly, $(\widehat{E}, \widehat{\Phi})$ corresponds to an element $L_2 \in \text{Prym}(\Sigma_q)$. We need to show that $L_2 = p^* L_1$. From the spectral correspondence we have an exact sequence

$$0 \rightarrow L_1 \otimes \check{\pi}^* \mathcal{O}(-3) \rightarrow \check{\pi}^* E \xrightarrow{\check{\pi}^* \Phi - \omega \text{id}_{\check{\pi}^* E}} \check{\pi}^* (E \otimes \mathcal{O}(2)).$$

Tensoring with $\mathcal{O}(3)$ and pulling back we obtain

$$0 \rightarrow p^* L_1 \rightarrow \pi^* \left(E' \otimes K_X^{\frac{1}{2}} \right) \xrightarrow{\pi^* \Phi' - \lambda \text{id}_{\pi^* E'}} \pi^* \left(E' \otimes K_X^{\frac{1}{2}} \right) \otimes (\check{\pi} \circ p)^* \mathcal{O}(2).$$

Here we used the commutativity of diagram 5. Finally, twisting by $K_X^{-\frac{1}{2}}$ we see that the line bundle associated to $\pi^*(E', \Phi')$ through the spectral correspondence is $p^* L_1$. The (pullback of the) Hecke modification $\widehat{E} \rightarrow E$ can potentially change this line bundle by

twisting it with a divisor supported at $\pi^{-1}R$. However, on the regular locus of \mathcal{M}_α the quasi-parabolic structure is uniquely determined through the Higgs field at each $p \in D$ (see [FMSW22, Prop. 8.1]). Hence, we can compute the eigen line bundle and the quasi-parabolic structure of $\pi^*(E', \Phi')$ with respect to the pullback of the local frame of (99).

Then it is easy to see that the eigen line bundle $p^*L_1 \otimes K_x^{-\frac{1}{2}}$ descends to a subbundle of $\pi^*(\widehat{E}, \widehat{\Phi})$. Hence $p^*L_1 = L_2$.

Solutions to Hitchin equation: Finally, we will show that a hermitian metric h on $(E, \Phi) \in \mathcal{M}_\alpha$ that solves the Hitchin equation is transformed to a solution to the Hitchin equation for $(\widehat{E}, \widehat{\Phi}) \in \mathcal{M}_{\text{SL}(2, \mathbb{C})}(X)$. First there is a section of $\mathcal{O}(4)$ with divisors D . Promoting this section to have norm 1 defines a singular hermitian metric on $\mathcal{O}(4)$ locally given by $|z|^{-2}$ at $y \in D$. This induces a singular hermitian metric on $\mathcal{O}(1)$ and hence on $\mathcal{O}(3)$. The latter is given by $|z|^{-\frac{3}{2}}$ at $y \in D$ and will be denoted by $h_{\mathcal{O}(3)}$. Similarly, $q \in H^0(X, K_X^2)$ induces a singular hermitian metric $h_{-\frac{1}{2}K_X}$ on $K_X^{-\frac{1}{2}}$ that is smooth away from $Z(q)$ and given by $|w|^{\frac{1}{2}}$ locally at $x \in Z(q)$. It is easy to see that singular hermitian metrics defined in this way are automatically flat.

We need to extend the morphism Θ to hermitian Higgs bundles (E, Φ, h) . In the first step we use the singular hermitian metrics defined above to obtain a hermitian metric on E' by

$$(E', h') = \pi^* \left((E, h) \otimes (\mathcal{O}(3), h_{\mathcal{O}(3)}) \right) \otimes (K_X^{-\frac{1}{2}}, h_{-\frac{1}{2}K_X}).$$

This hermitian metric is holomorphic on $X \setminus R$. The hermitian metric h' pulls back to a hermitian metric \widehat{h} on E through the Hecke modification $\widehat{E} \rightarrow E$ a priori holomorphic only on $X \setminus R$.

Now, we start with a poly-stable Higgs bundle $(E, \Phi) \in \mathcal{M}_\alpha$ and let h be a solution to Hitchin's equation. By the flatness of the hermitian metrics $h_{\mathcal{O}(3)}$ and $h_{-\frac{1}{2}K_X}$ the resulting hermitian metric h' on E will still be a solution to Hitchin's equation wherever it is smooth. By definition the Hecke modification $\widehat{E} \rightarrow E$ is an isomorphism on $X \setminus R$ and hence the induced metric \widehat{h} is a solution to Hitchin's equation on this locus. To show that it defines a solution to Hitchin's equation on X we are left with showing that it extends smoothly over R .

To do so we compute the local description of \widehat{h} at $x \in R$. The metric h is adapted to the parabolic structure. Hence, at $y \in D$ we can find a local frame s_1, s_2 of E , such that firstly

$$h = \begin{pmatrix} |z|^{2\alpha_1} & \\ & |z|^{2\alpha_2} \end{pmatrix},$$

secondly the Higgs field is given by (99) and thirdly the quasi-parabolic structure is the ascending flag $\langle s_1 \rangle \subset \langle s_1, s_2 \rangle$. We have to consider two cases depending on the ramification index of the fiber over $y \in D$. We will only give the details for $y \in D$ of ramification index 2, i.e. $\check{p}^{-1}y$ is made up from 3 : 1-ramification points.

In this case, the compatible parabolic weight are $\underline{\alpha}_y = (\frac{3}{4}, \frac{1}{4})$, so that $h = \text{diag}(|z|^{\frac{4}{3}}, |z|^{\frac{2}{3}})$. Tensoring with $(\mathcal{O}(3), h_{\mathcal{O}(3)})$, pulling back and then tensoring with $(K_X^{-\frac{1}{2}}, h_{-\frac{1}{2}K})$ we obtain a local description for h' at $x \in Z(q)$ with respect to the induced frame

$$h' = \check{p}^* \text{diag}(|z|^{\frac{4}{3}-\frac{3}{2}}, |z|^{\frac{2}{3}-\frac{3}{2}})|w|^{\frac{1}{2}} = \text{diag}(1, |w|^{-2}).$$

Here w is local coordinate at $x \in \check{p}^{-1}y$, such that $\check{p} : w \mapsto w^3 = z$. Finally, with respect to the frame s_1, s_2 the Hecke modification $\widehat{E} \rightarrow E$ is given by $\text{diag}(1, w)$ and hence the induced metric on \widehat{E} is indeed smooth at $x \in R$ of ramification index 2. The case of ramification index 1 follows along the same lines. Hence, \widehat{h} defines a smooth solution to the Hitchin equation for the Higgs bundle $(\widehat{E}, \widehat{\Phi})$.

In particular, for a poly-stable Higgs bundle $(E, \Phi) \in \mathcal{M}_\alpha$ the image $\Theta(E, \Phi)$ is poly-stable and hence indeed Θ defines a map of moduli spaces

$$\Theta : \mathcal{M}_\alpha \rightarrow \mathcal{M}_{\text{SL}(2, \mathbb{C})}(X).$$

This finishes the proof. \square

REMARK 6.4. If (X, q) is a uniform pillowcase cover, such that there is an odd number of $y \in D$ of ramification index 1, then the compatible weights are generic in the sense that semi-stability implies stability. In particular, $\mathcal{M}_{\underline{\alpha}}$ is the elliptic surface referred to as Hausel's toy model with the nilpotent cone being of Kodaira type I_0^* .

COROLLARY 6.5. *Let (X, q) be a uniform pillowcase cover. Then $\mathcal{I}_q \subset \mathcal{M}_{\mathrm{PGL}(2, \mathbb{C})}(X)$ is a hyperholomorphic subvariety.*

PROOF. By Theorem 6.3 Θ maps solutions to the Hitchin equation on (\mathbb{P}^1, D) to solutions to the Hitchin equation on X . Hence, it is holomorphic not only with respect to the holomorphic structure I , but also with respect to the holomorphic structure J and K on the moduli spaces $\mathcal{M}_{\underline{\alpha}}(\mathbb{P}^1, D)$ and $\mathcal{M}_{\mathrm{SL}(2, \mathbb{C})}(X)$. In particular, its image is a hyperholomorphic subvariety. However, Theorem 6.3 i) and ii) shows that Θ restricts to an isomorphism from the regular locus of $\mathcal{M}_{\underline{\alpha}}$ to the torsor $\mathcal{I}' \subset \mathcal{M}_{\mathrm{SL}(2, \mathbb{C})}(X)$ obtained by acting with the abelian scheme $E \subset \mathrm{Prym}(\Sigma/\mathbb{C}^\times q)$ over $\mathbb{C}^\times q$ on the Hitchin section. Hence, $\overline{\mathcal{I}'}$ the image of Θ is a hyperholomorphic subvariety. In particular, its image \mathcal{I}_q under the quotient map $\delta : \mathcal{M}_{\mathrm{SL}(2, \mathbb{C})}(X) \rightarrow \mathcal{M}_{\mathrm{PGL}(2, \mathbb{C})}(X)$ is hyperholomorphic. \square

Bibliography

- [ABBR15a] O. Amini, M. Baker, E. Brugallé, and J. Rabinoff. „Lifting harmonic morphisms I: metrized complexes and Berkovich skeleta“. In: *Research in the Mathematical Sciences* 2.7 (1 2015).
- [ABBR15b] O. Amini, M. Baker, E. Brugallé, and J. Rabinoff. „Lifting harmonic morphisms II: Tropical curves and metrized complexes“. In: *Algebra & Number Theory* 9.2 (2015), pp. 267–315.
- [ACP15] D. Abramovich, L. Caporaso, and S. Payne. „The tropicalization of the moduli space of curves“. In: *Annales Scientifiques de l'École Normale Supérieure* 48 (4 2015), pp. 765–809.
- [AEM17] A. Avila, A. Eskin, and M. Möller. In: *Journal für die reine und angewandte Mathematik (Crelles Journal)* 2017.732 (2017), pp. 1–20.
- [Bai07] M. Bainbridge. „Euler characteristics of Teichmüller curves in genus two“. In: *Geometry & Topology* 11.4 (2007), pp. 1887–2073.
- [Bak08] M. Baker. „Specialization of linear systems from curves to graphs“. In: *Algebra Number Theory* 2.6 (2008), pp. 613–653.
- [Bar16] D. Baraglia. „Classification of the automorphism and isometry groups of Higgs bundle moduli spaces“. English. In: *Proc. Lond. Math. Soc. (3)* 112.5 (2016), pp. 827–854.
- [BBCMMW21] M. Brandt, J. Bruce, M. Chan, M. Melo, G. Moreland, and C. Wolfe. *On the Top-Weight Rational Cohomology of \mathcal{A}_g* . 2021. arXiv: 2012.02892 [math.AG].
- [BCGGM18] M. Bainbridge, D. Chen, Q. Gendron, S. Grushevsky, and M. Möller. „Compactification of strata of Abelian differentials“. In: *Duke Math. J.* 167.12 (2018), pp. 2347–2416.
- [BCGGM19a] M. Bainbridge, D. Chen, Q. Gendron, S. Grushevsky, and M. Möller. „Strata of k -differentials“. In: *Algebr. Geom.* 6.2 (2019), pp. 196–233.
- [BCGGM19b] M. Bainbridge, D. Chen, Q. Gendron, S. Grushevsky, and M. Möller. *The moduli space of multi-scale differentials*. Preprint. 2019. arXiv: 1910.13492.
- [BDG22] F. Benirschke, B. Dozier, and S. Grushevsky. „Equations of linear subvarieties of strata of differentials“. In: *Geom. Topol.* 26.6 (2022), pp. 2773–2830.
- [Ben20] F. Benirschke. *The boundary of linear subvarieties*. 2020. arXiv: 2007.02502 [math.AG].
- [Ber90] V. Berkovich. *Spectral Theory and Analytic Geometry Over non-Archimedean Fields*. AMS Mathematical Surveys and Monographs 33. AMS, 1990.
- [BHPSS20] Y. Bae, D. Holmes, R. Pandharipande, J. Schmitt, and R. Schwarz. *Pixton's formula and Abel-Jacobi theory on the Picard stack*. 2020. arXiv: 2004.08676.
- [BJ16] M. Baker and D. Jensen. „Degeneration of linear series from the tropical point of view and applications“. In: *Nonarchimedean and Tropical Geometry* (2016), pp. 365–433.
- [BM10] I. I. Bouw and M. Möller. „Teichmüller curves, triangle groups, and Lyapunov exponents“. In: *Ann. Math. (2)* 172.1 (2010), pp. 139–185.
- [BN16] M. Baker and J. Nicaise. „Weight functions on Berkovich curves“. In: *Algebra & Number Theory* 10.10 (2016), pp. 2053–2079.
- [Boi15] C. Boissy. „Connected components of the strata of the moduli space of meromorphic differentials“. In: *Commentarii Mathematici Helvetici* 90.2 (2015), pp. 255–286.

- [BS16] D. Baraglia and L. P. Schaposnik. „Real structures on moduli spaces of Higgs bundles“. In: *Adv. Theor. Math. Phys.* 20.3 (2016), pp. 525–551.
- [BS19] D. Baraglia and L. P. Schaposnik. „Cayley and Langlands type correspondences for orthogonal Higgs bundles“. In: *Trans. Am. Math. Soc.* 371.10 (2019), pp. 7451–7492.
- [BU19] M. Brandt and M. Ulirsch. *Symmetric powers of algebraic and tropical curves: a non-Archimedean perspective*. 2019. arXiv: 1812.08740 [math.AG].
- [Cal04] K. Calta. „Veech surfaces and complete periodicity in genus two“. In: *J. Amer. Math. Soc.* 17.4 (2004), pp. 871–908.
- [Car15] D. Cartwright. „Lifting matroid divisors on tropical curves“. In: *Research in the Mathematical Sciences* 2 (1 2015).
- [CCM22] D. Chen, M. Costantini, and M. Möller. *On the Kodaira dimension of moduli spaces of Abelian differentials*. Preprint. 2022. arXiv: 2204.11943.
- [CGHMS22] D. Chen, S. Grushevsky, D. Holmes, M. Möller, and J. Schmitt. *A tale of two moduli spaces: logarithmic and multi-scale differentials*. 2022. arXiv: 2006.12815.
- [CGP21] M. Chan, S. Galatius, and S. Payne. „Tropical curves, graph complexes, and top weight cohomology of \mathcal{M}_g “. In: *Journal of the American Mathematical Society* 34 (2021), pp. 565–594.
- [Che22] D. Chen. „Dynamical invariants and intersection theory on the flex and gothic loci“. In: *European Journal of Mathematics* 8.1 (2022), pp. 42–52.
- [CJP15] D. Cartwright, D. Jensen, and S. Payne. „Lifting Divisors on a Generic Chain of Loops“. In: *Canadian Mathematical Bulletin* 58.2 (2015), pp. 250–262.
- [CMR16] R. Cavalieri, H. Markwig, and D. Ranganathan. „Tropicalizing the Space of Admissible Covers“. In: *Mathematische Annalen* 364 (2016), pp. 1275–1313.
- [CMS19] D. Chen, M. Möller, and A. Sauvaget. *Masur-Veech volumes and intersection theory: the principal strata of quadratic differentials*. With an appendix by G. Borot, A. Giacchetto and D. Lewanski. 2019. arXiv: 1912.02267. to appear in: *Duke Math J.*
- [CMS23] M. Costantini, M. Möller, and J. Schwab. *Chern classes of linear submanifolds with application to spaces of k -differentials and ball quotients*. 2023. arXiv: 2303.17929 [math.AG].
- [CMZ19] M. Costantini, M. Möller, and J. Zachhuber. *The area is a good enough metric*. 2019. arXiv: 1910.14151. to appear in: *Ann. Inst. Fourier*.
- [CMZ22] M. Costantini, M. Möller, and J. Zachhuber. „The Chern classes and Euler characteristic of the moduli spaces of Abelian differentials“. In: *Forum Math. Pi* 10 (2022), Paper No. e16, 55.
- [CMZ23] M. Costantini, M. Möller, and J. Zachhuber. „diffstrata—A Sage Package for Calculations in the Tautological Ring of the Moduli Space of Abelian Differentials“. In: *Experimental Mathematics* 32.3 (2023), pp. 545–565.
- [CT20] J. Chen and H. Tseng. „On the Bogomolov-Miyaoka-Yau inequality for stacky surfaces“. In: *Taiwanese J. Math.* 24.4 (2020), pp. 841–853.
- [Der20] M. Deraux. „A new nonarithmetic lattice in $\mathrm{PU}(3, 1)$ “. In: *Algebr. Geom. Topol.* 20.2 (2020), pp. 925–963.
- [DM86] P. Deligne and G. D. Mostow. „Monodromy of hypergeometric functions and nonlattice integral monodromy“. In: *Inst. Hautes Études Sci. Publ. Math.* 63 (1986), pp. 5–89.
- [DP12] R. Donagi and T. Pantev. „Langlands duality for Hitchin systems“. In: *Invent. Math.* 189.3 (2012), pp. 653–735.
- [DPP16] M. Deraux, J. Parker, and J. Paupert. „New non-arithmetic complex hyperbolic lattices“. In: *Invent. Math.* 203.3 (2016), pp. 681–771.
- [DR23] V. Delecroix and J. R uth. *A new orbit closure in genus 8?* 2023. arXiv: 2110.05407 [math.DS].

- [DSZ21] V. Delecroix, J. Schmitt, and J. van Zelm. „admcycles—a Sage package for calculations in the tautological ring of the moduli space of stable curves“. In: *J. Softw. Algebra Geom.* 11.1 (2021), pp. 89–112.
- [EM18] A. Eskin and M. Mirzakhani. „Invariant and stationary measures for the $SL(2, \mathbb{R})$ action on moduli space“. In: *Publ. Math. Inst. Hautes Études Sci.* 127 (2018), pp. 95–324.
- [EMM15] A. Eskin, M. Mirzakhani, and A. Mohammadi. „Isolation, equidistribution, and orbit closures for the $SL(2, \mathbb{R})$ action on moduli space“. In: *Ann. Math. (2)* 182.2 (2015), pp. 673–721.
- [EMMW20] A. Eskin, C. McMullen, R. Mukamel, and A. Wright. „Billiards, quadrilaterals, and Moduli spaces“. In: *J. Amer. Math. Soc.* 33.4 (2020), pp. 1039–1086.
- [Eva23] J. Evans. *Lectures on Lagrangian torus fibrations*. Vol. 105. Lond. Math. Soc. Stud. Texts. Cambridge: Cambridge University Press, 2023.
- [FGOP21] E. Franco, P. B. Gothen, A. Oliveira, and A. Peón-Nieto. „Unramified covers and branes on the Hitchin system“. In: *Adv. Math.* 377 (2021), p. 62.
- [FHHO23] E. Franco, R. Hanson, J. Horn, and A. Oliveira. *Lagrangians of Hecke cycles*. 2023. In preparation.
- [Fil16] S. Filip. „Splitting mixed Hodge structures over affine invariant manifolds“. In: *Ann. Math. (2)* 183.2 (2016), pp. 681–713.
- [Fil20] S. Filip. „Translation surfaces: Dynamics and Hodge theory“. In: (2020). eprint: https://math.uchicago.edu/~sfilip/public_files/surf_surv.pdf.
- [FJ22] E. Franco and M. Jardim. „Mirror symmetry for Nahm branes“. In: *Épjournal de Géom. Algèbr., EPIGA* 6 (2022), p. 29.
- [FMSW22] L. Fredrickson, R. Mazzeo, J. Swoboda, and H. Weiss. „Asymptotic geometry of the moduli space of parabolic $SL(2, \mathbb{C})$ -Higgs bundles“. In: *J. Lond. Math. Soc., II. Ser.* 106.2 (2022), pp. 590–661.
- [FP23] E. Franco and A. Peón-Nieto. „Branes on the singular locus of the Hitchin system via Borel and other parabolic subgroups“. In: *Mathematische Nachrichten* 296.5 (2023), pp. 1803–1841.
- [Fre99] D. S. Freed. „Special Kähler manifolds.“ English. In: *Commun. Math. Phys.* 203.1 (1999), pp. 31–52.
- [GAP22] *GAP – Groups, Algorithms, and Programming, Version 4.12.2*. The GAP Group, 2022.
- [Gee88] G. van der Geer. *Hilbert modular surfaces*. Vol. 16. Ergebnisse der Mathematik und ihrer Grenzgebiete (3) [Results in Mathematics and Related Areas (3)]. Springer-Verlag, Berlin, 1988, pp. x+291.
- [Gen20] Q. Gendron. „Les strates ne possèdent pas de variétés complètes“. In: *Comptes Rendus Mathématique* 358.2 (2020), pp. 197–200.
- [GK08] A. Gathmann and M. Kerber. „A Riemann-Roch Theorem in Tropical Geometry“. In: *Mathematische Zeitschrift* 259 (2008), pp. 217–230.
- [GKPT19] D. Greb, S. Kebekus, T. Peternell, and B. Taji. „The Miyaoka-Yau inequality and uniformisation of canonical models“. In: *Ann. Sci. Éc. Norm. Supér. (4)* 52.6 (2019), pp. 1487–1535.
- [Gou16] E. Goujard. „Volumes of strata of moduli spaces of quadratic differentials: getting explicit values“. In: *Ann. Inst. Fourier (Grenoble)* 66.6 (2016), pp. 2203–2251.
- [GP03] T. Graber and R. Pandharipande. „Constructions of nontautological classes on moduli spaces of curves“. In: *Michigan Math. J.* 51.1 (2003), pp. 93–109.
- [GP17] S. Ghazouani and L. Pirio. „Moduli spaces of flat tori with prescribed holonomy“. In: *Geom. Funct. Anal.* 27.6 (2017), pp. 1289–1366.
- [GP20] S. Ghazouani and L. Pirio. „Moduli spaces of flat tori and elliptic hypergeometric functions“. In: *Mém. Soc. Math. Fr. (N.S.)* 164 (2020), pp. viii+183.

- [GT21a] Q. Gendron and G. Tahar. „Différentielles abéliennes à singularités prescrites“. In: *Journal de l'École Polytechnique – Mathématiques* 8 (2021), pp. 1397–1428.
- [GT21b] Q. Gendron and G. Tahar. *Quadratic differentials with prescribed singularities*. 2021. arXiv: 2111.12653 [math.GT].
- [GT22a] Q. Gendron and G. Tahar. *K-differentials With Prescribed Singularities*. 2022. arXiv: 2208.11654 [math.CV].
- [GT22b] H. Guenancia and B. Taji. „Orbifold stability and Miyaoka-Yau inequality for minimal pairs“. In: *Geom. Topol.* 26.4 (2022), pp. 1435–1482.
- [Gub13] W. Gubler. „A guide to tropicalizations“. In: *Algebraic and combinatorial aspects of tropical geometry*. Vol. 589. Contemp. Math. Amer. Math. Soc., Providence, RI, 2013, pp. 125–189.
- [Hau98] T. Hausel. „Compactification of moduli of Higgs bundles“. In: *J. Reine Angew. Math.* 503 (1998), pp. 169–192.
- [HH09] B. Hassett and D. Hyeon. „Log canonical models for the moduli space of curves: the first divisorial contraction“. In: *Trans. Amer. Math. Soc.* 361.8 (2009), pp. 4471–4489.
- [HH22] T. Hausel and N. Hitchin. „Very stable Higgs bundles, equivariant multiplicity and mirror symmetry“. English. In: *Invent. Math.* 228.2 (2022), pp. 893–989.
- [Hit16] N. Hitchin. „Higgs bundles and characteristic classes“. In: *Arbeitstagung Bonn 2013. In memory of Friedrich Hirzebruch. Proceedings of the meeting, Bonn, Germany, May, 22–28, 2013*. Birkhäuser / Springer, 2016, pp. 247–264.
- [Hit17] N. Hitchin. „Spinors, Lagrangians and rank 2 Higgs bundles“. In: *Proc. Lond. Math. Soc. (3)* 115.1 (2017), pp. 33–54.
- [Hor22] J. Horn. „Semi-abelian spectral data for singular fibres of the $SL(2, \mathbb{C})$ -Hitchin system“. In: *Int. Math. Res. Not.* 2022.5 (2022), pp. 3860–3917.
- [HS01] P. Hubert and T. A. Schmidt. „Invariants of translation surfaces“. en. In: *Annales de l'Institut Fourier* 51.2 (2001), pp. 461–495.
- [HS18] S. Heller and L. P. Schaposnik. „Branes through finite group actions“. In: *J. Geom. Phys.* 129 (2018), pp. 279–293.
- [HS21] D. Holmes and J. Schmitt. „Infinitesimal structure of the pluricanonical double ramification locus“. In: *Compos. Math.* 157.10 (2021), pp. 2280–2337.
- [HS23] J. Horn and J. Schwab. *Visible Lagrangians for Hitchin systems and pillowcase covers*. 2023. arXiv: 2310.15716 [math.AG].
- [HT03] T. Hausel and M. Thaddeus. „Mirror symmetry, Langlands duality, and the Hitchin system“. In: *Invent. Math.* 153.1 (2003), pp. 197–229.
- [KKMS73] G. Kempf, F. Knudsen, D. Mumford, and B. Saint-Donat. *Toroidal embeddings. I*. Lecture Notes in Mathematics, Vol. 339. Springer-Verlag, Berlin-New York, 1973, pp. viii+209.
- [KM16] A. Kappes and M. Möller. „Lyapunov spectrum of ball quotients with applications to commensurability questions“. In: *Duke Math. J.* 165.1 (2016), pp. 1–66.
- [KM98] J. Kollár and S. Mori. *Birational geometry of algebraic varieties*. Vol. 134. Cambridge Tracts in Mathematics. With the collaboration of C. H. Clemens and A. Corti, Translated from the 1998 Japanese original. Cambridge University Press, Cambridge, 1998, pp. viii+254.
- [KN18] V. Koziarz and D.-M. Nguyen. „Complex hyperbolic volume and intersection of boundary divisors in moduli spaces of pointed genus zero curves“. In: *Ann. Sci. Éc. Norm. Supér. (4)* 51.6 (2018), pp. 1549–1597.
- [KNS89] R. Kobayashi, S. Nakamura, and F. Sakai. „A numerical characterization of ball quotients for normal surfaces with branch loci“. In: *Proc. Japan Acad. Ser. A Math. Sci.* 65.7 (1989), pp. 238–241.
- [KS00] R. Kenyon and J. Smillie. „Billiards on rational-angled triangles“. In: *Commentarii Mathematici Helvetici* 75.1 (2000), pp. 65–108.

- [KW07] A. Kapustin and E. Witten. „Electric-magnetic duality and the geometric Langlands program“. English. In: *Commun. Number Theory Phys.* 1.1 (2007), pp. 1–236.
- [Lia21] C. Lian. „The \mathcal{H} -tautological ring“. In: *Selecta Mathematica New Ser.* 27.96 (2021).
- [LU17] B. Lin and M. Ulirsch. „Towards a Tropical Hodge Bundle“. In: *Combinatorial Algebraic Geometry* (2017), pp. 353–368.
- [Mar94] E. Markman. „Spectral curves and integrable systems“. In: *Compos. Math.* 93.3 (1994), pp. 255–290.
- [McM03] C. McMullen. „Billiards and Teichmüller curves on Hilbert modular surfaces“. In: *J. Amer. Math. Soc.* 16.4 (2003), pp. 857–885.
- [McM06] C. McMullen. „Prym varieties and Teichmüller curves“. In: *Duke Math. J.* 133.3 (2006), pp. 569–590.
- [McM07a] C. McMullen. „Foliations of Hilbert modular surfaces“. In: *Amer. J. Math.* 129.1 (2007), pp. 183–215.
- [McM07b] C. McMullen. „Dynamics of $SL_2(\mathbb{R})$ over moduli space in genus two.“ English. In: *Ann. Math. (2)* 165.2 (2007), pp. 397–456.
- [McM17] C. McMullen. „The Gauss-Bonnet theorem for cone manifolds and volumes of moduli spaces“. In: *Amer. J. Math.* 139 (2017), pp. 261–291.
- [Mik05] G. Mikhalkin. „Enumerative tropical algebraic geometry in \mathbb{R}^2 “. In: *Journal of the American Mathematical Society* 18 (2005), pp. 313–377.
- [MM23] M. Möller and S. Mullane. *Teichmüller curves in hyperelliptic components of meromorphic strata*. 2023. arXiv: 2305.03309 [math.AG].
- [MMW17] C. McMullen, R. Mukamel, and A. Wright. „Cubic curves and totally geodesic subvarieties of moduli space“. In: *Ann. of Math. (2)* 185.3 (2017), pp. 957–990.
- [Mos86] G. D. Mostow. „Generalized Picard lattices arising from half-integral conditions“. In: *Inst. Hautes Études Sci. Publ. Math.* 63 (1986), pp. 91–106.
- [Mos88] G. D. Mostow. „On discontinuous action of monodromy groups on the complex n -ball“. In: *J. Amer. Math. Soc.* 1.3 (1988), pp. 555–586.
- [MT20] M. Möller and D. Torres-Teigell. „Euler characteristics of Gothic Teichmüller curves“. In: *Geometry & Topology* 24.3 (2020), pp. 1149–1210.
- [Mum74] D. Mumford. *Prym varieties. I*. English. *Contribut. to Analysis, Collect. of Papers dedicated to Lipman Bers*, 325–350. 1974.
- [Mum77] D. Mumford. „Hirzebruch’s proportionality theorem in the noncompact case“. In: *Invent. Math.* 42 (1977), pp. 239–272.
- [MUW21] M. Möller, M. Ulirsch, and A. Werner. „Realizability of tropical canonical divisors“. In: *Journal of the European Mathematical Society* 23 (2021), pp. 185–217.
- [MZ16] M. Möller and D. Zagier. „Modular embeddings of Teichmüller curves“. In: *Compos. Math.* 152.11 (2016), pp. 2269–2349.
- [Pet14] D. Petersen. „The structure of the tautological ring in genus one“. In: *Duke Mathematical Journal* 163.4 (2014).
- [PP11] G. Pareschi and M. Popa. „GV-sheaves, Fourier-Mukai transform, and generic vanishing“. In: *Am. J. Math.* 133.1 (2011), pp. 235–271.
- [PP19] C. Pauly and A. Peón-Nieto. „Very stable bundles and properness of the Hitchin map“. English. In: *Geom. Dedicata* 198 (2019), pp. 143–148.
- [PP21] R. Pandharipande and A. Pixton. „Relations in the tautological ring of the moduli space of curves“. In: *Pure Appl. Math. Q.* 17.2 (2021), pp. 717–771.
- [RS21] F. Röhrle and J. Schwab. *Realizability of tropical pluri-canonical divisors*. 2021. arXiv: 2109.03579 [math.AG].
- [Sau18] A. Sauvaget. „Volumes and Siegel–Veech constants of $\mathcal{H}(2g - 2)$ and Hodge integrals“. In: *Geom. Funct. Anal.* 28.6 (2018), pp. 1756–1779.
- [Sau21] A. Sauvaget. *Volumes of moduli spaces of flat surfaces*. 2021. arXiv: 2004.03198 [math.AG].

- [Sch22] C. Schnell. „The Fourier-Mukai transform made easy“. In: *Pure Appl. Math. Q.* 18.4 (2022), pp. 1749–1770.
- [Sie36] C. Siegel. „The volume of the fundamental domain for some infinite groups“. In: *Trans. Amer. Math. Soc.* 39.2 (1936), pp. 209–218.
- [SZ20] J. Schmitt and J. van Zelm. „Intersections of loci of admissible covers with tautological classes“. In: *Sel. Math. New Ser.* 26.79 (2020).
- [Thu07] A. Thuillier. „Géométrie toroïdale et géométrie analytique non archimédienne. Application au type d’homotopie de certain schémas formels“. In: *Manuscripta mathematica* 123 (2007), pp. 381–451.
- [Thu98] W. P. Thurston. „Shapes of polyhedra and triangulations of the sphere“. In: *The Epstein birthday schrift*. Vol. 1. Geom. Topol. Monogr. Geom. Topol. Publ., Coventry, 1998, 511–549 (electronic).
- [Uli15] M. Ulirsch. „Tropical compactification in log-regular varieties“. In: *Mathematische Zeitschrift* 280 (1-2 2015), pp. 195–210.
- [Uli17] M. Ulirsch. „Functorial tropicalization of logarithmic schemes: the case of constant coefficients“. In: *Proceedings of the London Mathematical Society* 114.6 (2017), pp. 1081–1113.
- [Vee89] W. Veech. „Teichmüller curves in moduli space, Eisenstein series and an application to triangular billiards.“ In: *Inventiones mathematicae* 97.3 (1989), pp. 553–584.
- [Vee91] W. Veech. „Erratum: Teichmüller curves in moduli space, Eisenstein series and an application to triangular billiards.“ In: *Inventiones mathematicae* 103.2 (1991), p. 447.
- [Vee93] W. Veech. „Flat surfaces“. In: *Amer. J. Math.* 115.3 (1993), pp. 589–689.
- [Viv13] F. Viviani. „Tropicalizing vs. compactifying the Torelli morphism“. In: *Tropical and non-Archimedean geometry*. Vol. 605. Contemp. Math. Amer. Math. Soc., Providence, RI, 2013, pp. 181–210.
- [War98] C. C. Ward. „Calculation of Fuchsian groups associated to billiards in a rational triangle“. In: *Ergodic Theory and Dynamical Systems* 18.4 (1998), pp. 1019–1042.
- [Wri15] A. Wright. „Cylinder deformations in orbit closures of translation surfaces“. In: *Geometry & Topology* 19.1 (2015), pp. 413–438.
- [Zor06] A. Zorich. „Flat surfaces“. In: *Frontiers in Number Theory, Physics and Geometry. Volume 1: On random matrices, zeta functions and dynamical systems*. Berlin: Springer-Verlag, 2006, pp. 439–586.

Zusammenfassung

1. \mathbb{R} -lineare Untermannigfaltigkeiten

Wenn Sie einen rechteckigen Billardtisch haben und eine Kugel in einer der Ecken starten, fragen Sie sich möglicherweise: „Wird diese Kugel jemals wieder in einer Ecke ankommen?“. Lassen Sie uns für den Moment annehmen, dass Ihr Billardtisch ganzzahlige Seitenlängen hat. Die Bahn der Kugel entlang aller Reflexionen im Rechteck nachzuverfolgen wird schnell mühsam. Statt jedes Mal die Kugel zu reflektieren, wenn sie eine Seite des Tisches berührt, ist es sehr viel angenehmer, stattdessen den Tisch zu reflektieren wie in Abbildung 6. Jetzt spannen die Ecken des Tischen ein Gitter in der Euklidischen Ebene auf. In diesem Setting übersetzt sich die Frage „Wird die Kugel jemals wieder in einer Ecke ankommen?“ zur Frage „Wird die Kugel jemals wieder einen Gitterpunkt treffen?“. Letztere Frage kann man sofort beantworten: Die Kugel wird einen Gitterpunkt treffen genau dann, wenn der Winkel zwischen der Trajektorie und einer Seite des Polygons ein rationales Vielfaches von π ist.

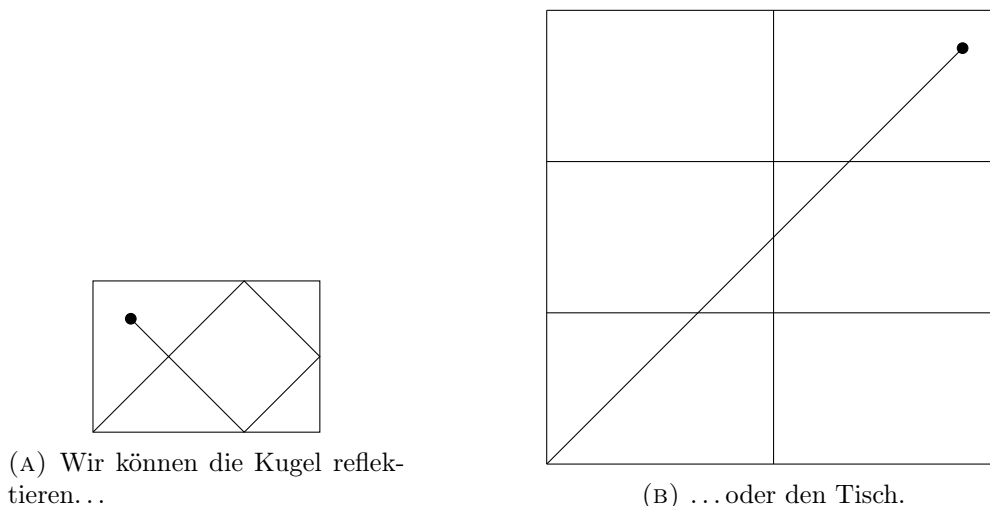


ABBILDUNG 6. Ein rechteckiger Billardtisch

Aber was passiert, wenn die Kugel keine Ecke trifft? In diesem Fall wird die Trajektorie offensichtlich unendlich sein. Wir können aber noch mehr sagen: Die Trajektorie wird dicht im Tisch liegen als Konsequenz des Dirichletschen Approximationssatz. Das motiviert die folgende Definition.

DEFINITION 1.1. *Ein Billardtisch in dem jede Trajektorie entweder geschlossen oder dicht ist hat optimale Dynamik.*

Wie wir oben gesehen haben, haben rechteckige Billardtische mit ganzzahligen (oder allgemeiner rationalen) Seitenlängen optimale Dynamik. Für einen komplizierteren Billardtisch wird unser Ansatz mit Gittern nicht mehr funktionieren. Stattdessen können wir die folgende Beobachtung benutzen. Den Tisch einmal zu reflektieren produziert einen Tisch mit einer anderen Orientierung im Vergleich zum ursprünglichem Tisch. Den neuen Tisch noch einmal in die gleiche Richtung zu reflektieren produziert einen Tisch mit der gleichen Orientierung wie der ursprüngliche Tisch. Anstatt also schon wieder einen neuen Tisch zu produzieren (wie wir es getan haben um das Gitter zu erhalten), können wir den zweiten Tisch mit dem ersten verkleben, um eine Mannigfaltigkeit zu erhalten. Wenn wir die ursprüngliche Fläche eingebettet in die Gaußschen Ebene denken, dann hat die so erhaltene

Mannigfaltigkeit in natürlicher Weise die Struktur einer Riemannschen Fläche X . Indem man die Differentialform dz von der Ebene auf die Fläche zurück zieht, erhält man eine Differentialform ω auf X . Dieser Prozess, um (X, ω) aus einem Polygon zu erhalten, heißt Entfaltung und wird uns in Kapitel III wieder begegnen. Das Paar (X, ω) heißt flache Fläche, da ω auf X eine flache Metrik induziert.

Umgekehrt kann jede flache Fläche (X, ω) mit Polygonen dargestellt werden: Wenn $\gamma_1, \dots, \gamma_n$ eine Basis der relativen Homologie $H_1(X, Z(\omega))$ ist, dann sind die Seiten der Polygone gegeben durch

$$(101) \quad \int_{\gamma_i} \omega.$$

Lassen Sie uns zu unserem Billardtisch zurückkehren. Die Trajektorien der Kugel korrespondieren zu Geodäten in X bezüglich der durch ω gegebenen Metrik. Der Begriff der optimalen Dynamik überträgt sich dadurch in eine Frage über die Geodäten auf (X, ω) . Das Hodgebündel $\Omega\mathcal{M}_g \rightarrow \mathcal{M}_g$ über dem Modulraum der kompakten Riemannschen Flächen in Geschlecht g ist ein Modulraum der flachen Flächen. Die Gruppe $\mathrm{GL}_2(\mathbb{R})^+$ operiert auf diesem Modulraum vermöge der Operation auf den Polygondarstellungen der Flächen. Veech [Vee89] [Vee91] hat beobachtet, dass die flache Fläche (X, ω) optimale Dynamik hat genau dann, wenn ihr $\mathrm{GL}_2(\mathbb{R})^+$ -Orbit abgeschlossen ist.

DEFINITION 1.2. *Falls der $\mathrm{GL}_2(\mathbb{R})^+$ -Orbit von $(X, \omega) \in \Omega\mathcal{M}_g$ abgeschlossen ist, dann heißt (X, ω) Veechfläche. Das Bild von $\mathrm{GL}_2(\mathbb{R})^+ \cdot (X, \omega)$ in \mathcal{M}_g heißt Teichmüllerkurve.*

Wir haben oben das einfachste Beispiel einer Veechfläche gesehen: Die Entfaltung unseres rechteckigen Billardtisches (X, ω) . Seit Veechs Beobachtung ist viel Arbeit in die Klassifikation von Teichmüllerkurven geflossen. Da man leicht neue Veechflächen aus bekannten mittels Überlagerungen konstruiert, ist man hauptsächlich an solchen Veechflächen interessiert, die nicht via Überlagerungen aus bekannten hervorgehen. Die zugehörigen Teichmüllerkurven heißen *primitiv*. In Geschlecht 2, 3 und 4 kennen wir unendlich viele primitive Teichmüllerkurven, die von Veech [Vee89], Ward [War98], Bouw-Möller [BM10], McMullen [McM03] [McM06], Calta [Cal04], Vorobets [HS01] und Kenyon-Smillie [KS00] entdeckt wurden. In jedem Geschlecht größer 4 kennen wir nur endlich viele primitive Teichmüllerkurven, die alle zur Bouw-Möller-Serie gehören. Die Existenz unendlich vieler primitiver Teichmüllerkurven in jedem Geschlecht ist eine offene Frage.

Für ein festes Geschlecht g sei $\mu = (m_1, \dots, m_n)$ eine ganzzahlige Partition von $2g - 2$. Für eine flache Fläche (X, ω) vom Geschlecht g sagen wir dass ω vom *Typ* μ ist, falls ω genau n Nullstellen mit Ordnungen m_1, \dots, m_n hat. Der Modulraum der flachen Flächen $\Omega\mathcal{M}_g$, auch bekannt als der Modulraum der abelschen Differentiale, besitzt eine natürliche Stratifikation bezüglich der Typen der Differentiale, und wir bezeichnen mit $\Omega\mathcal{M}_{g,n}(\mu)$ das Stratum der Differentiale vom Typ μ . Die Integrale in (101) sind lokale Koordinaten auf dem Stratum, genannt *Periodenkoordinaten*.

DEFINITION 1.3. *Ein Unterraum $\Omega\mathcal{H} \subseteq \Omega\mathcal{M}_{g,n}(\mu)$ heißt K -lineare Untermannigfaltigkeit, wenn er von linearen Gleichungen in Periodenkoordinaten mit Koeffizienten in einem Körper K ausgeschnitten wird.*

Da die $\mathrm{GL}_2(\mathbb{R})^+$ -Operation Gleichungen mit Koeffizienten in \mathbb{R} erhält, ist jede \mathbb{R} -lineare Untermannigfaltigkeit der Abschluss einer Vereinigung von $\mathrm{GL}_2(\mathbb{R})^+$ -Orbits. Nach dem bahnbrechenden Resultat von Eskin-Mirzakhani-Mohammadi ist auch die Umkehrung wahr.

THEOREM 1.4 ([EMM15]). *Jeder $\mathrm{GL}_2(\mathbb{R})^+$ -Bahnabschluss ist eine \mathbb{R} -lineare Untermannigfaltigkeit.*

Eine wichtige Invariante einer linearen Untermannigfaltigkeit ist ihr *Rang*, der wie folgt definiert werden kann. Über $\Omega\mathcal{H}$ betrachten wir das Bündel H^1 , dessen Faser über (X, ω) durch $H^1(X, \mathbb{C})$ gegeben ist und das Bündel H_{rel}^1 , dessen Faser über (X, ω) durch $H^1(X, Z(\omega), \mathbb{C})$ gegeben ist. Sei $p : H_{\mathrm{rel}}^1 \rightarrow H^1$ die natürliche Abbildung. Nach Avila-Eskin-Möller [AEM17] ist der Raum $p(T(\Omega\mathcal{H}))$ symplektisch, insbesondere von gerader Dimension, und wir definieren den Rang von $\Omega\mathcal{H}$ als $\frac{1}{2} \dim p(T(\Omega\mathcal{H}))$.

Teichmüllerkurven (oder genauer die Bahnabschlüsse der zugehörigen Veechflächen) sind \mathbb{R} -lineare Untermannigfaltigkeiten vom Rang 1. Es war eine Überraschung als McMullen-Mukamel-Wright [MMW17] die erste primitive \mathbb{R} -lineare Untermannigfaltigkeit vom Rang 2 gefunden haben, den sogenannten gotischen Lokus. Bis jetzt wurden nur 6 weitere primitive \mathbb{R} -lineare Untermannigfaltigkeiten vom Rang 2 von Eskin-McMullen-Mukamel-Wright [EMMW20] gefunden. Es gibt rechnerische Hinweise auf die Existenz mindestens einer weiteren solchen Untermannigfaltigkeit [DR23]. Die Existenz einer \mathbb{R} -linearen Untermannigfaltigkeit vom Rang mindestens 3 ist ein vollständig offenes Problem.

2. Die Chernklassen linearer Untermannigfaltigkeiten

Um mathematische Objekte (wie z.B. \mathbb{R} -lineare Untermannigfaltigkeiten, die wiederum komplexe Orbifaltigkeiten sind) zu klassifizieren, ist es oft eine gute Idee, Invarianten dieser Objekte zu bestimmen. Eine wichtige Invariante einer komplexen Orbifaltigkeit sind die Chernklassen. Hier gibt es ein kleines Problem: Auf einer linearen Untermannigfaltigkeit $\Omega\mathcal{H}$ operiert die Gruppe \mathbb{C}^\times via Skalierung des Differentials, so dass $\Omega\mathcal{H}$ eine triviales \mathbb{C}^\times -Bündel ist und damit alle Chernklassen 0 sind. Wenn wir nützliche Resultate erwarten wollen, sollten wir statt $\Omega\mathcal{H}$ die Projektifizierung $\mathcal{H} := \Omega\mathcal{H}/\mathbb{C}^\times$ betrachten. In Kapitel II, das in Zusammenarbeit mit Matteo Costantini und Martin Möller entstanden ist, beweisen wir eine Formel für den vollständigen Cherncharakter des logarithmischen Kotangentenbündels einer linearen Untermannigfaltigkeit in Theorem II.1.2. Das erlaubt es uns insbesondere, eine geschlossene Formel für die Eulercharakteristik einer linearen Untermannigfaltigkeit anzugeben. Für eine lineare Untermannigfaltigkeit \mathcal{H} sei $\xi_{\mathcal{H}} = c_1(\mathcal{O}(-1))$ die erste Chernklasse des tautologischen Bündels.

THEOREM 2.1 (Theorem II.1.3). *Sei $\mathcal{H} \rightarrow \mathbb{P}\Omega\mathcal{M}_{g,n}(\mu)$ eine projektifizierte lineare Untermannigfaltigkeit. Die orbifaltige Eulercharakteristik von \mathcal{H} ist gegeben durch*

$$\chi(\mathcal{H}) = (-1)^d \sum_{L=0}^d \sum_{\Gamma \in \text{LG}_L(\mathcal{H})} \frac{K_\Gamma^{\mathcal{H}} \cdot N_\Gamma^\top}{|\text{Aut}_{\mathcal{H}}(\Gamma)|} \cdot \prod_{i=0}^{-L} \int_{\mathcal{H}_\Gamma^{[i]}} \xi_{\mathcal{H}_\Gamma^{[i]}}^{d_{\mathcal{H}_\Gamma^{[i]}}},$$

wobei die Integrale über die Normalisierung der Abschlüsse $\overline{\mathcal{H}} \rightarrow \mathbb{P}\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ im Modulraum der Multiskalendifferentiale sind und ähnliche Integrale über Randstrata, wobei

- $\mathcal{H}_\Gamma^{[i]}$ die linearen Untermannigfaltigkeiten auf Level i von Γ wie in Abschnitt II.3.5 definiert,
- $d_\Gamma^{[i]} := \dim(\mathcal{H}_\Gamma^{[i]})$ die projektifizierten Dimensionen sind,
- $K_\Gamma^{\mathcal{H}}$ das Produkt der Anzahl der Zackenpaarungen für jede Kante von Γ , die tatsächlich in der linearen Untermannigfaltigkeit $\overline{\mathcal{H}}$ enthalten sind, ist,
- $\text{Aut}_{\mathcal{H}}(\Gamma)$ die Menge der Automorphismen des Graphen Γ , dessen induzierte Operationen auf einer Umgebung von $D_\Gamma^{\mathcal{H}}$ den Raum $\overline{\mathcal{H}}$ erhält, ist,
- $d := \dim(\mathcal{H})$ die projektifizierte Dimension ist.

Für die meisten Begriffe, die in obigem Theorem benutzt werden, verweisen wir den Leser auf Kapitel II. Als Einziges wollen wir an dieser Stelle hervorheben, dass das Theorem den Modulraum der Multiskalendifferentiale $\mathbb{P}\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ benutzt. Diese Kompaktifizierung des projektifizierten Stratum $\mathbb{P}\Omega\mathcal{M}_{g,n}(\mu)$ wurde von Bainbridge-Chen-Gendron-Grushevsky-Möller [BCGGM18] [BCGGM19b] konstruiert. Die Objekte im Rand bestehen grob gesprochen aus nodalen Kurven mit einem Differential auf jeder irreduziblen Komponente und einer Levelstruktur auf den irreduziblen Komponenten. Für mehr Details verweisen wir auf Abschnitt I.3. Der Rand dieser Kompaktifizierung lässt seinerseits wieder eine Stratifikation zu, wobei die Strata mit sogenannten angereicherten Levelgraphen indiziert sind: Diese sind die dualen Graphen der zugrundeliegenden nodalen Kurven zusammen mit einer Dekoration die Informationen über die Differentiale und die Levelstruktur festhält. Diese Kompaktifizierung wird im Folgenden mehrmals auftauchen.

3. Der gotische Lokus

Der gotische Lokus $\Omega G \subseteq \Omega\mathcal{M}_{4,6}(0^3, 2^3)$ ist der Bahnabschluss der Entfaltungsflächen aller Vierecke mit Winkeln $(\frac{1}{6}\pi, \frac{1}{6}\pi, \frac{1}{6}\pi, \frac{3}{2}\pi)$. Sein Name rührt von dem Fakt, dass er die Fläche in Abbildung 7 enthält, die dem Grundriss einer gotischen Kathedrale ähnelt. Wie

oben erwähnt, war der gotische Lokus ΩG die erste bekannte primitive \mathbb{R} -lineare Untermannigfaltigkeit von Rang 2. Dieser Lokus hat zusätzliche überraschende Eigenschaften: Er enthält eine dichte Teilmenge primitiver Teichmüllerkurven, und er kommt mit einer natürlichen Abbildung nach $\mathcal{M}_{1,3}$ und sein Bild unter dieser Abbildung ist eine geodätische Fläche bezüglich der Teichmüllermetrik.




ABBILDUNG 7. Die gotische Kathedrale (gegenüberliegende Seiten sind identifiziert soweit nicht anders angegeben)

Eine Teichmüllerkurve in einem Stratum meromorpher Differentiale heißt *offensichtlich*, falls sie der Schnitt einer Überlagerungskonstruktion mit einer Bedingung an die Residuen ist. In Kapitel III werden wir den Rand des Abschlusses $\mathbb{P}\Xi\overline{G} := \overline{\mathbb{P}\Omega G} \subseteq \overline{\mathbb{P}\Xi\mathcal{M}_{4,6}(0^3, 2^3)}$ analysieren. Diese Analyse ist noch nicht abgeschlossen, aber wir präsentieren im Folgenden einige Teilergebnisse. Als Teil des Randes finden wir ein Beispiel für eine nicht-offensichtliche Teichmüllerkurve.

THEOREM 3.1 (Theorem III.1.3). Sei $(X, \omega) \subseteq \Omega\mathcal{M}_{1,6}(-3^2, 2^3)$ die kanonische Überlagerung des 6-Differentials vom Typ $(-10, -5, 3)$. Das Differential (X, ω) erzeugt eine nicht-offensichtliche Teichmüllerkurve. In der Karte in Abbildung 8 ist diese Teichmüllerkurve ausgeschnitten durch die Gleichungen

$$w_i = -w_{i+3} \quad \text{für } i = 1, 2, 3 \quad \text{und} \quad w_1 + w_3 + w_5 = 0.$$




ABBILDUNG 8. Eine unendliche Fläche, die eine nicht-offensichtliche Teichmüllerkurve im Stratum $\Omega\mathcal{M}_{1,6}(-3^2, 2^3)$ erzeugt

Bis jetzt ist es uns nicht gelungen, genau zu bestimmen, welche Randstrata vom gotischen Lokus geschnitten werden, aber wir verfügen über partielle Informationen in diese Richtung. Für die horizontalen Strata zeigen wir:

PROPOSITION 3.2 (Proposition III.1.4). Der gotische Lokus $\mathbb{P}\Xi\overline{G}$ schneidet nur die horizontalen Strata die in Abbildung 9 gelistet sind.

Wir erinnern daran, dass der gotische Lokus eine unendliche Zahl an primitiven Teichmüllerkurven enthält. Diese Kurven sind nicht kompakt. Daher schneidet der Abschluss dieser Kurven den Rand von $\mathbb{P}\Xi\overline{G}$ in einigen Punkten, genannt Spitzen. Solche Spitzen können nur in rein horizontalen Randstrata enthalten sein. Für die Strata die zu den angereicherten Levelgraphen in Abbildung 9 korrespondieren zeigen wir:

PROPOSITION 3.3 (Proposition III.1.6). *Das Innere jedes der vier Strata $D_{\Gamma_1}^G$, $D_{\Gamma_2}^G$, $D_{\Gamma_3}^G$ und $D_{\Gamma_{20}}^G$ enthält Spitzen einer im gotischen Lokus ΩG enthalten Teichmüllerkurve. Das Innere des Stratums $D_{\Gamma_{19}}^G$ enthält Spitzen einer nicht-primitiven Teichmüllerkurve.*




ABBILDUNG 9. Die rein horizontalen Randstrata im gotischen Lokus

Man könnte hoffen, mittels Theorem 2.1 die Eulercharakteristik des gotischen Lokus $\mathbb{P}\Omega G$ zu bestimmen. Dieses Theorem kann so umformuliert werden, dass es reicht, die Fundamentalklasse des Bildes des gotischen Lokus in $\overline{\mathcal{M}}_{4,6}$ zu kennen. In Kapitel III werden wir einen möglichen Ansatz zur Bestimmung dieser Klasse skizzieren. Momentan fehlt uns das nötige Handwerkszeug, um die Rechnung tatsächlich auszuführen.

4. Strata von k -Differentialen

Wir betrachten eine Verallgemeinerung des obigen Settings: Anstelle von Strata von abelschen Differentialen können wir auch Strata von k -Differentialen $\Omega^k \mathcal{M}_{g,n}(\mu)$ betrachten, die Paare (X, η) parametrisieren, wobei η ein Schnitt von $\Omega^{\otimes k}(X)$ ist. Hierbei ist μ eine ganzzahlige Partition von $k(2g - 2)$. Wie im abelschen Fall, gibt es eine Kompaktifizierung dieser Strata, den Modulraum der k -Multiskalendifferentiale. Bainbridge-Chen-Gendron-Grushevsky-Möller [BCGGM19a] und Costantini-Möller-Zachhuber [CMZ19] haben in diese Richtung gearbeitet. In Kapitel II präzisieren wir die genau Struktur als Orbifaltigkeit dieser Räume. Diese Räume stehen, vermöge einer Überlagerungskonstruktion, in Beziehung mit mit linearen Untermannigfaltigkeiten (nicht notwendiger Weise \mathbb{R} -linear), und wir geben in Korollar II.1.5 eine geschlossene Formel für ihre Eulercharakteristik.

Diese Formel haben wir in einem Sage-Paket genannt `diffstrata` implementiert, das Teil des Paketes `admcycles` [DSZ21] ist. `Diffstrata` wurde ursprünglich von Costantini-Möller-Zachhuber [CMZ23] entwickelt, um deren Formel für die Chernklassen von Strata von abelschen Differentialen auszuwerten [CMZ22]. Wir haben `diffstrata` erweitert, so dass es nun mit allen Strata von k -differentialen arbeitet. Als Beispiel kann die Eulercharakteristik und das Masur-Veech-Volumen des Stratums $\mathbb{P}\Omega^2 \mathcal{M}_{2,2}(-1, 5)$ mit den folgenden Befehlen berechnet werden.

```

sage: from admcycles.diffstrata import Stratum
sage: X = Stratum((-1,5), k=2)
sage: X.euler_characteristic()
-7/15
sage: X.masur_veech_volume()
28/135*pi^4

```

Die Eulercharakteristiken der minimalen Strata in Geschlecht 2 sind, für kleine k , in Tabelle 1 gelistet. Wie das obige Beispiel schon zeigt, kann das Paket `diffstrata` mehr als nur die Eulercharakteristik auszurechnen. Zum Beispiel ist es möglich

- alle nicht-horizontalen Randstrata eines Stratums von k -Differentialen aufzulisten,
- beliebige Schnittprodukte im vertikalen tautologischen Ring (d.h. den Ring erzeugt von allen nicht-horizontalen Strata, ψ - und κ -Klassen) zu berechnen,
- das Pushforward von Klassen im Stratum zum Modulraum der markierten stabilen Kurven zu berechnen.

Die größte Einschränkung der `diffstrata` momentan unterliegt, ist die Tatsache, dass es nicht mit horizontalen Kanten umgehen kann. Dies würde deutlich allgemeinere Berechnungen erlauben, wie z.B. das Pullback von beliebigen Klassen aus dem tautologischen Ring des Modulraums der markierten stabilen Kurven zu einem Stratum von k -Differentialen.

k	1	3	4	5	6	7	8	9
$\chi(\mathbb{P}\Omega^k \mathcal{M}_{2,1}(2k))$	$-\frac{1}{40}$	$\frac{1}{3}$	$\frac{3}{2}$	$\frac{21}{5}$	9	18	30	51

TABELLE 1. Eulercharakteristiken einiger minimaler Strata von k -Differentialen, berechnet mit `diffstrata`

Als eine Anwendung unserer Resultate über die Chernklassen zeigen wir, dass für bestimmte Typen μ der Raum $\Omega^k \mathcal{M}_{g,n}(\mu)$ birational äquivalent zu einem Quotienten der komplexen Einheitskugel ist.

THEOREM 4.1 (Theorem II.1.7). *Angenommen $\mu = (-a_1, \dots, -a_5)$ ist ein Tupel mit $a_i \geq 0$ so dass*

$$\left(1 - \frac{a_i}{k} - \frac{a_j}{k}\right)^{-1} \in \mathbb{Z} \quad \text{falls } a_i + a_k < k$$

für all $i \neq j$. Dann gibt es einen birationalen Kontraktionsmorphismus $\mathbb{P}\Xi^k \overline{\mathcal{M}}_{0,5}(\mu) \rightarrow \overline{\mathfrak{B}}$ auf einen glatten eigentlichen DM-stack \mathfrak{B} für einen Ballquotienten \mathfrak{B} .

Diese Ballquotienten wurden in früheren Arbeiten von Deligne-Mostow [DM86] und Thurston [Thu98] mit anderen Methoden konstruiert.

5. Das tropische k -Hodgebündel

Analog zum abelschen Fall ist der Rand von $\mathbb{P}\Xi^k \overline{\mathcal{M}}_{g,n}(\mu)$ indiziert mit angereicherten k -Levelgraphen. Man kann hoffen, die Struktur des Randes zu verstehen, indem man die Graphen selbst als abstrakte Objekte studiert. Um über stetige (Un)degenerationen von Graphen reden zu können, versehen wir jede Kante mit einer reellen Länge. Das führt uns zur Definition einer tropischen Kurve.

DEFINITION 5.1. *Eine tropische Kurve ist ein zusammenhängender Graph mit reellen Kantenlängen und Gewichten $g: V \rightarrow \mathbb{N}$ an jedem Knoten.*

Die Levelstruktur eines angereicherten Levelgraphen können wir erfassen, indem wir jeder Kante eine ganzzahlige Steigung zuweisen. Die Nullstellen des Differential können mit zusätzlichen Halbkanten erfasst werden. In Kapitel I, das in Zusammenarbeit mit Felix Röhrlé entstanden ist, definieren wir das k -Hodgebündel $\mathbb{P}\Omega^k M_g^{\text{trop}}$ grob gesagt als den Raum aller tropischer Kurven mit Halbkanten und ganzzahligen Steigungen an den Kanten die bestimmte Kompatibilitätsbedingungen erfüllen. Diese Bedingungen sollen das

Verhalten von k -Differentialen widerspiegeln. Dieser Raum ist ein verallgemeinerter Kegelkomplex, aber nicht äquidimensional, siehe Theorem I.1.1.

Die Kluft zwischen der klassischen Welt, sprich $\mathbb{P}\Omega^k \mathcal{M}_g$, und der tropischen Welt, also $\mathbb{P}\Omega^k M_g^{\text{trop}}$, wird überbrückt von einem Prozess namens Tropikalisierung: Es gibt eine stetige Tropikalisierungsabbildung $\text{trop}_{\Omega^k} : \mathbb{P}\Omega^k \mathcal{M}_g^{\text{an}} \rightarrow \mathbb{P}\Omega^k M_g^{\text{trop}}$. Allerdings ist diese Abbildung nicht surjektiv. Das Bild von trop_{Ω^k} ist der *Realisierbarkeitslokus*. In Theorem I.1.4 geben wir ein kombinatorisches Kriterium um zu entscheiden, ob ein Element von $\mathbb{P}\Omega^k M_g^{\text{trop}}$ im Realisierbarkeitslokus enthalten ist. Das erlaubt es uns zu zeigen, dass der Realisierbarkeitslokus sehr viel schönere Eigenschaften hat als das tropische k -Hodgebündel selbst.

THEOREM 5.2 (Theorem I.1.5). *Für $k \geq 2$ hat der Realisierbarkeitslokus die Struktur eines verallgemeinerten Kegelkomplexes, dessen maximale Kegel alle von Dimension $(2 + 2k)(g - 1) - 1$ sind. Die Faser des Realisierbarkeitslokus über einem maximaldimensionalen Kegel in M_g^{trop} ist ein verallgemeinerter Kegelkomplex, dessen maximale Kegel alle von relativer Dimension $(2k - 1)(g - 1)$ sind.*

6. Pillowcase-Überlagerungen und sichtbare Lagrangesche

Zu einer Überlagerung von Riemannschen Flächen $f : X \rightarrow \mathbb{P}^1$, die über höchstens vier Punkten verzweigt ist, gibt es ein (bis auf Skalieren) eindeutiges quadratisches Differential η vom Typ (-1^4) auf \mathbb{P}^1 , so dass die einfachen Pole an den Verzweigungspunkten liegen. Wir können diese Differential zurückziehen, um ein Differential $q = f^* \eta$ auf X zu erhalten. Ein auf diese Art konstruiertes quadratisches Differential (X, q) heißt Pillowcase-Überlagerung. In Kapitel IV, welches in Zusammenarbeit mit Johannes Horn entstanden ist, untersuchen wir Riemannsche Flächen X , die mehrere quadratische Differentiale q_1, \dots, q_n zulassen, so dass

- der Verschwindungloki von q_1, \dots, q_n paarweise disjunkt sind,
- alle Paare (X, q_i) Pillowcase-Überlagerungen sind.

Wir nennen eine solche Riemannsche Fläche X eine *mehrfache Pillowcase-Überlagerung*. Wir sagen, dass eine Pillowcase-Überlagerung $f : X \rightarrow \mathbb{P}^1$ *uniform* ist, wenn jede Faser aus Verzweigungspunkten der gleichen Ordnung besteht.

THEOREM 6.1 (Theorem IV.5.2). *Für unendlich viele Geschlechter g gibt es mehrfache uniforme Pillowcase-Überlagerungen mit nur einfachen Nullstellen.*

Ein Beispiel einer mehrfachen Pillowcase-Überlagerung ist die Kleinsche Quartik. Man beachte, dass unserer Definition von mehrfacher Pillowcase-Überlagerung nicht verlangt, dass die q_i nicht-isomorph sind. Nichtsdestotrotz geben wir auch ein Beispiel für eine mehrfache Pillowcase-Überlagerung an, bei der die quadratischen Differentiale nicht alle isomorph sind.

Dies hat eine schöne Anwendung in der Theorie von Higgs-Bündeln. Für eine komplexe reductive Gruppe G sei \mathcal{M}_G der Modulraum der G -Higgsbündel und $\text{Hit} : \mathcal{M}_G \rightarrow \mathcal{B}_G$ die Hitchinabbildung. Eine komplexe Lagrangesche $\mathcal{L} \subseteq \mathcal{M}_G$ heißt *sichtbar*, wenn die Restriktion der Hitchinabbildung über eine eigentliche Untervarietät $\mathcal{B}' = \text{Hit}(\mathcal{L}) \subsetneq \mathcal{B}$ faktorisiert. Für den Spezialfall $G = \text{SL}(2, \mathbb{C})$ beweisen wir:

THEOREM 6.2 (Theorem IV.1.2). *Sei $q \in H^0(X, K_X^2)$ ein quadratisches Differential mit nur einfachen Nullstellen. Dann gibt es eine sichtbare Lagrangesche*

$$\mathcal{L} \rightarrow \mathcal{B}' = \{tq \mid t \in \mathbb{C}\} \subset \mathcal{B}_{\text{SL}(2, \mathbb{C})}(X)$$

genau dann, wenn (X, q) eine Pillowcase-Überlagerung ist.

Insbesondere geben unsere Beispiele von mehrfachen Pillowcase-Überlagerungen Beispiele von Riemannschen Flächen für die es mehrere Geraden in der $\text{SL}(2, \mathbb{C})$ -Hitchinbasis $\mathcal{B}_{\text{SL}(2, \mathbb{C})}$ gibt, die zu sichtbaren Lagrangeschen gehören.