



Universiteit
Leiden
The Netherlands

A tale of two moduli spaces: logarithmic and multi-scale differentials

Chen, D.; Grushevsky, S.; Holmes, D.S.T.; Möller, M.; Schmitt, J.

Citation

Chen, D., Grushevsky, S., Holmes, D. S. T., Möller, M., & Schmitt, J. (2025). A tale of two moduli spaces: logarithmic and multi-scale differentials. *Épijournal De Géométrie Algébrique*, 9. doi:10.46298/epiga.2025.11278

Version: Publisher's Version

License: [Creative Commons CC BY-SA 4.0 license](#)

Downloaded from: <https://hdl.handle.net/1887/4285498>

Note: To cite this publication please use the final published version (if applicable).

A tale of two moduli spaces: Logarithmic and multi-scale differentials

Dawei Chen, Samuel Grushevsky, David Holmes, Martin Möller, and Johannes Schmitt

Abstract. Multi-scale differentials were constructed by M. Bainbridge, D. Chen, Q. Gendron, S. Grushevsky, and M. Möller, from the viewpoint of flat and complex geometry, for the purpose of compactifying moduli spaces of curves together with a differential with prescribed orders of zeros and poles. Logarithmic differentials were constructed by S. Marcus and J. Wise, as a generalization of stable rubber maps from Gromov–Witten theory. Modulo the global residue condition that isolates the main components of the compactification, we show that these two kinds of differentials are equivalent, and establish an isomorphism of their (coarse) moduli stacks. Moreover, we describe the rubber and multi-scale spaces as an explicit blowup of the moduli space of stable pointed rational curves in the case of genus zero, and as a global blowup of the incidence variety compactification for arbitrary genera, which implies their projectivity. We also propose a refined double ramification cycle formula in the twisted Hodge bundle which interacts with the universal line bundle class.

Keywords. Multi-scale differentials, logarithmic rubber maps, moduli spaces

2020 Mathematics Subject Classification. 14H10, 14A21, 14D23, 14H15, 32G15

Received by the Editors on May 5, 2023, and in revised form on October 16, 2024. Accepted on May 11, 2025.

Dawei Chen

Department of Mathematics, Boston College, Chestnut Hill, MA 02467, USA

e-mail: dawei.chen@bc.edu

Samuel Grushevsky

Simons Center for Geometry and Physics, and Department of Mathematics, Stony Brook University, Stony Brook, NY 11794-3636, USA

e-mail: sam@math.stonybrook.edu

David Holmes

Mathematisch Instituut, Universiteit Leiden, Einsteinweg 55, 2333 CC Leiden, Netherlands

e-mail: holmesdst@math.leidenuniv.nl

Martin Möller

Institut für Mathematik, Goethe-Universität Frankfurt, Robert-Mayer-Str. 6-8, 60325 Frankfurt am Main, Germany

e-mail: moeller@math.uni-frankfurt.de

Johannes Schmitt

ETH Zurich, Department Mathematik, Raemistrasse 101, CH-8092 Zurich, Switzerland

e-mail: johannes.schmitt@math.uzh.ch

Research of D. C. is supported in part by the National Science Foundation under the grants DMS-20-01040 and DMS-23-01030, as well as by the Simons Foundation through Travel Support for Mathematicians and a Simons Fellowship. Research of S. G. is supported in part by the National Science Foundation under the grants DMS-18-02116 and DMS-21-01631. Research of D. H. is supported by grants 613.009.103 and VI.Vidi.193.006 of the Dutch Research Council (NWO). Research of M. M. is supported by the DFG-project MO 1884/2-1 and by the Collaborative Research Centre TRR 326 ‘Geometry and Arithmetic of Uniformized Structures’. Research of J. S. is supported by the Swiss National Science Foundation under the SNF Early Postdoc Mobility grant SNF-184245 and the grant SNF-184613. Moreover, part of the research was conducted at the Max-Planck Institute for Mathematics in Bonn, which the author wants to thank warmly.

© by the author(s) This work is licensed under https://www.google.com/url?q=http://creativecommons.org/licenses/by-sa/4.0/&source=gmail-imap&ust=1761768736000000&usg=A0vVaw0rw3lF_tnxer7_sxg0y8Vn

Contents

1. Introduction	3
2. Logarithmic rubber maps	8
3. Generalized multi-scale differentials	15
4. The underlying algebraic stack of Rub	26
5. From logarithmic to multi-scale	31
6. The Hodge DR conjecture	39
7. Blowup descriptions	42
Appendix. Sign conventions	49
References	49

1. Introduction

1.1. Background and main results

Let $\mu = (m_1, \dots, m_n)$ be a tuple of integers with $\sum_{i=1}^n m_i = 2g - 2$. The (projectivized) *stratum of differentials of type μ* is the moduli space of smooth curves X of genus g with distinct marked points $z_1, \dots, z_n \in X$ such that $\sum_{i=1}^n m_i z_i$ is a (possibly meromorphic) canonical divisor.

The study of differentials with prescribed zeros and poles is important for at least two reasons. On the one hand, a (holomorphic) differential induces a flat metric with conical singularities at its zeros, such that the underlying Riemann surface can be realized as a polygon with edges pairwise identified by translations. Varying the shape of the polygons by affine transformations of the plane induces an action on the strata of differentials (called Teichmüller dynamics), whose orbit closures (called affine invariant subvarieties) govern intrinsic properties of surface dynamics. On the other hand, a differential (up to multiplication by a scalar) corresponds to a canonical divisor in the underlying complex curve. Hence the union of the moduli spaces of differentials with all possible configurations of zeros stratifies the Hodge bundle over the moduli space of curves, thus producing a number of remarkable questions to investigate from the viewpoint of algebraic geometry, such as compactification, enumerative geometry, and cycle class calculations. The interplay of these aspects has brought the study of differentials to an exciting new stage (see, *e.g.*, [Zor06, Wri15, Che17] and the references therein for an introduction to this subject).

As in many other moduli problems, having a geometrically meaningful compactification plays a crucial role in the study of the strata of differentials. Extending the setup of canonical divisors with prescribed zeros and poles to (pre-)stable curves, we define an algebraic stack $G\Xi\overline{\mathcal{M}}_{g,n}(\mu)$, the moduli space of *generalized simple multi-scale differentials of type μ* . The relative coarse moduli space \mathcal{GMS}_μ over $\overline{\mathcal{M}}_{g,n}$ of this stack is defined the same way as the multi-scale differentials in [BCG⁺19], but dropping the global residue condition.⁽¹⁾ Compared to the multi-scale space, the key player in [BCG⁺19], the stack $G\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ has additional irreducible components whose generic elements parameterize differentials on (strictly) nodal curves. Indeed

⁽¹⁾Our definition thus solves a task left open in [BCG⁺19], namely to describe the smooth stack $\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ dominating the stack of multi-scale differentials \mathcal{MS}_μ without invoking Teichmüller markings.

$G\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ maps onto the space of twisted canonical divisors constructed by Farkas–Pandharipande [FP18]. A minimal logarithmic structure on the space of twisted canonical divisors was described in [CC19], which extracts the information of meromorphic differentials from lower levels, but does not specify the full level structure. The precise definitions on these related objects are recalled in Section 3.

On the logarithmic side, Marcus and Wise [MW20] defined, for any line bundle \mathcal{L} on the universal curve $X_{g,n}$ over $\overline{\mathcal{M}}_{g,n}$, a space $\mathbf{Rub}_{\mathcal{L}}$ over $\overline{\mathcal{M}}_{g,n}$. The fiber of $\mathbf{Rub}_{\mathcal{L}}$ over a curve X is the set of piecewise linear functions β on the tropicalization of X , together with an isomorphism of line bundles from $\mathcal{O}_X(\beta)$ to \mathcal{L} . The natural \mathbb{C}^* quotient, which forgets the data of the isomorphism, is denoted by $\mathbb{P}(\mathbf{Rub}_{\mathcal{L}})$. When $\mathcal{L} = \mathcal{O}_{X_{g,n}}(\sum_i m_i z_i)$, this space $\mathbb{P}(\mathbf{Rub}_{\mathcal{L}})$ is the space of rubber maps to \mathbb{P}^1 with zeros and poles prescribed by the m_i , giving an alternative definition to that of Li, Graber and Vakil [Li01, GV05]. This machinery gives an extremely clean and functional definition of the double ramification cycle, as well as its logarithmic, pluricanonical, universal, and iterated variants; see [BHP⁺23, HS21, MPS23, MR24, HMP⁺25].

To connect this space with moduli of differentials, we define the line bundle

$$\mathcal{L}_{\mu} = \omega_{X_{g,n}/\overline{\mathcal{M}}_{g,n}} \left(- \sum_{i=1}^n m_i z_i \right)$$

on $X_{g,n}$, leading to the stack $\mathbf{Rub}_{\mathcal{L}_{\mu}}$, together with its relative coarse moduli space $\mathbf{Rub}_{\mathcal{L}_{\mu}}^{\text{coarse}}$ over $\overline{\mathcal{M}}_{g,n}$.⁽²⁾ The virtual fundamental class of $\mathbb{P}(\mathbf{Rub}_{\mathcal{L}_{\mu}}^{\text{coarse}})$ is the ‘canonical’ double ramification cycle described in [HS21].

The definitions of the spaces $\mathbf{Rub}_{\mathcal{L}_{\mu}}$ and \mathcal{GMS}_{μ} look very different. They can be found in Sections 2 and 3, respectively. The main aim of this paper is to show that these definitions are in fact *essentially equivalent*. More precisely, we prove the following theorem.

Theorem 1.1. *For any tuple of integers $\mu = (m_1, \dots, m_n)$ with $\sum_{i=1}^n m_i = 2g - 2$, there is an isomorphism of algebraic stacks over $\overline{\mathcal{M}}_{g,n}$*

$$F: \mathbf{Rub}_{\mathcal{L}_{\mu}} \longrightarrow G\Xi\overline{\mathcal{M}}_{g,n}(\mu)$$

which induces an isomorphism of the corresponding relative coarse moduli spaces over $\overline{\mathcal{M}}_{g,n}$

$$\overline{F}: \mathbf{Rub}_{\mathcal{L}_{\mu}}^{\text{coarse}} \longrightarrow \mathcal{GMS}_{\mu}.$$

Note that the global residue condition described in [BCG⁺18] can isolate the main component of \mathcal{GMS}_{μ} , called the *multi-scale* space and denoted by MS_{μ} . In other words, a generalized multi-scale differential not satisfying the global residue condition is not smoothable while preserving the prescribed zero and pole orders. (This global residue condition arises from applying Stokes’ theorem to subcurves of the limiting nodal curve when differentials degenerate from nearby smooth curves, thereby imposing that certain sums of residues at the nodes vanish. See *op. cit.* for further details.) Moreover, in [BCG⁺19] the space of multi-scale differentials MS_{μ} was shown to possess nice geometric properties, such as smoothness (as a stack), normal crossings boundary, and extension of the $\mathrm{GL}_2(\mathbb{R})$ -action to the boundary (after a real oriented blowup). It would be interesting to see whether these properties can be obtained directly by using rubber differentials and logarithmic geometry.

1.2. Applications and related topics

In what follows we address several constructions, results, and conjectures related to the main result above.

⁽²⁾See [AOV11] for the definition of relative coarse moduli spaces. Moreover, note that one can replace ω with any power $\omega^{\otimes k}$ in the formula for \mathcal{L}_{μ} , extending the theory to k -canonical divisors.

1.2.1. A blowup description of the space of multi-scale differentials.— First, describing a modular compactification via *blowups* can be useful in many aspects, *e.g.*, for projectivity and intersection calculations. There is a natural action of \mathbb{C}^* on generalized multi-scale differentials by simultaneous rescaling of all differentials, and we denote the quotient, the space of ‘projectivized generalized multi-scale differentials’, by $\mathbb{P}(\mathcal{GMS}_\mu)$; Lemma 1.1 induces an isomorphism $\mathbb{P}(\mathbf{Rub}_{\mathcal{L}_\mu}^{\text{coarse}}) \xrightarrow{\sim} \mathbb{P}(\mathcal{GMS}_\mu)$.

In the case of genus zero, we can identify $\mathbb{P}(\mathbf{Rub}_{\mathcal{L}_\mu}^{\text{coarse}})$ with a blowup of $\overline{\mathcal{M}}_{0,n}$.

Theorem 1.2 (Lemma 7.5). *For $g = 0$ there exists an explicit sheaf of ideals in $\overline{\mathcal{M}}_{0,n}$ such that the normalization of its blowup is $\mathbb{P}(\mathbf{Rub}_{\mathcal{L}_\mu}^{\text{coarse}})$.*

We recall that the projectivized stratum of differentials can be compactified in different ways. Firstly, one can consider simply its closure in the Deligne–Mumford compactification $\overline{\mathcal{M}}_{g,n}$. Secondly, one can consider the closure of the stratum in the total space of the projectivized Hodge bundle over $\overline{\mathcal{M}}_{g,n}$ (twisted by the polar parts). This compactification is described completely in [BCG⁺18], and is called the incidence variety compactification (IVC). The IVC clearly admits a morphism onto the Deligne–Mumford closure of the stratum, while $\mathbb{P}(\mathcal{MS}_\mu)$ maps onto the IVC, and in general both these morphisms are ‘forgetful’, *i.e.* contract some loci in the compactifications. We further write NIVC for the normalization of the IVC.

In [Ngu24] Nguyen showed that, in the case of genus zero, the IVC can be described as an explicit blowup of $\overline{\mathcal{M}}_{0,n}$. From the above theorem, one can also retrieve Nguyen’s result, which we do in Lemma 7.7.

In arbitrary genus, recall that the multi-scale space \mathcal{MS}_μ is the main component of \mathcal{GMS}_μ , whose generic element parameterizes differentials with prescribed zero and pole orders on smooth curves.

Theorem 1.3 (Lemma 7.8). *For arbitrary genus there exists a global sheaf of ideals on the NIVC such that the normalization of the blowup of the NIVC along this ideal gives the projectivized multi-scale space $\mathbb{P}(\mathcal{MS}_\mu)$. Consequently, the coarse moduli space of the stack $\mathbb{P}(\mathcal{MS}_\mu)$ is a projective variety.*

In [BCG⁺19] a local blowup construction to obtain $\mathbb{P}(\mathcal{MS}_\mu)$ from the normalization of the IVC was described. That construction did not glue to a global sheaf of ideals on IVC, and hence did not yield the projectivity of $\mathbb{P}(\mathcal{MS}_\mu)$. In [CCM24] the projectivity of $\mathbb{P}(\mathcal{MS}_\mu)$ was established by constructing an explicit ample divisor class on it. Thus the above theorem provides a distinct conceptual understanding of the projectivity result.

1.2.2. A Hodge double ramification cycle.— Next we propose a refined version of the *double ramification (DR) cycle* in the twisted Hodge bundle and conjecture a Pixton-style formula for this class, involving coefficients of higher powers of the regularizing parameter ‘ r ’. For this purpose, we also generalize our considerations to k -differentials.

Let $A = (a_1, \dots, a_n) \in \mathbb{Z}^n$, where $|A| := \sum_{i=1}^n a_i = k(2g - 2 + n)$ for some $k > 0$, and denote by

$$\mathcal{L}_A := \omega^{\otimes k} \left(- \sum_{i=1}^n (a_i - k) z_i \right)$$

the associated degree zero line bundle on $X_{g,n}$, where $\pi: X_{g,n} \rightarrow \overline{\mathcal{M}}_{g,n}$ is the universal curve with sections z_i , and ω is the relative canonical bundle.⁽³⁾ Taking

$$\mathcal{H} := \omega^{\otimes k} \left(- \sum_{\substack{i \in \{1, \dots, n\}: \\ a_i < 0}} a_i z_i + \sum_{i=1}^n k z_i \right)$$

⁽³⁾Here we switch to the logarithmic version of indices to match the notation in [JPPZ17]. In other words, as a signature of k -differentials, each of the zero and pole orders is given by $a_i - k$. In particular, by slight abuse of notation, \mathcal{L}_A is simply the bundle we denoted by \mathcal{L}_μ in the previous convention.

to be the relative logarithmic k -canonical bundle twisted by the polar part, we obtain a natural diagram

$$(1.1) \quad \begin{array}{ccc} \mathbb{P}(\mathbf{Rub}_{\mathcal{L}_A}) & \xrightarrow{F} & \mathbb{P}(\pi_*\mathcal{H}) \\ & \searrow p & \downarrow q \\ & & \overline{\mathcal{M}}_{g,n} \end{array}$$

(see the discussion leading to (6.3) for more details). Pushing forward the virtual fundamental class of $\mathbb{P}(\mathbf{Rub}_{\mathcal{L}_A})$ gives a lift

$$\widetilde{\mathrm{DR}}_A^k = F_* [\mathbb{P}(\mathbf{Rub}_{\mathcal{L}_A})]^{\mathrm{vir}}$$

of the twisted DR cycles to $\mathbb{P}(\pi_*\mathcal{H})$, which we call the twisted *Hodge DR cycle*.

Let $H = c_1(\mathcal{O}(1))$ be the universal line bundle class on $\mathbb{P}(\pi_*\mathcal{H})$, and let $\eta = F^*H$ be its pullback to $\mathbb{P}(\mathbf{Rub}_{\mathcal{L}_A})$.⁽⁴⁾ By the projective bundle formula associated to the map q , to determine the class of $\widetilde{\mathrm{DR}}_A^k$ in the Chow ring $\mathrm{CH}^\bullet(\mathbb{P}(\pi_*\mathcal{H}))$, it suffices to determine the cycle class

$$(1.2) \quad q_* \left(\widetilde{\mathrm{DR}}_A^k \cdot H^u \right) = p_* \left([\mathbb{P}(\mathbf{Rub}_{\mathcal{L}_A})]^{\mathrm{vir}} \cdot \eta^u \right) \in \mathrm{CH}^{g+u}(\overline{\mathcal{M}}_{g,n})$$

for every u .

Before proceeding to give a conjectural formula for these cycles, let us make a remark about the case $k = 0$. When trying to follow the construction above, we encounter the issue that in general the higher cohomology of \mathcal{H} will not vanish, so that $\mathbb{P}(\pi_*\mathcal{H})$ is not a projective bundle. In Section 6 we explain how this can be remedied. However, there is also an alternative approach to defining η , which makes clearer a connection to relative Gromov–Witten theory: there both the space $\mathbb{P}(\mathbf{Rub}_{\mathcal{L}_A})$ and its forgetful map p to $\overline{\mathcal{M}}_{g,n}$ still make sense, and it was proven in [BHP⁺23, Proposition 50] that there is a natural isomorphism

$$\mathbb{P}(\mathbf{Rub}_{\mathcal{L}_A}) \cong \overline{\mathcal{M}}_{g,A}(\mathbb{P}^1, 0, \infty)^\sim$$

with the space of stable maps to *rubber* \mathbb{P}^1 relative to $0, \infty$, with contact orders specified by the vector A . This space of stable maps parameterizes maps from prestable curves to a chain of rational curves, with marked points $0, \infty$ at opposite ends of the chain (see [JPPZ17, Section 0.2.4] for details). What is important for us is that it still carries a natural divisor class $\eta = \Psi_\infty$ defined as the class of the cotangent line bundle at the marked point ∞ on the chain of rational curves.

Continuing in the general case $k \geq 0$, consider the space of twisted r -spin structures $\overline{\mathcal{M}}_{g,A}^{r,k}$ constructed in [Chi08, Jar00]. This is a compactification of the moduli space of smooth marked curves X and line bundles L on X with an isomorphism $L^{\otimes r} \cong \omega_X^{\otimes k}(-\sum(a_i - k)z_i)$. In the compactification, the curve X is allowed to acquire nodal singularities that are stacky points with some finite stabilizer group, making X a *twisted curve* in the sense of [AV02]. The moduli space then carries a universal curve $\pi: \mathcal{X} \rightarrow \overline{\mathcal{M}}_{g,A}^{k,r}$ with a line bundle \mathcal{L} and isomorphism $\mathcal{L}^{\otimes r} \cong \omega_\pi^{\otimes k}(-\sum(a_i - k)z_i)$. Here we follow the notation of [JPPZ17]. Forgetting the line bundle and the stacky structure on \mathcal{X} gives a map $\epsilon: \overline{\mathcal{M}}_{g,A}^{k,r} \rightarrow \overline{\mathcal{M}}_{g,n}$. Define the following Chiodo's class as the first cycle class given in [JPPZ17, Proposition 5]:

$$\mathrm{Ch}_{g,A}^{k,r,d} := r^{2d-2g+1} \epsilon_* c_d(-R^*\pi_*\mathcal{L}) \in \mathrm{CH}^d(\overline{\mathcal{M}}_{g,n}),$$

where R^* denotes derived push-forward. It is a polynomial in r (for r sufficiently large). Following computations of Chiodo [Chi08], the class $\mathrm{Ch}_{g,A}^{k,r,d}$ can be computed explicitly as a sum over stable graphs, decorated with polynomials in κ and ψ -classes (see [JPPZ17, Corollary 4]). We propose the following conjecture, giving a formula for the cycle classes (1.2).

⁽⁴⁾In the literature sometimes ξ denotes the universal line bundle class on the space of k -differentials and η denotes the tautological line bundle class $c_1(\mathcal{O}(-1))$.

Conjecture 1.4 (Hodge DR). *For every $g, k, u \geq 0$ and every $A \in \mathbb{Z}^n$ satisfying $|A| = k(2g - 2 + n)$, the following relation holds:*

$$p_* \left(\left[\mathbb{P}(\mathbf{Rub}_{\mathcal{L}_A}) \right]^{\text{vir}} \cdot \eta^u \right) = [r^u] \text{Ch}_{g,A}^{k,r,g+u} \in \text{CH}^{g+u}(\overline{\mathcal{M}}_{g,n}),$$

where $[r^u]$ means taking the coefficient of r^u .

For $u = 0$, by definition the left-hand side of this equation is the usual twisted DR cycle DR_A^k , and by [JPPZ17, Proposition 5] the right-hand side agrees with Pixton's formula for this cycle. Therefore, the conjecture is true for $u = 0$ by the results of [BHP⁺23].

For $u > 0$, the conjecture can be verified computationally in many examples for the special case $g = 0$. Indeed, in this case, the space $\mathbf{Rub}_{\mathcal{L}_A}$ agrees with the space of multi-scale k -differentials by Lemma 1.1 (since the global residue condition is automatically satisfied in the case $g = 0$). Then the software package `diffstrata`, see [CMZ23], can compute powers of η on this space using relations in its Picard group, and express the left-hand side of the conjecture in terms of tautological classes. On the other hand, the right-hand side of the conjecture can be computed in `admcycles`, see [DSvZ21], using the graph-sum formula from [JPPZ17]. Using this, the prediction of the conjecture has been verified for several example vectors A , giving many non-trivial equalities in the Chow group of $\overline{\mathcal{M}}_{0,n}$. The calculations in `diffstrata` for $k > 1$ rely on some code in development related to the forthcoming paper [CMS23].

On the other hand, for $k = 0$, the left-hand side of the conjecture has been computed in [FWY21, Corollary 4.3]. The formula given there is similar, but not equal, to the one above. However, using properties of the Chiodo class proven in [GLN23, Theorem 4.1(ii)], a short computation shows that the formula from [FWY21] can be simplified to the one we give above.⁽⁵⁾

Theorem 1.5. *Lemma 1.4 is true for $k = 0$: for any $g, u \geq 0$ and any vector $A \in \mathbb{Z}^n$ with sum $|A| = 0$, we have*

$$p_* \left(\left[\mathbb{P}(\mathbf{Rub}_{\mathcal{L}_A}) \right]^{\text{vir}} \cdot \eta^u \right) = p_* \left(\left[\overline{\mathcal{M}}_{g,A}(\mathbb{P}^1, 0, \infty) \right]^{\text{vir}} \cdot \Psi_\infty^u \right) = [r^u] \text{Ch}_{g,A}^{0,r,g+u}.$$

1.3. Sketch of the comparison

We hope that this paper will foster more communications between two groups of researchers, those working in logarithmic geometry for moduli spaces and those working in moduli of differentials for Teichmüller dynamics. With this in mind, we have written out definitions of objects on both sides of the story in a rather detailed way, in particular assuming minimal background knowledge about logarithmic structures. We now give an overview of the comparison in Lemma 1.1.

The definition of generalized multi-scale differentials on a stable curve X is geometrically very concrete but quite lengthy. The *level structure* (or *full order*) on the vertices of the dual graph Γ of X , corresponding to the irreducible components of X , encodes the vanishing orders of a differential in a family of differentials on smooth curves that degenerates to a given multi-scale differential on a nodal curve. One can twist differentials that vanish identically, on the irreducible components of the same level, by a rescaling parameter for that level, to obtain *twisted differentials* that are not identically zero on the components on that level. A multi-scale differential contains the combinatorial data of the zero and pole orders of twisted differentials at the nodes. Moreover, the *prong-matchings* of a multi-scale differential are combinatorial data that arise from choices of smoothing a nodal differential with matching zero and pole orders at the two branches at a node, under the flat metric induced by the differential. Lastly, a multi-scale differential stores the smoothing parameters of the nodes in a way consistent with the level structure, packaged in the notion of a *rescaling ensemble*. On all these parameters, a certain *level rotation torus* acts and induces a notion of equivalence that forgets the extra information due to various choices being made in the above process, e.g., how simultaneously scaling twisted differentials on the same level affects prong-matchings.

⁽⁵⁾Special thanks go to Longting Wu for patiently explaining their formula and to Danilo Lewański for informing us of the above property of the Chiodo class.

The definition of an element of $\mathbf{Rub}_{\mathcal{L}_\mu}$ is very concise; it is simply a piecewise linear function on the tropicalization subject to certain conditions (see Definitions 2.1 and 2.7). However, it may seem cryptic at a first reading. In particular, it may not be immediately apparent why the data of a log curve, a piecewise linear function, and an isomorphism of line bundles should yield up all the above data of an equivalence class of multi-scale differentials. Some parts of the comparison (such as the enhanced level graph) are obtained essentially by some bookkeeping, but extracting the *level rotation torus* and *prong-matchings* from the logarithmic data requires significantly more care.

Our first key insight about prong-matchings is Lemma 3.1, giving a new, coordinate-free characterization of prong-matching via the residue isomorphism. The second key insight exhibits the reason for the equivalence relation given by the level rotation torus in log language. We define a *log splitting* of a point in $\mathbf{Rub}_{\mathcal{L}_\mu}(B)$ essentially as a section of the quotient map $M_B \rightarrow \overline{M}_B$, from the sheaf of monoids M_B built into the log structure, to the ghost sheaf \overline{M}_B . The precise statement is given in Lemma 5.1. We show that the set of log splittings is closely related to the level rotation torus, and in particular changing the choice of log splitting corresponds to the action of the level rotation torus.

Finally, we remark that an analogue of Lemma 1.1 should also hold for rubber k -differentials and multi-scale k -differentials. Indeed, on the logarithmic side the generalization is straightforward, as noted earlier. Moreover, the space of multi-scale k -differentials was described similarly in [CMZ24]. Thus the arguments in the current paper can be adapted directly to compare the two versions of k -differentials. We leave the details to the interested reader.

Outline of the paper

In Section 2 we give the basic definitions of logarithmic rubber maps, and in Section 3 we do the same for generalized multi-scale differentials. In the somewhat technical Section 4, we describe the underlying algebraic stack that comes from the logarithmic definition in Section 2, which will be essential for what follows. Section 5 is the technical heart of our comparison theorem, where we show how to construct a multi-scale differential from a logarithmic one, and vice versa. In Section 6 we discuss several constructions of the universal line bundle class η that appears in the Hodge DR conjecture and prove the conjecture in the case of $k = 0$. In Section 7 we describe some of the moduli spaces concerned via blowup constructions. Finally, the sign conventions generally adopted in the logarithmic and multi-scale worlds are unfortunately opposite to one another; in the appendix we explain a small variation on the logarithmic definitions which makes the signs match.

Acknowledgments

The authors would like to thank Quentin Gendron for many helpful discussions. We also thank Qile Chen, Honglu Fan, Danilo Lewański, Rahul Pandharipande, Adrien Sauvaget, Johannes Schwab, and Longting Wu for sharing comments and insights on a preliminary version of the current paper. This project was initiated during an online seminar organized by D. H. and J. S.; we would like to thank all the participants for making the seminar a success. Finally, we thank the anonymous referees for carefully reading the paper and providing many helpful comments that helped clarify some arguments and presentation.

2. Logarithmic rubber maps

In the following we recall the relevant notions from logarithmic geometry that are needed to define logarithmic rubber maps and make the comparison to the space of generalized multi-scale differentials. Instead of a broader introduction, we focus on the precise parts of the theory that are needed. We refer the reader to [Kat89b, Ogul8] for a more detailed treatment of the basic notions of log geometry, and mention more specialized references where appropriate later in the text.

2.1. Overview of log divisors

A log scheme is a pair

$$(2.1) \quad (B, \alpha: M_B \rightarrow \mathcal{O}_B),$$

where B is a scheme, M_B is a sheaf of monoids on B , and α is a map of monoids, where \mathcal{O}_B is equipped with the multiplicative monoid structure, and where we assume that α induces an isomorphism $\alpha^{-1}(\mathcal{O}_B^\times) \rightarrow \mathcal{O}_B^\times$. We write $\overline{M}_B := M_B/\alpha^{-1}(\mathcal{O}_B^\times)$; this is called the *ghost sheaf* or *characteristic sheaf*. We write the monoid operation on M_B as multiplication, and that on \overline{M}_B as addition. Recall that a monoid M is called *saturated* if the natural map $M \rightarrow M^{\text{gp}}$ to its *groupification* is injective, and if, for every $n \in \mathbb{Z}_{\geq 1}$ and $g \in M^{\text{gp}}$ with $ng \in M$, we have $g \in M$. A log structure is called saturated if all its stalks are saturated. We work throughout only with fine saturated log structures (log structures admitting charts; see [Ogu18, Section III.1.2] by finitely generated saturated monoids).

If $\beta \in \Gamma(B, \overline{M}_B^{\text{gp}})$, then the preimage of β in the short exact sequence

$$(2.2) \quad 1 \longrightarrow \mathcal{O}_B^\times \longrightarrow M_B^{\text{gp}} \longrightarrow \overline{M}_B^{\text{gp}} \longrightarrow 1$$

is a \mathbb{G}_m -torsor, which we denote by $\mathcal{O}_B^\times(\beta)$. We write $\mathcal{O}_B(\beta)$ for the associated line bundle (see the appendix for our sign convention here).

Following [Kat00], the formal definition of a *log curve* is a morphism of log schemes⁽⁶⁾ $\pi: X \rightarrow B$ that is proper, saturated, log smooth, and has geometric fibers which are reduced and of pure dimension 1. This definition is rarely important to us, so rather than explicating the terms involved, we present a crucial structure result (to be found in [Kat00, Section 1.8]). If $\pi: X \rightarrow B$ is a log curve, then the underlying map of schemes is a prestable curve, and if x is a geometric point of X mapping to a geometric point b of B , then exactly one of the following three cases holds:

- (1) x is a smooth point of X , and the natural map $\overline{M}_{B,b} \rightarrow \overline{M}_{X,x}$ is an isomorphism;
- (2) x is a smooth point of X , and $\overline{M}_{X,x} \cong \overline{M}_{B,b} \oplus \mathbb{N}$ with the natural map $\overline{M}_{B,b} \rightarrow \overline{M}_{X,x}$ corresponding to the inclusion of the first summand (in this case we say x is a marked point, and we choose a total ordering on our markings to be compatible with the standard definition of marked prestable curves);
- (3) x is not a smooth point of the fiber X_b (i.e. x is a node), and there exist a unique element $\delta_x \in \overline{M}_{B,b}$ and an isomorphism

$$(2.3) \quad \overline{M}_{X,x} \cong \left\{ (u, v) \in \overline{M}_{B,b}^2 \text{ such that } \delta_x \text{ divides } u - v \right\}.$$

We warn the reader that the ghost sheaf \overline{M}_X does *not* fully determine the log structure; the units contain important additional information.

We write \mathfrak{N} for the fibered category over **LogSch** whose objects are log curves X/B , with the projection taking X/B to B . This is representable by an algebraic stack with log structure, see [GS13, Appendix A], generalizing the construction of [Kat00] in the stable case. As shown in those references, the underlying algebraic stack of \mathfrak{N} is naturally isomorphic to the stack of prestable curves. The stack \mathfrak{N} naturally contains all $\overline{\mathcal{M}}_{g,n}$ as open substacks, by equipping a stable curve X/B with its basic log structure (see [Kat00, GS13]), equivalently, with the log structure coming from the boundary divisor.

Given a log scheme, we define

$$\mathbb{G}_m^{\text{trop}}(B) := \Gamma(B, \overline{M}_B^{\text{gp}}),$$

which we call the tropical multiplicative group. It can naturally be extended to a presheaf $\mathbb{G}_{m,B}^{\text{trop}}$ on the category **LogSch** _{B} of log schemes over B , and admits a log smooth cover by log schemes (with subdivision $[\mathbb{P}^1/\mathbb{G}_m]$; see [MW20, Section 4.1].

⁽⁶⁾The reader concerned about the case $g = 1$, $n = 0$ should rather take log algebraic spaces.

Definition 2.1. We define **Rub** as the stack in groupoids over \mathfrak{M} , with objects being tuples

$$(\pi: X \rightarrow B, \beta: X \rightarrow \mathbb{G}_{m,B}^{\text{trop}})$$

with X/B a log curve, satisfying two conditions on each strict geometric fiber:

- (1) The image of β is fiberwise totally ordered,⁽⁷⁾ with largest element 0.
- (2) Writing R for the stack obtained from $\mathbb{G}_m^{\text{trop}}$ by subdividing at the image of β , we require that the fiber product $X \times_{\beta, \mathbb{G}_m^{\text{trop}}} R$ is a log curve.

The morphisms are defined by pullback.

Over a given geometric point of B , write $N + 1$ for the cardinality of the image of β ; since the latter is totally ordered, there is a unique isomorphism τ of totally ordered sets between the image of β and $\{0, -1, \dots, -N\}$. The composition

$$(2.4) \quad \ell := \tau \circ \beta$$

is then called the *normalized level function* associated with β .

Remark 2.2. This definition will be unpacked in Section 2.3, but for now we make a couple of remarks on how it differs from that given in Marcus–Wise [MW20]. Firstly, they declare the image of β to have *smallest* element 0; this makes no material difference, and the reason for our choice of conventions is explained in the appendix.

More significantly, condition (2) is not stated by Marcus and Wise. However, it is assumed, for example in datum (R1) in Section 5.5 of their paper. Most of their results go through without this condition, but it is necessary for making a connection to the spaces of rubber maps of Li, Graber–Vakil, *etc.*, and is also necessary for the comparison results in the present paper.

In fact, dropping condition (2) (and thus passing to the space of Marcus and Wise) is exactly the same as taking the coarse moduli space of **Rub** relative to \mathfrak{M} . We write \mathbf{Rub}^{MW} for this space. Combining Theorem 4.3.2 and Proposition 5.1.2 of [MW20] shows that the space \mathbf{Rub}^{MW} is an algebraic space over the relative Picard stack over \mathfrak{M} . However, because the line bundle $\mathcal{O}(\beta)$ is canonically trivial along the locus where β takes value 0, the map from \mathbf{Rub}^{MW} to the relative Picard stack factors via the relative Picard space, so that \mathbf{Rub}^{MW} is an algebraic space over \mathfrak{M} . On the other hand, $\mathbf{Rub} \rightarrow \mathbf{Rub}^{\text{MW}}$ is a root stack (see the proof of Lemma 2.4 for more details), and so the relative coarse space of **Rub** is exactly \mathbf{Rub}^{MW} .

Theorem 2.3 (*cf.* [MW20]). *The category **Rub** is a log algebraic stack locally of finite presentation.*

Marcus and Wise prove this in the absence of condition (2) above, but imposing this condition simply corresponds to a root stack construction, and does not affect the result. One benefit of imposing condition (2) is the following theorem, which did not hold for the version of **Rub** considered by Marcus and Wise.

Theorem 2.4. *The algebraic stack **Rub** is smooth.*

The proof of Lemma 2.4 will be given in Section 4.3.

Given $\beta \in \overline{M}_X^{\text{gp}}(X)$, then taking the preimage in the standard exact sequence (2.2) applied to X yields the line bundle $\mathcal{O}_X(\beta)$; in other words, it yields an Abel–Jacobi map

$$\text{aj}: \mathbf{Rub} \rightarrow \mathfrak{Pic}$$

to the Picard stack of the universal curve over \mathfrak{M} (the stack of pairs $(X/B, \mathcal{F})$, where X/B is a log curve and \mathcal{F} is a line bundle on X). One of the main results of [MW20] is that the composite of this Abel–Jacobi map with the forgetful map $\mathfrak{Pic} \rightarrow \text{Pic}$ to the relative Picard space is proper.

⁽⁷⁾Here we mean that for any two elements in the image of β , one of their differences is contained in \overline{M}_B .

Definition 2.5. Write n for the locally constant function on \mathfrak{M} giving the number of markings. Then there is an *outgoing slopes* map

$$\mathbf{Rub} \longrightarrow \mathbb{Z}^n$$

sending a point $(X/B, \beta)$ to the outgoing slopes of β , *i.e.* the values of β in the groupifications of the stalks $\overline{M}_{X/B}(z_i) := \overline{M}_X(z_i)/\overline{M}_B(\pi(z_i)) = \mathbb{N}$ at the markings.

Given a tuple $\mu = (m_1, \dots, m_n)$ of integers, we define \mathbf{Rub}_μ to be the open-and-closed substack of \mathbf{Rub} where the log curve has n markings and the outgoing slopes are given by μ .

Note that the forgetful map from \mathbf{Rub}_μ to the locus in \mathfrak{M} where the curves carry exactly n markings is birational (it is an isomorphism over the locus of smooth curves); in particular, if we fix a genus and a number of markings, then \mathbf{Rub}_μ is connected.

Writing $d := \sum_{i=1}^n m_i$, the image of \mathbf{Rub}_μ under the Abel–Jacobi map aj lands in the connected component Pic^d of Pic consisting of line bundles of (total) degree d .

Remark 2.6. In fact one can show that the map $\mathbf{Rub}_\mu \rightarrow \mathfrak{M}$ is not only birational onto the locus in \mathfrak{M} with n markings but also *log étale*. This is a type of map basically consisting of an iterated blowup of boundary strata, followed by root constructions⁽⁸⁾ on some of these strata, and then followed by taking an étale map. For the details, we refer the reader *e.g.* to the paper [HMP⁺25], where such morphisms are used extensively. An important point there is that they can be described uniquely by an (incomplete) subdivision of the tropicalization of \mathfrak{M} . While again we do not explain the details, one consequence is that one can obtain a smooth local model of the morphism $\mathbf{Rub}_\mu \rightarrow \mathfrak{M}$ by the toric map induced via some explicit subdivision of a cone.

In Figure 1 we use this to illustrate the importance of condition (2) in Lemma 2.1. For this, consider a point of \mathfrak{M} where the curve has genus zero and the stable graph Γ has three vertices and two edges e_1, e_2 as illustrated. Assume that each vertex carries one marking and that μ is chosen so that the unique slopes of a piecewise linear function on the edges are 1, 2 for e_1, e_2 (see Lemma 2.10 for a discussion of piecewise linear functions).

Then the tropicalization of \mathfrak{M} contains a cone $\sigma_\Gamma = \mathbb{R}_{\geq 0}^2$ parameterizing the ways of putting lengths ℓ_1, ℓ_2 on the two edges. Depending on which of the quantities ℓ_1 or $2\ell_2$ is greater, a piecewise linear function on Γ with the given slopes will take a larger value on either v_2 or v_1 . Then the smooth local picture of $\mathbf{Rub}_\mu \rightarrow \mathfrak{M}$ is given by the map of toric varieties associated to the subdivision of σ_Γ along the ray spanned by $(\ell_1, \ell_2) = (2, 1)$.

However, there is a subtlety: for the standard integral structure (black dots), the upper cone is simplicial, but not smooth. Indeed, the primitive generators $(0, 1), (2, 1)$ of its rays form a rational basis, but not an integral basis. Hence, the toric variety associated to this cone has a singularity, which would contradict Lemma 2.4. And indeed, this is precisely what happens for the variant of \mathbf{Rub} defined by omitting condition (2) from Lemma 2.1. Putting this condition forces us to adjoin the element $\ell_1/2$ to the dual of the lattice on the upper cone. Correspondingly, on that upper cone we take the sublattice of all points (ℓ_1, ℓ_2) with ℓ_1 an even integer (depicted by circled dots).⁽⁹⁾ Then the new ray generators are $(0, 1/2), (1, 1/2)$, which indeed form a basis of the integral structure $\mathbb{Z} \oplus (1/2)\mathbb{Z}$, so that \mathbf{Rub} is smooth as claimed.

2.2. Logarithmic rubber differentials

The stack \mathbf{Rub} is in some sense the universal space of logarithmic rubber maps. In this section we specialize to the case of logarithmic rubber differentials. For this we fix g, n and write $X_{g,n}/\overline{\mathcal{M}}_{g,n}$ for

⁽⁸⁾We cannot assume that the degrees of these roots are invertible in the base ring, so this map should probably not be considered to be log étale outside of characteristic zero (though conventions in the literature differ).

⁽⁹⁾Note that in contrast to the toric situation, not all cones in the tropicalization of \mathbf{Rub} lie in the same ambient vector space with integral structure, so that it is possible to change this integral structure on different cones of the tropicalization.




Figure 1. Subdivision associated to the drawn stable graph, with slope 1 at edge e_1 (of length ℓ_1) and slope 2 at edge e_2 (of length ℓ_2).

the universal curve, with markings $z = (z_1, \dots, z_n)$. Fix a tuple $\mu = (m_1, \dots, m_n)$ of integers such that $d = \sum_{i=1}^n m_i = 2g - 2$. We define a line bundle on the universal curve $X_{g,n}$ over $\overline{\mathcal{M}}_{g,n}$ by the formula

$$\mathcal{L} := \mathcal{L}_\mu := \omega_{X_{g,n}/\overline{\mathcal{M}}_{g,n}} \left(- \sum_{i=1}^n m_i z_i \right),$$

where $\omega = \omega_{X_{g,n}/\overline{\mathcal{M}}_{g,n}}$ is the relative dualizing sheaf of $X_{g,n} \rightarrow \overline{\mathcal{M}}_{g,n}$. Then \mathcal{L} induces a morphism

$$\varphi_{\mathcal{L}}: \overline{\mathcal{M}}_{g,n} \rightarrow \text{Pic}.$$

Definition 2.7. We define the *space of logarithmic rubber differentials* to be

$$(2.5) \quad \mathbf{Rub}_{\mathcal{L}} := \mathbf{Rub}_{\underline{0}} \times_{\text{Pic}, \varphi_{\mathcal{L}}} \overline{\mathcal{M}}_{g,n}.$$

Remark 2.8. If we had taken the fiber product over the relative Picard space (instead of the Picard stack), we would have obtained the projectivized space $\mathbb{P}(\mathbf{Rub}_{\mathcal{L}})$. This is the approach taken in [MW20, BHP⁺23], as the space $\mathbb{P}(\mathbf{Rub}_{\mathcal{L}})$ is what is needed for the study of the double ramification cycle.

Remark 2.9. There are two equivalent descriptions of the rubber differential space as

$$\mathbf{Rub}_{\mathcal{L}} = \mathbf{Rub}_{\underline{0}} \times_{\text{Pic}, \varphi_{\mathcal{L}}} \overline{\mathcal{M}}_{g,n} = \mathbf{Rub}_{\mu} \times_{\text{Pic}, \varphi_{\omega}} \overline{\mathcal{M}}_{g,n}.$$

2.3. Local description

In what follows we will make the definition of the space **Rub** more explicit for log curves over ‘sufficiently small’ bases, more precisely, for *nuclear* log curves as defined in [HMOP23]. This is a slight refinement of asking for the base to be atomic (in the sense of [AW18]), and is needed because a log curve even over a point does not have a well-defined dual graph unless the residue field is sufficiently large. We omit the details of the definition of a nuclear log curve, mentioning only the key properties we use:

- (1) For any family of log curves X/B with B locally of finite type, there exists a strict⁽¹⁰⁾ étale cover $\bigsqcup_{i \in I} B_i \rightarrow B$ such that each $X \times_B B_i \rightarrow B_i$ is nuclear.
- (2) For X/B a nuclear log curve and for any $b \in B$, the curve X_b has a well-defined dual graph Γ_b , with edges labelled by non-zero elements of $\overline{M}_{B,b}$; we denote the *label* (also called *length*) of e by δ_e ; this was denoted by δ_x in (2.3). If $\delta'_e \in M_B(B)$ is a lift of δ_e , then $\alpha(\delta'_e) \in \mathcal{O}_B(B)$ is a *smoothing parameter*

⁽¹⁰⁾A map $f: X \rightarrow Y$ of log schemes is *strict* if the log structure on X is the pullback of the log structure on Y . In particular, the strict étale topology on log schemes reflects very closely the usual étale topology on schemes.

for e , in the sense that X can be described locally around the corresponding point by an equation $uv = \alpha(\delta'_e)$. The stalk of \overline{M}_X at the corresponding node q of the fiber over $b \in B$ is given by

$$(2.6) \quad \overline{M}_{X,q} = \{(u, v) \in \overline{M}_{B,b} \oplus \overline{M}_{B,b} \text{ such that } \delta_e \mid (u - v)\}.$$

- (3) For X/B nuclear, the base B has a unique closed stratum,⁽¹¹⁾ and, for any b in that closed stratum, the restriction gives an isomorphism $\Gamma(B, \overline{M}_B) \xrightarrow{\sim} \overline{M}_{B,b}$.
- (4) If X/B is nuclear and $b, b' \in B$, with b in the closed stratum, there is a natural identification (of labelled graphs) of $\Gamma_{b'}$ with the graph obtained from Γ_b by mapping every label to $\overline{M}_{B,b'}$, and then contracting all edges that are labelled by 0. We often abuse notation by writing $\overline{M}_B := \overline{M}_{B,b}$ (for b in the closed stratum) in place of $\Gamma(B, \overline{M}_B)$. We often write Γ for the graph over any point in the closed stratum, which comes with an \overline{M}_B -metric.

If B is the spectrum of a noetherian strictly Henselian local ring with atomic log structure (for example, if B is the spectrum of a separably closed field), then by [HMOP23, Lemma 3.40] any log curve X/B is nuclear.

Let X/B be a nuclear log curve. Let $b \in B$ be a point in the closed stratum, with associated dual graph Γ with *vertex set* $V = V(\Gamma)$, *set of half-edges* $H = H(\Gamma)$ (including legs), and *set of non-leg half-edges* $H' = H'(\Gamma)$.

Definition 2.10. A *piecewise linear (PL) function* on X/B is an element of $\Gamma(X, \overline{M}_X^{\text{gp}})$.

A *combinatorial PL function* on X/B consists of the data:

- (1) a function $\beta': V(\Gamma) \rightarrow \overline{M}_{B,b}^{\text{gp}}$ (the *values* on the vertices), and
- (2) a function $\kappa: H'(\Gamma) \rightarrow \mathbb{Z}$ (the *slopes* on the non-leg⁽¹²⁾ half-edges)

such that if h_1 and h_2 are half-edges forming an edge e , with h_i attached to vertex v_i , we have

$$\kappa(h_2)\delta_e = \beta'(v_2) - \beta'(v_1)$$

(so that in particular $\kappa(h_1) + \kappa(h_2) = 0$). Edges of Γ with slope zero (that is, where both half-edges have slope zero) are called *horizontal*; all the other edges of Γ are called *vertical*.

We want to show that these two types of PL functions are in natural bijection. First, we construct a combinatorial PL function from any PL function. At generic points η of X_b , there is a natural isomorphism $\overline{M}_{B,b} = \overline{M}_{X,\eta}$, so the section $\beta \in H^0(X, \overline{M}_X^{\text{gp}})$ determines a function $\beta': V \rightarrow \overline{M}_{B,b}^{\text{gp}}$. To complete the definition of κ , we first show the following.

Lemma 2.11. *If h_1 and h_2 are half-edges forming an edge e , with h_i attached to vertex v_i , then for the function β' constructed from β as above, the value $\beta'(v_2) - \beta'(v_1)$ is an integer multiple of δ_e .*

Proof. This follows from (2.6) and the fact that the images of β under the two projections to $\overline{M}_{B,b}^{\text{gp}}$ are exactly given by $\beta'(v_1)$ and $\beta'(v_2)$. \square

In the notation of Lemma 2.11, we can then *define*

$$(2.7) \quad \kappa(h_2) := \frac{\beta'(v_2) - \beta'(v_1)}{\delta_e}$$

(which is unique because $\overline{M}_{B,b}$ is torsion-free). This accomplishes one direction of the following lemma.

Lemma 2.12. *The above construction induces a bijection between the set of PL functions and the set of combinatorial PL functions.*

⁽¹¹⁾Every log scheme comes with a decomposition into locally closed subschemes (called *strata*), where the ghost sheaf is locally constant.

⁽¹²⁾In this paper we do not include slopes on the legs, as we are interested only in the case where these slopes are equal to 0 (since we work throughout with **Rub**₀). Recall that, as discussed in Lemma 2.9, we have moved the data of the zeros and poles into the line bundle \mathcal{L}_μ .

Proof. Let β' be a combinatorial PL function; we build a PL function β giving the inverse image of β' under the construction above. If x is a smooth point of X_b , then $\overline{M}_{X,x} = \overline{M}_{B,b}$, and we define the value of β at x to be $\beta'(v)$, where v corresponds to the irreducible component of X_b containing x . The presentation (2.6) makes it clear that there is a unique way to extend this section to all non-smooth points $x \in X_b$. For any other point $b' \in B$, the combinatorial PL function can naturally be transferred (using property (4) of the definition of a nuclear log curve) to the fiber $X_{b'}$, and we repeat the above argument to give a PL function on $X_{b'}$. These then fit together to define a global PL function on X/B . \square

Our concrete local description of **Rub** is now given by the next proposition.

Proposition 2.13. *For X/B nuclear and $b \in B$ in the closed stratum, let V be the vertex set of the associated dual graph of b . Then there is a natural bijection between the set of X/B -points of \mathbf{Rub}_0 (i.e. the set of maps $B \rightarrow \mathbf{Rub}_0$ lying over X/B) and the set of maps*

$$(2.8) \quad \beta': V \longrightarrow \overline{M}_{B,b}^{\text{gp}}$$

satisfying the following conditions:

- (1) *The divisibility condition $\delta_e \mid \beta'(v_2) - \beta'(v_1)$ holds at every edge e in Γ_b connecting vertices $v_1, v_2 \in V$.*
- (2) *The image of β' is a totally ordered subset of $\overline{M}_{B,b}^{\text{gp}}$ with largest element being 0.*
- (3) *For every edge e connecting vertices v_1 and v_2 , with slope κ_e (defined as the absolute value of (2.7)), and for every $y \in \text{Image}(\beta')$ with $\beta'(v_1) < y < \beta'(v_2)$, the monoid $\overline{M}_{B,b}$ contains the element $\frac{y - \beta'(v_1)}{\kappa_e}$.*

Proof. Conditions (1) and (2) are translations of point (1) of Lemma 2.1. Condition (3) corresponds to point (2) of Lemma 2.1, as explained in [BHP⁺23, Section 6.2]. \square

Remark 2.14. If β'_1 and β'_2 are combinatorial PL functions with the same slopes κ_e , then there exists an element $c \in \overline{M}_{B,b}^{\text{gp}}$ such that $\beta'_1 = \beta'_2 + c$. In the definition of **Rub**, we restrict to PL functions whose values are totally ordered and take maximum value 0, and such functions are completely determined by the values of their slopes κ .

We would like to characterize in a similar spirit when a point of **Rub** lifts to **Rub** _{\mathcal{L}} . More concretely, this means describing explicitly the line bundle $\mathcal{O}_X(\beta)$ associated to a PL function. The next lemma describes the *restriction* of $\mathcal{O}_X(\beta)$ to the irreducible components of the curve X_b (in the case where β comes from **Rub**₀, i.e. has vanishing outgoing slopes). To describe the gluing between irreducible components would require us to get into quite a few more details of log geometry, and is not necessary for what we do in this paper.

Lemma 2.15 (cf. [RSPW19, Lemma 2.4.1]). *Let Y be the normalization of an irreducible component of X_b , corresponding to a vertex v . For each half-edge h attached to v , write κ_h for the slope (in the sense of (2.7)) and $z_h \in Y$ for the associated preimage of a node of X_b . Then there is a canonical isomorphism*

$$\mathcal{O}_X(\beta)|_Y = \pi^* \mathcal{O}_B(\beta'(v)) \otimes_{\mathcal{O}_Y} \mathcal{O}_Y \left(\sum_h \kappa_h z_h \right).$$

In particular, for a point $(X_b/b, \beta)$ of **Rub**₀ to lie in **Rub** _{\mathcal{L}} , it is necessary (though not in general sufficient) to require that, on the normalization Y of any irreducible component of X_b , there exists an isomorphism

$$\mathcal{L}_Y \cong \mathcal{O}_Y \left(\sum_h \kappa_h z_h \right),$$

where the sum runs over all half-edges h attached to v .

3. Generalized multi-scale differentials

We recall basic notions from [BCG⁺19], in order to define the groupoids $G\overline{\mathcal{M}}_{g,n}(\mu)$ of *simple generalized multi-scale differentials* and \mathcal{GMS}_μ of *generalized multi-scale differentials*, where $\mu = (m_1, \dots, m_n)$ is a tuple of integers with sum $2g - 2$. The adjective ‘generalized’ refers to the fact that we do not impose the global residue condition.

3.1. Enhanced level graphs

The boundary strata of the stack of generalized multi-scale differentials are indexed by *enhanced level graphs*. Such an enhanced level graph, typically denoted by Γ , is the dual graph of a stable curve, with legs corresponding to the marked points, with a level structure (*i.e.* a weak full order, equality being permitted) on the set of vertices $V(\Gamma)$, and with enhancements κ_e , which are non-negative integers attached to the edges. The edges $E(\Gamma)$ are grouped into the set of horizontal edges $E^h(\Gamma)$ joining vertices at the same level, and the set of vertical edges $E^v(\Gamma)$. The enhancements are required to be zero precisely for horizontal edges. We thus may consider an enhancement as a function

$$\kappa: H(\Gamma) \longrightarrow \mathbb{Z}$$

on the set of half-edges of Γ , assigning $\kappa_e > 0$ to the upper half and $-\kappa_e < 0$ to the lower half of a vertical edge, assigning zero to both halves of a horizontal edge, and letting κ agree with m_i at the legs of the graph. We normalize the set of levels so that the top level is zero, and let $L(\Gamma)$ be the set of levels below zero, usually given by consecutive negative integers $L(\Gamma) = \{-1, \dots, -N\}$, where $N := |L(\Gamma)|$, so that we typically use the *normalized level function*

$$(3.1) \quad \ell: V(\Gamma) \twoheadrightarrow \{0, -1, \dots, -N\}.$$

Occasionally, we use $L^\bullet(\Gamma)$ for the set of all levels including the zero level. In what follows we will only consider enhancements that are *admissible* in the sense that the degree equality

$$(3.2) \quad \deg(v) := \sum_{j \mapsto v} m_j + \sum_{e \in E^+(v)} (\kappa_e - 1) - \sum_{e \in E^-(v)} (1 + \kappa_e) - h(v) = 2g(v) - 2$$

holds, where $j \mapsto v$ means the leg of order m_j is attached to the vertex v ; *i.e.* the first sum goes over all legs attached to v , where $E^+(v)$ (resp. $E^-(v)$) is the set of vertical edges whose upper (resp. lower) end is the vertex v (we often write $e^+ = v$, resp. $e^- = v$, to express this adjacency), and $h(v)$ is the number of horizontal half-edges adjacent to v .

Enhanced level graphs come with two kinds of undegeneration maps. First, for any subset $I = \{i_1, \dots, i_n\}$ of $\{-1, \dots, -N\}$, there is the vertical undegeneration map δ_{i_1, \dots, i_n} , a map of graphs that contracts all vertical edges except those that go from the level at or above $i_k + 1$ to a level at or below i_k , for some $i_k \in I$. Especially important among those are the two-level undegenerations δ_i , which contract all vertical edges except those that cross a level passage above i , *i.e.* go from a vertex at level $i + 1$ or above to a vertex at level i or below. Second, for $S \subset E^h(\Gamma)$ there is the horizontal undegeneration maps δ_S^h that contract all the horizontal edges except those in S . An *undegeneration* of a level graph is a composition of a vertical and a horizontal undegeneration. Undegenerations determine the adjacency of boundary strata of the space of multi-scale differentials.

3.2. Prong-matchings

Let (X, ω) be a smooth complex curve with a meromorphic 1-form. If a differential ω has a zero of order $m \geq 0$ at $q \in X$, then there exists a local coordinate (and, in fact, there are $m + 1$ such choices) z on X centered at q such that locally in this coordinate $\omega = z^m dz$; similarly, for a pole of order $m \leq -2$ at $q \in X$, one can find a local coordinate such that $\omega = (z^m + r/z) dz$. Given such a local coordinate, a (*complex*) *prong*

is one of the $2|m+1|$ vectors in $T_q X$ of the form $\zeta^j \frac{\partial}{\partial z}$ in this local coordinate z , where ζ is a primitive root of unity of order $2|m+1|$; see [BCG⁺19, Definition 5.4]. In what follows we will mostly care about the set of P_q^{out} of *outgoing* prongs at the zeros of a differential, that is, the set of $m+1$ prongs there where the exponent j of ζ^j is even, and the set P_q^{inc} of *incoming* prongs at the poles of the differential, that is, those $|m+1|$ prongs there where the exponent j is odd.

Now let X be a stable curve with a node q corresponding to a vertical edge $e \in E^v(\Gamma)$ where two components of X meet, and suppose these components X_1 and X_2 come with differential forms ω_1 and ω_2 having a zero and a pole, respectively, at the respective preimages $q^+ \in X_1$ and $q^- \in X_2$ of q . A (*local*) *prong-matching* at the node q is a cyclic order-reversing bijection $\sigma_e: P_q^{\text{inc}} \rightarrow P_q^{\text{out}}$ between the incoming prongs at q^- and the outgoing prongs at q^+ .

Now let (X, z, Γ, ω) be a pointed stable curve with an enhanced level graph Γ , and let $\omega = (\omega_{(i)})_{i \in L(\Gamma)}$ be a *twisted differential of type μ* compatible with Γ , possibly except for the global residue condition. Following [BCG⁺18], this means a collection of meromorphic differentials ω_v for each vertex v , vanishing to order m_i at each of the marked points z_i , vanishing to order $\kappa(h) - 1$ at the preimages of nodes associated to the half-edges $h \in H'(\Gamma)$, and such that the residues at the two sides of a horizontal node add up to zero. Grouping objects level-wise, we denote by $\omega_{(i)}$ the tuple of differentials ω_v for all vertices v on level i .

Given a twisted differential, we have defined above local prong-matchings for each vertical edge. Packaging such a choice for each vertical edge $e \in E^v(\Gamma)$, we call the collection $\sigma = (\sigma_e)_{e \in E^v(\Gamma)}$ a *global prong-matching*.

There is an alternative viewpoint on prong-matchings, which can be generalized to germs of families $\mathcal{X} \rightarrow B$, where a node q corresponding to an edge e in the dual graph of the special fiber persists over the base. In the normalization of the family, there are two components X^\pm (as the edge is vertical, necessarily $X^+ \neq X^-$) that admit sections q^\pm that specify the two preimages of the node q . We let

$$(3.3) \quad \mathcal{N}_e^\vee := (q^+)^* \omega_{X^+} \otimes (q^-)^* \omega_{X^-}.$$

A *local prong-matching* (see [BCG⁺19, Definition 5.6]) is then a section σ_e of \mathcal{N}_e^\vee such that for any pair (v^+, v^-) of an incoming and an outgoing horizontal prong, the equation $\sigma_e(v^+ \otimes v^-)^{\kappa_e} = 1$ holds. To see the equivalence, given σ_e , we assign to v^- the prong v^+ given by the condition $\sigma_e(v^+ \otimes v^-) = 1$. Conversely, given a bijection $s: P_q^{\text{inc}} \rightarrow P_q^{\text{out}}$ of incoming and outgoing prongs, we define $\sigma_e \in \mathcal{N}_e^\vee$ by setting $\sigma_e(p \otimes s(p)) = 1$ for any $p \in P_q^{\text{inc}}$. The fact that s is order reversing implies that this $\sigma_e \in \mathcal{N}_e^\vee$ is well defined, which justifies using the same notation σ_e for both viewpoints on a prong-matching. A *global prong-matching* is a collection of local prong-matchings for each persistent node (as will be defined formally in Section 3.4) in the family.

We give another reformulation that eliminates the dependence on the choice of a preferred (“horizontal”) direction. Let U^\pm be neighbourhoods of the points q^\pm in the normalization of \mathcal{X} . Suppose the edge e joins level i to the lower level j . Then $\omega_{(i)}$ extends uniquely to a section of $\omega_{U^+/B}(-(\kappa_e - 1)q^+)$ and $\omega_{(j)}$ to a section of $\omega_{U^-/B}((\kappa_e + 1)q^-)$. Restricting these to q^+ and q^- , respectively, yields canonical elements

$$\tau^+ \in \omega_{U^+/B}(-(\kappa_e - 1)q^+)|_{q^+} = T_{q^+}^{\otimes -\kappa_e} \quad \text{and} \quad \tau^- \in \omega_{U^-/B}((\kappa_e + 1)q^-)|_{q^-} = T_{q^-}^{\otimes \kappa_e}$$

(where we use the residue isomorphism for the equalities). We define

$$\tau_e := (\tau^+)^{-1} \otimes (\tau^-) \in (T_{q^+} \otimes T_{q^-})^{\otimes \kappa_e} = \mathcal{N}_e^{\otimes \kappa_e}.$$

Lemma 3.1. *In the notation of the previous definition, let v^+ and v^- be some horizontal prongs at e . Then $(v^+ \otimes v^-)^{\otimes \kappa_e} \in \mathcal{N}_e^{\otimes \kappa_e}$ is independent of the choice of prongs and of the direction to be called horizontal, and we have*

$$(3.4) \quad \tau_e = (v^+ \otimes v^-)^{\otimes \kappa_e}.$$

Proof. For a fixed direction, the different choices of prongs v^+ differ by κ_e^{th} roots of unity, and likewise for v^- . Thus the formula for τ_e implies that it does not depend on these prong choices. On the other hand, changing the direction from horizontal to direction θ multiplies v^+ by $e^{2\pi i \theta}$ and v^- by $e^{-2\pi i \theta}$, and thus

preserves $v^+ \otimes v^-$. The equality is obvious, as can be seen by writing it out in any local coordinate that puts the differentials in normal form. \square

This implies that the earlier definitions of prong-matching agree with the following.

Definition 3.2. A *local prong-matching* is a section σ_e of \mathcal{N}_e^\vee such that $\sigma_e^{\kappa_e}(\tau_e) = 1$.

3.3. Level rotation tori

To an enhanced level graph, we associate some groups and algebraic tori. The *level rotation group* $R_\Gamma \cong \mathbb{Z}^{L(\Gamma)}$ acts on the set of all global prong-matchings, where the i^{th} factor twists by one (*i.e.* multiplies σ_e by $e^{2\pi i/\kappa_e}$) all prong-matchings associated to edges that cross the i^{th} *level passage*, a horizontal line above level i and below level $i+1$.⁽¹³⁾ The (*vertical*) *twist group* is the subgroup $\text{Tw}_\Gamma \subset R_\Gamma$ fixing all the prong-matchings under the above action. The level rotation group also acts (via its i^{th} component) on the set of prong-matchings of the two-level undegenerations $\delta_i(\Gamma)$. We define the *simple twist group* $\text{Tw}_\Gamma^s \subset \text{Tw}_\Gamma \subset R_\Gamma$ to be the subgroup that fixes each of the prong-matchings of each $\delta_i(\Gamma)$.

Let $\mathbb{C}^{L(\Gamma)} \rightarrow (\mathbb{C}^*)^{L(\Gamma)}$ be the universal covering of the algebraic torus $(\mathbb{C}^*)^{L(\Gamma)}$; we identify the level rotation group $R_\Gamma \subset \mathbb{C}^{L(\Gamma)}$ as the kernel of this covering. As a subgroup of the level rotation group, the (simple) twist group acts on $\mathbb{C}^{L(\Gamma)}$, and we define the *level rotation torus* $T_\Gamma := \mathbb{C}^{L(\Gamma)}/\text{Tw}_\Gamma$, together with its simple counterpart, the *simple level rotation torus* $T_\Gamma^s := \mathbb{C}^{L(\Gamma)}/\text{Tw}_\Gamma^s$. See Section 3.5 for an example when these two tori differ.

Next we define the data that provide the model for the toroidal embedding of the boundary inside the space of multi-scale differentials. Since $\text{Tw}_\Gamma^s = \bigoplus_i \text{Tw}_{\delta_i(\Gamma)}$ has by definition a direct sum decomposition level by level, the simple level rotation torus comes with a natural level-wise identification $T_\Gamma^s \cong (\mathbb{C}^*)^{L(\Gamma)}$. The embedding $\mathbb{C}^* \hookrightarrow \mathbb{C}$ with respect to these coordinates defines an embedding $T_\Gamma^s \hookrightarrow \overline{T}_\Gamma^s := \mathbb{C}^{L(\Gamma)}$. We let

$$(3.5) \quad a_i := a_{\delta_i(\Gamma)} := \text{lcm}_{e \in \delta_i(\Gamma)} \kappa_e$$

be the least common multiple of the enhancements of the edges of Γ that persist in the two-level undegeneration $\delta_i(\Gamma)$. Then $\text{Tw}_\Gamma^s \cong \bigoplus_i a_i \mathbb{Z} \subset R_\Gamma$. Consequently, T_Γ^s is a cover of the original torus $(\mathbb{C}^*)^{L(\Gamma)}$, of degree $\prod_i a_i$. Finally, we define the *quotient twist group* to be

$$(3.6) \quad K_\Gamma := \text{Tw}_\Gamma / \text{Tw}_\Gamma^s.$$

This group acts on T_Γ^s with quotient T_Γ . In coordinates the quotient map is given by

$$(3.7) \quad (\mathbb{C}^*)^{L(\Gamma)} \longrightarrow (\mathbb{C}^*)^{L(\Gamma)} \times (\mathbb{C}^*)^{E^v(\Gamma)}$$

$$(q_i) \longmapsto (\{r_i\}_{i \in L(\Gamma)}, \{\rho_e\}_{e \in E^v(\Gamma)}) := \left(\{q_i^{a_i}\}_{i \in L(\Gamma)}, \left\{ \prod_{i=\ell(e^-)}^{\ell(e^+)-1} q_i^{a_i/\kappa_e} \right\}_{e \in E^v(\Gamma)} \right),$$

where q_i, r_i, ρ_e denote the coordinates on the corresponding tori, and we view $T_\Gamma \subset (\mathbb{C}^*)^{L(\Gamma)} \times (\mathbb{C}^*)^{E^v(\Gamma)}$ as cut out by the equations

$$(3.8) \quad r_{\ell(e^-)} \cdots r_{\ell(e^+)-1} = \rho_e^{\kappa_e}$$

for each e . The action of K_Γ on T_Γ^s extends to an action on the closure \overline{T}_Γ^s , and we let $\overline{T}_\Gamma^n := \overline{T}_\Gamma^s / K_\Gamma$, which is the normalization of the closure of $T_\Gamma \subset (\mathbb{C}^*)^{L(\Gamma)} \times (\mathbb{C}^*)^{E^v(\Gamma)}$ inside $\mathbb{C}^{L(\Gamma)} \times \mathbb{C}^{E^v(\Gamma)}$.

All these tori come with their *extended versions*, denoted with an extra dot (*e.g.*, T_Γ^\bullet), that have an extra \mathbb{C}^* -factor. This factor will act on differentials of all levels simultaneously by multiplying all differentials by a common factor, and lead to the *projectivized* version of the corresponding quotient functor.

⁽¹³⁾In this paper we index levels and all quantities indexed by them, such as t_i, s_i, δ_i below, by negative integers, as in [BCG⁺19], but contrary to several subsequent papers that use this compactification.

3.4. Controlled families of generalized multi-scale differentials

We will now use these constructions to define families of generalized multi-scale differentials for families $\mathcal{X} \rightarrow B$ of curves over ‘small enough’ base schemes. The general case will then be treated by gluing. The notion of ‘small enough’ that we will use is that of *controlled curves* from [BH23, Section 2.6]. This deviates slightly from the corresponding definition in [BCG⁺19, Section 11], where such families are merely defined as germs, and allows to *a priori* say under which morphisms our families can be pulled back without having to pass to suitable representatives of germs.

We do not recall the definition of controlled curves in full, but recall the key properties that we will need. If $\mathcal{X} \rightarrow B$ is a controlled curve, then for every $p \in B$, the fiber X_p has a well-defined dual graph Γ_p . Moreover, there exists a *controlling point* $b \in B$ together with *smoothing parameters* $f_e \in \mathcal{O}_B(B)$ for every edge e of Γ_b , such that f_e vanishes exactly on the locus of $p \in B$ where the corresponding node persists in X_p , and such that the family has the local form $u_e v_e = f_e$ in a neighbourhood of the corresponding node. The dual graph of X_p is obtained from that of X_b by contracting exactly those e such that $f_e(p) \neq 0$. The function f_e is unique up to multiplication by units in $\mathcal{O}_B(B)$; we write $[f_e] \in \mathcal{O}_B(B)/\mathcal{O}_B^\times(B)$ for the equivalence class of the smoothing parameter. Given any family of stable curves $\mathcal{X} \rightarrow B$, over a locally noetherian base, there exists an étale cover $\bigsqcup_i B_i \rightarrow B$ with each $\mathcal{X} \times_B B_i \rightarrow B_i$ controlled (see [BH23, Lemma 2.6.9]). Our first step is to define ‘standard’ open subsets of controlled families of curves where the collection of rescaled differentials are defined. Let $\mu = (m_1, \dots, m_n)$ be a tuple of integers with $\sum_{i=1}^n m_i = 2g - 2$. If X is a stable n -marked curve with enhanced level graph structure on the dual graph, then for any level i , the subcurve $X^{[i]} := X_{(\leq i)} \setminus (X_{(> i)} \cup \mathbf{z}^\infty)$ is open in X , where $X_{(\leq i)}$ is the subcurve at and below level i , $X_{(> i)}$ is the subcurve above level i , and \mathbf{z}^∞ is the union of those markings with $m_j < 0$.

We now fix a stable n -marked controlled curve $\mathcal{X} \rightarrow B$ and a controlling point $b \in B$, and suppose we are given an enhanced level graph structure on the dual graph of X_b . If $p \in B$ is another point, we have an edge contraction $\Gamma_b \rightarrow \Gamma_p$, which naturally induces on Γ_p the structure of an enhanced level graph (see [BCG⁺19, Section 5.1] for details). In particular, the set L_b of levels of Γ_b naturally surjects onto the set L_p of levels of Γ_p : we have $l_p: L_b \twoheadrightarrow L_p$. The reader may check that the following sets $U_i \subset \mathcal{X}$ are indeed open.

Definition 3.3. Given a stable n -marked controlled curve $\mathcal{X} \rightarrow B$ and a level $i \in L_b$, we define $U_i \subseteq \mathcal{X}$, the *standard open set at level i* , to be the union over all $p \in B$ of $X_p^{[l_p(i)]}$.

We say that a node e is *persistent* in the family \mathcal{X} if $f_e = 0 \in \mathcal{O}_B$. If the dual graph Γ_b has been provided with an enhanced level graph structure, we say that a node e is *semi-persistent* if $f_e^{\kappa_e} = 0$. The notion of prong-matchings makes sense for a persistent node q .

For our families of multi-scale differentials, we need to include an explicit choice of smoothing parameters f_e into our data. This can be achieved via a section of the partial compactification \overline{T}_Γ^s of the simple level rotation torus. Indeed, given the coordinates (r_i, ρ_e) on the torus closure from (3.8), a morphism $R^s: B \rightarrow \overline{T}_\Gamma^s$ determines for each vertical edge e a function $f_e \in \mathcal{O}_B$, and for each level i a function $s_i \in \mathcal{O}_B$, defined as the compositions $f_e = \rho_e \circ P \circ R^s$ and $s_i = r_i \circ P \circ R^s$, where $P: \overline{T}_\Gamma^s \rightarrow \overline{T}_\Gamma^n$ is the canonical morphism. If an edge e joins levels $j < i$, then by (3.8) these functions satisfy

$$(3.9) \quad f_e^{\kappa_e} = s_j \cdots s_{i-1}.$$

The following definition makes precise the notion that a morphism R^s as above defines a compatible system of node-smoothing parameters.

Definition 3.4. A *simple rescaling ensemble* is a morphism $R^s: B \rightarrow \overline{T}_{\Gamma_b}^s$ such that the parameters $f_e \in \mathcal{O}_B(B)$ for each vertical edge e determined by R^s lie in the equivalence class $[f_e]$ determined by the family $\pi: \mathcal{X} \rightarrow B$. A *rescaling ensemble* is a morphism $R: B \rightarrow \overline{T}_{\Gamma_b}^n$ which arises as the composition $\pi \circ R^s$ for some simple rescaling ensemble R^s .

The s_i and f_e will be called the *rescaling parameters* and *smoothing parameters* determined by R or R^s . The composition of R^s with the coordinate projections gives functions t_i such that $s_i = t_i^{a_i}$. We refer to those t_i as the *level parameters*.

Recall that the adjective ‘generalized’ in the following definition refers to the fact that the global residue condition has been dropped, compared to [BCG⁺19]. For an illustration of some elements of the definition, see Figure 2. The well-definedness of the period in the following definition is checked (in any characteristic), e.g. in [Boj19, Lemma 1.8].

Definition 3.5. A collection of generalized rescaled differentials of type μ on the stable n -pointed controlled curve $(\pi: \mathcal{X} \rightarrow B, \mathbf{z})$ with a controlling point $b \in B$ is a collection of sections $\omega_{(i)}$ of $\omega_{\mathcal{X}/B}$ defined on the standard open subsets U_i of \mathcal{X} , indexed by the levels $i \in L^\bullet(\Gamma)$ of the enhanced level graph. The irreducible components of X_b on a level strictly below i are called *vertical zeros*, and those on a level strictly above i are called *vertical poles* of $\omega_{(i)}$. We require the collection to satisfy the following constraints:

- (1) There exist sections $s_i \in H^0(B, \mathcal{O}_B)$ with $s_i(b) = 0$ such that for any levels $j < i$, the differentials satisfy $\omega_{(i)} = s_j \cdots s_{i-1} \omega_{(j)}$ on $U_i \cap U_j$.
- (2) For any edge e joining levels $j < i$, for any $p \in B$ over which e persists, and for some (equivalently, any) choice of functions u_e, v_e on \mathcal{X} and f_e on B such that the family has local normal form $u_e v_e = f_e$, there exists a unit λ in the strict Henselization⁽¹⁴⁾ $(\mathcal{O}_B[u_e, v_e]/(u_e v_e - f_e))_{(p,0)}^{\text{sh}}$ of the local ring at $(p, u_e = 0, v_e = 0)$ such that

$$(3.10) \quad \omega_{(i)} = \lambda u_e^{\kappa_e} \frac{du_e}{u_e} \quad \text{and} \quad \omega_{(j)} = -\lambda v_e^{-\kappa_e} \frac{dv_e}{v_e}.$$

- (3) The $\omega_{(i)}$ have order m_k along the sections \mathcal{Z}_k of the k^{th} marked point that meet the level i subcurve of X_b ; these are called *horizontal zeros and poles* (where \mathcal{Z}^∞ records the horizontal poles). Moreover, $\omega_{(i)}$ is holomorphic and non-zero away from its horizontal and vertical zeros and poles.

If the rescaling and smoothing parameters s_i, f_e for the collection $\omega_{(i)}$ agree with those of a rescaling ensemble R^s or R , we call them *compatible*. We denote the collection by $\omega = (\omega_{(i)})_{i \in L^\bullet(\Gamma)}$.

The unit λ is unique since $\omega_{(i)}$ is a generating section of ω after inverting u_e and $\omega_{(j)}$ is a generating section of ω after inverting v_e . We then apply the following result from commutative algebra.

Lemma 3.6. Let R be a ring, $f \in R$, and $A = R[u, v]/(uv - f)$. Then $h: A \rightarrow A[u^{-1}] \times A[v^{-1}]$ is injective.

Proof. Assume $h(a) = 0$. Any element $a \in A$ has a unique normal form

$$a = e_0 + c_1 u + \cdots + c_n u^n + d_1 v + \cdots + d_m v^m.$$

Under the unique isomorphism $A[u^{-1}] \rightarrow R[u, u^{-1}]$ sending v to f/u , the element a is mapped to

$$e_0 + c_1 u + \cdots + c_n u^n + d_1 f u^{-1} + \cdots + d_m f^m u^{-m}.$$

Therefore, the vanishing of the first component of h implies $e_0 = c_1 = \cdots = c_n = 0$, as can be seen just by reading off the coefficients of the R -basis $\{u^t : t \in \mathbb{Z}\}$. Similarly, the vanishing of the second component implies $d_1 = \cdots = d_m = 0$. \square

Applying this to $R = \mathcal{O}_B$, $u = u_e$, $v = v_e$, and $f = f_e$, we first conclude that the map

$$\mathcal{O}_B[u_e, v_e]/(u_e v_e - f_e) \longrightarrow \mathcal{O}_B[u_e, v_e, u_e^{-1}]/(u_e v_e - f_e) \times \mathcal{O}_B[u_e, v_e, v_e^{-1}]/(u_e v_e - f_e)$$

is injective. Localization and strict Henselization are flat. Therefore, the above injectivity is preserved.

⁽¹⁴⁾In the complex analytic category, this would be the germ of a non-vanishing function around the node.

The reader comparing with the definition in [BCG⁺19] will realize that there in item (2) it is required that for any edge e joining levels $j < i$ of Γ , there are functions u_e, v_e on \mathcal{X} and f_e on B , such that the family has local normal form $u_e v_e = f_e$, and in these coordinates

$$(3.11) \quad \omega_{(i)} = \left(u_e^{\kappa_e} + f_e^{\kappa_e} r_{e,(j)} \right) \frac{du_e}{u_e} \quad \text{and} \quad \omega_{(j)} = - \left(v_e^{-\kappa_e} + r_{e,(j)} \right) \frac{dv_e}{v_e},$$

where κ_e is the enhancement of Γ_b at e . The two normal forms are equivalent at least if $f_e^{\kappa_e} \neq 0$: Equation (3.11) implies (3.10) by taking $\lambda = 1 + r_{e,(j)}(f_e/u_e)^{\kappa_e}$. Conversely, given (3.10) we may change v_e to a coordinate that is in the normal form (3.11) by [BCG⁺19, Theorem 4.3]. The form (3.10) is the one we can directly associate with a rubber differential; see Section 5.4.




Figure 2. The underlying curve for a family of generalized rescaled differentials of type $\mu = (4)$, with neighbourhoods U_0, U_{-1} (in red, green) and the vanishing cycle γ_e (in blue).

Remark 3.7. Let ω be a collection of generalized rescaled differentials with a compatible rescaling ensemble R^s or R . For any non-semi-persistent edge e , denote by $B_e \subset B$ the vanishing locus of f_e . Then there is a natural *induced prong-matching* σ_e over B_e , which is determined by the choice of the rescaled differentials $\omega_{(i)}$ and the rescaling ensemble. This prong-matching σ_e is defined explicitly in local coordinates by writing it as $\sigma_e = du_e \otimes dv_e$ when restricting to the nodal locus corresponding to e , where u_e and v_e are as in (3.11) with f_e prescribed by the rescaling ensemble. Any two possible choices of u_e and v_e are of the form $\alpha_e u_e$ and $\alpha_e^{-1} v_e$ for some unit $\alpha_e \in \mathcal{O}_B^*$ (see [BCG⁺19, Section 4]), so the induced prong-matching does not depend on this choice.

We can now package everything into our main notion.

Definition 3.8. Given a controlled family of pointed stable curves $(\pi: \mathcal{X} \rightarrow B, z)$, a (*controlled*) *family of generalized simple multi-scale differentials* of type μ over B consists of the following data:

- (1) the structure of an enhanced level graph on the dual graph Γ_b of the fiber X_b ;
- (2) a simple rescaling ensemble $R^s: B \rightarrow \overline{T}_{\Gamma_b}^s$, compatible with
 - (a) a collection of generalized rescaled differentials $\omega = (\omega_{(i)})_{i \in L \bullet(\Gamma_b)}$ of type μ , and
 - (b) a collection of prong-matchings $\sigma = (\sigma_e)_{e \in E^v(\Gamma)}$, where σ_e is a section of \mathcal{N}_e^v over B_e , the vanishing locus of f_e . If e is a non-semi-persistent node, σ_e is required to agree with the induced prong-matching defined in Lemma 3.7.

A section of the simple level rotation torus $T_{\Gamma_b}^s(\mathcal{O}_B)$, that is, a morphism $\xi: B \rightarrow T_{\Gamma_b}^s$, acts on all of the above data via

$$\xi \cdot (\omega_{(i)}, R^s, \sigma_e) = (\xi \cdot \omega_{(i)}, \xi^{-1} \cdot R^s, \xi \cdot \sigma_e).$$

Here, for $\xi \in T_{\Gamma_b}^s(\mathcal{O}_B)$ mapping to $((r_i)_{i \in L(\Gamma_b)}, (\rho_e)_{e \in E^v(\Gamma_b)})$ under the quotient map (3.7), the action is defined by

$$\xi \cdot \omega_{(i)} = \left(\prod_{\ell \geq i} r_\ell \right) \omega_{(i)}, \quad \xi \cdot \sigma_e = \rho_e \sigma_e,$$

and $\xi^{-1} \cdot R^s$ denotes the post-composition of R^s with the multiplication by ξ^{-1} .⁽¹⁵⁾

An isomorphism between two controlled families of generalized simple multi-scale differentials over the same base B and with the same controlling point b

$$(3.12) \quad (\pi': \mathcal{X}' \rightarrow B, z', \Gamma_b', (R^s)', \omega', \sigma') \rightarrow (\pi: \mathcal{X} \rightarrow B, z, \Gamma_b, R^s, \omega, \sigma)$$

consists of an isomorphism $\varphi: \mathcal{X}' \rightarrow \mathcal{X}$ and an element $\xi \in T_{\Gamma_b}^s(\mathcal{O}_B)$ such that

- i) φ defines an isomorphism of families of pointed stable curves,
- ii) the induced isomorphism of dual graphs $\Gamma_b' \rightarrow \Gamma_b$ is also an isomorphism of enhanced level graphs,
- iii) the action of ξ sends $((R^s)', \omega', \sigma')$ to $\varphi^*(R^s, \omega, \sigma)$.

Pullbacks of families of controlled generalized multi-scale differentials are defined as in [BCG⁺19, Section 11.2]. This step requires some care, since the number of levels, the nodes where the prong-matching is an induced prong-matching, and the target of the rescaling ensemble map change. With this in hand, we can define a family of generalized multi-scale differentials over any scheme locally of finite type over \mathbb{C} by sheafifying the notion already defined for controlled families, using that controlled families form a base for the étale topology in the sense of [BH23, Definition A.3]. This definition is analogous to that in [BCG⁺19, Section 11.3], and can be seen as a groupoid-version of the constructions worked out in [BH23, Appendix A].

Definition 3.9. We let $G\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ be the groupoid of families of generalized simple multi-scale differentials.

There are two variants of this definition. First, replacing $T_{\Gamma_b}^s(\mathcal{O}_{B,b})$ with the extended level rotation torus $T_{\Gamma_b}^\bullet(\mathcal{O}_{B,b})$ in the definition of a morphism, we obtain projectivized generalized simple multi-scale differentials. Here the additional torus factor acts by scaling the differential on all levels simultaneously, including level 0. These are relevant to get compact spaces. Here we compare the unprojectivized definitions and will not elaborate further on this.

Second, there is a ‘non-simple’ variant that we need to compare to the relative coarse moduli space. The remarks above about pullback and sheafification apply here as well.

Definition 3.10. A family of controlled generalized multi-scale differentials of type μ is defined as in Lemma 3.8, replacing (2) by a rescaling ensemble $R: B \rightarrow \overline{T}_{\Gamma_b}^n$. A morphism of such controlled families consists of $(\varphi, \tilde{\varphi}, \xi)$ as above, except that now we allow $\xi \in T_{\Gamma_b'}(\mathcal{O}_{B'})$. We let \mathcal{GMS}_μ be the resulting groupoid of families of generalized multi-scale differentials.

Modifying Lemma 3.5 by additionally imposing the global residue condition gives a groupoid that we denote by $\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ for the simple version (Lemma 3.8) and by \mathcal{MS}_μ for the non-simple version (Lemma 3.10). We state the comparison to the objects defined in [BCG⁺19].

Proposition 3.11. The stack $\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ is a smooth DM-stack. The stack \mathcal{MS}_μ is a stack with finite quotient singularities and agrees with the normalization of the orderly blowup of the normalized incidence variety compactification; see [BCG⁺19, Section 14].

⁽¹⁵⁾Most of the checks that this action is well defined are straightforward. To verify part (2) of Lemma 3.8, assume we are given local coordinates u, v around a node associated to $e \in E^v(\Gamma_b)$ such that the differentials have the normal form (3.11). Then the rescaled differential is put in the required normal form using the new coordinates $\widehat{u} = (\prod_{\ell \geq i} r_\ell)^{1/\kappa_e} u$ and $\widehat{v} = (\prod_{\ell \geq j} r_\ell)^{-1/\kappa_e} v$.

The paper [BCG⁺19, Section 14.2] also defines a smooth stack denoted by $\Xi\overline{\mathcal{M}}_{g,n}(\mu)$, patched locally from quotients of stacks with a Teichmüller marking. The full proof that this stack is isomorphic to the stack with the same symbol defined here would require recalling the lengthy definitions of level-wise real blowup and Teichmüller marking from [BCG⁺19, Section 12]. This identification directly implies the second statement of the proposition. The proof given here provides the main content of the proposition, the smoothness of this stack, without using the smoothness results from [BCG⁺19].

Proof. Recall from [BCG⁺19, Section 8.1 and Theorem 10.1] that a versal deformation space B of \mathcal{MS}_μ is given by a product $B = \overline{T}_{\Gamma_b}^n \times B_0$, where $\overline{T}_{\Gamma_b}^n$ gives a parameterization of possible rescaling ensembles R (which have values in $\overline{T}_{\Gamma_b}^n$), and where B_0 parameterizes the remaining data (deformations of the components \mathcal{X}_v for $v \in V(\Gamma_b)$ and twisted differentials on these components). In fact, this local structure is given for the model space in [BCG⁺19, Section 8.1]. This model space is locally isomorphic to the Dehn space by the plumbing construction given in [BCG⁺19, Theorem 10.1], and [BCG⁺19, Proposition 12.5] shows that every family can locally be lifted to the Dehn space.

Consider the fiber product

$$\begin{array}{ccc} \widehat{B} := B \times_{\mathcal{MS}_\mu} \Xi\overline{\mathcal{M}}_{g,n}(\mu) & \longrightarrow & \Xi\overline{\mathcal{M}}_{g,n}(\mu) \\ \downarrow & & \downarrow \\ B & \longrightarrow & \mathcal{MS}_\mu. \end{array}$$

We claim that \widehat{B} is equal to the stack quotient $[\overline{T}_\Gamma^s/K_\Gamma]$ times the product B_0 of the other factors. Then the maps $\widehat{B} \rightarrow \Xi\overline{\mathcal{M}}_{g,n}(\mu)$ provide a smooth cover by spaces which are smooth themselves, which implies the claimed smoothness of $\Xi\overline{\mathcal{M}}_{g,n}(\mu)$.

To show that \widehat{B} is equal to $[\overline{T}_\Gamma^s/K_\Gamma] \times B_0$, we write down explicitly the maps $\widetilde{B} \rightarrow \widehat{B}$, where \widetilde{B} carries a controlled curve. For this, recall⁽¹⁶⁾ that a morphism to a fiber product as above is given by a triple

$$(\widetilde{B} \rightarrow \Xi\overline{\mathcal{M}}_{g,n}(\mu), \widetilde{B} \rightarrow B, G),$$

where G is a 2-isomorphism between the compositions

$$\widetilde{B} \longrightarrow \Xi\overline{\mathcal{M}}_{g,n}(\mu) \longrightarrow \mathcal{MS}_\mu \quad \text{and} \quad \widetilde{B} \longrightarrow B \longrightarrow \mathcal{MS}_\mu.$$

Inserting the definitions of the moduli stacks, this data above is equivalent to a triple of

- a controlled family $(\pi: \mathcal{X} \rightarrow \widetilde{B}, z, \Gamma_b, R^s: \widetilde{B} \rightarrow \overline{T}_{\Gamma_b}^s, \omega, \sigma)$ of generalized simple multi-scale differentials,
- morphisms $s_T: \widetilde{B} \rightarrow \overline{T}_{\Gamma_b}^n$ and $s_0: \widetilde{B} \rightarrow B_0$ (which together can be thought of as a morphism $(s_T, s_0): \widetilde{B} \rightarrow \overline{T}_{\Gamma_b}^n \times B_0 = B$),
- an isomorphism $(\mathcal{X} \cong \mathcal{X}', \xi \in T_{\Gamma_b}(\mathcal{O}_{\widetilde{B}}))$ of generalized (non-simple) multi-scale differentials, sending the family $(\pi: \mathcal{X} \rightarrow \widetilde{B}, z, \Gamma_b, R, \omega, \sigma)$ to the family $(\pi': \mathcal{X}' \rightarrow \widetilde{B}, z', \Gamma_b, R', \omega', \sigma')$ induced by $(s_T, s_0): \widetilde{B} \rightarrow B$.

By identifying the families of curves $\mathcal{X} \cong \mathcal{X}'$, we can act on the pair (s_T, s_0) with the section ξ of the level rotation torus. Replacing (s_T, s_0) by this modified pair, we obtain a new, equivalent, triple of data, where the isomorphism in the last bullet point is taken as the identity. But then we see that such a triple is uniquely determined by the pair

$$(R^s: \widetilde{B} \longrightarrow \overline{T}_{\Gamma_b}^s, s_0: \widetilde{B} \longrightarrow B_0),$$

by taking s_T in the second bullet point as the composition $\widetilde{B} \rightarrow \overline{T}_{\Gamma_b}^s \rightarrow \overline{T}_{\Gamma_b}^n$ and taking the data $(\pi, z, \Gamma_b, \omega, \sigma)$ in the first bullet point that is determined by the non-simple generalized multi-scale differential from $(s_T, s_0): \widetilde{B} \rightarrow B$.

⁽¹⁶⁾For a reminder on fiber products of stacks, we recommend the excellent paper [Fan01].

Above we have found that any morphism $\widetilde{B} \rightarrow \widehat{B}$ can be described by a morphism $(R^s, s_0): \widetilde{B} \rightarrow \overline{T}_{\Gamma_b}^s \times B_0$. Two such morphisms are 2-isomorphic if they can be related by compatible isomorphisms for the stacks B and $\Xi \overline{\mathcal{M}}_{g,n}(\mu)$ in the fiber product. Since B is a scheme, the only such isomorphisms come from sections $\xi: \widetilde{B} \rightarrow \overline{T}_{\Gamma_b}^s$ leaving the underlying non-simple generalized multi-scale differential fixed. These are exactly identified with sections $\xi: \widetilde{B} \rightarrow K_{\Gamma_b}$, which act in a natural way on the first morphism $R^s: \widetilde{B} \rightarrow \overline{T}_{\Gamma_b}^s$. Since \widetilde{B} is connected, the section ξ is necessarily constant, so that we have identified⁽¹⁷⁾

$$\mathrm{Mor}(\widetilde{B}, \widehat{B}) = \mathrm{Mor}(\widetilde{B}, \overline{T}_{\Gamma_b}^s \times B_0) / K_{\Gamma} = \mathrm{Mor}(\widetilde{B}, [\overline{T}_{\Gamma_b}^s / K_{\Gamma}] \times B_0).$$

This proves the isomorphism $\widehat{B} \cong [\overline{T}_{\Gamma_b}^s / K_{\Gamma}] \times B_0$. Since both the quotient stack $[\overline{T}_{\Gamma_b}^s / K_{\Gamma}]$ and B_0 are smooth, this finishes the proof. \square

Proof of Lemma 1.1, second part. Assuming the first part of the theorem, the proof of the second part is completed by showing that the map $G\Xi \overline{\mathcal{M}}_{g,n}(\mu) \rightarrow \mathcal{GMS}_{\mu}$ is the relative coarse moduli space over $\overline{\mathcal{M}}_{g,n}$. First, we observe that the map $\mathcal{GMS}_{\mu} \rightarrow \overline{\mathcal{M}}_{g,n}$ is representable. Indeed, the stabilizers $(\varphi, \widetilde{\varphi}, \xi)$ of a family of generalized multi-scale differentials lying over the identity morphism $\varphi = \mathrm{id}_B$, $\widetilde{\varphi} = \mathrm{id}_X$ of the underlying stable curves are those $\xi \in T_{\Gamma_b}(\mathcal{O}_B)$ fixing both the differentials ω and the prong-matchings σ . By the definition of the level rotation torus, this forces ξ to be trivial, so that indeed the stabilizers of \mathcal{GMS}_{μ} inject to the stabilizers of $\overline{\mathcal{M}}_{g,n}$.

By the definition of the relative coarse space, we then have a factorization

$$G\Xi \overline{\mathcal{M}}_{g,n}(\mu) \longrightarrow G\Xi \overline{\mathcal{M}}_{g,n}(\mu)^{\mathrm{coarse}} \longrightarrow \mathcal{GMS}_{\mu},$$

and we show that the second map is an isomorphism. For this, let $B \rightarrow \mathcal{GMS}_{\mu}$ be associated to a controlled family of generalized multi-scale differentials. Then we have a commutative diagram, where we *define* the diagrams on the right to be cartesian:

$$\begin{array}{ccccc} [\overline{T}_{\Gamma_b}^s / K_{\Gamma}] & \longleftarrow \cdots & G\Xi \overline{\mathcal{M}}_{g,n}(\mu)_B & \longrightarrow & G\Xi \overline{\mathcal{M}}_{g,n}(\mu) \\ \downarrow & & \downarrow & & \downarrow \\ & & G\Xi \overline{\mathcal{M}}_{g,n}(\mu)_B^{\mathrm{coarse}} & \longrightarrow & G\Xi \overline{\mathcal{M}}_{g,n}(\mu)^{\mathrm{coarse}} \\ & & \downarrow & & \downarrow \\ \overline{T}_{\Gamma_b}^n & \xleftarrow{R} & B & \longrightarrow & \mathcal{GMS}_{\mu}. \end{array}$$

For the family $B \rightarrow \mathcal{GMS}_{\mu}$ given by a tuple $(\pi: \mathcal{X} \rightarrow B, z, \Gamma_b, R, \omega, \sigma)$ and a map $s: \widetilde{B} \rightarrow B$, we claim that after we shrink \widetilde{B} in the étale topology, the morphisms $\widetilde{B} \rightarrow G\Xi \overline{\mathcal{M}}_{g,n}(\mu)_B$ lying over s are precisely given by the data of

$$(3.13) \quad R^s: \widetilde{B} \longrightarrow \overline{T}_{\Gamma_b}^s \quad \text{such that} \quad (\overline{T}_{\Gamma_b}^s \rightarrow \overline{T}_{\Gamma_b}^n) \circ R^s = R \circ s$$

with the automorphisms of this data described by the (necessarily locally constant) sections $\xi: \widetilde{B} \rightarrow K_{\Gamma_b}$. To see this, one repeats the analysis of the stacky fiber product from Lemma 3.11 for $G\Xi \overline{\mathcal{M}}_{g,n}(\mu)_B$.⁽¹⁸⁾ From this description, we see that $G\Xi \overline{\mathcal{M}}_{g,n}(\mu)_B$ is also the fiber product of B with the map $[\overline{T}_{\Gamma_b}^s / K_{\Gamma}] \rightarrow \overline{T}_{\Gamma_b}^n$, so that the dotted arrow on the top left makes the left diagram cartesian.

⁽¹⁷⁾For the second equality below, we use that for a finite group K acting on a scheme \overline{T} , the morphisms $\widetilde{B} \rightarrow [\overline{T}/K]$ can be identified with the groupoid $\{\widetilde{B} \rightarrow \overline{T}\}/K$, perhaps after replacing \widetilde{B} by an étale cover, which is harmless for the argument. This itself uses the definition of the quotient stack together with the fact that all K -torsors over a scheme \widetilde{B} as above are étale-locally trivial.

⁽¹⁸⁾In that proof we did use the normal form $B = \overline{T}_{\Gamma_b}^n \times B_0$ to split the map $s: \widetilde{B} \rightarrow B$ as (s_T, s_0) , where $s_T = R \circ s$. However, the only place where this was actually used was in observing that the data $(s = (s_T, s_0), R^s)$ satisfying the condition in (3.13) is equivalent to (R^s, s_0) by setting $s_T = (\overline{T}_{\Gamma_b}^s \rightarrow \overline{T}_{\Gamma_b}^n) \circ R^s$.

To conclude, we first note that by [AOVII, Proposition 3.4] the space $G\Xi\overline{\mathcal{M}}_{g,n}(\mu)_B^{\text{coarse}}$, which was defined as a fiber product, is in fact a relative coarse space for $G\Xi\overline{\mathcal{M}}_{g,n}(\mu)_B$ over $\overline{\mathcal{M}}_{g,n}$. But since the map $G\Xi\overline{\mathcal{M}}_{g,n}(\mu)_B \rightarrow \overline{\mathcal{M}}_{g,n}$ factors through the representable map $B \rightarrow \overline{\mathcal{M}}_{g,n}$, the space $G\Xi\overline{\mathcal{M}}_{g,n}(\mu)_B^{\text{coarse}}$ is *also* a coarse space of $G\Xi\overline{\mathcal{M}}_{g,n}(\mu)_B$ over B , by an application of Lemma 3.12 below to $\mathcal{X} = G\Xi\overline{\mathcal{M}}_{g,n}(\mu)_B$, $\mathcal{Y}' = B$, and $\mathcal{Y} = \overline{\mathcal{M}}_{g,n}$.

On the other hand, since $\overline{T}_{\Gamma_b}^n$ is the coarse space of $[\overline{T}_{\Gamma_b}^s/K_{\Gamma}]$ (over $\text{Spec}(\mathbb{C})$), applying [AOVII, Proposition 3.4] again shows that B *itself* is the coarse space of $G\Xi\overline{\mathcal{M}}_{g,n}(\mu)_B$. This proves that the map $G\Xi\overline{\mathcal{M}}_{g,n}(\mu)_B^{\text{coarse}} \rightarrow B$ is an isomorphism. Since we prove this for any $B \rightarrow \mathcal{GMS}_{\mu}$, we conclude that $G\Xi\overline{\mathcal{M}}_{g,n}(\mu)^{\text{coarse}} \rightarrow \mathcal{GMS}_{\mu}$ is an isomorphism, as desired. \square

Lemma 3.12. *Consider a sequence of morphisms $\mathcal{X} \rightarrow \mathcal{Y}' \rightarrow \mathcal{Y}$ of algebraic stacks, locally of finite presentation, and assume the relative inertia $I(\mathcal{X}/\mathcal{Y}) \rightarrow \mathcal{X}$ is finite. Then if $\mathcal{Y}' \rightarrow \mathcal{Y}$ is representable, the relative coarse space $\mathcal{X}^{\text{coarse}, \mathcal{Y}}$ of \mathcal{X} over \mathcal{Y} is isomorphic to the relative coarse space $\mathcal{X}^{\text{coarse}, \mathcal{Y}'}$ of \mathcal{X} over \mathcal{Y}' .*

Proof. Consider the solid diagram of morphisms

$$\begin{array}{ccccc} \mathcal{X} & \longrightarrow & \mathcal{X}^{\text{coarse}, \mathcal{Y}'} & \longrightarrow & \mathcal{Y}' \\ \parallel & & \uparrow f \quad \downarrow g & & \downarrow \\ \mathcal{X} & \longrightarrow & \mathcal{X}^{\text{coarse}, \mathcal{Y}} & \longrightarrow & \mathcal{Y}. \end{array}$$

Then by the properties of relative coarse spaces (see [AOVII, Theorem 3.1(2)]), there exists a morphism f as above, since $\mathcal{X}^{\text{coarse}, \mathcal{Y}}$ is initial among factorizations of $\mathcal{X} \rightarrow \mathcal{Y}$ via a representable map. But then the induced map $\mathcal{X}^{\text{coarse}, \mathcal{Y}} \rightarrow \mathcal{Y}'$ via f is representable, so by the same universal property we obtain the map g , and one verifies that f, g are inverse isomorphisms. \square

Remark 3.13. In practice it is often relevant to determine the number of projectivized multi-scale differentials on a given pointed curve with twisted differential (X, z, Γ, ω) . By the definition of the above equivalence relation, this is the number of *prong-matching equivalence classes*, i.e. the number of orbits of the set of global prong-matchings under the action of the level rotation group R_{Γ} . Determining this number is complicated in general, but, for a two-level graph with prongs $\kappa_1, \dots, \kappa_s$, there are $\prod \kappa_i / \text{lcm}(\kappa_i)$ prong-matching equivalence classes.

3.5. Quotient twist group and rescaling ensembles in a worked example

Consider the triangle graph Γ with three levels, each containing one vertex, and three edges forming a triangle, as illustrated in Figure 3 (to which we also refer for the labelling of the edges). For simplicity we




Figure 3. The triangle graph (the generic fiber X , left) and its subdivision (the special fiber X_L , right) where the extra vertex stands for a semistable rational component.

restrict to the case $\kappa_1 = 1 = \kappa_2$ and $\kappa_3 = n$. In this case the simple twist group is $\text{Tw}_{\Gamma}^s = n\mathbb{Z} \oplus n\mathbb{Z}$. The full

twist group is generated by the simple twist group and the element $(1, -1)$. In particular, we note that the quotient twist group is

$$(3.14) \quad K_\Gamma = \text{Tw}_\Gamma / \text{Tw}_\Gamma^s = \mathbb{Z}/n\mathbb{Z}.$$

To work explicitly with invariants, we specialize to the case $n = 3$ in what follows.

The simple level rotation torus is isomorphic to $(\mathbb{C}^*)^2$, hence $\overline{T}_\Gamma^s \cong \mathbb{C}^2$, and a morphism $R^s: B \rightarrow \overline{T}_\Gamma^s$ is given by two functions (t_{-1}, t_{-2}) . Consequently,

$$\begin{aligned} \overline{T}_\Gamma^n &= \overline{T}_\Gamma^s / K_\Gamma = \left\{ (s_{-1}, s_{-2}, f_1, f_2, f_3) : f_1^1 = s_{-1}, f_2^1 = s_{-2}, f_3^3 = s_{-1}s_{-2} \right\} \\ &= \left\{ (f_1, f_2, f_3) : f_3^3 = f_1 f_2 \right\}, \end{aligned}$$

where $s_{-1} = f_1 = t_{-1}^3$, $s_{-2} = f_2 = t_{-2}^3$, and $f_3 = t_{-1}t_{-2}$. In this case the rescaling ensemble R induced by R^s is given by the composition of R^s with the quotient map $\overline{T}_\Gamma^s \rightarrow \overline{T}_\Gamma^s / K_\Gamma$, and has coordinates

$$(s_{-1}, s_{-2}, f_1, f_2, f_3) = (t_{-1}^3, t_{-2}^3, t_{-1}^3, t_{-2}^3, t_{-1}t_{-2}).$$

Let $\omega = (\omega_0, \omega_{-1}, \omega_{-2})$ be a twisted differential on some pointed stable automorphism-free curve (X, z) compatible with the Γ discussed here. By plumbing (see [BCG⁺19, Section 12]), we get a family of curves $\mathcal{X} \rightarrow \overline{T}_\Gamma^n$ with an underlying collection of rescaled differentials

$$\omega_{(0)} = \omega_0, \quad \omega_{(-1)} = s_{-1}\omega_{-1}, \quad \omega_{(-2)} = s_{-1}s_{-2}\omega_{-2}$$

and the rescaling ensemble R .⁽¹⁹⁾

To summarize: *near (X, z, Γ, ω) as above, $\mathcal{GMS}_\mu = \mathcal{MS}_\mu$, since there are no global residue conditions; both functors are representable by an algebraic variety; this algebraic variety is singular with a quotient singularity given by the group K_Γ .*

Finally, we remark that as illustrated in Figure 3, a geometric way to think of the $[\mathbb{P}^1/\mathbb{G}_m]$ subdivision is to modify the definition of level graphs by eliminating all long edges (*i.e.* edges crossing more than one level passage), and instead inserting semistable rational vertices at each level crossed by a long edge, with the same number of prongs. Then the corresponding twist group, level rotation torus, and rescaling ensemble have only their ‘simple’ versions. To see this concretely, suppose $uv = f$ is the local model of a node corresponding to a long edge crossing k level passages, where $f^\kappa = s_{i-k} \cdots s_{i-1}$ as in (3.9). Introduce new parameters u_j, v_j, f_j for $i-k \leq j \leq i-1$ satisfying $u_j v_j = f_j$, $f_j^\kappa = s_j$, $v_j u_{j-1} = 1$, $u_{i-1} = u$, and $v_{i-k} = v$. Then v_j and u_{j-1} are coordinates on the inserted semistable rational vertex at each intermediate level j that can subdivide the long edge into k short edges, where the differential on the semistable rational vertex is $u_{j-1}^\kappa (du_{j-1}/u_{j-1}) = -v_j^{-\kappa} (dv_j/v_j)$ and the number of prongs at each node in the inserted semistable rational curves remains equal to κ .

A logarithmic version of this construction (the replacement of an edge crossing k levels by a chain of $k-1$ projective lines) can be described in terms of the *divided tropical lines* of [MW20]. We will not use this in what follows, but we give a brief sketch. The PL function β with maximum value 0 can be seen as a map from X to the ‘tropical line’ $[\mathbb{P}^1/\mathbb{G}_m]$. Write \mathcal{P} for the subfunctor of $\text{Hom}_{\text{LogSch}}(-, [\mathbb{P}^1/\mathbb{G}_m])$ consisting of log maps to $[\mathbb{P}^1/\mathbb{G}_m]$ such that the images of the vertices of X form a totally ordered set in the characteristic monoid. Then $\mathcal{P} \rightarrow [\mathbb{P}^1/\mathbb{G}_m]$ is a quotient by \mathbb{G}_m of a contraction map from a chain of rational curves to a single rational curve. The fiber product $X \times_{[\mathbb{P}^1/\mathbb{G}_m]} \mathcal{P}$ is then the ‘divided’ curve constructed above, all of whose edges cross at most one level.

⁽¹⁹⁾In fact, replacing the initial datum ω by the universal equisingular deformation inside the appropriate stratum of differentials and taking as new base \overline{T}_Γ^n times the base of the equisingular deformation, we obtain the universal family.

4. The underlying algebraic stack of \mathbf{Rub}

The category \mathbf{Rub} is naturally fibered over \mathbf{LogSch} . Our goal in this section is to understand its underlying algebraic stack (a fibered category over \mathbf{Sch}). We use the notion of minimal log structures from [Gil12] and [Wis16, Appendix B]. We describe here the minimal log structures on points of \mathbf{Rub} , as a variation on the description of the minimal log structures on \mathbf{Div} given in the proof of [MW20, Theorem 4.2.4].

Throughout this section we work with \mathbf{Rub}_0 in place of \mathbf{Rub} , as it is notationally slightly simpler, and fits better with what we do in the rest of the paper. The interested reader will check that the results go through essentially unchanged for \mathbf{Rub} .

4.1. Brief recap on minimal log structures

This is taken from [Wis16, Appendix B], based on [Gil12]. The purpose of minimal log structures is to understand how to pass from a category fibered in groupoids (CFG) X over \mathbf{LogSch} to a CFG over \mathbf{Sch} . The wide subcategory of \mathbf{LogSch} with morphisms the strict morphisms is a CFG over \mathbf{Sch} , so one could just take the corresponding wide subcategory of X and composite. However, this is the ‘wrong’ way to extract the underlying CFG of X over \mathbf{Sch} . For an elementary example, let $X := (\mathrm{pt}, \mathbb{N}^2)$ be a point with log structure \mathbb{N}^2 . Then there are very many maps from $Y := (\mathrm{pt}, \mathbb{N})$ to X : one can choose both the underlying monoid map $\mathbb{N}^2 \rightarrow \mathbb{N}$ and the lift to the log structure giving a \mathbb{C}^* parameter. Hence if we take the CFG over \mathbf{LogSch} associated to X and view it as a CFG over \mathbf{Sch} via taking strict morphisms and then the forgetful functor, we will get a very large and complicated object,⁽²⁰⁾ when what we really wanted was a point!

However, given a map $T \rightarrow \mathrm{pt}$ of schemes, there exist a unique log structure M on T and a morphism $(T, M) \rightarrow X = (\mathrm{pt}, \mathbb{N}^2)$ such that any other log morphism $(T, M') \rightarrow X$ factors through $(T, M) \rightarrow X$. Namely, M is simply the pullback of the log structure \mathbb{N}^2 on pt under $T \rightarrow \mathrm{pt}$. Such a log structure is called *minimal*, and if we take the full subcategory of log schemes over X given by minimal objects, then view it as a CFG over \mathbf{Sch} via the forgetful functor, we recover exactly what we wanted, namely a point.

In the next two subsections, we will apply the same machinery to the CFG \mathbf{Rub}_0 over \mathbf{LogSch} . An object $(X/B, \beta)$ of \mathbf{Rub}_0 is called *minimal* if every solid diagram in \mathbf{Rub}_0

$$(4.1) \quad \begin{array}{ccc} (X'/B', \beta') & \xrightarrow{\quad} & (X/B, \beta), \\ & \searrow & \nearrow \text{dashed} \\ & (X''/B'', \beta'') & \end{array}$$

with the induced maps $\underline{B}' \rightarrow \underline{B}$ and $\underline{B}' \rightarrow \underline{B}''$ on underlying schemes being isomorphisms, admits a unique dashed arrow.

Gillam’s results imply that the full subcategory of \mathbf{Rub}_0 consisting of minimal objects, together with its natural forgetful functor to \mathbf{Sch} , is (equivalent to) the underlying algebraic stack of \mathbf{Rub}_0 . Thus, objects are those log points of \mathbf{Rub}_0 for which the log structure is minimal, and morphisms are simply the usual morphisms of log objects.

As such, if we want to understand the relative inertia of \mathbf{Rub}_0 over \mathfrak{M} , we need to understand not only the minimal objects and their morphisms, but also all possible ways of equipping a schematic object of \mathbf{Rub}_0 with minimal log structure.

Remark 4.1. A warning: suppose that one starts with a CFG over \mathbf{LogSch} which is equivalent to a category fibered in setoids (CFS), and which has enough minimal objects. It is then representable by an algebraic stack with log structure, but this *need not* be equivalent to a category fibered in setoids over schemes (in other words, it can still have non-trivial stacky structure). The most elementary example of this is perhaps the

⁽²⁰⁾For example, the fiber over $\mathrm{pt} \in \mathbf{Sch}$ is the category of pairs of a log structure M on pt and an associated log morphism $(\mathrm{pt}, M) \rightarrow X$.

subdivision of $\mathbb{G}_m^{\text{trop}}$ at 0, which is certainly a category fibered in setoids over **LogSch**, but whose underlying algebraic stack is $[\mathbb{P}^1/\mathbb{G}_m]$. This is because a given schematic point can admit two (or more) different minimal log structures, which can have several isomorphisms between them even if we have a CFS over **LogSch**; the fiber over *any given* log scheme can still have no non-trivial automorphisms.

4.2. Minimal log structures for Rub_0

Let $(X/B, \beta)$ be a point of Rub_0 with X/B nuclear, where \mathbb{M}_B is the sheaf of monoids on B . Recall that from this family and any $b \in B$, we obtain

- the stable graph Γ describing the shape of X_b ;
- the length maps $\delta: E(\Gamma) \rightarrow \overline{\mathbb{M}}_{B,b}$, which we extend to a monoid homomorphism

$$\delta: \mathbb{N}\langle E(\Gamma) \rangle \longrightarrow \overline{\mathbb{M}}_{B,b};$$

- the value map $\beta: V(\Gamma) \rightarrow \overline{\mathbb{M}}_{B,b}^{\text{gp}}$ at vertices, whose image is totally ordered, inducing the level map $\ell: V(\Gamma) \longrightarrow \{0, -1, \dots, -N\} = \{0\} \sqcup L(\Gamma)$;
- the slopes $\kappa: H(\Gamma) \rightarrow \mathbb{Z}$ at half-edges, where, given an edge $e \in E(\Gamma)$ consisting of half-edges h, h' , we set $\kappa_e = |\kappa(h)| = |-\kappa(h')|$ and denote by $E^v = \{e \in E(\Gamma) : \kappa_e > 0\}$ the set of vertical edges and by $E^h = \{e \in E(\Gamma) : \kappa_e = 0\}$ the set of horizontal edges.

For $i \in L(\Gamma)$, we define with (3.5)

$$a_i := \text{lcm}_e \kappa_e,$$

where the lcm runs over the set of all edges e such that $\ell(e^-) \leq i < \ell(e^+)$ (we say such an edge e *crosses level* i). We let $\widetilde{P} := \mathbb{N}\langle p_{-1}, \dots, p_{-N} \rangle$ be the free monoid on $N = |L(\Gamma)|$ generators. Then we can define a map $g: E^v \rightarrow \widetilde{P}$ by

$$(4.2) \quad g(e) := \sum_{i=\ell(e^-)}^{\ell(e^+)-1} \frac{a_i}{\kappa_e} p_i,$$

and extend this map additively to a map $g: \mathbb{N}\langle E^v \rangle \rightarrow \widetilde{P}$. Finally, we let

$$\sigma_i := \beta(v_i) - \beta(v_{i-1}) \in \overline{\mathbb{M}}_{B,b},$$

where v_i is any vertex of level i .

Lemma 4.2. *The difference σ_i is divisible by a_i in $\overline{\mathbb{M}}_{B,b}$.*

Proof. Showing that σ_i is divisible by a_i is exactly equivalent to showing that it is divisible by κ_e for every edge e crossing level i (since we work with saturated monoids, if an element is divisible by two integers, then it is also divisible by their least common multiple). But this is exactly condition (3) in Lemma 2.13. \square

Set $\tau_i := \sigma_i/a_i \in \overline{\mathbb{M}}_{B,b}$ (noting that division in $\overline{\mathbb{M}}_{B,b}$ is unique since $\overline{\mathbb{M}}_{B,b}$ is sharp, integral, and saturated), and define a monoid homomorphism

$$(4.3) \quad \psi: \widetilde{P} \longrightarrow \overline{\mathbb{M}}_{B,b}, \quad \psi: p_i \longmapsto \tau_i.$$

Lemma 4.3. *The triangle*

$$(4.4) \quad \begin{array}{ccc} \mathbb{N}\langle E^v \rangle & \xrightarrow{g} & \widetilde{P} \\ & \searrow \delta & \downarrow \psi \\ & & \overline{\mathbb{M}}_{B,b} \end{array}$$

commutes.

Proof. We compute

$$(4.5) \quad \psi(g(e)) = \psi\left(\sum_i \frac{a_i}{\kappa_e} p_i\right) = \sum_i \frac{a_i}{\kappa_e} \tau_i = \frac{1}{\kappa_e} \sum_i \sigma_i = \frac{1}{\kappa_e} (\beta(v_+) - \beta(v_-)) = \delta_e,$$

where the last equality comes from the fact that β is a PL function. \square

Definition 4.4. We say $(X/B, \beta)$ for B nuclear is *basic* if the natural map

$$\psi \oplus \delta|_{E^h}: \widetilde{P} \oplus \mathbb{N}\langle E^h \rangle \longrightarrow \overline{M}_{B,b}$$

is an isomorphism. In general we say a point of \mathbf{Rub}_0 is *basic* if it is so on a nuclear cover.

Our motivation for introducing this definition lies in Lemma 4.6. The intuition behind the definition is that $\overline{M}_{B,b}$ is precisely big enough to contain the elements that are necessary to accommodate the images of the maps δ , the differences of images of β , and roots of these differences whose existence is required by condition (3) in Lemma 2.13.

Lemma 4.5. *Every point of \mathbf{Rub}_0 comes with a map to a basic object.*

Proof. For $(X/B, \beta)$ a nuclear point of \mathbf{Rub}_0 , we define a sheaf of monoids P on B as the fiber product

$$(4.6) \quad P := (\widetilde{P} \oplus \mathbb{N}\langle E^h \rangle) \times_{\overline{M}_B} M_B.$$

This P comes with a map $P \rightarrow \mathcal{O}_B$, namely the composition of the projection to the second factor M_B and the old log structure map $M_B \rightarrow \mathcal{O}_B$, making it into a log structure.

This uses that, for *any* nuclear point $(X/B, \beta)$, the map $\psi \oplus \delta|_{E^h}$ from the definition above satisfies that the preimage of $0 \in \overline{M}_B$ is $0 \in \widetilde{P} \oplus \mathbb{N}\langle E^h \rangle$. From this it also follows that, at any point b in the closed stratum of B , the stalk at b of the ghost sheaf \overline{P} of P is given by

$$\overline{P}_b = \widetilde{P} \oplus \mathbb{N}\langle E^h \rangle.$$

Now we make (\underline{B}, P) into a point of \mathbf{Rub}_0 : we take the underlying family $\underline{X}/\underline{B}$ of curves, and equip \underline{X} with a log structure making it a log curve over (\underline{B}, P) with length map

$$\widetilde{\delta}: E(\Gamma) \longrightarrow \widetilde{P} \oplus \mathbb{N}\langle E^h \rangle, \quad e \longmapsto \begin{cases} \left(\sum_{i=\ell(e^-)}^{\ell(e^+)-1} \frac{a_i}{\kappa_e} p_i, 0 \right) & \text{for } e \in E^v, \\ (0, e) & \text{for } e \in E^h. \end{cases}$$

With this we obtain a family of log curves $(\widetilde{X}/(\underline{B}, P))$. Using Lemma 2.13 we then lift to a (\underline{B}, P) -point of \mathbf{Rub}_0 by specifying the combinatorial PL function

$$(4.7) \quad \beta: V(\Gamma) \longrightarrow (\widetilde{P} \oplus \mathbb{N}\langle E^h \rangle)^{\text{gp}}, \quad v \longmapsto - \sum_{j=\ell(v)}^{-1} a_j p_j.$$

The construction gives a map from $(X/B, \beta)$ to this basic object $(\underline{B}, P) \rightarrow \mathbf{Rub}_0$. \square

Lemma 4.6. *The \mathbf{Rub}_0 -point $(X/B, \beta)$ is basic if and only if it is minimal.*

Proof. Our proof mimics closely that of [MW20, Theorem 4.2.4]. To show that basic objects are minimal, consider a diagram as in (4.1); the problem comes down to showing there is a unique dashed arrow making the diagram

$$(4.8) \quad \begin{array}{ccc} \overline{M}_{B'} & \xleftarrow{\quad} & \widetilde{P} \oplus \mathbb{N}\langle E^h \rangle \\ & \swarrow \quad \searrow & \\ & \overline{M}_{B''} & \end{array}$$

commute. The existence of this arrow comes from the fact that we can apply the same formula (4.3) to define the map, and the arrow is unique since the image of $\mathbb{N}\langle E \rangle$ has finite index in $\tilde{P} \oplus \mathbb{N}\langle E^h \rangle$, and division is unique in sharp, integral, saturated monoids. Conversely, applying (4.3) shows that every minimal monoid admits a map from one which is basic (and hence minimal), and the definition of minimality furnishes an inverse to this map. \square

Definition 4.7. Let \mathbf{Rub}'_0 be the full subcategory of \mathbf{Rub}_0 whose objects have minimal log structure, viewed as a fibered category over \mathbf{Sch} via forgetting the log structure and the curve.

As explained in Section 4.1, Gillam's minimality machinery immediately yields the main theorem of this section, slightly refining the results of [MW20].

Theorem 4.8. *The underlying algebraic stack of \mathbf{Rub}_0 is given by \mathbf{Rub}'_0 .*

4.3. Smoothness of \mathbf{Rub}_0

With the preparations above, we can now prove Lemma 2.4, stating that the algebraic stack \mathbf{Rub}_0 is smooth.

Proof of Lemma 2.4. We write k for the base ring (which the reader may take to be \mathbb{C} , but we are careful to make this proof work in any characteristic). We equip $\mathrm{Spec}(k)$ with trivial log structure.

We write $\mathbf{Rub}_0^{\mathrm{MW}}$ for the variant of \mathbf{Rub}_0 considered by Marcus and Wise [MW20]; this is the same as our space except that they drop condition (2) of the definition. The map $\mathbf{Rub}_0^{\mathrm{MW}} \rightarrow \mathfrak{N}$ is proven in [MW20, Lemma 4.2.5 and Corollary 5.3.5] to be log étale; hence $\mathbf{Rub}_0^{\mathrm{MW}}$ is log smooth (over $\mathrm{Spec}(k)$ with trivial log structure).

Let p be a geometric point of \mathbf{Rub}_0 mapping to a geometric point p' of $\mathbf{Rub}_0^{\mathrm{MW}}$, and write \bar{P} and \bar{Q} for the stalks of the respective characteristic monoids. The map $\bar{Q} \rightarrow \bar{P}$ is injective and has finite cokernel; it corresponds to taking roots of suitable parameters in order to make condition (2) in the definition of \mathbf{Rub}_0 be satisfied.

The log smoothness of $\mathbf{Rub}_0^{\mathrm{MW}}$ implies that there exist a scheme U and smooth strict morphisms $f: U \rightarrow \mathbf{Rub}_0^{\mathrm{MW}}$ and $g: U \rightarrow \mathrm{Spec}(k[\bar{Q}])$ such that p' lies in the image of f . Since \mathbf{Rub}_0 is obtained from $\mathbf{Rub}_0^{\mathrm{MW}}$ by taking the roots that transform \bar{Q} into \bar{P} , we have a diagram of pullback squares

$$(4.9) \quad \begin{array}{ccccc} \mathbf{Rub}_0 & \longleftarrow & V & \longrightarrow & \mathrm{Spec}(k[\bar{P}]) \\ \downarrow & & \downarrow & & \downarrow \\ \mathbf{Rub}_0^{\mathrm{MW}} & \longleftarrow & U & \longrightarrow & \mathrm{Spec}(k[\bar{Q}]). \end{array}$$

Now Lemma 4.4, Lemma 4.6, and Lemma 4.8 together imply that the stalks of the characteristic monoid of \mathbf{Rub}_0 are free monoids of finite rank; in other words, $\mathrm{Spec}(k[\bar{P}])$ is an affine space over k , in particular is smooth over k . \square

Note that the base-change $\mathbf{Rub}_{\mathcal{L}}$ is *not* in general smooth, except in genus zero (when the map $\mathfrak{N} \rightarrow \mathrm{Pic}$ is an open immersion, hence smooth). For example, it can have many non-reduced irreducible components; see [HS21]. In particular, the smoothness of the main component of $\mathbf{Rub}_{\mathcal{L}_\mu}$ (proven in [BCG⁺19] granting the verification that the spaces named $\Xi \overline{\mathcal{M}}_{g,n}(\mu)$ there and in Lemma 3.11 indeed agree) does not follow directly from Lemma 2.4 outside of genus zero.

4.4. Relative automorphisms

As a log stack, \mathbf{Rub}_0 has trivial automorphisms relative to the stack of log curves. But (as discussed in Lemma 4.1) this does not mean that the underlying algebraic stack of minimal objects has trivial

automorphisms. Rather, they come from automorphisms of the log structure; the following remark makes this precise.

Remark 4.9. In general, given a map $\mathcal{X} \rightarrow \mathcal{Y}$ of log stacks with underlying stacks $\underline{\mathcal{X}}, \underline{\mathcal{Y}}$ and a point $\underline{x}: \text{Spec}(\mathbb{C}) \rightarrow \underline{\mathcal{X}}$, we can ask: what is the relative inertia of \underline{x} over $\underline{y} = (\underline{\mathcal{X}} \rightarrow \underline{\mathcal{Y}}) \circ \underline{x}$? For this, let $(\text{Spec}(\mathbb{C}), M_x) \rightarrow \mathcal{X}$ and $(\text{Spec}(\mathbb{C}), M_y) \rightarrow \mathcal{Y}$ be the minimal log structures lifting $\underline{x}, \underline{y}$. Then by minimality of M_y , the composition $(\text{Spec}(\mathbb{C}), M_x) \rightarrow \mathcal{X} \rightarrow \mathcal{Y}$ must factor through a map

$$f: (\text{Spec}(\mathbb{C}), M_x) \longrightarrow (\text{Spec}(\mathbb{C}), M_y).$$

Such a map is uniquely described by a monoid map $M_y \rightarrow M_x$ over $(\mathbb{C}, \times) = (\mathcal{O}_{\text{Spec}(\mathbb{C})}, \times)$. Then the desired group of automorphisms is just the group of those automorphisms of M_x that are constant on the image of M_y and commute with the map to (\mathbb{C}, \times) .

Returning to our situation, the ‘tropical’ part of the log structure (the ghost sheaf \overline{M}) has no non-trivial automorphisms relative to the stack of log curves. Thus the automorphisms all arise from automorphisms of the log structure M that are trivial on \overline{M} and trivial on the structure sheaf. So they are really automorphisms of the extension structure of M .

4.5. The worked example again

Let $(X/\mathbb{C}, \beta \in \overline{M}_X^{\text{gp}})$ be a point of \mathbf{Rub}_0 with the underlying enhanced level graph given by Figure 3, still restricting to the case $\kappa_1 = \kappa_2 = 1$ and $\kappa_3 = n$. We would like to understand the relative inertia of this point of \mathbf{Rub}_0 over \mathfrak{M} .

The minimal monoid on \mathbb{C} for the curve X/\mathbb{C} is just $\mathbb{N}\langle E \rangle = \mathbb{N}\langle e_1, e_2, e_3 \rangle$, and the minimal monoid as a point in \mathbf{Rub}_0 is given by $\widetilde{P} = \mathbb{N}\langle p_{-1}, p_{-2} \rangle$, with one generator p_i for each *level* i (there are no horizontal edges in this example; otherwise, they should also appear in this minimal monoid). The natural map is then given by

$$g: \mathbb{N}\langle E \rangle \longrightarrow \widetilde{P}, \quad e_1 \mapsto np_{-1}, \quad e_2 \mapsto np_{-2}, \quad e_3 \mapsto p_{-1} + p_{-2}.$$

To see this, note that $a_1 = a_2 = n$, and then apply Formula (4.2). Note that there are no non-trivial automorphisms of \widetilde{P} that act as the identity on the image of g . The map g extends in the obvious manner to a map on the stalks of the log structures

$$\mathbb{N}\langle E \rangle \oplus \mathbb{C}^* \longrightarrow P = \widetilde{P} \oplus \mathbb{C}^*,$$

and the relative inertia is then given by the automorphisms of $\widetilde{P} \oplus \mathbb{C}^*$ which act as the identity on the image of $\mathbb{N}\langle E \rangle \oplus \mathbb{C}^*$, and which lie over the identity map on \widetilde{P} (since any automorphism of \widetilde{P} constant on the image of g must be the identity). Such an automorphism is defined on $((1, 0), 1)$ and $((0, 1), 1)$ by

$$((1, 0), 1) \mapsto ((1, 0), u) \quad \text{and} \quad ((0, 1), 1) \mapsto ((0, 1), v)$$

for some $u, v \in \mathbb{C}^*$ satisfying

- (1) $u^n = 1$, because $n((1, 0), 1) = ((n, 0), 1^n)$ lies in the image of $\mathbb{N}\langle E \rangle \oplus \mathbb{C}^*$ and is thus fixed;
- (2) $v^n = 1$ for the analogous reason;
- (3) $uv = 1$, because $((1, 1), 1)$ lies in the image of $\mathbb{N}\langle E \rangle \oplus \mathbb{C}^*$ and is thus fixed.

Such a choice of u, v evidently determines such an automorphism. (Or, more precisely, there are two canonical isomorphisms with the roots of unity, one coming from ‘above’ and the other from ‘below’ on the graph, and the composite of these isomorphisms is the inversion map on the group of roots of unity).

We conclude that the *relative inertia for this triangle graph is equal to the group K_Γ computed in (3.14)*.

5. From logarithmic to multi-scale

In this section we construct the morphism of functors $F: \mathbf{Rub}_{\mathcal{L}_\mu} \rightarrow \mathcal{GMS}_\mu$ whose existence is claimed in Lemma 1.1, and then prove that theorem. At the end of the section, we include two related results about describing the multi-scale space as a Zariski closure and describing a morphism from the rubber space to the Hodge bundle, which can be of independent interest.

Let $(X/B, \beta \in \Gamma(X, \overline{M}_X^{\text{gp}}), \varphi: \mathcal{O}_X(\beta) \xrightarrow{\sim} \mathcal{L}_\mu) \in \mathbf{Rub}'_{\mathcal{L}_\mu}$. Recall that the prime on \mathbf{Rub} indicates that we are working with the minimal log structure as described in Section 4, and that we always work with *saturated* log structures.

We assume moreover for now, and for most of this section, that X/B is nuclear, and explain at the end why the functor glues to general families.

5.1. The enhanced level graph

The first item to build the F -image of $(X/B, \beta, \varphi)$ is an enhanced level graph. As the underlying graph Γ , we simply take the dual graph of the curve fiber over the closed stratum of B . The level structure, given in terms of a normalized level function, comes from $\beta \in \overline{M}_X^{\text{gp}}(X)$ as explained in (2.4). The definition of the enhancement κ is given in (2.7), where the divisibility required for this definition has been proven in Lemma 2.11. The stability condition just comes from the fact that we work with stable curves.

Given a vertex v and the corresponding component X_v of the central fiber, the admissibility of κ comes down to the equalities

$$(5.1) \quad 2g(v) - 2 + \#H'(v) - \sum_{j \mapsto v} m_j = \deg(\mathcal{L}_\mu|_{X_v}) = \sum_{h \mapsto v} \kappa_h$$

(recall that $H'(v)$ denotes the set of non-leg half-edges attached to a vertex v). The first equality is immediate from the definition of \mathcal{L}_μ , and the second comes from the isomorphism φ and a computation of the degree of $\mathcal{O}_X(\beta)$ on the component X_v presented in Lemma 2.15.

The dual graph $\Gamma_{b'}$ of the fiber over a general b' (possibly outside the closed stratum) comes with a level structure obtained from Γ by undegeneration (as defined in Section 3.1), by the Key Property (4) of nuclear log curves from Section 2.3. Constructing the rest of the data of a generalized multi-scale differential requires more work, which we now begin.

5.2. Logarithmic splittings and rotations

We write $\widetilde{P} = \mathbb{N}\langle p_{-1}, \dots, p_{-N} \rangle$ as in Section 4.

Definition 5.1. A *log splitting* is a map

$$(5.2) \quad \widetilde{\psi}: \widetilde{P} \longrightarrow M_B$$

whose composition with the canonical map $M_B \rightarrow \overline{M}_{B,b}$ is the map $\psi: \widetilde{P} \hookrightarrow \overline{M}_{B,b}$ from (4.3) (recall that we work throughout this section with minimal objects).

The *simple log level rotation torus* T_{\log}^s , abbreviated as *simple LLRT*, is the set of log splittings.

Remark 5.2. Let us unpack the simple log level rotation torus definition a bit. Recall our key exact sequence (2.2). The presence of the gp is not so important, as we always work with integral monoids (*i.e.* monoids which inject into their groupifications). Consequently, a choice of a splitting is essentially a choice of an invertible function on B (which we think of as a scalar) for every level below 0. Pre-composing $\widetilde{\psi}$ with the map g from (4.2) and using Lemma 4.3, we then also obtain a lift of the map δ , *i.e.* a choice of a scalar for each edge. These must satisfy appropriate compatibility equations, and the saturation condition also imposes the existence of certain roots.

Definition 5.3. The *simple log rotation group* is the group

$$\mathrm{Hom}_{\mathrm{mon}}(\widetilde{P}, \mathcal{O}_B^\times(B)) = \mathrm{Hom}_{\mathrm{gp}}(\widetilde{P}^{\mathrm{gp}}, \mathcal{O}_B^\times(B)),$$

where the identification stems from the universal property of the groupification.⁽²¹⁾

We define an action of an element φ of the simple log rotation group on the simple log rotation torus by the formula

$$(5.3) \quad (\varphi \cdot \widetilde{\psi})(p) := \varphi(p)\widetilde{\psi}(p) \quad \text{for } p \in \widetilde{P}.$$

Lemma 5.4. *Via the action (5.3), the simple LLRT is either empty, or a torsor for the simple log rotation group. After possibly shrinking B , we can ensure the existence of a log splitting.*

Recall that a pseudo-torsor is a space with a free transitive action, but unlike a torsor, it may be empty (here, if the base B is too large to support the appropriate sections). Thus Lemma 5.4 says that the simple LLRT is a pseudo-torsor.

Proof. In the exact sequence

$$1 \longrightarrow H^0(B, \mathcal{O}_B^\times) \longrightarrow H^0(B, \mathcal{M}_B^{\mathrm{gp}}) \longrightarrow \underbrace{H^0(B, \overline{\mathcal{M}}_B^{\mathrm{gp}})}_{= \overline{\mathcal{M}}_{B,b}^{\mathrm{gp}}} \longrightarrow H^1(B, \mathcal{O}_B^\times) \longrightarrow \cdots,$$

if all elements $\psi(p_i) = \tau_i \in \overline{\mathcal{M}}_{B,b}$ map to zero in $H^1(B, \mathcal{O}_B^\times)$, then they have preimages in $H^0(B, \mathcal{M}_B)$; by the freeness of $\overline{\mathcal{M}}_{B,b}$, this implies the existence of a log splitting. Any such choices of preimages differ precisely by elements in $H^0(B, \mathcal{O}_B^\times)$, which together define an element of the simple log rotation group. Thus the action of this group is free and transitive.

Finally, if the elements $\tau_i \in \overline{\mathcal{M}}_{B,b}$ do *not* map to zero in $H^1(B, \mathcal{O}_B^\times) = \mathrm{Pic}(B)$, we can always find an open neighbourhood B_0 of $b \in B$ on which these line bundles are trivial after all. Then on B_0 , the long exact sequence and the argument above show the existence of a lift, finishing the proof. \square

5.3. Log viewpoint on smoothing and rescaling parameters

In this subsection we construct the rescaling ensemble from the choice of a log splitting, and provide auxiliary statements about the smoothing and rescaling functions contained in the ensemble.

Let $\widetilde{\psi}: \widetilde{P} \rightarrow \mathcal{M}_B$ be a log splitting. Recall the definitions of the maps $g: \mathbb{N}\langle E^v \rangle \rightarrow \widetilde{P}$ from (4.2) and $\alpha: \mathcal{M}_B \rightarrow \mathcal{O}_B$ from the definition of a log scheme.

Definition 5.5. The *smoothing parameter associated to a vertical edge $e \in E^v(\Gamma)$ by the log splitting $\widetilde{\psi}$* is

$$(5.4) \quad f_e := (\alpha \circ \widetilde{\psi} \circ g)(e).$$

Fix a level $i \in L(\Gamma)$. The *level parameter* and *rescaling parameter* associated to i by $\widetilde{\psi}$ are

$$(5.5) \quad t_i := (\alpha \circ \widetilde{\psi})(p_i) \quad \text{and} \quad s_i := (\alpha \circ \widetilde{\psi})(a_i p_i).$$

The collection of functions $\mathbf{t} = (t_i)_{i \in L(\Gamma)}$ defines a map $R^s: B \rightarrow \overline{T}_\Gamma^s$ to the closure of the simple level rotation torus, which is just \mathbb{C}^N , and a rescaling parameter $s_i = r_i \circ \pi \circ R^s$ in the notation of Section 3.4.

Lemma 5.6. *The morphism $R^s: B \rightarrow \overline{T}_\Gamma^s$ defined above is a simple rescaling ensemble.*

⁽²¹⁾Note that there is also a (non-simple) log rotation group, consisting of the set of compatible choices of elements in $\mathcal{O}_B^\times(B)$ for all $e \in E^v$ and the elements $\sigma_i = \beta(v_i) - \beta(v_{i-1})$. Since this non-simple group will not be needed in the following, we do not give a formal definition.

Proof. By Lemma 3.4 we must verify that the functions f_e from (5.4) are indeed smoothing parameters for their respective nodes, lying in the correct equivalence class in $\mathcal{O}_B/\mathcal{O}_B^\times$. To see this, consider the following diagram:

$$\begin{array}{ccccc}
 \mathbb{N}\langle E^v \rangle & \xrightarrow{g} & \tilde{P} & & \\
 \downarrow \delta & & \downarrow \tilde{\psi} & \swarrow & \swarrow \\
 & & M_B & \xrightarrow{\alpha} & \mathcal{O}_B \\
 & \swarrow & \nwarrow & \nwarrow & \nwarrow \\
 \overline{M}_{B,b} & \xrightarrow{\bar{\alpha}} & \mathcal{O}_B/\mathcal{O}_B^\times & & \mathcal{O}_B^\times
 \end{array}$$

What we must show is that $f_e = (\alpha \circ \tilde{\psi} \circ g)(e) \in \mathcal{O}_B$ maps to the class of a smoothing parameter in $\mathcal{O}_B/\mathcal{O}_B^\times$. Now the commutativity of the upper-left quadrilateral follows from Lemma 4.3 and the assumption that $\tilde{\psi}$ lifts the map $\psi: \tilde{P} \rightarrow \overline{M}_{B,b}$. On the other hand, the map $\bar{\alpha}$ is just *defined* to make the lower quadrilateral commute. Then we have

$$[f_e] = (\alpha \circ \tilde{\psi} \circ g)(e) = \bar{\alpha}(\delta(e)) \in \mathcal{O}_B/\mathcal{O}_B^\times.$$

The fact that $\bar{\alpha}$ maps $\delta(e)$ to a smoothing parameter for the node associated to e is then just a basic property of families of log curves; see Key Property (2) of Section 2.3. \square

5.4. The collection of rescaled differentials

By the definition of lying in $\mathbf{Rub}_{\mathcal{L}_\mu}$, we are given an isomorphism

$$(5.6) \quad \varphi: \omega_{X/B} \left(- \sum_{k=1}^n m_k z_k \right) \xrightarrow{\sim} \mathcal{O}_X(\beta).$$

On the other hand, it follows from the definition of ψ that the element $-\sum_{j=i}^{-1} a_j p_j \in \tilde{P}^{\text{gp}}$ maps to $\beta(v_i) \in \overline{M}_{B,b}^{\text{gp}}$ under ψ , where $v_i \in V(\Gamma)$ is any vertex on level i . Using the log splitting $\tilde{\psi}$, we obtain the elements

$$o_i := \tilde{\psi} \left(- \sum_{m=i}^{-1} a_m p_m \right) \in M_B^{\text{gp}}$$

in the preimage of $\beta(v_i)$. Since this preimage can be identified as the complement of the zero section in $\mathcal{O}_B(\beta(v_i))$, we can see o_i as a nowhere-vanishing section of $\mathcal{O}_B(\beta(v_i))$.

Adapting the convention from Lemma 3.5 to the family X/B , we define $X_{(>i)} \subseteq X$ to be the closed subset of components of fibers $X_{b'}$ whose closure in X does not intersect any component of the central fiber X_b at a level less than or equal to i . Then we define $U_i = X \setminus (X_{(>i)} \cup Z^\infty)$, where $Z^\infty \subseteq X$ is the image of sections associated to marked poles. Then we claim that there is a well-defined map

$$(5.7) \quad w_i: \pi^* \mathcal{O}_B(\beta(v_i))|_{U_i} \longrightarrow \mathcal{O}_X(\beta)|_{U_i}.$$

Indeed, the left-hand side is the line bundle on U_i associated to the piecewise linear function which is *constant*, equal to $\beta(v_i)$. Since we removed $X_{(>i)}$, this function dominates the function β on the right, so we have a map as desired. Thus $w_i(\pi^* o_i)$ gives a section of $\mathcal{O}_X(\beta)$ on U_i , and we define

$$(5.8) \quad \omega_{(i)} := \varphi^{-1} w_i(\pi^* o_i) \in H^0 \left(U_i, \omega_{X/B} \left(- \sum_{k=1}^n m_k z_k \right) \right).$$

We check that $\omega_{(i)}$ satisfies the conditions in Lemma 3.5 and that the smoothing and rescaling parameters f_e and s_i defined in (5.4) and (5.5) (and thus the simple rescaling ensemble R^s) are compatible with these generalized rescaled differentials.

- (1) For any levels $j < i < 0$, there is a natural map of line bundles $\mathcal{O}_B(\beta(v_i)) \rightarrow \mathcal{O}_B(\beta(v_j))$. On the level of sections, we then have

$$o_i = \tilde{\psi}\left(-\sum_{m=i}^{-1} a_m p_m\right) = \tilde{\psi}\left(-\sum_{m=j}^{-1} a_m p_m\right) \cdot \prod_{m=j}^{i-1} \tilde{\psi}(a_m p_m) \mapsto o_j \cdot \prod_{m=j}^{i-1} \tilde{\psi}(a_m p_m).$$

Via the isomorphism φ^* , and using that $s_m = \alpha(\tilde{\psi}(a_m p_m))$, this becomes the desired equality $\omega_{(i)} = \omega_{(j)} \cdot \prod_{m=j}^{i-1} s_m$. The fact that s_i vanishes at the closed point of B comes from the fact that the map of line bundles is the zero map when restricted to the fibers over the closed point of B .

- (2) Choose local coordinates u_e and v_e so that X is locally given by $u_e v_e = f_e$. Then the isomorphism φ corresponds near the node to a non-vanishing section of $\omega_{X/B}(-\beta)$. Now ω has a local generating section which is given after inverting u_e by $\frac{du_e}{u_e}$ and after inverting v_e by $-\frac{dv_e}{v_e}$, and (perhaps after adjusting the choices of coordinate u_e and v_e) $\mathcal{O}_X(-\beta)$ has a local generating section which is given after inverting u_e by $u_e^{\kappa_e} o_i$ and after inverting v_e by $o_j v_e^{-\kappa_e}$. By dividing the just-described local generating section by the one given by the isomorphism φ , we obtain a local unit λ such that

$$(5.9) \quad \omega_{(i)} = \lambda u_e^{\kappa_e} \frac{du_e}{u_e} \quad \text{and} \quad \omega_{(j)} = -\lambda v_e^{-\kappa_e} \frac{dv_e}{v_e}.$$

- (3) On the normalization Y_i of the union of all components of the special fiber X_b sitting at level i , we have (see Lemma 2.15)

$$\mathcal{O}_X(\beta)|_{Y_i} = \pi^* \mathcal{O}_B(\beta(v_i)) \otimes_{\mathcal{O}_{Y_i}} \mathcal{O}_{Y_i}\left(\sum_h \kappa_h h\right),$$

where the sum is taken over all non-leg half-edges h attached to the vertices at level i . Pulling back via φ^* , the line bundle on the left becomes

$$\omega_{X_b}\left(-\sum_{k=1}^n m_k z_k\right)|_{Y_i} = \omega_{Y_i}\left(-\sum_{k=1}^n m_k z_k + \sum_h h\right).$$

Tensoring with $\mathcal{O}_{Y_i}(-\sum_h \kappa_h h)$ on both sides, we then get

$$\omega_{Y_i}\left(-\sum_{k=1}^n m_k z_k - \sum_h (\kappa_h - 1)h\right) \cong \pi^* \mathcal{O}_B(\beta(v_i)).$$

Seeing $\omega_{(i)}$ as a meromorphic section on the left, it then corresponds to the nowhere-vanishing section $\pi^* o_i$ on the right. Thus it extends to all of Y_i on the left-hand side. But then this extension seen as a meromorphic section of ω_{Y_i} has the desired order m_k at the marked points z_k and $\kappa_h - 1$ at the preimage of the node associated to h .

5.5. Prong-matchings

To recall the notion of a prong-matching, consider a vertical edge $e \in E^v$, and let $B_e \hookrightarrow B$ be the closed subscheme of B over which the node e persists, *i.e.* the vanishing locus of the smoothing parameter f_e .

The sections q^\pm of the two preimages of the node identify B_e as a subscheme of the blowup of $X \times_B B_e$ along the section corresponding to e . Recalling (3.3), we let $\mathcal{N}_e^\vee = (q^+)^* \omega_{X_+} \otimes (q^-)^* \omega_{X_-}$ be the corresponding line bundle on B_e . Then a local prong-matching at e is a section σ_e of \mathcal{N}_e^\vee such that $\sigma_e^{\kappa_e}(\tau_e) = 1$ for the section $\tau_e \in \mathcal{N}_e^{\kappa_e}$ defined in Lemma 3.1.

To identify this notion in the logarithmic context, recall that we have the element $\delta(e) \in \overline{M}_{B,b}$. Then the bundle \mathcal{N}_e^\vee has an interpretation as follows.

Lemma 5.7. *There are canonical isomorphisms of line bundles*

$$(5.10) \quad \mathcal{O}_B(\delta(e))|_{B_e} = \mathcal{N}_e^\vee$$

for each edge e . Moreover, let $\mathfrak{f} \in M_B$ be an element mapping to $\delta(e) \in \overline{M}_{B,b}$, so that we can see it as a section of $\mathcal{O}_B(\delta(e))$. Then the function $f = \alpha(\mathfrak{f}) \in \mathcal{O}_B$ is a smoothing parameter for the node associated to e . Let u, v be local coordinates around the node on X such that the Henselized local ring at the node is the Henselization of $\mathcal{O}_B[u, v]/(uv - f)$. Then the isomorphism (5.10) sends the section $\mathfrak{f}|_{B_e} \in \mathcal{O}_B(\delta(e))|_{B_e}$ to

$$du \otimes dv \in \mathcal{N}_e^\vee = (q^+)^* \omega_{X_+} \otimes (q^-)^* \omega_{X_-}.$$

Proof. Since both sides commute with base-change, it is enough to check this in the universal case, in which the log structure is divisorial coming from the boundary (and the map α of the log structure is injective, so there are no non-trivial automorphisms of the log structure). Over a versal deformation R , the local equation of the node is given by $R[u, v]/(uv - f)$, where $f \in R$ is an element corresponding to $\delta(e)$. So $\mathcal{O}_B(\delta(e))$ is canonically identified with the ideal sheaf generated by f in R (cf. the appendix). On the other hand, \mathcal{N}_e^\vee is canonically identified with the conormal bundle in R to the locus $f = 0$ (see [ACG11, Section XIII.3]) and thus agrees with $\mathcal{O}_B(\delta(e))|_{Z_e}$. Tracing through the constructions of these canonical identifications yields the second part of the lemma; alternatively, this can be seen as a very slight generalization of [Edi98, Section 4], where his $c(x)$ corresponds to the element $du \otimes dv$ and his $\pi^{x(e)}$ to the element f . \square

Let $\tilde{\psi}: \tilde{P} \rightarrow M_B$ be a log splitting, and let e be a vertical edge. By Lemma 4.3 the element $(\tilde{\psi} \circ g)(e) \in M_B$ maps to $\delta(e) \in \overline{M}_{B,b}$ and hence lies in $\mathcal{O}_B^\times(\delta(e)) \subseteq M_B$ (by the definition of this bundle via (2.2)). Applying the isomorphism of (5.10), we thus obtain a section of \mathcal{N}_e^\vee .

Definition 5.8. We call the section $\sigma_e = (\tilde{\psi} \circ g)(e)|_{B_e} \in H^0(B_e, \mathcal{N}_e^\vee)$ the *local prong-matching* $\sigma_e = \sigma_e(\tilde{\psi})$ at e determined by the log splitting. The collection $\sigma = (\sigma_e)_{e \in E^v(\Gamma)}$ of these is called the *global prong-matching determined by the log splitting*.

There are two compatibility statements to check for this definition: to get a prong-matching, see the discussion after (3.3), and to make this part of a multi-scale differential, see Lemma 3.8(2b).

Lemma 5.9. *The prong-matching σ determined by any log splitting is indeed a prong-matching in the sense of Lemma 3.2.*

Proof. Assume that the vertical edge e connects levels $i > j$ in Γ . From Lemma 3.2, we need to show that $\sigma_e^{\kappa_e}(\tau_e) = 1$, where τ_e is the section of $\mathcal{N}_e^{\kappa_e}$ defined as $\tau_e = (q^+)^* \omega_{(i)}^{-1} \otimes (q^-)^* \omega_{(j)}$.

On the other hand, the differentials $w_{(i)}$ and $w_{(j)}$ are also determined in (5.8) by the formulae

$$\omega_{(i)} = \varphi^* w_i \left(\pi^* \tilde{\psi} \left(- \sum_{m=i}^{-1} a_m p_m \right) \right) \quad \text{and} \quad \omega_{(j)} = \varphi^* w_j \left(\pi^* \tilde{\psi} \left(- \sum_{m=j}^{-1} a_m p_m \right) \right).$$

We put this into the formula for τ_e ; the pullbacks $(q^+)^*, (q^-)^*$ cancel the pullback π^* . Interpreting τ_e as a section of $\mathcal{O}_B(-\kappa_e \delta(e))$ via (5.10), we thus have

$$\tau_e = \tilde{\psi} \left(\sum_{m=i}^{-1} a_m p_m - \sum_{m=j}^{-1} a_m p_m \right) = \tilde{\psi} \left(- \sum_{m=j}^{i-1} a_m p_m \right) = \tilde{\psi}(-\kappa_e g(e)) = \sigma_e^{-\kappa_e}.$$

Here in the second to last equality, we used the definition of g from (4.2). This finishes the proof that $\sigma_e^{\kappa_e}(\tau_e) = 1$, and thus that σ_e is a local prong-matching. \square

Lemma 5.10. *Let $\tilde{\psi}: \tilde{P} \rightarrow M_B$ be a log splitting and e a non-semi-persistent vertical node (i.e. $f_e^{K_e} \neq 0$). Then the local prong-matching determined by $\tilde{\psi}$ is equal to that induced in Lemma 3.7.*

Proof. The local prong-matching σ_e of Lemma 3.7 is constructed by writing the local equation of the node as $uv = f_e$ and setting

$$\sigma_e := du \otimes dv \in \mathcal{N}_e^\vee = (q^+)^* \omega_{X_+} \otimes (q^-)^* \omega_{X_-}.$$

On the other hand, the local prong-matching $\tilde{\sigma}_e$ associated to e by $\tilde{\psi}$ is given by applying the isomorphism in Lemma 5.7 to the element $(\tilde{\psi} \circ g)(e)$.

Recalling that $f_e = (\alpha \circ \tilde{\psi} \circ g)(e)$, we see that the desired equality $\sigma_e = \tilde{\sigma}_e$ is then the second part of Lemma 5.7. \square

5.6. Morphism of functors from rubber to multi-scale

We put the above together to build a morphism of functors after restricting to base schemes which are locally of finite type over \mathbb{C} (this restriction is harmless since $\mathbf{Rub}_{\mathcal{L}_\mu}$ is representable by a locally finite-type Deligne–Mumford stack). We first make this construction in a local situation, then globalize. We start with a family $(X/B, \beta \in \overline{M}_X(X), \varphi)$, which we take to have minimal saturated log structure, and which is both nuclear and controlled. This immediately gives us the structure of an enhanced level graph. We *choose* a log splitting $\tilde{\psi}: \tilde{P} \rightarrow M_B$ (perhaps after shrinking B). This determines a simple rescaling ensemble, a collection of rescaled differentials, and induces local prong-matchings at each node. Hence we have a simple multi-scale differential.

We next claim that a different choice of log splitting yields an isomorphic simple multi-scale differential, together with a choice of isomorphism. Indeed, by Lemma 5.4 any two log splittings differ by the action of the simple LLRT, and one checks easily that the action of the simple LLRT corresponds to the action of the simple level rotation torus.

For general B locally of finite type, a family of multi-scale differentials is defined as a family of multi-scale differentials on a nuclear controlled cover, compatible on overlaps.

It is clear from the constructions that the above map is independent of choices and is compatible with shrinking the base B ; more precisely, given a map $B' \rightarrow B$ and a family of log differentials on B , we can either first apply our map (obtaining a family of multi-scale differentials on B) and then pull back to B' , or first pull back and then apply our map; unravelling the definitions yields that the results are canonically isomorphic. By descent we obtain a global morphism of functors $F: \mathbf{Rub}_{\mathcal{L}_\mu} \rightarrow G\Xi \overline{\mathcal{M}}_{g,n}(\mu)$.

5.7. Showing the map of functors induces an isomorphism

The above construction gives a morphism from the logarithmic space to the multi-scale space. In this section we complete the proof of Lemma 1.1 by showing that this functor induces an isomorphism.

Theorem 5.11. *The morphism*

$$(5.11) \quad F: \mathbf{Rub}_{\mathcal{L}_\mu} \longrightarrow G\Xi \overline{\mathcal{M}}_{g,n}(\mu).$$

is an isomorphism.

Proof. Given a map $B \rightarrow G\Xi \overline{\mathcal{M}}_{g,n}(\mu)$ with the implicit stable curve over B being controlled, we show that there exists a unique map $B \rightarrow \mathbf{Rub}_{\mathcal{L}_\mu}$ making the following diagram commute:

$$(5.12) \quad \begin{array}{ccc} B & & \\ \downarrow & \searrow & \\ \mathbf{Rub}_{\mathcal{L}_\mu} & \longrightarrow & G\Xi \overline{\mathcal{M}}_{g,n}(\mu). \end{array}$$

The claimed isomorphism in the global situation then follows by descent. Let $(\pi: X \rightarrow B, z, \Gamma, R^s, \omega, \sigma)$ be the simple multi-scale differential corresponding to $B \rightarrow G\Xi\overline{\mathcal{M}}_{g,n}(\mu)$. Given $i \in L(\Gamma)$, we write $t_i \in \mathcal{O}_B(B)$ for the composition with the appropriate coordinate projections $B \rightarrow \overline{T}^s \rightarrow \mathbb{C}$.

Let M_B be the minimal log structure making X/B into a log curve; in particular, its characteristic monoid $\overline{M}_{B,b}$ is canonically identified with the free monoid $\mathbb{N}\langle E \rangle$ on the edges of Γ . For each edge e , the local prong-matching σ_e at e is by Lemma 5.7 a section of the fiber of the bundle $\mathcal{O}(\delta_e)$ over the locus B_e where the node persists, and we choose a section in $M_B(B)$ lifting f_e and restricting over B_e to σ_e , yielding a splitting

$$\mathfrak{f}: \overline{M}_{B,b} \longrightarrow M_B.$$

Denote by $\widetilde{P} := \langle p_{-1}, \dots, p_{-N} \rangle$ the free monoid on the levels, as usual, and define

$$(5.13) \quad t: \widetilde{P} \longrightarrow \mathcal{O}_B, \quad p_i \longmapsto t_i,$$

and

$$(5.14) \quad t': \widetilde{P} \oplus \mathbb{N}\langle E^h \rangle \longrightarrow \mathcal{O}_B,$$

acting as t on the first summand and as \mathfrak{f} on the second.

Then let

$$(5.15) \quad g': \mathbb{N}\langle E \rangle \longrightarrow \widetilde{P} \oplus \mathbb{N}\langle E^h \rangle$$

be the map given by g on the vertical edges and by the identity on the horizontal edges.

The equalities

$$(5.16) \quad f_e = \prod_{i=\ell(e^-)}^{\ell(e^+)-1} t_i^{\frac{a_i}{k_e}},$$

which come from (3.7) (where coordinates were denoted by $f_e = \rho_e$ and $t_i = q_i$), imply that the diagram

$$(5.17) \quad \begin{array}{ccc} M_B & \xrightarrow{\alpha} & \mathcal{O}_B \\ \mathfrak{f} \uparrow & & \uparrow t' \\ \overline{M}_B & \xrightarrow{g'} & \widetilde{P} \oplus \mathbb{N}\langle E^h \rangle \end{array}$$

commutes.

Now we define a sheaf of monoids P as the pushout

$$(5.18) \quad \begin{array}{ccc} M_B & \longrightarrow & P \\ \mathfrak{f} \uparrow & & \uparrow \\ \overline{M}_B & \xrightarrow{g'} & \widetilde{P} \oplus \mathbb{N}\langle E^h \rangle, \end{array}$$

which by the commutativity of the previous diagram comes with a map $\alpha_P: P \rightarrow \mathcal{O}_B$. One checks easily that P is in fact a log structure on B , with characteristic sheaf $\overline{P} = \widetilde{P} \oplus \mathbb{N}\langle E^h \rangle$ at a point $b \in B$ in the closed stratum. The map $M_B \rightarrow P$ gives $X/(B, P)$ the structure of a log curve, and mapping a vertex v of level i to the element

$$(5.19) \quad \left(-\sum_{j=i}^{-1} a_j p_j, 0 \right) \in (\widetilde{P} \oplus \mathbb{N}\langle E^h \rangle)^{\text{gp}}$$

defines a map $\beta: V \rightarrow \overline{P}^{\text{gp}}$ so that the pair $(X/B, \beta)$ is a (minimal) point of **Rub**.

To lift this point to a point of $\mathbf{Rub}_{\mathcal{L}_\mu}$, we need to build an isomorphism of line bundles

$$(5.20) \quad \mathcal{O}_X(\beta) \xrightarrow{\sim} \omega_{X/B} \left(- \sum_{i=1}^n m_i z_i \right).$$

We first define this map on the smooth locus; let $p \in B$ and let $x \in X_p$ be a smooth point of X_p , lying in the component associated to a vertex $v \in \Gamma$. Then the image of β in $\overline{M}_{X,x}^{\text{gp}} = \overline{P}_p^{\text{gp}}$ is given by $\beta(v)$. Our splitting $\overline{P} \rightarrow P$ from (5.18) extends to $\overline{P}^{\text{gp}} \rightarrow P^{\text{gp}}$, and thus $\beta(v)$ maps to a unique section of $\mathcal{O}_B(\beta(v)) \subseteq P$. Then we define

$$(5.21) \quad \mathcal{O}_X(\beta)_x \xrightarrow{\sim} \omega_{X/B,x}$$

to be the unique map sending this section to the differential $\omega_{(\ell(v))}$. Next we check that this isomorphism extends over the nodes. We treat only the case of a vertical node e , say passing from a vertex v_i of level i to a vertex v_j of level j with $j < i$; then the map near a horizontal node can be constructed just as in the smooth case. Write \mathfrak{f}' for the natural map $\widetilde{P} \rightarrow P$, so that

$$(5.22) \quad \beta(v_i) = \mathfrak{f}' \left(- \sum_{m=i}^{-1} a_m p_m \right) \quad \text{and} \quad \beta(v_j) = \mathfrak{f}' \left(- \sum_{m=j}^{-1} a_m p_m \right).$$

Setting $\mathfrak{f}_e := \mathfrak{f}'(g'(\delta_e))$, we have

$$(5.23) \quad \beta(v_i) = \beta(v_j) \cdot \mathfrak{f}_e^{\kappa_e} \in P.$$

We choose local coordinates u and v near the node, say v vanishes on the upper level component corresponding to v_i . The line bundle ω has a local generating section which is given after inverting u by $\frac{du}{u}$ and after inverting v by $-\frac{dv}{v}$. The line bundle $\mathcal{O}(-\beta)$ has (perhaps after adjusting the local coordinates u and v) a local generating section which is given after inverting u by $u^{\kappa_e} \beta(v_i)^{-1}$ and after inverting v by $\mathfrak{f}_e^{\kappa_e} v^{-\kappa_e} \beta(v_j)^{-1}$. As such, the bundle $\omega(-\beta)$ has a local generating section that is given after inverting u by $u^{\kappa_e} \beta(v_i)^{-1} \frac{du}{u}$ and after inverting v by $-\mathfrak{f}_e^{\kappa_e} v^{-\kappa_e} \beta(v_j)^{-1} \frac{dv}{v}$. The isomorphism (5.21) then corresponds to the section of $\omega(-\beta)$ that is given after inverting u by $\omega_{(i)} \beta(v_i)^{-1}$ and after inverting v by $\beta(v_j)^{-1} \omega_{(j)} = \beta(v_i)^{-1} \mathfrak{f}_e^{\kappa_e} \omega_{(j)}$. To conclude, we need to show that there exists an invertible function λ near the node such that after inverting u we have $\omega_{(i)} = \lambda u^{\kappa_e} \beta(v_i)^{-1} \frac{du}{u}$, and after inverting v we have $\mathfrak{f}_e^{\kappa_e} \omega_{(j)} = -\lambda \mathfrak{f}_e^{\kappa_e} v^{-\kappa_e} \beta(v_i)^{-1} \frac{dv}{v}$. But by condition (2) of Lemma 3.5, we know that there exists an invertible function λ such that

$$(5.24) \quad \omega_{(i)} = \lambda u^{\kappa_e} \frac{du}{u} \quad \text{and} \quad \omega_{(j)} = -\lambda v^{-\kappa_e} \frac{dv}{v},$$

and this λ clearly suffices.

Unravelling the constructions earlier in this section verifies that the constructed point of $\mathbf{Rub}_{\mathcal{L}_\mu}$ indeed maps to our starting point in $G\Xi \overline{\mathcal{M}}_{g,n}(\mu)$ under F .

To show that we have constructed an isomorphism of fibered categories, we must finally check that the composites

$$(5.25) \quad \mathbf{Rub}_{\mathcal{L}_\mu}(B) \longrightarrow G\Xi \overline{\mathcal{M}}_{g,n}(\mu)(B) \longrightarrow \mathbf{Rub}_{\mathcal{L}_\mu}(B)$$

and

$$(5.26) \quad G\Xi \overline{\mathcal{M}}_{g,n}(\mu)(B) \longrightarrow \mathbf{Rub}_{\mathcal{L}_\mu}(B) \longrightarrow G\Xi \overline{\mathcal{M}}_{g,n}(\mu)(B)$$

are isomorphic to the respective identities. This can be done by comparing the actions of the simple LLRT and the simple level rotation torus on the respective spaces; we omit the details. \square

5.8. The multi-scale space as a Zariski closure

Fix g, n , and define \mathcal{L}_μ on the universal curve over $\overline{\mathcal{M}}_{g,n}$ as before.

Definition 5.12. We define $\mathbf{Rub}_{\mathcal{L}_\mu}^{\text{trop}}$ to be the fibered category of $\mathbf{LogSch}_{\overline{\mathcal{M}}_{g,n}}$ whose objects are pairs $(X/B, \beta)$, where $X/B \in \overline{\mathcal{M}}_{g,n}$ and β is a PL function satisfying condition (1) of Lemma 2.1, and such that the line bundle $\mathcal{L}_\mu(-\beta)$ has multi-degree 0 on each geometric fiber.

This is a slight variant on $\mathbb{P}(\mathbf{Rub}_{\mathcal{L}_\mu})$. By ignoring the divisibility condition in Lemma 2.1, we are effectively taking the coarse moduli space relative to $\overline{\mathcal{M}}_{g,n}$, and we only require that $\mathcal{L}_\mu(-\beta)$ has multi-degree 0, rather than requiring it to be trivial. Since we in particular do not record the data of an isomorphism, we are effectively also taking a \mathbb{C}^* -quotient.

The map $\mathbf{Rub}_{\mathcal{L}_\mu}^{\text{trop}} \rightarrow \overline{\mathcal{M}}_{g,n}$ is birational and representable, but not in general proper. Using stability conditions as in [HMP⁺25], we can construct a compactification

$$\mathbf{Rub}_{\mathcal{L}_\mu}^{\text{trop}} \longrightarrow \mathbb{P}(\mathbf{Rub}_{\mathcal{L}_\mu}^\theta) \longrightarrow \overline{\mathcal{M}}_{g,n},$$

where $\mathbb{P}(\mathbf{Rub}_{\mathcal{L}_\mu}^\theta) \rightarrow \overline{\mathcal{M}}_{g,n}$ is proper, birational, and representable, and $\mathbf{Rub}_{\mathcal{L}_\mu}^{\text{trop}} \rightarrow \mathbb{P}(\mathbf{Rub}_{\mathcal{L}_\mu}^\theta)$ is an open immersion; but we do not pursue this here as it would require substantial additional notation.

Let $\mathbb{P}(\mathcal{MS}^0) \subseteq \mathcal{M}_{g,n}$ be the locus of smooth curves over which \mathcal{L}_μ admits a non-zero global section; this can be seen as the interior of the locus of (projectivized, generalized) multi-scale differentials.

Theorem 5.13. *The Zariski closure of \mathcal{MS}^0 in $\mathbf{Rub}_{\mathcal{L}_\mu}^{\text{trop}}$ (or, equivalently, in $\mathbb{P}(\mathbf{Rub}_{\mathcal{L}_\mu}^\theta)$) is equal to $\mathbb{P}(\mathcal{MS}_\mu)$, the projectivized space of (non-generalized) multi-scale differentials.*

Proof. There is a natural closed immersion $\mathbb{P}(\mathbf{Rub}_{\mathcal{L}_\mu}^{\text{coarse}}) \rightarrow \mathbf{Rub}_{\mathcal{L}_\mu}^{\text{trop}}$, and the main component of the space $\mathbb{P}(\mathbf{Rub}_{\mathcal{L}_\mu}^{\text{coarse}})$ is $\mathbb{P}(\mathcal{MS}_\mu)$. \square

One can obtain the stacky version $\Xi \overline{\mathcal{M}}_{g,n}(\mu)$ (of which \mathcal{MS}_μ is the relative coarse moduli space) in a similar fashion, replacing $\mathbb{P}(\mathbf{Rub}_{\mathcal{L}_\mu}^\theta)$ with a stacky modification; we leave the details to the interested reader.

6. The Hodge DR conjecture

In this section we present several equivalent constructions of the universal line bundle introduced in Section 1.2.2, discuss its various properties, and prove Lemma 1.5.

As explained in Section 1.2.2, the projectivized space of (generalized) multi-scale differentials comes with a map to the projectivized Hodge bundle, obtained by taking the differential at the top level, and allowing it to vanish at all lower levels. Pulling back $\mathcal{O}(1)$ from the Hodge bundle gives a line bundle on the generalized multi-scale space. We begin by giving several equivalent versions of this construction.

First we write out explicitly the objects of the fibered category $\mathbb{P}(\mathbf{Rub})$:

$$\mathbb{P}(\mathbf{Rub}) = \{(\pi: X \longrightarrow B, \beta, \mathcal{F})\},$$

where $(X/B, \beta)$ is a point of \mathbf{Rub} as in Lemma 2.1, and \mathcal{F} is a line bundle on B . The Abel-Jacobi map sends such an object to $\pi^* \mathcal{F}(\beta)$, giving a proper, see [MW20, Theorem 4.3.2], map $\mathbb{P}(\mathbf{Rub}) \rightarrow \mathbb{P}\text{ic}$.

Now fix a line bundle \mathcal{L} on $X_{g,n}/\overline{\mathcal{M}}_{g,n}$, which is of total degree 0 on each fiber. Then we can write explicitly the fibered category of $\mathbb{P}(\mathbf{Rub}_{\mathcal{L}})$ as

$$\mathbb{P}(\mathbf{Rub}_{\mathcal{L}}) = \{(X/B, \beta, \mathcal{F}, \varphi)\},$$

where $(X/B, \beta, \mathcal{F})$ is an object of $\mathbb{P}(\mathbf{Rub})$ with X/B stable of genus g , and $\varphi: \pi^* \mathcal{F}(\beta) \rightarrow \mathcal{L}$ is an isomorphism.

Construction 1: Tautological bundle

This is just the bundle \mathcal{F} on $\mathbb{P}(\mathbf{Rub})$, or its pullback to \mathcal{F} on $\mathbb{P}(\mathbf{Rub}_{\mathcal{L}})$ along the tautological map. We denote the *dual* of this line bundle by η .

Construction 2: Projective embedding

Let D be an effective divisor on $X_{g,n}$ such that $R^1\pi_*\mathcal{L}(D)$ vanishes, so in particular $\pi_*\mathcal{L}(D)$ is a vector bundle on $\overline{\mathcal{M}}_{g,n}$. Such a D can always be found as an element of the linear system of a sufficiently relatively ample sheaf on $X_{g,n}$ over $\overline{\mathcal{M}}_{g,n}$. Then over $\mathbb{P}(\mathbf{Rub}_{\mathcal{L}})$ we have natural maps

$$(6.1) \quad \pi^*\mathcal{F} \xrightarrow{\sim} \mathcal{L}(-\beta) \longrightarrow \mathcal{L} \longrightarrow \mathcal{L}(D),$$

where the first map is induced by φ , the second is induced by the natural map $\mathcal{O}(-\beta) \rightarrow \mathcal{O}$, and the third by the natural map $\mathcal{O} \rightarrow \mathcal{O}(D)$. Adjunction yields a map

$$(6.2) \quad \mathcal{F} = \pi_*\pi^*\mathcal{F} \longrightarrow \pi_*\mathcal{L}(D),$$

which is by definition a map

$$(6.3) \quad F: \mathbb{P}(\mathbf{Rub}_{\mathcal{L}}) \longrightarrow \mathbb{P}_{\mathbb{P}(\mathbf{Rub}_{\mathcal{L}})}(\pi_*\mathcal{L}(D)).$$

Indeed, note that our projectivizations are moduli of subbundles, not quotient bundles, so it is enough to check that this map is universally injective. But the formation of both sides commutes with arbitrary base-change, so it is enough to check this over a point, where it is clear.

Lemma 6.1. $F^*\mathcal{O}(1) = \eta$.

Proof. The equality $F^*\mathcal{O}(1) = \mathcal{F}^\vee$ is immediate from [Sta25, Example 0FCY]; we obtain \mathcal{F}^\vee instead of \mathcal{F} because we define the projectivization to be the moduli of rank 1 subbundles, not rank 1 quotient bundles. \square

In particular, we observe that the line bundle $F^*\mathcal{O}(1)$ turns out to be independent of the choice of the sufficiently relatively ample divisor D . In the case considered in the introduction, we take

$$(6.4) \quad \mathcal{L} = \omega_{X_{g,n}/\overline{\mathcal{M}}_{g,n}}^{\otimes k} \left(- \sum_{i=1}^n (a_i - k)z_i \right)$$

and $D = \sum_{i:a_i>0} a_i z_i$.

Construction 3: Pullback from rubber target

For this construction we restrict to the case where $\mathcal{L} = \mathcal{O}_X(\sum_i a_i z_i)$ for $k = 0$; put another way, we choose a rational section of \mathcal{L} whose locus of zeros and poles is contained in a union of disjoint sections of $X \rightarrow B$.

We write

$$D = \sum_{i:a_i>0} a_i z_i \quad \text{and} \quad E = - \sum_{i:a_i<0} a_i z_i.$$

Since these are effective divisors, we have natural maps

$$\mathcal{O}_X \longrightarrow \mathcal{O}_X(D) \quad \text{and} \quad \mathcal{O}_X \longrightarrow \mathcal{O}_X(E),$$

and combining with the natural map $\mathcal{O}_X \rightarrow \mathcal{O}_X(\beta)$ and the isomorphism $\varphi: \pi^*\mathcal{F}(\beta) \xrightarrow{\sim} \mathcal{O}_X(D - E)$ yields maps

$$\mathcal{O}_X(-E)(-\beta) \longrightarrow \mathcal{O}_X \quad \text{and} \quad \mathcal{O}_X(-E)(-\beta) \longrightarrow \mathcal{O}_X(D - E)(-\beta) \xrightarrow{\sim} \pi^*\mathcal{F}.$$

The induced map

$$\mathcal{O}_X(-E)(-\beta) \longrightarrow \mathcal{O}_X \oplus \pi^*\mathcal{F}$$

is universally injective since the first map is injective around the support of D and the second is injective away from the support of D . This induces a map

$$X \longrightarrow \mathbb{P}(\mathcal{O}_B \oplus \mathcal{F}).$$

The cotangent line at ∞ to this rubber target is then given by

$$(6.5) \quad \Psi_\infty = \mathcal{F}^\vee.$$

We have deduced the following result.

Lemma 6.2. $\Psi_\infty = \eta$.

Remark 6.3. Above we have constructed a rubber target of length 1 (*i.e.* with no expansions). This is because we are only interested in what happens near the infinity section, so we do not need to construct the whole expanded chain. The reader who is more comfortable with expansions may verify that the length 1 target we construct here is exactly what is obtained by following through the proof of the expanded target in [BHP⁺23, Proposition 50], and then contracting all except the top component.

6.1. Computation of η for $k = 0$

Here we prove Lemma 1.5, which we restate for the convenience of the reader.

Theorem 6.4. *Lemma 1.4 is true for $k = 0$: for any $g, u \geq 0$ and any vector $A \in \mathbb{Z}^n$ with sum $|A| = 0$, we have*

$$p_* \left(\left[\mathbb{P}(\mathbf{Rub}_{\mathcal{L}_A}) \right]^{\text{vir}} \cdot \eta^u \right) = p_* \left(\left[\overline{\mathcal{M}}_{g,A}(\mathbb{P}^1, 0, \infty)^\sim \right]^{\text{vir}} \cdot \Psi_\infty^u \right) = [r^u] \text{Ch}_{g,A}^{0,r,g+u}.$$

Proof. The first equality follows from Lemmas 6.1 and 6.2. For the second equality, we note that the term on the left has been computed in [FWY21, Corollary 4.3] in terms of a slightly modified Chiodo class. Indeed, we define an r -shifted version $A(r)$ of A by

$$A(r)_i = \begin{cases} a_i & \text{for } a_i \geq 0, \\ r + a_i & \text{for } a_i < 0. \end{cases}$$

In other words, for all indices i with $a_i < 0$ (which form a subset $I_\infty \subseteq \{1, \dots, n\}$), we shift the vector A by r in the i^{th} entry. Then the Chiodo class $\text{Ch}_{g,A(r)}^{0,r,d}$ is a polynomial in r , for r sufficiently large. Denote by

$$\text{Ch}_{g,A(r)}^{0,r,\bullet} = \sum_{d \geq 0} \text{Ch}_{g,A(r)}^{0,r,d}$$

the associated mixed-degree class. Then in this notation, the formula from [FWY21, Corollary 4.4] reads as follows:

$$\begin{aligned} p_* \left(\left[\overline{\mathcal{M}}_{g,A}(\mathbb{P}^1, 0, \infty)^\sim \right]^{\text{vir}} \cdot \Psi_\infty^u \right) &= \sum_{\vec{e} \in \mathbb{Z}_{\geq 0}^{I_\infty}} \prod_{i \in I_\infty} (-a_i \psi_i)^{e_i} \cdot [r^{u-|\vec{e}|}] \text{Ch}_{g,A(r)}^{0,r,u+g-|\vec{e}|} \\ &= [r^u] \left[\sum_{\vec{e} \in \mathbb{Z}_{\geq 0}^{I_\infty}} \prod_{i \in I_\infty} (-a_i r \psi_i)^{e_i} \cdot \text{Ch}_{g,A(r)}^{0,r,\bullet} \right]_{\text{codim } g+u} \\ &= [r^u] \left[\prod_{i \in I_\infty} \frac{1}{1 + a_i r \psi_i} \cdot \text{Ch}_{g,A(r)}^{0,r,\bullet} \right]_{\text{codim } g+u} \\ &= [r^u] \left[\text{Ch}_{g,A}^{0,r,\bullet} \right]_{\text{codim } g+u}. \end{aligned}$$

Here the last step uses [GLN23, Theorem 4.1(ii)]. □

6.2. (A)symmetry

Above we gave three constructions of the line bundle $\eta = \eta(\mathcal{L})$ on $\mathbb{P}(\mathbf{Rub}_{\mathcal{L}})$. We know that the push-forwards to $\overline{\mathcal{M}}_{g,n}$ of $[\mathbb{P}(\mathbf{Rub}_{\mathcal{L}})]^{\text{vir}}$ and $[\mathbb{P}(\mathbf{Rub}_{\mathcal{L}^\vee})]^{\text{vir}}$ agree. However, once we intersect with the class η , things are a little more subtle. The universal curve over $\mathbb{P}(\mathbf{Rub})$ carries a PL function β , totally ordered and with maximum value 0. The *minimum* value of β we denote by β^{\min} ; this is a PL function on $\mathbb{P}(\mathbf{Rub})$. We set $\eta(\beta^{\min}) = \eta \otimes_{\mathcal{O}} \mathcal{O}(\beta^{\min})$.

Lemma 6.5. *We have*

$$(6.6) \quad p_*([\mathbb{P}(\mathbf{Rub}_{\mathcal{L}^\vee})]^{\text{vir}} \cdot c_1(\eta)^u) = p_*([\mathbb{P}(\mathbf{Rub}_{\mathcal{L}})]^{\text{vir}} \cdot (-c_1(\eta(\beta^{\min})))^u).$$

Proof. There is a natural isomorphism (compatible with the virtual fundamental classes) over $\overline{\mathcal{M}}_{g,n}$ from $\mathbb{P}(\mathbf{Rub}_{\mathcal{L}})$ to $\mathbb{P}(\mathbf{Rub}_{\mathcal{L}^\vee})$, given by

$$(6.7) \quad (X/B, \beta, \mathcal{F}, \varphi) \mapsto (X/B, \beta^{\min} - \beta, (\mathcal{F}(\beta^{\min}))^\vee, \varphi'),$$

where φ' is the composite

$$(6.8) \quad \pi^*(\mathcal{F}(\beta^{\min}))^\vee (\beta^{\min} - \beta) = \pi^* \mathcal{F}^\vee(-\beta) \xrightarrow{(\varphi^\vee)^{-1}} \mathcal{L}^\vee. \quad \square$$

7. Blowup descriptions

In this section we give a description of $\mathbb{P}(\mathbf{Rub}_{\mathcal{L}}^{\text{coarse}})$ as a global blowup.

First, in genus zero, we construct an explicit sheaf of ideals on $\overline{\mathcal{M}}_{0,n}$, such that blowing up $\overline{\mathcal{M}}_{0,n}$ along this sheaf gives $\mathbb{P}(\mathbf{Rub}_{\mathcal{L}}^{\text{coarse}})$. In [Ngu24] Nguyen described the incidence variety compactification (IVC) in the case of genus zero as an explicit blowup of $\overline{\mathcal{M}}_{0,n}$. Note that in genus zero there are no global residue conditions (because any top-level vertex must have a marked pole), and hence in genus zero the rubber space and the space of generalized multi-scale differentials coincide with the space of multi-scale differentials. Our blowup description can thus recover Nguyen's result about the IVC of the strata of meromorphic 1-forms in genus zero as a blowup of $\overline{\mathcal{M}}_{0,n}$. We also provide an example demonstrating the difference between the rubber space and the IVC in genus zero.

Next, for arbitrary genus, we construct a globally defined sheaf of ideals on the normalization of the incidence variety compactification (NIVC) whose blowup gives the (projectivized) multi-scale moduli space (*i.e.* the main component of $\mathbf{Rub}_{\mathcal{L}}^{\text{coarse}}$). Consequently, it follows that the (coarse) space of projectivized (non-generalized) multi-scale differentials is a projective variety for all g . Recall that in [BCG⁺19, Section 14.1] the moduli space of multi-scale differentials was described as a local blowup, where the ideals locally defining the center of the blowup can differ by principal ideals on the overlaps of local charts. In particular, the description of [BCG⁺19] did not yield projectivity of the space of multi-scale differentials. By constructing an explicit ample divisor class, the projectivity of the moduli space of multi-scale differentials was later established in [CCM24, Section 3]. Our global blowup description thus provides a direct conceptual understanding of this projectivity result.

Besides projectivity, knowing a blowup description of compactified strata of differentials can be helpful for obtaining geometric invariants, such as volumes of the strata, by using intersection theory; see [Ngu24]. We also provide a tropical interpretation of our blowup, which sheds further light on the geometry of the construction.

7.1. The sheaf of ideals in genus zero

Let Γ be the dual graph of a boundary stratum $D_\Gamma \subset \overline{\mathcal{M}}_{0,n}$. For each vertex $v \in V(\Gamma)$, let $d(v)$ be the degree of \mathcal{L}_μ restricted to v (so $\sum_{v \in V(\Gamma)} d(v) = 0$ by definition). Since Γ is a tree, there exists a unique ‘slope’⁽²²⁾ function $\kappa: H \rightarrow \mathbb{Z}$ from the set $H = H(\Gamma)$ of half-edges of Γ such that

- (1) κ agrees with m_i at the leg corresponding to a marked point z_i ;
- (2) $\kappa(h) + \kappa(h') = 0$ for any h and h' that are opposite halves of an edge;
- (3) for all vertices v , we have $\sum_{h \in H(v)} \kappa(h) = d(v)$, where we sum over all half-edges attached to v .

For every pair of vertices v and v' , let γ be the unique path from v to v' in the tree Γ . We view this (directed) path as a sequence of half-edges, where if an edge $e = (h, h') \in E(\Gamma)$ appears in γ in the direction going *from* h *to* h' , meaning that along the path γ in the direction from v to v' the half-edge h appears first, followed by h' , then we put (only) h in our sequence of half-edges. We define an ideal locally around the boundary stratum $D_\Gamma \subseteq \overline{\mathcal{M}}_{0,n}$ by

$$I(v, v') := \prod_{h \in \gamma} \delta(h)^{\max(\kappa(h), 0)},$$

where we write $\delta(h)$ for the ideal associated to the edge containing h (that is, for the defining equation of the boundary divisor of $\overline{\mathcal{M}}_{0,n}$ where the corresponding node exists). Define

$$J(v, v') := I(v, v') + I(v', v);$$

this evidently satisfies $J(v, v') = J(v', v)$ and $J(v, v) = (1)$. Finally, we set

$$w(v) := \text{valence}(v) - 2,$$

which is a positive integer by the stability of the curve, and define

$$J(\Gamma) := \prod_{(v, v') \in V \times V} J(v, v')^{w(v)w(v')}.$$

A concrete example of this ideal is given in Lemma 7.6 below.⁽²³⁾

Remark 7.1. Note that it is possible to define *locally* an ideal that is simpler than J and whose blowup produces the same space. Indeed, taking the product of J and any principal ideal works. Nevertheless, such local ideals may not always glue to form a *global* sheaf of ideals. Blowing up a globally defined sheaf of ideals (from a projective variety) can directly imply the projectivity of the resulting space, while gluing local blowups together does not. Therefore, the ideal J in the above was designed with some delicate exponents to make it a globally defined sheaf of ideals, as we will check in the next section. This idea will be used in Section 7.6 to define a global ideal sheaf on NIVC to conclude the projectivity of the multi-scale space for arbitrary genus.

7.2. Compatibility under degeneration in genus zero

To show that the ideals $J(\Gamma)$ defined in the neighbourhood of each stratum $D_\Gamma \subseteq \overline{\mathcal{M}}_{0,n}$ glue to a global ideal sheaf over $\overline{\mathcal{M}}_{0,n}$, we need to show that they behave well under degeneration. As any dual graph Γ can be obtained from any other Γ' by a series of operations of inserting and contracting edges, it is enough to check that the ideals glue under contracting a single edge of the graph.

Lemma 7.2. *Let e be an edge of Γ , and let Γ' be the graph obtained from Γ by contracting e . Then $J(\Gamma') = J(\Gamma)$, after inverting the ideal $\delta(e)$.*

⁽²²⁾The justification for this terminology is given by (5.1), which shows that the slopes of points $\mathbf{Rub}_{\mathcal{L}_\mu}$ satisfy the same conditions.

⁽²³⁾If the reader prefers not splitting up the half-edges into ones with increasing and decreasing slopes, alternatively we can define the *fractional* ideal $J'(v, v') = \prod_{h \in \gamma} (\delta(h)^{\kappa(h)}, 1)$, and define $J'(\Gamma) := \prod_{(v, v') \in V \times V} J'(v, v')^{w(v)w(v')}$. Then J' induces a globally defined fractional ideal whose blowup is the same as that of J (as we will justify for J in the following subsections.)

Note that inverting $\delta(e)$ geometrically corresponds to restricting to the locus where the edge e is contracted, *i.e.* where the corresponding node of the curve is smoothed out.

Proof. We denote by $c: \Gamma \rightarrow \Gamma'$ the contraction map, let v_1 and v_2 be the endpoints of e , and let v' be the vertex of Γ' to which e is contracted, so that $d(v') = d(v_1) + d(v_2)$.

If v is any vertex of Γ different from v_1 and v_2 , then clearly $w(v) = w(c(v))$. Furthermore, the slope function on Γ clearly restricts to the slope function on Γ' . Thus for any two vertices u_1 and u_2 of Γ distinct from v_1 and v_2 , we have

$$J_\Gamma(u_1, u_2) \sim J_{\Gamma'}(u_1, u_2),$$

where to simplify notation we write $I \sim J$ if the ideal sheaves I and J become equal after inverting $\delta(e)$. Similarly, $J_\Gamma(v_1, v_2) \sim (1)$.

Based on the above analysis, we only need to consider the pairs of vertices in Γ' and in Γ that involve v' and v_1 or v_2 , respectively. It therefore suffices to show that

$$(7.1) \quad \prod_{v \in V(\Gamma')} J(v', v)^{2w(v')w(v)} \sim \prod_{v \in V(\Gamma)} J(v_1, v)^{2w(v)w(v_1)} J(v_2, v)^{2w(v)w(v_2)}.$$

Let $V^\circ := V(\Gamma) \setminus \{v_1, v_2\} = V(\Gamma') \setminus \{v'\}$. Then (7.1) reduces to showing

$$\prod_{v \in V^\circ} J(v', v)^{w(v')w(v)} \sim \prod_{v \in V^\circ} J(v_1, v)^{w(v_1)w(v)} J(v_2, v)^{w(v_2)w(v)}.$$

This follows from $w(v') = w(v_1) + w(v_2)$ and the relations

$$J(v', v) \sim J(v_1, v) \sim J(v_2, v)$$

for all $v \in V'$. □

Definition 7.3. Define $J(\mathcal{L}_\mu)$ to be the (global) ideal sheaf on $\overline{\mathcal{M}}_{0,n}$ that for any boundary stratum D_Γ restricts to the ideal $J(\Gamma)$ on a neighbourhood of D_Γ .

The existence of $J(\mathcal{L})$ follows from Lemma 7.2.

7.3. A tropical picture in genus zero

The normalized blowup in the ideal $J(\Gamma)$ corresponds tropically to a subdivision of the positive orthant in the vector space $\mathbb{Q}\langle E \rangle$, where $E = E(\Gamma)$ is the edge set. This subdivision is built by taking a hyperplane (or sometimes the whole space) for every pair of vertices in Γ : if γ is the path from v to v' as above, then the corresponding hyperplane $L(v, v')$ is cut out by the equation

$$\sum_{h \in \gamma} \kappa(h) e(h) = 0,$$

where $e(h)$ is the edge containing the half-edge h , viewed as an element of the group $\mathbb{N}\langle E \rangle$ (and we recall that a half-edge h is said to be contained in a directed path γ if γ goes via h before going through the complementary half-edge of the same edge).

These local subdivisions glue to a global subdivision of the tropicalization of $\overline{\mathcal{M}}_{0,n}$, inducing a proper birational map $\widetilde{\mathcal{M}}_{0,n} \rightarrow \overline{\mathcal{M}}_{0,n}$.

Lemma 7.4. *The normalization of the blowup of $\overline{\mathcal{M}}_{0,n}$ in the ideal $J(\mathcal{L}_\mu)$ is equal to the proper birational map $\widetilde{\mathcal{M}}_{0,n} \rightarrow \overline{\mathcal{M}}_{0,n}$ induced by the subdivision above.*

Proof. The standard dictionary (see [Kat89a, Section 1.3.3, p. 14]) between toric blowups and subdivisions implies that the normalized blowup in $J(v, v')$ is equal to that induced by the subdivision in $L(v, v')$. Since $w(v) \geq 1$ (by stability), blowing up in $J(v, v')$ is the same as blowing up in $J(v, v')^{w(v)w(v')}$. Normalized blowup in a product of ideals corresponds to superimposing their subdivisions. □

7.4. Comparing blowups and rubber maps in genus zero

We are ready to prove our main statement in genus zero.

Theorem 7.5. *The normalization of the blowup $\widetilde{\mathcal{M}}_{0,n}$ of $\overline{\mathcal{M}}_{0,n}$ along the ideal sheaf $J(\mathcal{L}_\mu)$ is the projectivized coarse moduli space of rubber differentials $\mathbb{P}(\mathbf{Rub}_\mathcal{L}^{\text{coarse}})$.*

Proof. Let X/B be a nuclear log curve of genus zero.

Claim: There exists a PL function β on X such that $\mathcal{L}_\mu \cong \mathcal{O}(\beta)$, and moreover such a β is unique up to translation by an element of $\overline{M}_B(B)^{\text{gp}}$.

To prove the claim, we use the fact that the graph is a tree to deduce that there is a unique collection of admissible slopes κ_e . We pick a vertex v_0 and let β be the unique PL function vanishing on v_0 and with slopes given by the κ_e . The line bundle $\mathcal{L}_\mu(-\beta)$ has multi-degree zero, and is hence trivial since X has genus zero. This proves the claim.

Now recall that $\mathbf{Rub}_\mathcal{L}^{\text{coarse}}$ can be obtained by omitting the divisibility condition (2) from Lemma 2.1. In other words, the point X/B lies in $\mathbf{Rub}_\mathcal{L}^{\text{coarse}}$ if and only if the values of β on the vertices of Γ form a totally ordered set. It therefore remains to check that this is equivalent to the map $B \rightarrow \overline{\mathcal{M}}_{0,n}$ factoring via the subdivision described in Section 7.3.

If γ is a directed path in Γ , we define

$$\varphi(\gamma) := \sum_{h \in \gamma} \kappa_h \delta_h.$$

Since the difference of values of β at the two ends of an edge is the slope κ_e of that edge (with the appropriate sign), the values of β at the two ends of a path γ differ by $\varphi(\gamma)$.

Fix a vertex v_0 , and write γ_v for the unique path from v_0 to v . Then the set $\{\beta(v) : v \in V(\Gamma)\}$ is totally ordered if and only if the set

$$\{\varphi(\gamma_v) : v \in V(\Gamma)\}$$

is totally ordered. This is in turn equivalent to requiring that for every path $\gamma \subset \Gamma$ (not necessarily a path from v_0), the element $\varphi(\gamma)$ is comparable to 0, *i.e.* either $\varphi(\gamma) \in \overline{M}_S$ or $-\varphi(\gamma) \in \overline{M}_S$. Imposing this condition is equivalent to subdividing $\mathbb{N}\langle E \rangle$ in the hyperplane $L(v, v')$ of Section 7.3, where v and v' are the endpoints of γ . \square

7.5. Comparison to Nguyen's blowup in genus zero

As mentioned, in genus zero Nguyen [Ngu24] described the IVC as an explicit blowup of $\overline{\mathcal{M}}_{0,n}$ (also for the more general case of k -differentials in genus zero). Since the rubber/multi-scale space is the normalization of a blowup of the normalization of the IVC, our blowup described in Lemma 7.5 must dominate the blowup defined by Nguyen. In this subsection we recall Nguyen's construction, provide a viewpoint of his blowup from our setup, and give an alternative proof for Nguyen's result that blowing up $\overline{\mathcal{M}}_{0,n}$ in his ideal gives the IVC.

We begin by recalling Nguyen's construction of a sheaf of ideals on $\overline{\mathcal{M}}_{0,n}$. Let X/B be a nuclear log curve of genus zero with graph Γ , and let κ be the slope function on the edges of Γ , *i.e.* the PL function constructed in the proof of Lemma 7.5. For a given vertex $v \in V(\Gamma)$ and an edge $e \in E(\Gamma)$, let $h_v(e)$ be the half-edge of e such that the path from the end of $h_v(e)$ to v passes through e . For a vertex $v \in V(\Gamma)$, we define

$$(7.2) \quad \delta_v := \prod_{e \in E(\Gamma)} \delta_e^{\kappa_{v,e}},$$

where $\kappa_{v,e} := \max(\kappa(h_v(e)), 0)$. Let $N(\Gamma)$ be the (local) ideal (in the variables δ_e , as in our setup) generated by the set of elements δ_v for all vertices $v \in V(\Gamma)$. It was shown in [Ngu24] that these $N(\Gamma)$ can be patched




Figure 4. The graph Γ of a stratum in $\overline{\mathcal{M}}_{0,7}$; the desired slopes κ can be obtained, *e.g.*, by using the signature $\mu = (-1^6, 4)$ with the six markings associated to simple poles attached to the vertices v_1 , v_2 , and v_3 .

together to a sheaf of ideals N globally defined on $\overline{\mathcal{M}}_{0,n}$. This can be seen the same way as Lemma 7.2, and we will discuss this in more generality in Lemma 7.9 for arbitrary genera.

Before proceeding, we illustrate Nguyen’s ideal and our ideal in the following example.

Example 7.6. Consider a (partially ordered) dual graph Γ as illustrated in Figure 4, with all slopes $\kappa = 1$. Recalling the definition of δ_v in (7.2) for a vertex $v \in V(\Gamma)$ and writing $\delta_i := \delta_{e_i}$ to lighten notation, we obtain $\delta_{v_0} = \delta_1 \delta_2 \delta_3$, $\delta_{v_1} = \delta_2 \delta_3$, $\delta_{v_2} = \delta_1 \delta_3$, and $\delta_{v_3} = \delta_1 \delta_2$. In this case Nguyen’s ideal $N(\Gamma)$ is given by

$$N(\Gamma) = (\delta_1 \delta_2, \delta_1 \delta_3, \delta_2 \delta_3, \delta_1 \delta_2 \delta_3) = (\delta_1 \delta_2, \delta_1 \delta_3, \delta_2 \delta_3).$$

In contrast, our ideal $J(\Gamma)$ is given by

$$J(\Gamma) = (\delta_1, \delta_2)^2 (\delta_1, \delta_3)^2 (\delta_2, \delta_3)^2 (\delta_1)^4 (\delta_2)^4 (\delta_3)^4.$$

When we blow up $J(\Gamma)$, each ideal generated by a pair (δ_i, δ_j) for $1 \leq i < j \leq 3$ becomes principal, and so does the ideal $N(\Gamma)$. Therefore, the blowup in $J(\Gamma)$ dominates the blowup in $N(\Gamma)$.

Nguyen [Ngu24] proved that blowing up $\overline{\mathcal{M}}_{0,n}$ along the globally defined sheaf of ideals N gives the IVC. Indeed, in the example above we see explicitly that locally around the boundary stratum with the dual graph Γ , the rubber/multi-scale space obtained by blowing up along J is a further blowup of the IVC.

The situation of this example can also be understood in general, from our viewpoint, which gives an alternative proof of the result of Nguyen.

Proposition 7.7. *The local blowup of $\overline{\mathcal{M}}_{0,n}$ near D_Γ along the ideal $J(\Gamma)$ makes the ideal $N(\Gamma)$ become principal. Moreover, in genus zero the blowup of $\overline{\mathcal{M}}_{0,n}$ along the ideal sheaf N is the IVC.*

Before giving the proof, we first reinterpret $N(\Gamma)$ geometrically as follows. If two vertices v and v' are joined by an edge e , and if $\ell(v) > \ell(v')$, then δ_v divides $\delta_{v'}$. Therefore, the ideal $N(\Gamma)$ is the same as the ideal generated only by the elements δ_v where v ranges over all vertices that are local maxima of Γ (in the sense that all edges from v go down – recall that this is a partial order on the graph, and the datum of a multi-scale differential upgrades this to a full order). A vertex v that is a local maximum of Γ , such that the corresponding δ_v generates the ideal $N(\Gamma)$ after the blowup, becomes a global top-level vertex. On the other hand, those local maxima v whose δ_v terms do *not* generate the principal ideal after blowing up $N(\Gamma)$ may not divide each other, and thus remain unordered. This corresponds to the fact that a point in the IVC records actual differentials merely on top-level vertices where the stable differential is not identically zero, while on any lower vertex the stable differential is identically zero (though the underlying marked zeros and poles of the twisted differential are still remembered).

Proof. For the first claim, note that the edge parameter δ_e appears with the same exponent in the expressions of δ_v and $\delta_{v'}$ unless e lies in the unique path from v to v' , in which case the exponents of δ_e in δ_v and $\delta_{v'}$ are the same as those in $I(v, v')$ and $I(v', v)$, respectively. Since blowing up along $J(\Gamma)$ makes the ideal

$(I(v, v'), I(v', v))$ principal, it follows that each ideal $(\delta_v, \delta_{v'})$ becomes principal under that blowup. This is to say that after blowing up in $J(\Gamma)$, one of δ_v and $\delta_{v'}$ must divide the other. Doing this for all v and v' shows that after the blowup along $J(\Gamma)$, there is a set of elements $\{\delta_{v_1}, \dots, \delta_{v_k}\}$ such that for every i and each $v \in V(\Gamma)$, δ_{v_i} divides δ_v . In particular, such δ_{v_i} and δ_{v_j} divide each other and thus differ by multiplication by a unit, and the ideal $N(\Gamma)$, after the blowup along $J(\Gamma)$, is generated by any one of these δ_{v_i} , and hence it becomes principal.

For the second claim, we will construct the desired morphisms between the blowup and the IVC in both directions that are inverses of each other. Since the blowup and the IVC both admit natural maps onto $\overline{\mathcal{M}}_{0,n}$, these morphisms will be constructed locally over each boundary stratum D_Γ of $\overline{\mathcal{M}}_{0,n}$.

The upshot underneath the constructions is that δ_v for $v \in V(\Gamma)$ is an *adjusting parameter* in the sense of [BCG⁺19, Proposition 11.13], which means that multiplying by δ_v^{-1} makes the limiting differential become not identically zero on the component corresponding to v . To see this, let D_{e_i} be the boundary divisor of $\overline{\mathcal{M}}_{0,n}$ corresponding to a given edge e_i of Γ . Contracting all edges of Γ except e_i produces a graph with two vertices connected by the edge e_i , and the family of stable differentials over it vanishes on the irreducible component corresponding to the lower-level vertex, with generic vanishing order $|\kappa_{e_i}|$. Given a vertex v of Γ , if the image of v under this contraction is the lower of the two resulting vertices, then over D_Γ the family of stable differentials vanishes identically on the irreducible component corresponding to v . Applying this observation to all edges e_i in Γ where v becomes lower after the above edge contractions, it follows that the expression of δ_v records exactly the total vanishing order over D_Γ of the stable differentials on the irreducible component corresponding to v . Therefore, multiplying by δ_v^{-1} makes the limiting differential become not identically zero on v . By definition, this implies that δ_v is an adjusting parameter for v .

Now we construct a morphism from the IVC to the blowup of $\overline{\mathcal{M}}_{0,n}$ along N by using the universal property of the blowup. More precisely, as we blow up (in a neighbourhood of D_Γ) the ideal generated by all δ_v , it suffices to check that this ideal becomes principal on the IVC. Recall that the IVC parameterizes pointed stable differentials (of prescribed type) that are not identically zero, where a stable differential is a section of the dualizing sheaf over the stable curve, considered up to an overall scaling by a non-zero constant factor. If a vertex v is not a local maximum of Γ , *i.e.* if there exists an edge e going up from v , then the (stable) differential on the irreducible component corresponding to v is identically zero. Thus given a (not identically zero) stable differential, we can declare a local maximum vertex v of Γ to be a global maximum if and only if the stable differential on the corresponding irreducible component of the curve is not identically zero. By the preceding discussion, this is precisely to say that all adjusting parameters δ_v for the global maxima vertices v differ by units, and divide all the other δ_v . Hence the ideal $N(\Gamma)$ pulls back a principal ideal on the IVC, which induces the map (locally) from the IVC to the blowup of $\overline{\mathcal{M}}_{0,n}$ along $N(\Gamma)$.

Next we construct a morphism in the opposite direction, locally near D_Γ from the blowup of $\overline{\mathcal{M}}_{0,n}$ along $N(\Gamma)$ to the IVC, by using the universal property of the Hodge bundle over $\overline{\mathcal{M}}_{0,n}$ (twisted by the polar part of the differentials, and projectivized as always).

Consider the universal family of differentials with prescribed zeros and poles over a punctured neighbourhood of D_Γ in $\mathcal{M}_{0,n}$. We claim that this family of differentials extends to a family of stable differentials over the local blowup of $\overline{\mathcal{M}}_{0,n}$ along $N(\Gamma)$. Indeed, for each point in the preimage of D_Γ in the blowup, we know the set of global maxima v_1, \dots, v_k of the graph (with $k \geq 1$), such that any other δ_v is divisible by one of the adjusting parameters $\delta_{v_1}, \dots, \delta_{v_k}$. It follows that the limiting stable differential will be not identically zero precisely on the irreducible components corresponding to v_1, \dots, v_k , and thus in particular not identically zero on the stable curve as a whole. By the universal property of the projectivized Hodge bundle, the blowup along $N(\Gamma)$ carrying a family of (not identically zero) stable differentials admits locally a morphism to this bundle. Moreover, since, over the locus of smooth curves, this family of differentials coincides with the family of differentials in a given stratum, it implies that the image of the morphism from the blowup to the Hodge bundle is the closure of the stratum, *i.e.* the IVC. By construction, it is clear that

this map is the inverse of the local morphism in the other direction. As both the blowup of $\overline{\mathcal{M}}_{0,n}$ along the global ideal sheaf N and the IVC admit forgetful surjective maps to $\overline{\mathcal{M}}_{0,n}$, and as we proved that they are locally isomorphic near every stratum $D_\Gamma \subset \overline{\mathcal{M}}_{0,n}$, it follows that they are globally isomorphic. \square

7.6. A blowup description for arbitrary genus

Recall that the NIVC denotes the normalization of the incidence variety compactification (*i.e.* of the closure of the stratum in the Hodge bundle), and let Γ be a partially ordered level graph of a boundary stratum in the NIVC. For every vertex $v \in V(\Gamma)$, by normality an *adjusting parameter* h_v exists by [BCG⁺19, Proposition 11.13]. Recall that by definition this means that multiplying by h_v^{-1} makes the limiting differential in a degenerating family not identically zero on the irreducible component of the stable curve corresponding to v . Define an ideal locally around the boundary stratum of the NIVC corresponding to Γ by

$$J(\Gamma) := \prod_{(v,v') \in V(\Gamma) \times V(\Gamma)} (h_v, h_{v'})^{w(v)w(v')},$$

where the product runs over *all ordered* pairs of vertices (including the case $v = v'$) and where $w(v) := 2g(v) - 2 + \text{valence}(v)$. Since the blowup in $J(\Gamma)$ makes the adjusting parameters comparable for any two vertices, the (local) blowup of the NIVC along $J(\Gamma)$ is orderly (recall that this means that after the blowup the divisibility relation induces a full order on the set of adjusting parameters; see [BCG⁺19, Definition 11.15]). By the same argument as in the proof of [BCG⁺19, Theorem 14.8], it follows that the normalization of this blowup is isomorphic to the moduli space of multi-scale differentials.

Finally, we show that the locally defined ideals $J(\Gamma)$ are compatible under degeneration, so that they glue to form a global sheaf of ideals J on the NIVC. For this, again it is enough to check compatibility under an edge contraction (recalling that unlike in the genus zero case, the edge can be a loop). First, in the case of a loop, by the formula for $w(v)$, we see that contracting a loop does not change $J(\Gamma)$. Now suppose that two distinct vertices v_1, v_2 of Γ connected by an edge e are merged, when e is contracted, to a vertex v' in the resulting graph Γ' . Note that this contraction makes $h_{v_1} \sim h_{v_2} \sim h_{v'}$ modulo units. Moreover, $w(v') = w(v_1) + w(v_2)$. Then for any vertex u different from v_1, v_2, v' , we have

$$\begin{aligned} (h_{v_1}, h_u)^{2w(v_1)w(u)} (h_{v_2}, h_u)^{2w(v_2)w(u)} &\sim (h_{v'}, h_u)^{2w(v')w(u)}, \\ (h_{v_1}, h_{v_1})^{w(v_1)^2} (h_{v_2}, h_{v_2})^{w(v_2)^2} (h_{v_1}, h_{v_2})^{2w(v_1)w(v_2)} &\sim (h_{v'}, h_{v'})^{w(v')^2}. \end{aligned}$$

It follows that $J(\Gamma')$ specializes to $J(\Gamma)$. Therefore, the local ideals $J(\Gamma)$ can be glued to a global sheaf of ideals J . In summary, we have proven the following theorem.

Theorem 7.8. *The main component $\mathbb{P}(\mathcal{MS}_\mu)$ of $\mathbb{P}(\mathcal{GMS}_\mu)$ is the normalization of the blowup of the NIVC in the ideal sheaf J ; in particular, its coarse moduli space is a projective variety.*

Remark 7.9. For arbitrary genera, one can describe the IVC (and then also the rubber and multi-scale spaces) by blowing up the normalization of the closure of the stratum in the Deligne–Mumford compactification $\overline{\mathcal{M}}_{g,n}$, which we denote by NDM. The argument is similar to the one in the proof of Lemma 7.7. Since the NDM is normal, for every vertex v of Γ , an adjusting parameter h_v for v exists as in [BCG⁺19, Proposition 11.13]. Then the blowup of the NDM along the (local) ideals $(h_{v_1}, \dots, h_{v_k})$, where v_1, \dots, v_k are local maximum vertices of Γ , carries a family of stable differentials, and hence it maps to the IVC by the universal property of the Hodge bundle. The inverse map from the IVC to this blowup is similarly obtained by using the universal property of the blowup.

To see that these local ideals patch together to form a global sheaf of ideals, suppose that a local maximum vertex v_1 joins a lower vertex v_0 via an edge e . Suppose further that e is contracted so that v_1 and v_0 merge as one vertex v'_1 , which makes $h_{v_1} \sim h_{v'_1}$ modulo units. If v'_1 remains a local maximum, then we have $(h_{v_1}, h_{v_2}, \dots, h_{v_k}) = (h_{v'_1}, h_{v_2}, \dots, h_{v_k})$ after contracting e , so these ideals match. If v'_1 is not a local maximum,

then there exists another local maximum vertex, say v_2 , that goes along a path downward to v'_1 (in terms of the partial order of Γ). It follows that h_{v_2} divides $h_{v_1} \sim h_{v'_1}$ and hence $(h_{v_1}, h_{v_2}, \dots, h_{v_k}) = (h_{v_2}, \dots, h_{v_k})$ after contracting e , so these ideals still match.

Appendix. Sign conventions

The sign conventions adopted in [MW20] and in [BCG⁺19] are opposite to one another; as this sign plays a more prominent role in [BCG⁺19], we follow that sign convention, which we now explain in the logarithmic language.

Let (X, M_X) be a log scheme and $\beta \in \overline{M}_X^{\text{gp}}(X)$. The preimage of β in the exact sequence

$$1 \longrightarrow \mathcal{O}_X^\times \longrightarrow M_X^{\text{gp}} \longrightarrow \overline{M}_X^{\text{gp}} \longrightarrow 1$$

is a \mathbb{G}_m -torsor that we denote by $\mathcal{O}_X^\times(\beta)$, from which we construct a line bundle $\mathcal{O}_X(\beta)$ by gluing in the *zero*⁽²⁴⁾ section. In particular,

- (1) if X has divisorial log structure and $\beta \in \overline{M}_X(X)$, then $\mathcal{O}_X(\beta)$ is naturally an ideal sheaf on X ;
- (2) if (X, x) is a DVR with divisorial log structure at x , then the stalk of \overline{M}_X at x is naturally identified with \mathbb{N} , and the association $\beta \mapsto \mathcal{O}_X(\beta)$ sends n to $\mathcal{O}_X(-nx)$;
- (3) if $a \leq b \in \overline{M}_X(X)^{\text{gp}}$, then we have a natural map $\mathcal{O}_X(b) \rightarrow \mathcal{O}_X(a)$.

If $e: u \rightarrow v$ is a directed edge of a graph Γ of length δ_e , and β is a function on the vertices of γ with slope κ along e , then $\beta(v) = \beta(u) + \kappa \cdot \delta_e$. We identify a half-edge h attached to a vertex e with an *outgoing* edge at e .

If $(X/B, \beta)$ is a nuclear object of **Rub**, then the image of β is totally ordered with *largest* element 0. If the image of β has cardinality $N + 1$, then there is a unique isomorphism of totally ordered sets between $\text{Im}(\beta)$ and $\{0, -1, \dots, -N\}$ (the latter having largest element 0). We denote by $\ell: V \rightarrow \{0, -1, \dots, -N\}$ the induced map, in accordance with (3.1).

If e is an edge between vertices u and v , we define $\ell^+(e)$ and $\ell^-(e)$ to be the unique pair of elements of $\{0, -1, \dots, -N\}$ such that $\ell^+(e) \geq \ell^-(e)$ and $\{\ell^+(e), \ell^-(e)\} = \{\ell(u), \ell(v)\}$.

References

- [AOV11] D. Abramovich, M. Olsson, and A. Vistoli, *Twisted stable maps to tame Artin stacks*, J. Algebraic Geom. **20** (2011), no. 3, 399–477, doi:10.1090/S1056-3911-2010-00569-3.
- [AV02] D. Abramovich and A. Vistoli, *Compactifying the space of stable maps*, J. Amer. Math. Soc. **15** (2002), no. 1, 27–75, doi:10.1090/S0894-0347-01-00380-0.
- [AW18] D. Abramovich and J. Wise, *Birational invariance in logarithmic Gromov–Witten theory*, Compos. Math. **154** (2018), no. 3, 595–620, doi:10.1112/S0010437X17007667.
- [ACG11] E. Arbarello, M. Cornalba, and P. Griffiths, *Geometry of algebraic curves. Volume II* (with a contribution by J.D. Harris), Grundlehren math. Wiss., Vol. 268, Springer, Heidelberg, 2011, doi:10.1007/978-3-540-69392-5.
- [BHP⁺23] Y. Bae, D. Holmes, R. Pandharipande, J. Schmitt, and R. Schwarz, *Pixton’s formula and Abel–Jacobi theory on the Picard stack*, Acta Math. **230** (2023), no. 2, 205–319, doi:10.4310/acta.2023.v230.n2.a1.

⁽²⁴⁾Here Marcus and Wise glue in the ∞ -section.

- [BCG⁺18] M. Bainbridge, D. Chen, Q. Gendron, S. Grushevsky, and M. Möller, *Compactification of strata of Abelian differentials*, Duke Math. J. **167** (2018), no. 12, 2347–2416, doi:10.1215/00127094-2018-0012.
- [BCG⁺19] ———, *The moduli space of multi-scale differentials*, preprint arXiv:1910.13492 (2019).
- [BH23] O. Biesel and D. Holmes, *Fine compactified moduli of enriched structures on stable curves*, Mem. Amer. Math. Soc. **285** (2023), no. 1416, doi:10.1090/memo/1416.
- [Boj19] V. Bojković, *Mittag-leffler problems on Berkovich curves*, preprint arXiv:1901.07650 (2019).
- [Che17] D. Chen, *Teichmüller dynamics in the eyes of an algebraic geometer*, In: *Surveys on recent developments in algebraic geometry*, pp. 171–197, Proc. Sympos. Pure Math., Vol. 95, Amer. Math. Soc., Providence, RI, 2017, doi:10.1090/pspum/095/01626.
- [CC19] D. Chen and Q. Chen, *Spin and hyperelliptic structures of log twisted differentials*, Selecta Math. (N.S.) **25** (2019), no. 2, Paper No. 20, doi:10.1007/s00029-019-0467-x.
- [CCM24] D. Chen, M. Costantini, and M. Möller, *On the Kodaira dimension of moduli spaces of Abelian differentials*, Camb. J. Math. **12** (2024), no. 3, 623–752, doi:10.4310/cjm.241001000503.
- [Chi08] A. Chiodo, *Towards an enumerative geometry of the moduli space of twisted curves and r th roots*, Compos. Math. **144** (2008), no. 6, 1461–1496, doi:10.1112/S0010437X08003709.
- [CMS23] M. Costantini, M. Möller, and J. Schwab, *Chern classes of linear submanifolds with application to spaces of k -differentials and ball quotients*, preprint arXiv:2303.17929 (2023).
- [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*, Exp. Math. **32** (2023), no. 3, 545–565, doi:10.1080/10586458.2021.1980750.
- [CMZ24] ———, *The area is a good enough metric*, Ann. Inst. Fourier (Grenoble) **74** (2024), no. 3, 1017–1059, doi:10.5802/aif.3592.
- [DSvZ21] 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*, J. Softw. Algebra Geom. **11** (2021), no. 1, 89–112, doi:10.2140/jsag.2021.11.89.
- [Edi98] B. Edixhoven, *On Néron models, divisors and modular curves*, J. Ramanujan Math. Soc. **13** (1998), no. 2, 157–194.
- [FWY21] H. Fan, L. Wu, and F. You, *Higher genus relative Gromov–Witten theory and double ramification cycles*, J. Lond. Math. Soc. (2) **103** (2021), no. 4, 1547–1576, doi:10.1112/jlms.12417.
- [Fan01] B. Fantechi, *Stacks for everybody*, In: *European Congress of Mathematics, Vol. I* (Barcelona, 2000), pp. 349–359, Progr. Math., vol. 201, Birkhäuser Verlag, Basel, 2001, doi:10.1007/978-3-0348-8268-2_20.
- [FP18] G. Farkas and R. Pandharipande, *The moduli space of twisted canonical divisors*, J. Inst. Math. Jussieu **17** (2018), no. 3, 615–672, doi:10.1017/S1474748016000128.
- [GLN23] A. Giacchetto, D. Lewański, and P. Norbury, *An intersection-theoretic proof of the Harer–Zagier formula*, Algebr. Geom. **10** (2023), no. 2, 130–147, doi:10.14231/AG-2023-004.
- [Gill12] W. D. Gillam, *Logarithmic stacks and minimality*, Internat. J. Math. **23** (2012), no. 7, 1250069, doi:10.1142/S0129167X12500693.
- [GV05] T. Graber and R. Vakil, *Relative virtual localization and vanishing of tautological classes on moduli spaces of curves*, Duke Math. J. **130** (2005), no. 1, 1–37, doi:10.1215/S0012-7094-05-13011-3.

- [GS13] M. Gross and B. Siebert, *Logarithmic Gromov-Witten invariants*, J. Amer. Math. Soc. **26** (2013), no. 2, 451–510, doi:10.1090/S0894-0347-2012-00757-7.
- [HMOP23] D. Holmes, S. Molcho, G. Orecchia, and T. Poirer, *Models of Jacobians of curves*, J. reine angew. Math. **801** (2023), 115–159, doi:10.1515/crelle-2023-0031.
- [HMP⁺25] D. Holmes, S. Molcho, R. Pandharipande, A. Pixton, and J. Schmitt, *Logarithmic double ramification cycles*, Invent. Math. **240** (2025), no. 1, 35–121, doi:10.1007/s00222-025-01318-z.
- [HS21] D. Holmes and J. Schmitt, *Infinitesimal structure of the pluricanonical double ramification locus*, Compos. Math. **157** (2021), no. 10, 2280–2337, doi:10.1112/s0010437x21007557.
- [JPPZ17] F. Janda, R. Pandharipande, A. Pixton, and D. Zvonkine, *Double ramification cycles on the moduli spaces of curves*, Publ. Math. Inst. Hautes Études Sci. **125** (2017), 221–266, doi:10.1007/s10240-017-0088-x.
- [Jar00] T.J. Jarvis, *Geometry of the moduli of higher spin curves*, Internat. J. Math. **11** (2000), no. 5, 637–663, doi:10.1142/S0129167X00000325.
- [Kat00] F. Kato, *Log smooth deformation and moduli of log smooth curves*, Internat. J. Math. **11** (2000), no. 2, 215–232, doi:10.1142/S0129167X0000012X.
- [Kat89a] K. Kato, *Logarithmic degeneration and dieudonné theory*, preprint, available at <https://www.math.brown.edu/dabramov/LOGGEOM/Kato-Dieudonne.pdf> (1989).
- [Kat89b] ———, *Logarithmic structures of Fontaine-Illusie*, In: *Algebraic analysis, geometry, and number theory* (Baltimore, MD, 1988), pp. 191–224, Johns Hopkins Univ. Press, Baltimore, MD, 1989.
- [Li01] J. Li, *Stable morphisms to singular schemes and relative stable morphisms*, J. Differential Geom. **57** (2001), no. 3, 509–578, doi:10.4310/jdg/1090348132.
- [MW20] S. Marcus and J. Wise, *Logarithmic compactification of the Abel-Jacobi section*, Proc. Lond. Math. Soc. (3) **121** (2020), no. 5, 1207–1250, doi:10.1112/plms.12365.
- [MPS23] S. Molcho, R. Pandharipande, and J. Schmitt, *The Hodge bundle, the universal \mathcal{O} -section, and the log Chow ring of the moduli space of curves*, Compos. Math. **159** (2023), no. 2, 306–354, doi:10.1112/S0010437X22007874.
- [MR24] S. Molcho and D. Ranganathan, *A case study of intersections on blowups of the moduli of curves*, Algebra Number Theory **18** (2024), no. 10, 1767–1816, doi:10.2140/ant.2024.18.1767.
- [Ngu24] D.M. Nguyen, *The incidence variety compactification of strata of d -differentials in genus 0*, Int. Math. Res. Not. IMRN (2024), no. 10, 8734–8757, doi:10.1093/imrn/rnad326.
- [Ogu18] A. Ogus, *Lectures on logarithmic algebraic geometry*, Cambridge Stud. Adv. Math., Vol. 178, Cambridge Univ. Press, 2018, doi:10.1017/9781316941614.
- [RSPW19] D. Ranganathan, K. Santos-Parker, and J. Wise, *Moduli of stable maps in genus one and logarithmic geometry, I*, Geom. Topol. **23** (2019), no. 7, 3315–3366, doi:10.2140/gt.2019.23.3315.
- [Sta25] The Stacks Project Authors, *Stacks project*, <http://stacks.math.columbia.edu>, 2025.
- [Wis16] J. Wise, *Moduli of morphisms of logarithmic schemes*, Algebra Number Theory **10** (2016), no. 4, 695–735, doi:10.2140/ant.2016.10.695.
- [Wri15] A. Wright, *Translation surfaces and their orbit closures: an introduction for a broad audience*, EMS Surv. Math. Sci. **2** (2015), no. 1, 63–108, doi:10.4171/EMSS/9.
- [Zor06] A. Zorich, *Flat surfaces*, In: *Frontiers in number theory, physics, and geometry. I*, pp. 437–583, Springer-Verlag, Berlin, 2006.