

THE MODULI SPACE OF MULTI-SCALE DIFFERENTIALS

MATT BAINBRIDGE, DAWEI CHEN, QUENTIN GENDRON, SAMUEL GRUSHEVSKY,
AND MARTIN MÖLLER

ABSTRACT. We construct a compactification $\mathbb{P}\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ of the moduli spaces of abelian differentials on Riemann surfaces with prescribed zeroes and poles. This compactification, called the moduli space of multi-scale differentials, is a complex orbifold with normal crossing boundary. Locally, $\mathbb{P}\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ can be described as the normalization of an explicit blowup of the incidence variety compactification, which was defined in [BCGGM18] as the closure of the stratum of abelian differentials in the closure of the Hodge bundle. We also define families of projectivized multi-scale differentials, which gives a proper Deligne-Mumford stack, and $\mathbb{P}\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ is the orbifold corresponding to it. Moreover, we perform a real oriented blowup of the unprojectivized space $\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ such that the $\mathrm{SL}_2(\mathbb{R})$ -action in the interior of the moduli space extends continuously to the boundary.

A multi-scale differential on a pointed stable curve is the data of an enhanced level structure on the dual graph, prescribing the orders of poles and zeroes at the nodes, together with a collection of meromorphic differentials on the irreducible components satisfying certain conditions. Additionally, the multi-scale differential encodes the data of a prong-matching at the nodes, matching the incoming and outgoing horizontal trajectories in the flat structure. The construction of $\mathbb{P}\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ furthermore requires defining families of multi-scale differentials, where the underlying curve can degenerate, and understanding the notion of equivalence of multi-scale differentials under various rescalings.

Our construction of the compactification proceeds via first constructing an augmented Teichmüller space of flat surfaces, and then taking its suitable quotient. Along the way, we give a complete proof of the fact that the conformal and quasiconformal topologies on the (usual) augmented Teichmüller space agree.

Date: October 29, 2019.

Research of the second author was supported in part by the National Science Foundation under the CAREER grant DMS-13-50396 and a von Neumann Fellowship at the Institute for Advanced Study in Spring 2019.

Research of the fourth author was supported in part by the National Science Foundation under the grant DMS-18-02116.

Research of the fifth author is partially supported by the DFG-project MO 1884/2-1 and by the LOEWE-Schwerpunkt “Uniformisierte Strukturen in Arithmetik und Geometrie”.

CONTENTS

1. Introduction	2
2. Notation and background	9
3. The topology on (classical) augmented Teichmüller space	14
4. Normal forms for differentials on families that acquire a node	27
5. Prong-matched differentials	34
6. Twist groups and level rotation tori	43
7. Families of multi-scale differentials	52
8. Real oriented blowups	61
9. Augmented Teichmüller space of flat surfaces	65
10. The model domain	69
11. Modifying differentials and perturbed period coordinates	74
12. The Dehn space and the complex structure	83
13. The universal property of the Dehn space	101
14. The moduli space of multi-scale differentials	104
15. Extending the $\mathrm{SL}_2(\mathbb{R})$ -action to the boundary	114
Index of notation	116
References	119

1. INTRODUCTION

The goal of this paper is to construct a compactification of the (projectivized) moduli spaces of abelian differentials $\mathbb{P}\Omega\mathcal{M}_{g,n}(\mu)$ of type $\mu = (m_1, \dots, m_n)$ with zeros and poles of order m_i at the marked points. Our compactification shares almost all of the useful properties of the Deligne-Mumford compactification $\overline{\mathcal{M}}_g$ of the moduli space of curves \mathcal{M}_g . These properties include a normal crossing boundary divisor, natural coordinates near the boundary, and representing a natural moduli functor. Applications of the compactification include justification for intersection theory computations, a notion of the tautological ring, an algorithm to compute Euler characteristics of $\mathbb{P}\Omega\mathcal{M}_{g,n}(\mu)$, and potentially contributions to the classification of $\mathrm{SL}_2(\mathbb{R})$ -orbit closures. Throughout this paper the zeroes and poles are labeled. The reader may quotient by a symmetric group action as discussed in Section 2 to obtain the (unmarked) strata of abelian differentials.

The description of our compactification as a moduli space of what we call multi-scale differentials should be compared with the objects characterizing the naive compactification, the *incidence variety compactification (IVC)* we studied in [BCGGM18]. The IVC is defined as the closure of the moduli space $\Omega\mathcal{M}_{g,n}(\mu)$ in the extension of the Hodge bundle $\Omega\overline{\mathcal{M}}_{g,n}$ over $\overline{\mathcal{M}}_{g,n}$ in the holomorphic case, and as the closure in a suitable twist in the meromorphic case. The IVC can have bad singularities near the boundary, e.g. they can fail to be \mathbb{Q} -factorial (see Section 14.3 and Example 14.7), and we are not aware of a good coordinate system near the boundary. Points in the IVC can be described by twisted differentials, whose definition we now briefly recall.

The *dual graph* of a stable curve X has vertices $v \in V(\Gamma)$ corresponding to irreducible components X_v of the stable curve, and edges $e \in E(\Gamma)$ corresponding to nodes q_e . A *level graph* endows Γ with a level function $\ell: V(\Gamma) \rightarrow \mathbb{R}$, and we may assume that its image, called the set of levels $L^\bullet(\Gamma)$, is the set $\{0, -1, \dots, -N\}$ for some $N \in \mathbb{Z}_{\geq 0}$. We write $X_{(i)}$ for the union of all irreducible components of X that are at level i . A *twisted differential of type μ* compatible with a level graph is a collection $(\eta_{(i)})_{i \in L^\bullet(\Gamma)}$ of non-zero meromorphic differentials on the subcurves $X_{(i)}$, having order prescribed by μ at the marked points and satisfying the matching order condition, the matching residue condition, and the global residue condition (GRC) that we restate in detail in Section 2.4.

The top level $X_{(0)}$ is the subcurve on which, in a one-parameter family over a complex disc with parameter t , the limit of differentials ω_t is a non-zero differential $\eta_{(0)}$, while this limit is zero on all lower levels. By rescaling with appropriate powers of t , we obtain the non-zero limits on the lower levels. The order of the levels here reflects the exponents of t . Note that a point in the IVC determines a twisted differential only up to rescaling individually on each irreducible component of the limiting curve.

The notion of a multi-scale differential refines the notion of a twisted differential in three ways. First, the equivalence relation is a rescaling level-by-level, by the level rotation torus (defined below, see also Section 6.3) instead of component-by-component. Second, the graph records besides the level structure an enhancement prescribing the vanishing order at the nodes, see Section 2.5. Third, we additionally record in a prong-matching (defined below, see also Section 5.4) a finite amount of extra data at every node, a matching of horizontal directions for the flat structure at the two preimages of the node.

Definition 1.1. A *multi-scale differential of type μ* on a stable curve X consists of

- (i) an enhanced level structure on the dual graph Γ of X ,
- (ii) a twisted differential of type μ compatible with the enhanced level structure,
- (iii) and a prong-matching for each node of X joining components of non-equal level.

Two multi-scale differentials are considered equivalent if they differ by the action of the level rotation torus. \triangle

The notion of a family of multi-scale differentials requires to deal with the subtleties of the enhanced level graph varying, with vanishing rescaling parameters, and also with the presence of nilpotent functions on the base space. The complete definition of a family of multi-scale differentials, the corresponding functor \mathbf{MS}_μ on the category of complex spaces, and the groupoid \mathcal{MS}_μ will be given in Section 7. They come with projectivized versions, denoted by \mathbb{PMS}_μ and $\mathbb{P}\mathcal{MS}_\mu$.

Theorem 1.2 (Main theorem). *There is a complex orbifold $\Xi\overline{\mathcal{M}}_{g,n}(\mu)$, the moduli space of multi-scale differentials, with the following properties:*

- (1) *The moduli space $\Omega\mathcal{M}_{g,n}(\mu)$ is open and dense within $\Xi\overline{\mathcal{M}}_{g,n}(\mu)$.*
- (2) *The boundary $\Xi\overline{\mathcal{M}}_{g,n}(\mu) \setminus \Omega\mathcal{M}_{g,n}(\mu)$ is a normal crossing divisor.*
- (3) *$\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ admits a \mathbb{C}^* -action, and the projectivization $\mathbb{P}\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ is compact.*
- (4) *$\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ and $\mathbb{P}\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ are algebraic varieties.*

- (5) The complex space underlying $\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ is a coarse moduli space for \mathbf{MS}_μ .
- (6) The complex space underlying $\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ admits a forgetful map to the normalization of the IVC.

In fact, the codimension of a boundary stratum of multi-scale differentials compatible with an enhanced level graph Γ is equal to the number of levels below zero plus the number of horizontal nodes, i.e., nodes joining components on the same level.

Our proof of algebraicity requires us to recast this theorem in the language of stacks.

Theorem 1.3 (Functorial viewpoint). *The groupoid $\mathbb{P}\mathcal{MS}_\mu$ of projectivized multi-scale differentials is a proper Deligne-Mumford stack. Moreover, there is a map of stacks $\mathbb{P}\Xi\overline{\mathcal{M}}_{g,n}(\mu) \rightarrow \mathbb{P}\mathcal{MS}_\mu$, which is an isomorphism over the open substack $\mathbb{P}\Omega\mathcal{M}_{g,n}(\mu)$.*

The map could be made globally an isomorphism by using a hybrid object that is an orbifold where curve automorphisms are present, and a singular analytic space where the toroidal compactification forces such singularities. We elaborate on this in Section 14.3 after the proof of the theorem.

With a similar construction one can obtain a compactification of the space of k -differentials for all $k \geq 1$ with the same good properties, see [CMZ19a] for details.

Other compactifications. We briefly mention the relation with other compactifications in the literature. The space constructed in [FP18] is not a compactification, but rather a reducible space that contains the IVC as one of its components. As emphasized in that paper, the moduli spaces of meromorphic k -differentials can be viewed as generalizations of the double ramification cycles. There are several (partial) compactifications of the (k -twisted version of the) double ramification cycle, see e.g. [HKP18] and [HS19], mostly with focus on extending the Abel-Jacobi maps.

Mirzakhani-Wright [MW17] considered the compactification of holomorphic strata that simply forgets all irreducible components of the stable curve on which the limit differential is identically zero. This is called the WYSIWYG (“what you see is what you get”) compactification. Since this compactification reflects much of the tangent space of an $\mathrm{SL}_2(\mathbb{R})$ -orbit closure, it has proven useful to their classification. This compactification is however not even a complex analytic space, see [CW19].

Applications. Many applications of our compactification are based on the normal crossing boundary divisor and a good coordinate system, given by the perturbed period coordinates (see Section 11.2) near the boundary. The first application in [CMZ19a] shows that the area form is a good metric on the tautological bundle. This is required in [Sau18] for a direct computation of Masur-Veech volumes, and in [CMSZ19] to justify the formula for the spin components.

A second application in [CMZ19b] is the construction of an analog of the Euler sequence for projective space on $\Xi\overline{\mathcal{M}}_{g,n}(\mu)$. This allows to recursively compute all Chern classes of the (logarithmic) cotangent bundle to $\Xi\overline{\mathcal{M}}_{g,n}(\mu)$. In particular this gives a recursive way to compute the orbifold Euler characteristic of the moduli spaces $\mathbb{P}\Omega\mathcal{M}_{g,n}(\mu)$. Moreover, it gives a formula for the canonical bundle. As in the case of the moduli space of curves, this opens the gate towards determining the Kodaira dimension of $\mathbb{P}\Omega\mathcal{M}_{g,n}(\mu)$. This has been solved so far only for some series of cases, see [Gen18], [Bar18] and [FV14].

A third large circle of potential applications concerns the dynamics of the action of $\mathrm{SL}_2(\mathbb{R})$ on $\Omega\mathcal{M}_{g,n}(\mu)$, in particular in the case when the type μ corresponds to holomorphic differentials. In this case the results of Eskin-Mirzakhani-Mohammadi [EM18; EMM15] and Filip [Fil16] show that the closure of every orbit is an algebraic variety. The classification of those orbit closures is an important goal towards which tremendous progress has been made recently, see e.g. constraints found by Eskin-Filip-Wright [EFW18] and the constructions of orbit closures by Eskin-McMullen-Mukamel-Wright in [EMMW18]. Our compactification provides a natural bordification of $\Omega\mathcal{M}_{g,n}(\mu)$.

Theorem 1.4. *There exists an orbifold with corners $\Xi\widehat{\mathcal{M}}_{g,n}(\mu)$ containing $\Omega\mathcal{M}_{g,n}(\mu)$ as open and dense subspace with the following properties.*

- (1) *There is a continuous map $\Xi\widehat{\mathcal{M}}_{g,n}(\mu) \rightarrow \Xi\overline{\mathcal{M}}_{g,n}(\mu)$ whose fiber over a multi-scale differential with N levels below zero is isomorphic to the real torus $(S^1)^N$.*
- (2) *$\Xi\widehat{\mathcal{M}}_{g,n}(\mu)$ admits an $\mathbb{R}_{>0}$ -action, and the quotient $\Xi\widehat{\mathcal{M}}_{g,n}(\mu)/\mathbb{R}_{>0}$ is compact.*
- (3) *The action of $\mathrm{SL}_2(\mathbb{R})$ extends continuously to $\Xi\widehat{\mathcal{M}}_{g,n}(\mu)$.*
- (4) *Points in $\Xi\widehat{\mathcal{M}}_{g,n}(\mu)$ are in bijection with real multi-scale differentials.*

These real multi-scale differentials are similar to multi-scale differentials, with a coarser equivalence relation, see Definition 15.2. This bordification will be constructed in Section 15 as a special case of our construction of level-wise real blowup. This blowup is an instance of the classical real oriented blowup construction, where the blowup is triggered by the level structure underlying a family of multi-scale differentials, see Section 8.

We hope that the study of orbit closures in $\Xi\widehat{\mathcal{M}}_{g,n}(\mu)$ will provide new insights on the classification problem.

New notions and techniques. We next give intuitive explanations of the new objects and techniques used to construct the moduli space of multi-scale differentials.

Prong-matchings. This is simply a choice of how to match the horizontal directions at the pole of the differential at one preimage of a node to the horizontal directions at the zero of the differential at the other preimage of the same node. To motivate that recording this data is necessary to construct a space dominating the normalization of the IVC, consider two differentials in standard form, locally given by $\eta_1 = u^\kappa(du/u)$ and $\eta_2 = C \cdot v^{-\kappa}(dv/v)$ in local coordinates u, v around two preimages of a node given by $uv = 0$ of X , where $C \in \mathbb{C}^*$ is some constant. Then plumbing these differentials on the plumbing fixture $uv = t$ is possible if and only if $\eta_1 = \eta_2$ after the change of coordinates $v = t/u$, which is equivalent to $t^\kappa = -C$. Thus for a given C the different choices of t differ by multiplication by κ -th roots of unity, and the prong-matching is used to record this ambiguity, in the limit of a degenerating family. The notion of a prong-matching will be introduced formally in Section 5.

As the above motivation already indicates, this requires locally choosing coordinates such that the differential takes the standard form in these coordinates. Pointwise, this is a classical result of Strebel. These normal forms for a family of differentials are a technical underpinning of much of the current paper. The relevant analytic results

are proven in Section 4, by solving the suitable differential equations and applying the Implicit Function Theorem in the suitable Banach space.

Level rotation torus. This algebraic torus has one copy of \mathbb{C}^* for each level below zero. Its action makes the intuition of rescaling level by level precise. As indicated above, differentials on lower level in degenerating families are obtained by rescaling by a power of t . As such a family can also be reparameterized by multiplying t by a constant, such a scaled limit on a given component is only well-defined up to multiplication by a non-zero complex number. Suppose now that while keeping differential at one side of the node fixed, we start multiplying the differential on the other side by $e^{i\theta}$. If we start with a given prong-matching, which is just some fixed choice of $(-C)^{1/\kappa}$, this choice of the root is then being multiplied by $e^{i\theta/\kappa}$. Consequently, varying θ from 0 to 2π ends up with the same differential, but with a different prong-matching.

Thus the equivalence relation among multi-scale differentials that we consider records simultaneously all possible rescalings of the differentials on the levels of the stable curve and the action on the prong-matchings. This leads to the notion of the level rotation torus T_Γ , which will be defined as a finite cover of $(\mathbb{C}^*)^N$ in Section 6. See in particular (6.11) for its action on multi-scale differentials.

Twist groups and the singularities at the boundary. The twist group Tw_Γ can be considered as the subgroup of T_Γ fixing all prongs under the rotation action. The rank of this group equals the number N of levels below zero, but the decomposition into levels does in general not induce an isomorphism of the twist group with \mathbb{Z}^N . Instead, there is a subgroup $\text{Tw}_\Gamma^s \subset \text{Tw}_\Gamma$ of finite index that is generated by rotations of one level at a time. We comment below in connection with the model domain why this subgroup naturally appears from the toroidal aspects of our compactification. The quotient group $K_\Gamma = \text{Tw}_\Gamma / \text{Tw}_\Gamma^s$ is responsible for the orbifold structure at the boundary of our compactification. These groups are of course always abelian. Our running Example in Section 2.6, a triangle graph, provides a simple instance where this group K_Γ is non-trivial (see Section 6.4).

From Teichmüller space down to the moduli space. Even though the result of our construction is an algebraic moduli space, our construction of $\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ starts via Teichmüller theory and produces intermediate results relevant for the geometry of moduli spaces of marked meromorphic differentials.

To give the context, recall that recently Hubbard-Koch [HK14] completed a program of Bers to provide the quotient of Abikoff's augmented Teichmüller space $\overline{\mathcal{T}}_{g,n}$ by the mapping class group with a complex structure such that this quotient is isomorphic to the Deligne-Mumford compactification $\overline{\mathcal{M}}_{g,n}$. As an intermediate step they also provided, for any multicurve Λ , the classical Dehn space D_Λ (which Bers in [Ber74] called "deformation space"), the quotient of $\overline{\mathcal{T}}_{g,n}$ by Dehn twists along Λ , with a complex structure. Our proof proceeds along similar lines, taking care at each step of the extra challenges due to the degenerating differential.

As a first step, recall that there are several natural topologies on $\overline{\mathcal{M}}_{g,n}$. One can define the *conformal topology* where roughly a sequence X_n of pointed curves converges to X if there exist diffeomorphisms $g_n: X \rightarrow X_n$ that are conformal on compact subsets that exhaust the complement of nodes and punctures. In the *quasiconformal*

topology one relaxes from conformal to quasiconformal, but requires that the modulus of quasiconformality tends to one. Conformal maps are convenient, since they pull back holomorphic differentials to holomorphic differentials. On the other hand, quasiconformal maps are easier to glue when a surface is constructed from several subsurfaces. We therefore need both topologies, see Section 3 for precise definitions. The following is an abridged version of Theorem 3.2, which was announced in [Mar87] and [EM12].

Theorem 1.5. *If $n \geq 1$, then the conformal and quasiconformal topologies on the augmented Teichmüller space $\overline{\mathcal{T}}_{g,n}$ are equivalent.*

We upgrade this result in Section 3.3 to provide also the universal bundle of one-forms with a conformal topology that coincides with the usual vector bundle topology.

An outline for the construction of $\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ is then the following. We start with a construction of the *augmented Teichmüller space* $\Omega\overline{\mathcal{T}}_{(\Sigma,s)}(\mu)$ of *flat surfaces of type μ* . As a set, this is the union over all multicurves Λ of the moduli spaces $\Omega\mathcal{T}_{\Lambda}^{pm}(\mu)$ of marked prong-matched twisted differentials as defined in Section 5.5. This mimics the classical case, with marked prong-matched twisted differentials taking the role of Λ -marked stable curves. We then provide $\Omega\overline{\mathcal{T}}_{(\Sigma,s)}(\mu)$ with a topology that makes it a Hausdorff topological space in Theorem 9.4. For each multicurve Λ , the subspace of $\Omega\overline{\mathcal{T}}_{(\Sigma,s)}(\mu)$ of strata less degenerate than Λ admits an action of the twist group Tw_{Λ} and the quotient is the *Dehn space* $\Xi\mathcal{D}_{\Lambda}$. Providing $\Xi\mathcal{D}_{\Lambda}$ with a complex structure is the goal of the lengthy plumbing construction in Section 12. As a topological space, $\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ is the quotient of the augmented Teichmüller space $\Omega\overline{\mathcal{T}}_{(\Sigma,s)}(\mu)$ by the action of the mapping class group, and its structure as complex orbifold stems from its covering by the images of Dehn spaces for all Λ .

We next elaborate on two technical concepts in this construction.

Welded surfaces and nearly turning-number-preserving maps. In order to provide the augmented Teichmüller space $\Omega\overline{\mathcal{T}}_{(\Sigma,s)}(\mu)$ with a topology, we roughly declare a sequence $(X_n, \omega_{n,(i)})$ to be convergent if the curves converge in the conformal topology, there exist rescaling parameters $c_{n,(i)}$ such that the rescaled differentials pull back to nearly the limit differential, and such that the rescaling parameters reflect the relative sizes determined by the level graph of the limit nodal curve.

To get a Hausdorff topological space one has to rule out unbounded twisting of the diffeomorphism near the developing node in a degenerating family. The literature contains formulations in the conformal topology that are not convincing and notions based on Fenchel-Nielsen coordinates (see e.g. [ACGH85, Section 15.8]) that do not work well conformally. Our solution is the following.

Take a nodal curve with a twisted differential and perform a real blowup of the nodes, i.e. replacing each preimage of each node with an S^1 . A prong-matching uniquely determines a way to identify the boundary circles at the two preimages of each node to form a seam, thereby obtaining a smooth *welded surface*. On such a surface we have a notion of turning number of arcs non-tangent to the seams and we require the diffeomorphisms g_n exhibiting convergence to be nearly turning-number-preserving. Details are given in Section 5.6.

Level-wise real oriented blowups. Usually, in the definition of the classical Dehn space D_Λ , markings are considered isomorphic if they agree up to twists along Λ . Such a definition however loses their interaction with the marking. We are forced to mark the welded surfaces instead. This, in turn, is not possible over the base B of a family, since the welded surface depends on the choice of the twisted differential in its T_Γ -orbit. As a consequence we define a functorial construction of a level-wise real oriented blowup $\widehat{B} \rightarrow B$ and define markings using the pullback of the family to \widehat{B} , see Section 8. Our construction is similar in spirit to several blowup constructions in the literature, e.g. the Kato-Nakayama blowup [KN99].

The model domain and toroidal aspects of the compactification. The augmented Teichmüller space of flat surfaces parameterizes (marked) multi-scale differentials, and in particular admits families in which the underlying Riemann surfaces can degenerate. In contrast, the model domain only parameterizes equisingular families, where the topology of the underlying (nodal) Riemann surfaces remains constant, and only the scaling of the differential on the components varies, while remaining non-zero. Families of such objects are called model differentials, which serve as auxiliary objects for our construction. The open model domain MD_Λ is a finite cover of a suitable product of (quotients of) Teichmüller spaces and thus automatically comes with a complex structure and a universal family.

We define a toroidal compactification \overline{MD}_Λ of MD_Λ roughly by allowing the scalings to attain the value zero. The actual definition given in Section 10 is not as simple as locally embedding $(\Delta^*)^N \hookrightarrow \Delta^N$, but rather a quotient of this embedding by the group K_Γ defined above. As a result, \overline{MD}_Λ is a smooth orbifold and the underlying singular space is a fine moduli space for families of model differentials, called the model domain.

The plumbing construction and perturbed period coordinates. We use the model domain to induce a complex structure on $\Xi\mathcal{D}_\Lambda$. In order to do this, we define a plumbing construction that starts with a family of model differentials and constructs a family of multi-scale differentials. The point of this construction is that starting with an equisingular family of curves with variable scales for differentials, which may in particular be zero, the plumbing constructs a family of curves of variable topology with a family of non-zero differentials on the smooth fibers. Whenever the scale is non-zero, the plumbing “plumbs” the node, i.e. smoothes it in a controlled way. The goal of our elaborate plumbing construction is to establish the local homeomorphism of the Dehn space with the model domain. As in [BCGGM18], to be able to plumb one needs to match the residues of the differentials at the two preimages of every node, and thus in particular one needs to add a small modification differential. We then argue that the resulting map will still be a homeomorphism of moduli spaces, and to this end we use the perturbed period coordinates introduced in Section 11.

Perturbed period coordinates are coordinates at the boundary of our compactification. They consist of periods of the twisted differential, parameters for the level-wise rescaling, and a classical additional plumbing parameter for each node joining components on the same level. These periods are close, but not actually equal, to the periods

of the plumbed differential, whence the name. See (11.8) for the precise amount of perturbation.

Finally, in Section 12 we complete this setup and define the plumbing map in full generality and prove in Theorem 12.1 that plumbing is a local diffeomorphism. This is used in Theorem 12.2 to show that $\Xi\mathcal{D}_\Lambda$ is a complex orbifold and that the underlying singular complex space is a fine moduli space for (Teichmüller) marked multi-scale differentials.

Algebraicity, families, and the orderly blowup construction. To prove the algebraicity in the main theorem and to prove the precise relation of $\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ to the IVC we need to encounter some of the details of families of multi-scale differentials. First, since $\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ is normal, the forgetful map factors through the normalization of the IVC. This corresponds to memorizing the extra datum of enhancement of the dual graph and the prong-matching. Second, a family of multi-scale differentials admits level-by-level rescaling, while twisted differentials a priori do not. While for twisted differentials there exists a rescaling parameter for each irreducible component, they might be mutually incomparable or, as we say, disorderly. We thus design in Section 14.3 locally a blowup, the *orderly blowup* of the base of a family such that the rescaling parameters can be put in order, i.e. a divisibility relation according to the level structure. However, the resulting blowup is in general not even normal. The third step is thus geometrically the normalization of the resulting space. In families of multi-scale differentials this is reflected by including the notion of a *rescaling ensemble* given in Definition 7.1. It ultimately reflects the normality of the toroidal compactifications by Δ^N/K_Γ used above. This procedure, culminating in Theorem 14.14, is summarized as follows.

Theorem 1.6. *The moduli stack of projectivized multi-scale differentials \mathbb{PMS}_μ is the normalization of the orderly blowup of the normalization of the IVC.*

Algebraicity and the remaining properties of the main theorems above follow from this result. The zoo of notations is summarized in a table at the end of the paper.

Acknowledgments. We are very grateful to the American Institute of Mathematics (AIM) for supporting our research as SQuARE meetings in 2017–2019, where we made much progress on this project. We are also grateful to Institut für Algebraische Geometrie of the Leibniz Universität Hannover, Mathematisches Forschungsinstitut Oberwolfach (MFO), Max Planck Institut für Mathematik (MPIM, Bonn), Casa Matemática Oaxaca (CMO), and the Mathematical Sciences Research Institute (MSRI, Berkeley) as well as the organizers of various workshops there, where various subsets of us met and collaborated on this project. Finally we thank Qile Chen, Matteo Costantini, John Smillie, Jakob Stix, Scott Wolpert, Alex Wright, Jonathan Zachhuber, and Anton Zorich for inspiring discussions on related topics.

2. NOTATION AND BACKGROUND

The purpose of this section is to recall notation and the main result from [BCGGM18]. Along the way we introduce the notion of *enhanced level graphs* that records the extra data of orders of zeros and poles that compatible twisted differentials should have.

2.1. Flat surfaces with marked points and their strata. A *type* of a (possibly meromorphic) abelian differential on a Riemann surface is a tuple of integers $\mu = (m_1, \dots, m_n) \in \mathbb{Z}^n$ with $m_j \geq m_{j+1}$, such that $\sum_{j=1}^n m_j = 2g - 2$. We assume that there are r positive m 's, s zeroes, and l negative m 's, with $r + s + l = n$, i.e., that we have $m_1 \geq \dots \geq m_r > m_{r+1} = \dots = m_{r+s} = 0 > m_{r+s+1} \geq \dots \geq m_n$. Note that $m_j = 0$ is allowed, representing an ordinary marked point. We use the abbreviation $\bar{n} = \{1, \dots, n\}$.

A (*pointed*) *flat surface* or equivalently a (*pointed*) *abelian differential* is a triple (X, \mathbf{z}, ω) , where X is a (smooth and connected) compact genus g Riemann surface, ω is a non-zero meromorphic one-form on X , and $\mathbf{z}: \bar{n} \hookrightarrow X$ is an injective function such that $\mathbf{z}(j)$ is a singularity of degree m_j of ω . We also denote by z_j the marked point $\mathbf{z}(j)$.

The rank g Hodge bundle of holomorphic (stable) differentials on n -pointed stable genus g curves, denoted by $\Omega\overline{\mathcal{M}}_{g,n} \rightarrow \overline{\mathcal{M}}_{g,n}$, is the total space of the relative dualizing sheaf $\pi_*\omega_{\mathcal{X}/\overline{\mathcal{M}}_{g,n}}$, where $\pi: \mathcal{X} \rightarrow \overline{\mathcal{M}}_{g,n}$ is the universal curve. We denote the polar part of μ by $\tilde{\mu} = (m_{r+s+1}, \dots, m_n)$. We then define the (*pointed*) *Hodge bundle twisted by $\tilde{\mu}$* to be the bundle

$$K\overline{\mathcal{M}}_{g,n}(\tilde{\mu}) = \pi_*\left(\omega_{\mathcal{X}/\overline{\mathcal{M}}_{g,n}}\left(-\sum_{j=r+s+1}^n m_j \mathcal{Z}_j\right)\right)$$

over $\overline{\mathcal{M}}_{g,n}$, where we have denoted by \mathcal{Z}_j the image of the section z_j of the universal family π given by the j -th marked point. The formal sums

$$(2.1) \quad \mathcal{Z}^0 = \sum_{j=1}^{r+s} m_j \mathcal{Z}_j \quad \text{and} \quad \mathcal{Z}^\infty = \sum_{j=r+s+1}^n m_j \mathcal{Z}_j$$

are called the (*prescribed*) *horizontal zero divisor* and (*prescribed*) *horizontal polar divisor* respectively.

The *moduli space of abelian differentials of type μ* is denoted (still) by $\Omega\mathcal{M}_{g,n}(\mu) \subseteq K\overline{\mathcal{M}}_{g,n}(\tilde{\mu})$, and consists of those pointed flat surfaces where the divisor of ω is equal to $\sum_{j=1}^n m_j z_j$. We denote by adding \mathbb{P} to the Hodge bundle (resp. to the strata) the projectivization, i.e., when we want to parameterize differentials up to scale. The (*ordered*) *incidence variety compactification* (IVC for short) is then defined to be the closure $\mathbb{P}\Omega\overline{\mathcal{M}}_{g,n}^{\text{inc}}(\mu)$ inside $\mathbb{P}K\overline{\mathcal{M}}_{g,n}(\tilde{\mu})$ of the (projectivized) moduli space of abelian differentials of type μ . A point $(X, \omega, z_1, \dots, z_n) \in \mathbb{P}\Omega\overline{\mathcal{M}}_{g,n}(\tilde{\mu})$ is called a *pointed stable differential*. The main result of [BCGGM18] is to precisely describe this closure, as we recall below.

2.2. Removing the labelling by the $\text{Sym}(\mu)$ -action. We emphasize again that throughout the paper and in particular in the moduli space $\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ in our main theorem the points are labelled. We let $\text{Sym}(\mu) \subset S_n$ be the subgroup of permutations that permutes only points with the same prescribed order m_i . This group acts on the moduli space $\Omega\mathcal{M}_{g,n}(\mu)$ with quotient the moduli space $\Omega\mathcal{M}_g(\mu)$, which gives the usual strata of the Hodge bundle if μ is the zero type of holomorphic differentials. The reader

is invited to check along the whole paper that $\text{Sym}(\mu)$ acts everywhere and in particular on $\Xi\overline{\mathcal{M}}_{g,n}(\mu)$. The quotient $\Xi\overline{\mathcal{M}}_{g,n}(\mu)/\text{Sym}(\mu)$ is a compactification of $\Omega\mathcal{M}_g(\mu)$.

2.3. Graphs, level graphs and ordered stable curves. Throughout this paper Γ will be a graph, connected unless explicitly stated otherwise, with the set of vertices denoted by $V(\Gamma)$, the set of edges by $E(\Gamma)$ and with possibly half-edges at the vertices. We denote by $\text{val}(v)$ the *valence* of a vertex $v \in V$ (a self-loop is counted twice).

A (*weak*) *full order* \succcurlyeq on the graph Γ is an order \succcurlyeq on the set of vertices $V(\Gamma)$ that is reflexive, transitive, and such that for any $v_1, v_2 \in V$ at least one of the statements $v_1 \succcurlyeq v_2$ or $v_2 \succcurlyeq v_1$ holds. The pair $\overline{\Gamma} = (\Gamma, \succcurlyeq)$ is called a *level graph*. In what follows it will be convenient to assume that the full order on $\overline{\Gamma}$ is induced by a *level function* $\ell: V(\Gamma) \rightarrow \mathbb{Z}_{\leq 0}$ such that the vertices of top level are elements of the set $\ell^{-1}(0) \neq \emptyset$, and the comparison between vertices is by comparing their ℓ -values. Any full order can be induced by a level function, but not by a unique one. We thus use the words level graph and a full order on a graph interchangeably. We let $L^\bullet(\overline{\Gamma}) = \{a \in \mathbb{Z} : \ell^{-1}(a) \neq \emptyset\}$ be the set of levels and let $L(\overline{\Gamma})$ be the set of all but the top level. We usually use the *normalized level function*

$$(2.2) \quad \ell: \Gamma \rightarrow \underline{N} = \{0, -1, \dots, -N\},$$

where $N = |L^\bullet(\overline{\Gamma})| - 1 = |L(\overline{\Gamma})| \in \mathbb{Z}_{\geq 0}$ is the number of levels strictly below 0.

For a given level i we call the subgraph of $\overline{\Gamma}$ that consists of all vertices v with $\ell(v) > i$, along with edges between them, the *graph above level i* of $\overline{\Gamma}$, and denote it by $\overline{\Gamma}_{>i}$. We similarly define the graph $\overline{\Gamma}_{\geq i}$ *above or at level i* , and use $\overline{\Gamma}_{(i)}$ to denote the graph *at level i* . Note that these graphs are usually disconnected.

If Γ_X is the dual graph of a stable curve with pointed differential of type μ , we denote by μ_v for $v \in V(\Gamma)$ the subset of the type corresponding to the marked points on the component X_v . We also let $n_v = \text{val}(v) + |\mu_v|$ be the total number of special points (marked points and nodes) of such a component X_v of a stable curve.

The dual graph Γ_X of a pointed stable curve (X, \mathbf{z}) is allowed to have half-edges. These half-edges at a vertex v correspond to the marked points z_i contained in the component X_v .

Definition 2.1. An edge $e \in E(\overline{\Gamma})$ of a level graph $\overline{\Gamma}$ is called *horizontal* if it connects two vertices of the same level, and is called *vertical* otherwise. Given a vertical edge e , we denote by e^+ (resp. e^-) the vertex that is its endpoint of higher (resp. lower) order. \triangle

We denote the sets of vertical and horizontal edges by $E(\overline{\Gamma})^v$ and by $E(\overline{\Gamma})^h$ respectively. Implicit in this terminology is our convention that we draw level graphs so that the map ℓ is given by the projection to the vertical axis.

We call a stable curve X equipped with a full order on its dual graph Γ_X an *ordered stable curve*. We will write X_v for the irreducible component of X associated to a vertex v , and $X_{(i)}$ for the (possibly disconnected) union of the irreducible components X_v such that $\ell(v) = i$. We write q_e for the node associated to an edge e . We call such a node *vertical* or *horizontal* accordingly. The set of nodes of X is denoted by N_X , the set of vertical nodes by N_X^v and the set of horizontal nodes by N_X^h .

For a vertical node q_e of X corresponding to an edge e we write $q_e^+ \in X_{(\ell(e^+)})$ and $q_e^- \in X_{(\ell(e^-))}$ for the two points lying above q_e in the normalization, and for the irreducible components in which they lie, ordered so that $X_{(\ell(e^-))} \prec X_{(\ell(e^+)})$. Moreover we denote the levels of q_e^\pm by $\ell(e^\pm)$, respectively. We use the same notation for horizontal nodes, making an arbitrary choice of label \pm .

2.4. Twisted differentials and the IVC. Recall from [BCGGM18] that a *twisted differential* η of type μ on a stable n -pointed curve (X, \mathbf{z}) is a collection of (possibly meromorphic) differentials η_v on the irreducible components X_v of X , such that no η_v is identically zero, with the following properties:

- (0) (**Vanishing as prescribed**) Each differential η_v is holomorphic and non-zero outside of the nodes and marked points of X_v . Moreover, if a marked point z_j lies on X_v , then $\text{ord}_{z_j} \eta_v = m_j$.
- (1) (**Matching orders**) For any node of X that identifies $q_1 \in X_{v_1}$ with $q_2 \in X_{v_2}$,

$$\text{ord}_{q_1} \eta_{v_1} + \text{ord}_{q_2} \eta_{v_2} = -2.$$
- (2) (**Matching residues at simple poles**) If at a node of X that identifies $q_1 \in X_{v_1}$ with $q_2 \in X_{v_2}$ the condition $\text{ord}_{q_1} \eta_{v_1} = \text{ord}_{q_2} \eta_{v_2} = -1$ holds, then $\text{Res}_{q_1} \eta_{v_1} + \text{Res}_{q_2} \eta_{v_2} = 0$.

Let $\bar{\Gamma} = (\Gamma_X, \succ)$ be a level graph where Γ_X is the dual graph of X . A twisted differential η of type μ on X is called *compatible with $\bar{\Gamma}$* if in addition it also satisfies the following two conditions:

- (3) (**Partial order**) If a node of X identifies $q_1 \in X_{v_1}$ with $q_2 \in X_{v_2}$, then $v_1 \succ v_2$ if and only if $\text{ord}_{q_1} \eta_{v_1} \geq -1$. Moreover, $v_1 \prec v_2$ if and only if $\text{ord}_{q_1} \eta_{v_1} = -1$.

We remark that this condition only uses the partial order induced by $\bar{\Gamma}$ on the vertices that are connected by an edge, while the most subtle condition, which uses the full order, is the following.

- (4) (**Global residue condition**) For every level i and every connected component Y of $X_{>i}$ that does not contain a marked point with a prescribed pole (i.e., there is no $z_i \in Y$ with $m_i < 0$) the following condition holds. Let q_1, \dots, q_b denote the set of all nodes where Y intersects $X_{(i)}$. Then

$$\sum_{j=1}^b \text{Res}_{q_j^-} \eta = 0,$$

where by definition $q_j^- \in X_{(i)}$.

For brevity, we write GRC for the global residue condition. We denote a twisted differential compatible with a level \preceq by $(X, \mathbf{z}, \eta, \preceq)$. Moreover, we will usually group the restrictions of the twisted differential η according to the levels of ℓ . We will denote the restriction of η to the subsurface $X_{(i)}$ by $\eta_{(i)}$.

We have shown in [BCGGM18, Theorem 1.5]:

Theorem 2.2. *A pointed stable differential (X, ω, \mathbf{z}) is contained in the incidence variety compactification of $\mathbb{P}\Omega\mathcal{M}_{g,n}(\mu)$ if and only if the following conditions hold:*

- (i) *There exists an order \succ on the dual graph Γ_X of X such that its maxima are the irreducible components X_v of X on which ω is not identically zero.*
- (ii) *There exists a twisted differential η of type μ on X , compatible with the level graph $\bar{\Gamma} = (\Gamma_X, \succ)$.*
- (iii) *On every irreducible component X_v where ω is not identically zero, $\eta_v = \omega|_{X_v}$.*

2.5. Enhanced level graphs. Note that a boundary point of the IVC does not necessarily determine a twisted differential uniquely, see [BCGGM18, Examples 3.4 and 3.5]. The full combinatorics of a twisted differential is encoded by the following notion.

An *enhanced level graph* Γ^+ of type $\mu = (m_1, \dots, m_n)$ is a level graph $\bar{\Gamma}$ together with a numbering of the half-edges by \bar{n} and with an assignment of a positive number $\kappa_e \in \mathbb{N}$ for each vertical edge $e \in E(\bar{\Gamma})^v$. The *degree* of a vertex v in Γ^+ is defined to be

$$\deg(v) = \sum_{j \rightarrow v} m_j + \sum_{e \in E^+(v)} (\kappa_e - 1) - \sum_{e \in E^-(v)} (\kappa_e + 1),$$

where the first sum is over all half-edges incident to v , and the remaining sums are over the edges $E^+(v)$ and $E^-(v)$ incident to v that are going from v to a respectively lower and upper vertex. In terms of the notation in Definition 2.1, the set $E^+(v)$ is the set of edges $\{e \in E(\bar{\Gamma}) : e^+ = v\}$. We require that

- (i) **(Admissible degrees)** the degree of each vertex is even and at least -2 , and
- (ii) **(Stability)** the valence of each vertex of degree -2 is at least three.

Our notion of enhancement is equivalent to the notion of twist used e.g. in [FP18] or [CMSZ19]. The main example is the enhanced level graph Γ_X^+ of a twisted differential (X, \mathbf{z}, η) , obtained by assigning to each vertical node q the weight

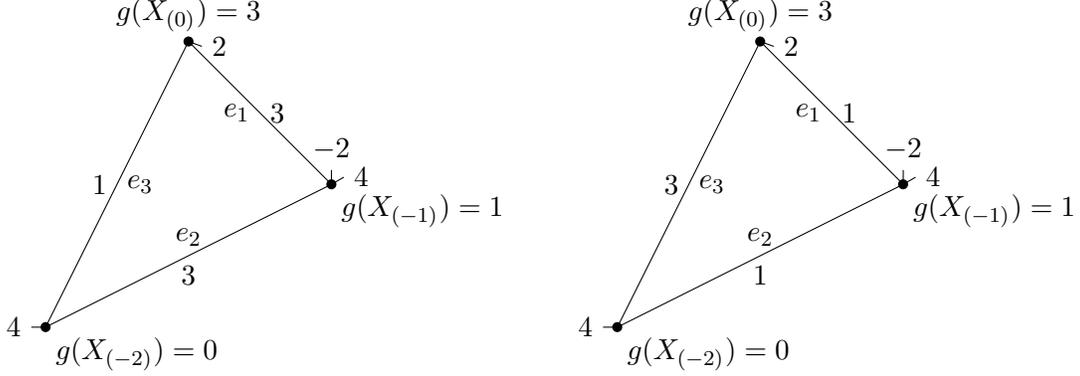
$$(2.3) \quad \kappa_q = \text{ord}_{q^+} \eta + 1.$$

In these terms, the above stability condition is equivalent to stability of (X, \mathbf{z}) . The degree of a vertex v is the degree of η_v . The admissible degrees condition ensures that such a Γ^+ can be realized as the enhanced level graph of some twisted differential. We also say that a twisted differential (X, η) is compatible with Γ^+ if it is compatible with the underlying level graph $\bar{\Gamma}$ and if the markings of Γ^+ are the weights of η just defined.

In order to keep notation concise, we will denote by Γ the dual graph Γ_X of a curve X , a level graph $\bar{\Gamma}$ and write for an enhanced graph Γ^+ , or simply Γ , as appropriate.

2.6. The running example. In order to illustrate the notions that were introduced, we will describe an example. This example will be used throughout the text to exemplify the different notions that we will introduce. We will refer to it as the *running example*.

The example is for the moduli space $\Omega\mathcal{M}_{5,4}(4, 4, 2, -2)$. We fix the curve whose dual graph is a triangle, with the level function taking three different values $0, -1, -2$ on it, so that the level graph is fixed. The irreducible components are of genus 3 (at top level), genus 1 (at the intermediate level) and genus 0 (at the bottom level). This level graph admits two different enhanced structures, which we denote Γ_1 and Γ_2 , as pictured in Figure 1.

FIGURE 1. Two different enhanced orders Γ_1 and Γ_2 on Γ .

We denote the twisted differentials compatible with the level graphs Γ_i by (X, \mathbf{z}, η^i) . The enhanced structure Γ_1 tells us that the differential $\eta_{(-2)}^1$ is in $\Omega\mathcal{M}_{0,3}(4, -2, -4)$, $\eta_{(-1)}^1$ is in $\Omega\mathcal{M}_{1,4}(4, 2, -2, -4)$ and $\eta_{(0)}^1$ is in $\Omega\mathcal{M}_{3,3}(2, 2, 0)$. Similarly $\eta_{(-2)}^2$ is in $\Omega\mathcal{M}_{0,3}(4, -2, -4)$, $\eta_{(-1)}^2$ is in $\Omega\mathcal{M}_{1,4}(4, 0, -2, -2)$ and $\eta_{(0)}^2$ is in $\Omega\mathcal{M}_{3,3}(2, 2, 0)$. The global residue condition in both cases says that the differential $\eta_{(-1)}^i$ has no residue at its pole at the point $q_{e_1}^-$. Note that it follows from [GT17, Theorem 1.2] that this locus is not empty.

3. THE TOPOLOGY ON (CLASSICAL) AUGMENTED TEICHMÜLLER SPACE

The classical augmented Teichmüller space contains the Teichmüller space as a dense subset such that the action of the mapping class group extends continuously and such that the quotient by the mapping class group is the Deligne-Mumford compactification of the moduli space of curves. In this section we compare various topologies on the augmented Teichmüller space, and on the related spaces of one-forms.

3.1. Augmented Teichmüller space. To give the precise definition of the augmented Teichmüller space, we fix a “base” compact n -pointed oriented differentiable surface (Σ, \mathbf{s}) of genus g . We regard \mathbf{s} as a set of $n \geq 0$ distinct labeled points $\{s_1, \dots, s_n\} \subset \Sigma$, or alternatively as an injective function $\mathbf{s}: \bar{n} \hookrightarrow \Sigma$. Let $\mathcal{T}_{g,n} = \mathcal{T}_{(\Sigma, \mathbf{s})}$ be the Teichmüller space of (Σ, \mathbf{s}) . Next, recall that a *multicurve* Λ on $\Sigma \setminus \mathbf{s}$ is a collection of disjoint simple closed curves, such that no two curves are isotopic on $\Sigma \setminus \mathbf{s}$, and no curve in Λ is isotopic to any puncture s_i . Two multicurves are equivalent if the curves they consist of are pairwise isotopic. To ease notation, we will speak of curves of a multicurve Λ both when we mean the actual curves or their isotopy equivalence classes, as should be clear from the context.

Definition 3.1. A *marked pointed stable curve* (X, \mathbf{z}, f) is a pointed stable curve (X, \mathbf{z}) together with a marking $f: (\Sigma, \mathbf{s}) \rightarrow (X, \mathbf{z})$, where a *marking* of a pointed stable curve is a continuous map $f: \Sigma \rightarrow X$ such that

- (i) the inverse image of every node $q \in X$ is a simple closed curve on $\Sigma \setminus \mathbf{s}$,

- (ii) if we denote by $\Lambda \subset \Sigma$ the set of the preimages of the set of nodes N_X of X , which is a multicurve on Σ that we call the *pinched multicurve*, then the restriction of f to $\Sigma \setminus \Lambda$ is an orientation-preserving diffeomorphism $\Sigma \setminus \Lambda \rightarrow X \setminus N_X$,
- (iii) the map f preserves the marked points, that is, $f \circ \mathbf{s} = \mathbf{z}$.

Two marked pointed stable curves are equivalent if there is an isomorphism of pointed stable curves that identifies the markings up to isotopy rel \mathbf{s} . \triangle

The *augmented Teichmüller space* $\overline{\mathcal{T}}_{g,n} = \overline{\mathcal{T}}_{(\Sigma, \mathbf{s})}$ is the set of all equivalence classes of pointed stable curves marked by (Σ, \mathbf{s}) . We caution the reader that $\overline{\mathcal{T}}_{g,n}$ is not a manifold, and is not even locally compact at the boundary, in the topology which we define below.

The *mapping class group* $\text{Mod}_{g,n}$ acts properly discontinuously on the classical Teichmüller space $\mathcal{T}_{g,n}$ and this action (by pre-composition of the marking) extends to a continuous action on the augmented Teichmüller space (whose topology is defined below).

The augmented Teichmüller space is stratified according to the pinched multicurve. Given a multicurve $\Lambda \subset \Sigma \setminus \mathbf{s}$, we define $\mathcal{T}_\Lambda \subset \overline{\mathcal{T}}_{g,n}$ to be the stratum consisting of stable curves where exactly the curves in Λ have been pinched to nodes. In particular, the empty multicurve recovers the interior $\mathcal{T}_\emptyset = \mathcal{T}_{g,n}$. Each \mathcal{T}_Λ is itself a finite unramified cover of the product of the Teichmüller spaces of the components of $(\Sigma, \mathbf{s}) \setminus \Lambda$ that takes into account the identification of the branches of the nodes. In particular each \mathcal{T}_Λ is smooth.

The topology on the augmented Teichmüller space can be described in several ways. For us, the conformal topology (introduced by [Mar87], see also Earle-Marden [EM12]) will be most useful. Abikoff [Abi77] described several equivalent topologies on the augmented Teichmüller space. We recall the definition of his *quasiconformal topology* below (somewhat confusingly, he called this the conformal topology). The equivalence of the two topologies is claimed in [EM12, Theorem 6.1]. We include a complete proof of this equivalence here. We mention [Mon09] for several other viewpoints of the topology, mainly based on hyperbolic length functions.

We define an *exhaustion* of a (possibly open) Riemann surface X to be a sequence of compact subsurfaces with boundary, $K_m \subset X$, such that each K_m is a deformation retract of X , and such that the union $\cup_{m=1}^\infty K_m$ is all of X . An important example of an exhaustion that is used throughout this article is the following. For any sequence ϵ_m of positive numbers (smaller than the Margulis constant) converging to zero, the ϵ_m -thick parts of $X \setminus \mathbf{z}$, denoted by $(X, \mathbf{z})_{\epsilon_m}$, form an exhaustion of $X \setminus \mathbf{z}$. Note that the fact that the ϵ_m are smaller than the Margulis constant ensures that the thin part is a union of annular neighborhoods of short geodesics or cusps.

Let (X, \mathbf{z}, f) be a marked pointed stable curve in $\overline{\mathcal{T}}_{(\Sigma, \mathbf{s})}$, and let $X^s = X \setminus N_X$ denote the smooth part of X , that is, the complement of its nodes. A sequence of marked pointed stable curves (X_m, \mathbf{z}_m, f_m) in $\overline{\mathcal{T}}_{g,n}$ *converges quasiconformally* to (X, \mathbf{z}, f) if for some exhaustion $\{K_m\}$ of X^s , there exists a sequence of quasiconformal maps $g_m: K_m \rightarrow X_m$ such that for each m the maps $f_m \circ f^{-1}$ and g_m are homotopic on K_m , the map g_m respects the marked points (i.e. $g_m \circ \mathbf{z} = \mathbf{z}_m$), and the quasiconformal dilatations $\|\bar{\partial}g_m/\partial g_m\|_\infty$ tend to 0 as $m \rightarrow \infty$. The sequence *converges conformally* if the K_m are instead an exhaustion of $X^s \setminus \mathbf{z}$ and the g_m can be taken to be conformal. We

call the topologies on $\overline{\mathcal{T}}_{(\Sigma, \mathbf{s})}$ induced by these notions of convergence the *quasiconformal topology* and the *conformal topology*, respectively.

Note that for conformal convergence it no longer makes sense to require that the g_m respect the marked points, since they are not in the domain. However, each marked point of X is contained in a unique connected component of $X^s \setminus K_m$, and the hypothesis that $f_m \circ f^{-1} \simeq g_m$ forces g_m to respect these complementary components.

We will sometimes need the conformal maps g_m to respect the marked points. We say that (X_m, \mathbf{z}_m, f_m) converges *strongly conformally* to (X, \mathbf{z}, f) if the conformal maps g_m can be defined on an exhaustion $\{K_m\}$ of X^s and the g_m respect the marked points (i.e. $g_m \circ \mathbf{z} = \mathbf{z}_m$).

The idea of the proof of the equivalence of these topologies is that given a quasiconformal map on X with small dilatation and an open set U (generally a neighborhood of a node or marked point), one can find a nearby quasiconformal map which pushes all of the “quasiconformality” into U . Since strong conformal convergence requires the maps to be conformal near the marked points, we see that it should only be equivalent to the other types of convergence in the presence of nodes, as we will need U to be a neighborhood of the nodes in this case.

Theorem 3.2. *If $n \geq 1$, then the conformal and quasiconformal topologies on $\overline{\mathcal{T}}_{g,n}$ are equivalent. For any n , if $X \in \overline{\mathcal{T}}_{g,n}$ has any nodes, then quasiconformal, conformal, and strong conformal convergence of a sequence to X are all equivalent.*

Given a measurable subset E of a Riemann surface X , we denote by $\mathcal{M}(E)$ the Banach space of measurable L^∞ -Beltrami differentials supported on E , and we denote by $\mathcal{M}^r(E) \subset \mathcal{M}(E)$ the open ball of radius r .

The proof of Theorem 3.2 is based on the following Lemma.

Lemma 3.3. *Let (X, \mathbf{z}) be a compact pointed Riemann surface and $K \subset U \subset X$ subsets such that K is compact with positive Lebesgue measure and U is open. Then there is a constant $0 < k < 1$ such that for every Beltrami differential ν on $X \setminus K$ with $\|\nu\|_\infty < k$, there exists a quasiconformal homeomorphism $f_\nu: X \rightarrow X$, preserving the marked points, such that the Beltrami differential of f_ν restricted to $X \setminus K$ agrees with ν , and $f_\nu(K) \subset U$.*

Moreover, the collection of such maps f_ν may be regarded as a holomorphic map $\mathcal{M}^k(X \setminus K) \rightarrow \text{QC}^0(X)$, to the space of quasiconformal homeomorphisms of X isotopic to the identity, equipped with the compact-open topology.

Proof. A Beltrami differential $\nu \in \mathcal{M}^1(X)$ induces a conformal structure on X which we denote by X_ν . This defines a holomorphic map

$$\Phi: \mathcal{M}^1(X) = \mathcal{M}^1(K) \oplus \mathcal{M}^1(X \setminus K) \rightarrow \mathcal{T}_{g,n}.$$

Consider the derivative operator defined as

$$D = D_1 \Phi_{(0,0)}: \mathcal{M}(K) \rightarrow T_{(X,\mathbf{z})} \mathcal{T}_{g,n}$$

of Φ restricted to the tangent space of the first factor of the splitting. We claim that D is surjective. This is equivalent to show that the dual operator $D^*: T_{(X,\mathbf{z})}^* \mathcal{T}_{g,n} \rightarrow$

$\mathcal{M}(K)^*$ is injective. Under the usual identification of the cotangent space to Teichmüller space at (X, \mathbf{z}) with $Q(X, \mathbf{z})$, the space of quadratic differentials q on X with at worst simple poles contained in \mathbf{z} , the dual D^* is given explicitly by the pairing

$$D^*(q)(\nu) = \int_K q\nu.$$

Taking ν_q to be the restriction of $\bar{q}/|q|$ to K , we obtain

$$D^*(q)(\nu_q) = \int_K |q| > 0,$$

so injectivity follows.

Since $T_{(X, \mathbf{z})}\mathcal{T}_{g, n}$ is finite-dimensional, the kernel of D is closed and of finite codimension, so it has a complementary closed subspace. Thus D is a split surjection, and the Implicit Function Theorem applies.

By the Implicit Function Theorem, there is for some $0 < k < 1$ a holomorphic map $\psi: \mathcal{M}^k(X \setminus K) \rightarrow \mathcal{M}^1(K)$ such that $\Phi(\psi(\nu), \nu) = (X, \mathbf{z})$. In other words for each $\nu \in \mathcal{M}^k(X \setminus K)$, there is a quasiconformal map $f_\nu: X \rightarrow X$ with Beltrami differential given by $\psi(\nu) + \nu$.

The map $\nu \mapsto f_\nu$ can be regarded as a map $\Psi: \mathcal{M}^k(X \setminus K) \rightarrow \text{QC}^0(X)$. By holomorphic dependence of solutions to the Beltrami equation on parameters (see e.g. [Hub06]), this map Ψ is holomorphic, and in particular continuous, as desired. Therefore, by the definition of the compact-open topology, by possibly decreasing the constant k , we can make $f_\nu(K) \subset U$. \square

Proof of Theorem 3.2. We first show that quasiconformal convergence implies conformal convergence and, if nodes are present, also strong conformal convergence. Suppose a sequence of marked pointed curves (X_m, \mathbf{z}_m, f_m) converges to (X, \mathbf{z}, f) in the quasiconformal topology, so that there is an exhaustion of X^s by compact sets K_m and quasiconformal maps $g_m: K_m \rightarrow X_m$ isotopic to f_m , whose dilatation tends to 0. Let $U \subset X$ be an (arbitrarily small) open neighborhood of the nodes and the marked points. To show convergence in the conformal topology, we must produce, for m sufficiently large, a conformal map $h_m: X \setminus U \rightarrow X_m$ isotopic to f_m .

Let $K \subset U$ be compact with positive Lebesgue measure. By Lemma 3.3, for m sufficiently large, there is a quasiconformal map $k_m: X \rightarrow X$ sending K into U and whose Beltrami differential restricted to $X \setminus U$ agrees with the Beltrami differential of g_m . The composition $h_m = g_m \circ k_m^{-1}$ is then conformal outside U as desired.

If X has nodes, this argument works just as well to get strong conformal convergence by taking U to be a neighborhood of the nodes only.

We now show that conformal convergence implies quasiconformal convergence. We choose an exhaustion $\{K_m\}$ of $X^s \setminus \mathbf{z}$ so that the inclusion $K_m \hookrightarrow X \setminus \mathbf{z}$ is a homotopy equivalence, and let $g_m: K_m \rightarrow X_m$ be the conformal maps which exhibit conformal convergence. Let $\{K_m^f\}$ be the exhaustion of X^s obtained by filling in the disks containing the marked points z_i (in this proof, the superscript f will always mean that we fill in the disks around the marked points). We must show that we can replace the g_m with quasiconformal maps g_m^f on K_m^f in the same isotopy class which respect

the marked points and whose dilatation tends to 0. For concreteness, we fill in the disk containing z_1 .

Let Y be the component of X containing z_1 , and let $J_m = K_m \cap Y$ and note that $J_m \hookrightarrow Y$ is a homotopy equivalence. We first represent Y as \mathbb{H}/Γ for some Fuchsian group Γ . The fundamental group of the subsurface $L_m^f = g_m(J_m)^f \subset X_m$ is a subgroup of the fundamental group of the component of X_m containing L_m^f , so it determines a cover of X_m , which we represent as \mathbb{H}/Γ_m for some Fuchsian group Γ_m . Let $\tilde{J}_m^f \subset \mathbb{H}$ and $\tilde{L}_m^f \subset \mathbb{H}$ be the unique connected subsurfaces, invariant under the Fuchsian groups, with $\tilde{J}_m^f/\Gamma = J_m^f$ and $\tilde{L}_m^f/\Gamma_m = L_m^f$. The conformal map g_m then lifts to a conformal map $\tilde{g}_m: \tilde{J}_m^f \rightarrow \tilde{L}_m^f$ which is equivariant in the sense that

$$(3.1) \quad \rho_m(\gamma) \cdot \tilde{g}_m(z) = \tilde{g}_m(\gamma \cdot z)$$

for some isomorphism $\rho_m: \Gamma \rightarrow \Gamma_m$ and for each $\gamma \in \Gamma$. Note that the Fuchsian groups are really only defined up to conjugacy. We normalize the Γ, Γ_m and all related objects by requiring that $0, 1, \infty$ belong to the limit set of Γ and the extension of \tilde{g}_m to this limit set fixes these three points.

We claim now that \tilde{g}_m converges uniformly on compact sets to the identity and that the Fuchsian groups Γ_m converge to Γ algebraically (meaning that for each $\gamma \in \Gamma$, the limit of $\rho_m(\gamma)$ is γ). By Montel's Theorem, any subsequence of \tilde{g}_m has a further subsequence which converges uniformly on compact sets to some $G: \mathbb{H} \rightarrow \overline{\mathbb{H}}$. Since each \tilde{g}_m is conformal and fixes three points on the boundary of \mathbb{H} , in fact G must be the identity map. Since every subsequence of \tilde{g}_m converges to the identity, we see that g_m converges uniformly on compact sets to the identity. Algebraic convergence of Γ_m to Γ then follows immediately from (3.1).

Now choose a conformal map $p: \Delta \rightarrow \tilde{J}_m^f$ whose image is an open disk U which covers the complementary disk containing z_1 , which sends 0 to z_1 , and which maps $\partial\Delta$ onto a smooth curve γ which is eventually contained in \tilde{J}_m . The composition $\tilde{g}_m \circ p$ sends the boundary circle to a smooth curve $\gamma_m \subset \tilde{L}_m^f$ which bounds a disk U_m containing the marked point $z_{1,m}$. Choose two points $a_1, a_2 \in \partial\Delta$, and let $p_m: \Delta \rightarrow \tilde{L}_m^f$ be the Riemann mapping of Δ onto U_m which is normalized so that $p_m(w) = \tilde{g}_m \circ p(w)$ for the points $w = a_1, a_2$, and 0. Since \tilde{g}_m converges to the identity uniformly on compact sets, the sets U_m converge to U in the Carathéodory topology on disks (see [McM94, Section 5.1]). In fact, they converge uniformly on $\overline{\Delta}$, since the closed sets $\mathbb{H} \setminus U_m$ are uniformly locally connected (see [Pom92, Corollary 2.4]). Let $\alpha_m: \Delta \rightarrow \Delta$ be the Douady-Earle extension of $p_m^{-1} \circ \tilde{g}_m \circ p|_{\partial\Delta}$. The boundary map is uniformly close to the identity, so $\alpha_m(0)$ is close to 0, and we may construct a quasiconformal map $\beta_m: \Delta \rightarrow \Delta$ which is the identity on the boundary, sends $\alpha_m(0)$ back to 0, and has small quasiconformal dilatation. Finally, we define our extension g_m^f of g_m as before on the complement of U , and we define it to be $p_m \circ \beta_m \circ \alpha_m \circ p^{-1}$ on \overline{U} . This is the desired quasiconformal extension of g_m sending z_1 to $z_{1,m}$. \square

Another reformulation of the same idea allows to assume, for X smooth and with at least one marked point, g_m to be conformal on an exhaustion K_m of X minus a single marked point.

These ideas allow a similar definition of the “universal curve” over the augmented Teichmüller space. While we are not interested directly in this object as it is not an honest flat family of curves, it is useful for defining universal curves over other spaces.

Convergence of sequences in the universal curve is defined analogously to convergence in $\overline{\mathcal{T}}_{g,n}$. Given (X, \mathbf{z}, p, f) such that p is not a node, we say that a sequence $(X_m, \mathbf{z}_m, p_m, f_m)$ converges to (X, \mathbf{z}, p, f) if $(X_m, \mathbf{z}_m, f_m) \rightarrow (X, \mathbf{z}, f)$ as pointed stable curves; and moreover, if $g_m: K_m \rightarrow X_m$ are the (conformal or quasiconformal) maps which exhibit this convergence, then $g_m^{-1}(p_m)$ converges to p .

This definition does not work in the case when p is a node, as then $p \in X \setminus K_m$, and thus the map g_m is never defined at p . Instead we require, for m sufficiently large, the point p_m to lie in the end of $X_m \setminus g_m(K_m)$ that corresponds to the end of $X \setminus K_m$ containing p . This is well-defined, since g_m eventually induces a bijection of the components of $X \setminus K_m$ and $X_m \setminus g_m(K_m)$, as remarked above.

3.2. The Dehn space and the Deligne-Mumford compactification. We briefly recall the construction of the Deligne-Mumford compactification $\overline{\mathcal{M}}_{g,n}$ of $\mathcal{M}_{g,n}$, as well as the closely related Dehn spaces \mathcal{D}_Λ , which give simple models for $\overline{\mathcal{M}}_{g,n}$ near its boundary. For more details and proofs of all of these statements, we refer the reader to [HK14], see also [ACG11] for some of the statements.

Given a multicurve $\Lambda \subset \Sigma \setminus \mathbf{s}$, the *full Λ -twist group* $\text{Tw}_\Lambda^{\text{full}} \subset \text{Mod}_{g,n}$ is the free abelian subgroup generated by Dehn twists around the curves of Λ . The *Dehn space* \mathcal{D}_Λ is the space obtained by adjoining to $\mathcal{T}_{g,n}$ the stable curves where $f(\Lambda')$ for some subset Λ' of Λ has been pinched, and then taking the quotient by $\text{Tw}_\Lambda^{\text{full}}$. That is,

$$\mathcal{D}_\Lambda = \bigcup_{\Lambda' \subset \Lambda} \mathcal{T}_{\Lambda'} / \text{Tw}_\Lambda^{\text{full}}.$$

(Bers [Ber74] called this space the “deformation space”.) Each \mathcal{D}_Λ is a contractible complex manifold. It has a unique complex structure which agrees with the complex structure induced by $\mathcal{T}_{g,n}$ in the interior, and such that the boundary is a normal crossing divisor.

The universal curve $\pi: \mathcal{X}_\Lambda \rightarrow \mathcal{D}_\Lambda$ is the quotient

$$\mathcal{X}_\Lambda = \bigcup_{\Lambda' \subset \Lambda} \overline{\mathcal{X}}_{g,n}|_{\mathcal{T}_{\Lambda'}} / \text{Tw}_\Lambda^{\text{full}},$$

where $\overline{\mathcal{X}}_{g,n}$ is the universal family over $\mathcal{T}_{\Lambda'}$ and where the full twist group acts trivially on each fiber. It is a flat family of stable curves over \mathcal{D}_Λ , as can be seen using the plumbing construction of [HK14].

The *Deligne-Mumford compactification* $\overline{\mathcal{M}}_{g,n}$ of $\mathcal{M}_{g,n}$ is the quotient $\overline{\mathcal{T}}_{g,n} / \text{Mod}_{g,n}$. For each multicurve Λ , the natural map $\mathcal{D}_\Lambda \rightarrow \overline{\mathcal{M}}_{g,n}$ is a local homeomorphism. The image of \mathcal{D}_Λ is the complement of the locus of stable curves with a node not arising from pinching Λ . These local homeomorphisms provide an atlas of charts for $\overline{\mathcal{M}}_{g,n}$ which give it the structure of a compact complex orbifold such that the boundary is a normal crossing divisor.

One may also compactify $\mathcal{M}_{g,n}$ as a projective variety $\overline{\mathcal{M}}_{g,n}^{\text{alg}}$ (see [DM69] or [ACG11]). Hubbard-Koch [HK14] showed that $\overline{\mathcal{M}}_{g,n}^{\text{alg}} \cong \overline{\mathcal{M}}_{g,n}$ as complex orbifolds, so the natural topology of the algebraic variety $\overline{\mathcal{M}}_{g,n}^{\text{alg}}$ gives yet another equivalent topology on $\overline{\mathcal{M}}_{g,n}$.

3.3. Spaces of one-forms. We now consider topologies on various spaces of surfaces with holomorphic one-forms. For surfaces with one-forms, the conformal topology is much more convenient than the quasiconformal topology, as pullbacks of holomorphic one-forms by quasiconformal maps are in general only measurable. On the other hand, these spaces already have topologies coming from algebraic geometry, and we will show that these topologies coincide.

Consider the universal curve over the Dehn space $\pi: \mathcal{X}_\Lambda \rightarrow \mathcal{D}_\Lambda$, with its relative cotangent sheaf $\omega_{\mathcal{X}_\Lambda/\mathcal{D}_\Lambda}$. The pushforward $\pi_*\omega_{\mathcal{X}_\Lambda/\mathcal{D}_\Lambda}$ is the sheaf of sections of the *Hodge bundle* $\Omega\mathcal{D}_\Lambda \rightarrow \mathcal{D}_\Lambda$, a (trivial) rank g vector bundle whose fiber over a point X is the space $\Omega(X)$ of stable forms on X . As $\Omega\mathcal{D}_\Lambda$ is a vector bundle, it comes with a natural topology, which we call the *vector bundle topology*.

On the other hand, the conformal topology on \mathcal{D}_Λ gives a second natural topology on $\Omega\mathcal{D}_\Lambda$. A sequence $(X_m, \mathbf{z}_m, \omega_m, f_m)$ of marked pointed stable forms converges to $(X, \mathbf{z}, \omega, f)$ in the *conformal topology* if for some exhaustion K_m of $X^s \setminus \mathbf{z}$, there is a sequence of conformal maps $g_m: K_m \rightarrow X_m$ such that $f_m \simeq g_m \circ f$ and $g_m^*\omega_m$ converges to ω uniformly on compact sets. Again, we say that such a sequence converges *strongly conformally* if the g_m are moreover defined on an exhaustion K_m of X^s and respect the marked points. We show below that these topologies agree.

Lemma 3.4. *Suppose $3g - 3 + n > 0$. Let (X_m, \mathbf{z}_m, f_m) be a sequence in \mathcal{D}_Λ converging to (X, \mathbf{z}, f) , and let $g_m: K_m \rightarrow X_m$ be a sequence of conformal maps exhibiting this convergence, where K_m is an exhaustion of X . Then the maps g_m (regarded as maps into the universal curve \mathcal{X}_Λ) converge uniformly on compact sets to the identity map on X .*

Proof. First, we claim that there is a subsequence which converges uniformly on compact sets. We show via the Arzelá-Ascoli Theorem that a subsequence converges uniformly on $K = (X, \mathbf{z})_{\epsilon_0}$, and convergence on compact sets follows from the usual diagonal trick.

Choose a Riemannian metric ρ' on \mathcal{X}_Λ^s , the complement of the nodes and marked points in \mathcal{X}_Λ , whose restriction to the fibers is the vertical hyperbolic metric ρ . The map g_m is eventually defined on $W = (X, \mathbf{z})_{\epsilon_0/2}$, and by the Schwarz lemma, g_m is contracting for ρ' and the hyperbolic metric ρ_W on W , so g_m is an equicontinuous family of maps.

To apply Arzelá-Ascoli, we just need that the g_m are contained in a compact subset of \mathcal{X}_Λ^s . By Mumford's compactness criterion, the ϵ -thick part in the vertical hyperbolic metric $(\mathcal{X}_\Lambda^s)_\epsilon$ is compact, so it suffices to show that g_m maps K into $(X_m, \mathbf{z}_m)_\epsilon$ for some uniform ϵ .

Suppose that $\|g'_m\|$ is small at some point of K . By [McM94, Corollary 2.29], there is a constant C , depending only on the injectivity radius and diameter of K in the

ρ_W -metric (in particular, independent of m), such that for all $x_1, x_2 \in K$,

$$\frac{1}{C} \leq \frac{\|g'_m(x_1)\|}{\|g'_m(x_2)\|} \leq C.$$

It follows that $\|g'_m\|$ is uniformly small on all of K , so g_m sends K into the thin part of X_m , which is a union of annuli by the Margulis lemma. Since g_m is compatible with the markings, g_m is π_1 -injective on K , but K is not an annulus as $X \setminus \mathbf{z}$ is finite type and hyperbolic, so this is a contradiction. Therefore $\|g'_m\|$ is uniformly bounded below (independent of m) on K , and it follows immediately that $g_m(K)$ is contained in $(X_m, \mathbf{z}_m)_\epsilon$ for some uniform ϵ .

Second, we show convergence to the identity. The preceding argument in fact shows that every subsequence of g_m has a further uniformly convergent subsequence. As the g_m are injective, any subsequential limit is a conformal automorphism of $X \setminus \mathbf{z}$. As the g_m are compatible with the markings, this map is homotopic to the identity, so must in fact be the identity, since $X \setminus \mathbf{z}$ is finite type and hyperbolic. Thus any subsequence of g_m has a further subsequence which converges to the identity map, and it follows that g_m converges to the identity. \square

Proposition 3.5. *The vector bundle and conformal topologies on $\Omega\mathcal{D}_\Lambda$ coincide.*

Proof. On the base surface (Σ, \mathbf{s}) , choose g disjoint homologically independent “ α -curves” $\alpha_1, \dots, \alpha_g$, such that each α_i is either part of the multicurve Λ or disjoint from each curve in Λ . As these curves are fixed by the twist group $\text{Tw}_\Lambda^{\text{full}}$, they are dual to a basis of relative forms. Then there are relative one-forms η_1, \dots, η_g on the universal curve \mathcal{X}_Λ over \mathcal{D}_Λ such that in each fiber,

$$\int_{\alpha_i} \eta_j = \delta_{ij}.$$

Now suppose a sequence $(X_m, \mathbf{z}_m, \omega_m, f_m)$ converges to $(X, \mathbf{z}, \omega, f)$ in the conformal topology, and let $g_m: K_m \subset X \rightarrow X_m$ be the sequence of conformal maps exhibiting this convergence. We may write each ω_m and ω as a linear combination of the η_i :

$$(3.2) \quad \omega_m = \sum_{i=1}^g c_{mi} \eta_i|_{X_m} \quad \text{and} \quad \omega = \sum_{i=1}^g c_i \eta_i|_X.$$

Convergence in the vector bundle topology is then equivalent to convergence of each c_{mi} to c_i . Since each c_{mi} can be recovered as the integral of $g_m^* \omega_m$ over $f(\alpha_i)$, this follows from the uniform convergence of $g_m^* \omega_m$ to ω .

Conversely, suppose the sequence converges in the vector bundle topology. Writing the form ω_m in the basis η_i as in (3.2), this means that c_{mi} converge to c_i for each i . Then (X_m, \mathbf{z}_m, f_m) converge to (X, \mathbf{z}, f) as marked pointed surfaces, and by Theorem 3.2 there is a sequence of maps $g_m: K_m \rightarrow X_m$, defined on an exhaustion of X , which exhibit convergence in the conformal topology. By Lemma 3.4, these g_m converge uniformly on compact sets to the identity, and so do their derivatives. It follows that $g_m^* \eta_i \rightarrow \eta_i|_X$ uniformly on compact sets, so $g_m^* \omega_m \rightarrow \omega$ as well. \square

These notions of convergence will appear in several similar contexts. We will often need convergence of one-forms on part of X only. A sequence of stable differentials

(X_m, z_m, ω_m) converges to (X, z, ω) on an irreducible component $Y \subset X$ if there are conformal maps $g_m: K_m \rightarrow X_m$ so that $g_m^* \omega_m$ converge to ω uniformly on compact sets, where K_m is an exhaustion of Y . In another direction, one may allow the ω_m to have poles of prescribed order at the marked points. These notions of convergence may be formalized similarly to the vector bundle topology described above by twisting the relative cotangent bundle, giving a notion of convergence equivalent to conformal convergence.

We will occasionally need a more flexible topology which allows us to pull back one-forms by maps which are only C^1 . We say that a sequence $(X_m, z_m, \omega_m, f_m)$ of marked pointed stable forms converges to (X, z, ω, f) in the C^1 topology if for some exhaustion K_m of X^s there is a sequence of C^1 maps $h_m: K_m \rightarrow X_m$ such that

- $f_m \simeq h_m \circ f$,
- h_m is L_m -quasiconformal with $L_m \rightarrow 1$, and such that
- $h_m^* \omega_m$ converges to ω uniformly on compact sets.

Proposition 3.6. *The C^1 and conformal topologies on $\Omega\mathcal{D}_\Lambda$ coincide.*

Proof. Suppose we have a convergent sequence in the C^1 topology, exhibited by a sequence of C^1 maps $h_m: K_m \rightarrow X_m$, so that the forms $\eta_m = h_m^* \omega_m$ converge to ω uniformly on compact sets. Let μ_m be the Beltrami differential of h_m . Since $L_m \rightarrow 1$, we have $\mu_m \rightarrow 0$ uniformly on compact sets.

Fix a compact subsurface K of X^s . We wish to construct conformal maps $g_m: K \rightarrow X_m$ so that $\nu_m = g_m^* \omega_m$ converges uniformly on K to ω . Since K is arbitrary, conformal convergence follows.

Fix also compact subsurfaces $K'' \supset K' \supset K$, each containing the next in its interior. By Lemma 3.3, there is a sequence of quasiconformal maps $k_m: X \rightarrow X$ whose Beltrami differentials agree with μ_m on K' . By the same Lemma, these k_m converge uniformly to the identity. Then for m sufficiently large $k_m(K'') \supset K'$, so $j_m = k_m^{-1}$ is defined on K' and sends K' into K'' . The composition $g_m = h_m \circ j_m$ is then conformal on K' .

If we show that $\nu_m = g_m^* \omega_m$ converges to ω uniformly on compact sets, conformal convergence follows by applying the diagonal trick to some exhaustion of $X^s \setminus z$. This convergence would follow from L^2 convergence of ν_m to ω on K . We know η_m converges to ω uniformly, hence in L^2 , on K' . It would therefore be enough to show $\|\eta_m - \nu_m\|_2 = \|\eta_m - j_m^* \eta_m\|_2 \rightarrow 0$ on K . This convergence is straightforward to see using that

- η_m converges to ω uniformly on K' ,
- j_m converges uniformly to the identity on K , and
- the derivatives Dj_m are L^2 and converge in L^2 to the identity (see [LV73, Theorems 4.1.2 and 5.5.3]). □

3.4. Strengthening conformal convergence. We have defined conformal convergence of one-forms as uniform convergence of the pullbacks $g_m^* \omega_m$ to ω on compact sets. A natural strengthening is to require the pullbacks to be equal to ω . This is not always possible: if ω_m and ω have different relative periods, then they cannot be identified by any conformal map. It turns out that relative periods are the only obstruction.

Theorem 3.7. *Let X be a closed Riemann surface, containing open subsurfaces U and W with $\overline{U} \subset W \subsetneq X$, and let $Z, P \subset U$ be disjoint discrete sets. Suppose moreover that the boundaries of U and W are smooth and that U is a deformation retract of W . Let ν_m and η_m be two sequences of meromorphic one-forms on W converging uniformly on compact sets to a single non-zero meromorphic form ω . Suppose moreover that*

(i) *all of the forms ν_m , η_m and ω have the same set of poles P and the same set of zeroes Z , and moreover*

(ii) *the orders $\text{ord}_z \nu_m$, $\text{ord}_z \eta_m$ and $\text{ord}_z \omega$ coincide for every m and $z \in U$, and*

(iii) *for each m , the classes $[\nu_m]$ and $[\eta_m]$ in $H^1(U \setminus P, Z; \mathbb{C})$ are equal.*

Then for m sufficiently large, there exists a conformal map $h_m: U \rightarrow W$ fixing each point of $Z \cup P$, and such that $h_m^(\nu_m) = \eta_m$. Moreover one can choose h_m to converge uniformly to the identity as $m \rightarrow \infty$.*

The proof will follow from applying the Implicit Function Theorem to a suitable holomorphic map on an open subset of $\mathcal{H} \times E$, where \mathcal{H} is a space of holomorphic maps $U \rightarrow X$, and E is a Banach space parameterizing one-forms on W . In the next Lemma, we give \mathcal{H} the structure of a Banach manifold modeled on a space of vector fields on U .

Given an open set V in some Banach space, we denote by \mathcal{O}_V the Banach space of bounded holomorphic functions on V equipped with the sup norm. More generally, if E is a normed vector space, $\mathcal{O}_V(E)$ will denote the Banach space of bounded holomorphic functions $V \rightarrow E$, equipped with the sup norm. We use the following notation for derivatives of maps between Banach spaces. We denote by $D_i^n F_z$ the n -th partial derivative with respect to the i -th variable at z and we let $D^n F_z$ denote the derivative of F at z . We use several times that standard results from calculus and complex analysis hold in the context of holomorphic maps on Banach spaces. See [Muj86; Nac69] for details.

Lemma 3.8. *Let $Y \subset \mathbb{C}^3$ be a smooth analytic curve, $U \subset Y$ a relatively compact open set, and $S \subset U$ a finite subset. In the space $\mathcal{O}_U(\mathbb{C}^3)^S$ of bounded holomorphic functions $g: U \rightarrow \mathbb{C}^3$ which fix S pointwise, let \mathcal{H} be the locus of those functions sending U into Y , and let \mathcal{B}_ϵ be the ϵ -ball centered at the identity map id . Then for some $\epsilon > 0$ the intersection $\mathcal{H}_\epsilon = \mathcal{B}_\epsilon \cap \mathcal{H}$ has the structure of a Banach manifold isomorphic to an open ball in $\mathcal{V}(U)^S$, the space of bounded holomorphic tangent vector fields to U which vanish at each point of S .*

Proof. By [BF82, Corollary 1.5], every analytic curve Y in \mathbb{C}^3 is an ideal-theoretic complete intersection, meaning there are holomorphic functions $F_1, F_2: \mathbb{C}^3 \rightarrow \mathbb{C}$ so that Y is defined by the equations $F_1 = F_2 = 0$ and the derivative $DF: \mathbb{C}^3 \rightarrow \mathbb{C}^2$ (where $F = (F_1, F_2)$) is surjective at each point of Y .

Now $\mathcal{O}_U(\mathbb{C}^3)^S \subset \mathcal{O}_U(\mathbb{C}^3)$ is a finite-codimension affine subspace, which may be identified with the Banach space $\mathcal{O}_U(\mathbb{C}^3)_S$ of functions which vanish on S . We define $\Phi: \mathcal{O}_U(\mathbb{C}^3)_S \rightarrow \mathcal{O}_U(\mathbb{C}^2)_S$ by $\Phi(g) = F \circ (\text{id} + g) - \text{id}$, a holomorphic map with derivative $D\Phi_{\text{id}}(g) = DF \cdot g$. The space \mathcal{H} is then the fiber of Φ over 0. If we could show that $D\Phi_{\text{id}}$ is a split surjection by constructing a right-inverse to DF , it would then follow that \mathcal{H}_ϵ is a Banach manifold modeled on the kernel of DF , which is clearly $\mathcal{V}(U)^S$, as claimed.

The derivative DF is explicitly the 2×3 matrix whose ij th entry is the entire function $\frac{\partial F_i}{\partial z_j}$. Let M_i be the 3×2 matrix obtained by replacing the i th row of DF^T by zeros, and let μ_i be the i th minor of DF , so that

$$(3.3) \quad DF \cdot M_i = \mu_i I.$$

Since DF is surjective, the minors μ_i have no common zero on Y . In other words the functions $F_1, F_2, \mu_1, \mu_2, \mu_3$ have no common zero in \mathbb{C}^3 . Let \mathfrak{a} be the ideal generated by these functions in the ring $\mathcal{O}_{\mathbb{C}^3}$ of entire holomorphic functions. By a version of Forster's analytic Nullstellensatz (see [ABF16]), the radical ideal $\sqrt{\mathfrak{a}}$ is dense in $\mathcal{O}_{\mathbb{C}^3}$ (in the topology of normal convergence). There are then entire functions α_k, β_k, h and an integer n so that

$$h^n = \alpha_1 \mu_1 + \alpha_2 \mu_2 + \alpha_3 \mu_3 + \beta_1 F_1 + \beta_2 F_2,$$

with h nonzero on U . Using (3.3), we then have that $h^{-n} \sum_k \alpha_k M_k$ is the desired right-inverse to DF on U . \square

Remark 3.9. When this Lemma is applied below, Y is an algebraic curve. In this case, by the Ferrand-Szpiro Theorem Y is a set-theoretic complete intersection (see [Szp79]) which may not be an ideal-theoretic complete intersection. So even when Y is algebraic, we are forced to use analytic equations defining Y .

Proof of Theorem 3.7. Choose $Q \in X \setminus W$ and fix an embedding of $Y = X \setminus Q$ in \mathbb{C}^3 as an affine space curve. By the previous Lemma, the space of holomorphic maps $U \rightarrow Y$ which fix the subset S and are sufficiently close to the identity may be identified with a δ -ball $\mathcal{V}(U)_\delta^S$.

Let \mathcal{O}_W^0 be the closed subspace of \mathcal{O}_W consisting of those f such that $f\omega$ has trivial periods in $W \setminus P$. We can thus write $\nu_m = (1 + f_m)\omega$ and $\omega_m = (1 + f_m + g_m)\omega$ with $f_m \in \mathcal{O}_W$ and $g_m \in \mathcal{O}_W^0$ both converging to zero as $m \rightarrow \infty$.

Let \mathcal{B}_ϵ denote the ϵ -ball in $\mathcal{V}(U)^S \times \mathcal{O}_W \times \mathcal{O}_W^0$ centered at $(id, 0, 0)$. Consider the map $\Psi: \mathcal{B}_\epsilon \rightarrow \mathcal{O}_U^0$ defined by

$$\Psi(\phi, f, g) = \frac{\phi^*((1 + f + g)\omega) - (1 + f)\omega}{\omega}.$$

Once we have shown that Ψ is well-defined, holomorphic and that the tangent map $D_1\Psi_{(id,0,0)}$ is a split surjection, the Implicit Function Theorem allows to construct a family of maps $\phi(f, g)$, parameterized by f and g in some ϵ -ball, so that $\phi(0, 0)$ is the identity map and $\Psi(\phi(f, g), f, g) = 0$. We can then set $h_m = \phi(f_m, g_m)$.

To check that Ψ is well-defined, that is, the image is contained in \mathcal{O}_U^0 , note first that the numerator and denominator have the same zeros and poles, since they are fixed by ϕ . Moreover, the right hand is bounded on U for ϵ sufficiently small, as it extends to a holomorphic function on a neighborhood of \overline{U} , so it does indeed belong to \mathcal{O}_U^0 . Note that for ϵ sufficiently small, the maps ϕ are sufficiently close to the identity map U to W , so that the pullback in the definition of Ψ is defined.

We claim that Ψ is holomorphic. We write \mathcal{X} for $X \times \mathcal{B}_\epsilon$, and similarly \mathcal{U} and \mathcal{W} for the trivial families of subsets of X over \mathcal{B}_ϵ . We have a universal holomorphic map $\Phi: \mathcal{U} \rightarrow \mathcal{X}$ whose fiber over a point in $\mathcal{V}(U)^S$ is the map which that point represents.

Similarly, there are universal bounded holomorphic functions $F, G: \mathcal{W} \rightarrow \mathbb{C}$ associated to the factors \mathcal{O}_W and \mathcal{O}_W^0 of \mathcal{B}_ϵ . The form ω can be regarded as a relative one-form Ω on \mathcal{W} . The function

$$H = \frac{\Phi^*((1 + F + G)\Omega) - (1 + F)\Omega}{\Omega}$$

is holomorphic on \mathcal{U} and uniformly bounded on $\partial\mathcal{U}$. Here we use the Cauchy Integral formula to bound the ‘‘vertical’’ derivatives of Φ on $\partial\mathcal{U}$. By Lemma 3.10, this induces a holomorphic map into \mathcal{O}_U^0 which is none other than Ψ , and moreover $D\Psi$ can be computed with (3.4).

The derivative operator $D_1\Psi_{(id,0,0)}: \mathcal{V}(U)^S \rightarrow \mathcal{O}_U^0$ is

$$D_1\Psi_{(id,0,0)}(v) = \mathcal{L}_v\omega/\omega,$$

where \mathcal{L}_v is the Lie derivative. We now show that this map is a split surjection by constructing a right inverse Υ to $D_1\Psi$. We define

$$\Upsilon: \mathcal{O}_U^0 \rightarrow \mathcal{V}(U)^S, \quad \Upsilon(f) = \frac{1}{\omega} \int_{z_0} f\omega$$

and argue now that this is well-defined. The integral is over any path starting at z_0 which is either an arbitrary choice of basepoint in Z , or an arbitrary basepoint if Z is empty. The integral depends only on the endpoints of the path, since $f\omega$ has trivial absolute periods, and moreover since it has trivial relative periods, it vanishes at each point in Z to order one larger than ω . It follows that $\Upsilon(f)$ is a holomorphic vector field on U which vanishes at $Z \cup P$.

This defines an operator $V: \mathcal{O}_U^0 \rightarrow \mathcal{V}(U)^S$ which is evidently bounded. It is a left inverse to $D_1\Psi$ by Cartan’s equality $\mathcal{L}_v\omega = d(\omega(v))$ (for closed ω). This completes the verification of the hypothesis of the Implicit Function Theorem. \square

To complete the proof of Theorem 3.7 it remains to verify the following result.

Lemma 3.10. *Let E and F be complex Banach spaces containing open sets U and V respectively, and let $f: U \times V \rightarrow \mathbb{C}$ a bounded holomorphic function. Then the map $F: U \rightarrow \mathcal{O}_V$ defined by $F(z)(w) = f(z, w)$ is a holomorphic function with*

$$(3.4) \quad DF_z(w) = D_1f_{(z,w)}.$$

Proof. Given $(z, w) \in U \times V$, suppose $B_R(z)$ is contained in U . By the Cauchy integral formula, we then have

$$(3.5) \quad \|D_1^p f_{(z,w)}\| \leq p! \frac{M}{R^p},$$

where M is a uniform bound for $|f|$ on $U \times V$. Given $z \in U$, let $D_z: E \rightarrow \mathcal{O}_V$ be the bounded operator $D_z(h)(w) = D_1f_{(z,w)}(h)$. We claim that F is differentiable at z with first derivative D_z . Since D_z is complex linear, it implies that F is holomorphic. This follows immediately from the bound

$$\begin{aligned} |F(z+h) - F(z) - D_z(h)| &= \sup_{w \in V} |f(z+h, w) - f(z, w) - D_1f_{(z,w)}(h)| \\ &\leq \frac{M}{(R-|h|)^2} |h|^2, \end{aligned}$$

where the last inequality follows from the bound (3.5) for the second derivative and Taylor's Theorem. \square

3.5. Compactness for meromorphic differentials. In this subsection, we study convergence for sequences of curves equipped with a meromorphic differential, establishing a compactness result which we later use in Section 14 to obtain compactness of the moduli space of multi-scale differentials.

Given a pointed stable curve (X, \mathbf{z}) we denote the punctured surface $X^s \setminus \mathbf{z}$ by X' , which will always be equipped with its Poincaré hyperbolic metric ρ . Recall that X_ϵ denotes the ϵ -thick part of X (with ϵ smaller than the Margulis constant).

Consider a degenerating sequence of pointed meromorphic differentials $(X_m, \mathbf{z}_m, \omega_m)$ in $\Omega\mathcal{M}_{g,n}(\mu)$ such that the underlying pointed curves converge to some pointed stable curve (X, \mathbf{z}) . It may happen that on some components of the thick part of X'_m the flat metric $|\omega_m|$ is much smaller than on other components. As a result the limit of ω_m may be non-zero on some components of X'_m , and vanish identically on others. In order to get non-zero limits everywhere, we allow ourselves to rescale the differential on different components at different rates. These rescaling parameters arise from a notion of size for the thick parts of the X'_m which we now define.

Given a meromorphic differential $(X, \omega) \in \Omega\mathcal{M}_{g,n}(\mu)$, for any $p \in X'$, let $|\omega|_p$ be its norm at p with respect to the hyperbolic metric. If Y is a component of the thick part X'_ϵ , we define the *size* of Y by

$$(3.6) \quad \lambda(Y) = \sup_{p \in Y} |\omega|_p.$$

A similar notion of size is defined in [Raf07].

Theorem 3.11. *Suppose $(X_m, \mathbf{z}_m, \omega_m)$ is a sequence of meromorphic differentials in $\Omega\mathcal{M}_{g,n}(\mu)$ such that (X_m, \mathbf{z}_m) converges to some (X, \mathbf{z}) as a sequence of pointed stable curves. Let $Y \subset X$ be a component and choose ϵ small enough so that Y_ϵ is connected. For large m , choose $(Y_m)_\epsilon$ to be the sequence of components of $(X_m)_\epsilon$ such that $(Y_m)_\epsilon$ converges to Y_ϵ . Let $\lambda_m = \lambda((Y_m)_\epsilon)$. Then we may pass to a subsequence so that the sequence of rescaled differentials ω_m/λ_m has a non-zero limit on Y .*

Note that if we only wanted a limiting differential defined on Y_ϵ , since $|\omega_m/\lambda_m|$ is bounded on $(Y_m)_\epsilon$, this would be a trivial consequence of Montel's Theorem. To get convergence on all of Y , we establish *a priori* bounds (depending only on ϵ and μ) for the size of any component of the ϵ -thick part of X' , in terms of the norm $|\omega|_p$ at any point of Y .

To this end, we introduce the *Poincaré distortion function* of a pointed meromorphic differential (X, \mathbf{z}, ω) as the function $\mathfrak{T}: X' \rightarrow \mathbb{R}$ defined by

$$\mathfrak{T}(p) = |\beta|_p \quad \text{where} \quad \beta = d \log |\omega/\rho|.$$

This function measures how quickly the flat metric $|\omega|$ varies with respect to the hyperbolic metric ρ . Note that \mathfrak{T} is independent of the scale of ω , so can be regarded as a function on the punctured universal curve over $\mathbb{P}\Omega\mathcal{M}_{g,n}(\mu)$.

Lemma 3.12. *There is a constant C depending only on μ and ϵ so that for any $(X, \mathbf{z}, \omega) \in \Omega\mathcal{M}_{g,n}(\mu)$, the distortion function \mathfrak{T} is bounded by C on the ϵ -thick part of X' .*

Proof. We wish to define a compactification of $\mathbb{P}\Omega\mathcal{M}_{g,n}(\mu)$ so that \mathfrak{T} extends continuously to the universal curve over the compactification. To this end, let $\mathbb{P}\Omega\overline{\mathcal{M}}_{g,n}^{\text{ninc}}(\mu)$ be the normalization of the Incidence Variety Compactification, with the universal curve $\pi: \tilde{\mathcal{X}} \rightarrow \mathbb{P}\Omega\overline{\mathcal{M}}_{g,n}^{\text{ninc}}(\mu)$. The universal curve is equipped with a family of one-forms ω , defined up to scale. Its divisor consists of horizontal components (whose π -image is $\mathbb{P}\Omega\overline{\mathcal{M}}_{g,n}^{\text{ninc}}(\mu)$) along the marked zeros and poles, and also some vertical components (whose π -image is a boundary divisor of $\mathbb{P}\Omega\overline{\mathcal{M}}_{g,n}^{\text{ninc}}(\mu)$).

Suppose $D \subset \tilde{\mathcal{X}}$ is an irreducible vertical component of the zero divisor of ω lying over D' . Since the base is normal, by Proposition 7.13 below near any point $p \in D'$ there is a regular function f defined near p so that ω/f is regular on D near the fiber over p , and moreover such an f is unique up to multiplication by a regular function which does not vanish at p . The family of one-forms

$$\tilde{\beta} = d \log |\omega/\rho f|$$

is then a continuous extension of β which is defined in a neighborhood of the fiber over p in the punctured universal curve \mathcal{X}' and depends neither on the choice of f nor on the scale of ω . Here we are using the fact that the vertical hyperbolic metric is C^1 on $\tilde{\mathcal{X}}$ by [Wol90]. Since we may extend β on a neighborhood of any such vertical zero divisor, this gives a continuous extension of $\tilde{\beta}$ over all of \mathcal{X}' . The function $\tilde{\mathfrak{T}}(p) = |\tilde{\beta}|_p$ is then the desired continuous extension of \mathfrak{T} to $\tilde{\mathcal{X}}$. Since the ϵ -thick part of $\tilde{\mathcal{X}}$ is compact, we see that \mathfrak{T} is bounded on the ϵ -thick part. \square

Corollary 3.13. *There exists a constant L that depends only on μ and ϵ , such that for any pointed meromorphic differential $(X, \mathbf{z}, \omega) \in \Omega\mathcal{M}_{g,n}(\mu)$, for any points p and q in the ϵ -thick part of X , we have*

$$|\omega|_q \leq L|\omega|_p.$$

Proof. By Lemma 3.12, $\log |\omega/\rho|$ is C -Lipschitz on X_ϵ for a uniform constant C . The diameter of X_ϵ is bounded by a uniform constant E , so we can take $L = Ce^E$. \square

Proof of Theorem 3.11. Let $f_m: K_m \rightarrow X_m$ be conformal maps on an exhaustion $\{K_m\}$ of Y' that exhibit the convergence of the (X_m, \mathbf{z}_m) . By Corollary 3.13, and convergence of the Poincaré metrics of X'_m to that of Y' , the differentials $f_m^*(\omega_m/\lambda_m)$ are uniformly bounded on the $(1/k)$ -thick part of Y for every k . By Montel's Theorem, there is a subsequence which converges uniformly on $Y_{1/k}$. The diagonal trick gives a sequence converging uniformly on compact subsets of Y . \square

4. NORMAL FORMS FOR DIFFERENTIALS ON FAMILIES THAT ACQUIRE A NODE

This section provides auxiliary statements for normal forms for differentials on Riemann surfaces, and for families degenerating to a nodal Riemann surface. There are two types of statements. The first is for a fixed Riemann surface, in fact a disk or an

annulus. If moreover the differential is fixed, this goes back to Strebel. For a varying differential we proved such a normal form statement in [BCGGM18, §4.2] and we give a slight generalization below. The second type of normal form theorem is for differentials on a family of surfaces whose topology changes. This statement has two subcases corresponding to the local situation at vertical nodes and horizontal nodes, respectively.

We first recall Strebel's standard local coordinates for meromorphic differentials in the complex plane. A meromorphic differential ω defined on a neighborhood of 0 in \mathbb{C} has two local conformal invariants, its order of vanishing $k = \text{ord}_0 \omega$ and its residue $r = \text{Res}_0 \omega$. Strebel constructed a standard normal form for ω , which depends only on k and r .

Theorem 4.1 (Normal form on a disk, [Str84]). *Consider a meromorphic differential ω on the δ -disk $\Delta_\delta \subset \mathbb{C}$ with $k = \text{ord}_0 \omega$ and $r = \text{Res}_0 \omega$. Then for some $\epsilon > 0$, there exists a conformal map $\phi: (\Delta_\epsilon, 0) \rightarrow (\Delta_\delta, 0)$ such that*

$$(4.1) \quad \phi^* \omega = \begin{cases} z^k dz & \text{if } k \geq 0, \\ r \frac{dz}{z} & \text{if } k = -1, \\ (z^{k+1} + r) \frac{dz}{z} & \text{if } k < -1. \end{cases}$$

The germ of ϕ is unique up to multiplication by a $(k+1)$ -st root of unity when $k \geq 0$, and up to multiplication by a non-zero constant if $k = -1$. For $k < -1$ the map ϕ is uniquely determined by (4.1) and the specification of the image of some point p in $\Delta_\epsilon \setminus \{0\}$. Moreover, if ϕ satisfies (4.1) and $\phi(p) = q$, then there exists a neighborhood U of q such that for every $\tilde{q} \in U$ there exists a map $\tilde{\phi}$ satisfying (4.1) and with $\tilde{\phi}(p) = \tilde{q}$.

This statement also holds for families of differentials ω_t on families of disks, as long as the order $\text{ord}_0 \omega_t = k$ is the same for all t . For families of differentials such that the order $\text{ord}_0 \omega_t$ is not constant, the situation is more complicated. This has essentially been dealt with in [BCGGM18, §4.2], and we implement here two minor generalizations. First, the differential is given a priori only over an annulus, and second, the locus where the differential is assumed to be in standard form is an arbitrary closed subvariety of some open ball $U \subset \mathbb{C}^N$. Let $A_{\delta_1, \delta_2} := \{z : \delta_1 < |z| < \delta_2\} \subset \Delta_{\delta_2}$ be an annulus and let $\zeta_j := e(j/(k+1))$ be a $(k+1)$ -st root of unity (where we denote $e(z) = \exp(2\pi\sqrt{-1}z)$).

Theorem 4.2 (Normal form of a deformation on an annulus). *Let ω_t be a holomorphic family of nowhere vanishing holomorphic differentials on $U \times A_{\delta_1, \delta_2}$ such that its restriction over a closed complex subspace $Y \subset U$ is in standard form (4.1). Choose a basepoint $p \in A_{\delta_1, \delta_2}$ and a holomorphic map $\varsigma: U \rightarrow A_{\delta_1, \delta_2}$ such that $\varsigma(Y) = \zeta_j p$.*

Then there exists a neighborhood $U_0 \subset U$ of Y , together with $\delta_1 < \epsilon_1 < \epsilon_2 < \delta_2$ and a holomorphic map $\phi: U_0 \times A_{\epsilon_1, \epsilon_2} \rightarrow A_{\delta_1, \delta_2}$ such that $\phi_t^(\omega_t)$ is in standard form (4.1), and such that $\phi|_{Y \times A_{\epsilon_1, \epsilon_2}}$ is the inclusion of annuli composed with multiplication by ζ_j , and such that $\phi_t(p) = \varsigma(t)$ for all $t \in U_0$.*

We now pass to families where the topology of the underlying Riemann surfaces changes. Fix some arbitrary complex (base) space B , possibly singular and possibly

non-reduced, with a base point $p \in B$. Any family of Riemann surfaces over B with at worst nodal singularities can be locally embedded in $\tilde{V} = \tilde{V}_\delta = \Delta_\delta^2 \times B$, for some radius δ , where the family is given by $V(f, \delta) = \{uv = f\}$, where f is a holomorphic function on B and where u and v are the two coordinates on the disk. For simplicity we sometimes write $V(f)$ or V for $V(f, \delta)$ when there is no confusion. We denote the “upper” component of the nodal fibers by $X^+ = \{f = 0, v = 0\}$, and the “lower” component by $X^- = \{f = 0, u = 0\}$ respectively. The next statement gives a local normal form for a family of differentials on V near the nodal locus $X^+ \cap X^-$.

Theorem 4.3 (Normal form near vertical nodes). *Let ω be a family of holomorphic differentials on V , not identically zero on every irreducible component of V , which does not vanish at a generic point of X^+ and vanishes to order exactly $k = \kappa - 1 \geq 0$ at the nodal locus $X^+ \cap X^-$. Suppose that f^κ is not identically zero and that there exists an adjusting function h on B such that $\omega = h\eta$ for some family of meromorphic differentials η on V , which is holomorphic away from X^+ and nowhere zero.*

Then for some $\epsilon > 0$, after restricting B to a sufficiently small neighborhood of p , there exists an $r \in \mathcal{O}_{B,p}$ divisible by f^κ , and a change of coordinates $\phi: V(f, \epsilon) \rightarrow V(f, \delta)$, which lifts the identity map of B to itself, such that

$$(4.2) \quad \phi^*\omega = (u^\kappa + r)\frac{du}{u}.$$

Moreover, given a section $\varsigma_0: B \rightarrow V$ and an (initial) section ς that both map to X^- along $f = 0$ and with ς sufficiently close to ς_0 , there exists a unique change of coordinates ϕ as above that further satisfies $\phi \circ \varsigma = \varsigma_0$.

The notion of adjusting function will be formally defined and used later, see Definition 7.11. We split the proof in several steps.

Lemma 4.4. *Under the assumption of Theorem 4.3, the following statements hold:*

- (i) *There exists a holomorphic function g on V such that we can write $\omega = u^\kappa g(u, v)\frac{du}{u}$. Moreover, g can be taken with constant term 1 after rescaling u by a unit.*
- (ii) *Up to multiplying η by a unit, we can assume that $h = f^\kappa$.*
- (iii) *We have $f^\kappa \mid r$.*

Proof. We will see that the second and third statements follow from the proof of the first one. Using the defining equation of V and the fact that ω is holomorphic, we can expand ω in series as

$$(4.3) \quad \omega = \left(\sum_{i \geq 0} c_i u^i + \sum_{i > 0} c_{-i} v^i \right) \frac{du}{u}$$

for some local functions c_i, c_{-i} on B . An arbitrary holomorphic function g on V can be uniquely written, possibly after shrinking the neighborhood to guarantee convergence, as a series $g = \sum_{i \geq 0} a_i u^i + \sum_{i > 0} b_{-i} v^i$, so that our goal is to write ω as

$$(4.4) \quad \omega = u^\kappa g(u, v) \frac{du}{u} = \left(\sum_{i \geq \kappa} a_{i-\kappa} u^i + \sum_{0 \leq i < \kappa} b_{i-\kappa} f^{\kappa-i} u^i + \sum_{i > 0} b_{-\kappa-i} f^\kappa v^i \right) \frac{du}{u}.$$

Since η is holomorphic outside the locus $v = 0$, we can also expand it as

$$(4.5) \quad \eta = \left(\sum_{i \geq 0} e_i v^{-i} + \sum_{i > 0} e_{-i} v^i \right) \frac{du}{u}.$$

The hypothesis on the vanishing order of ω implies that $c_i(p) = 0$ for $0 \leq i < \kappa$, but $c_\kappa(p) \neq 0$. We consider the equation $\omega = h\eta$ near X^- and write $u^i = f^i v^{-i}$ in the defining power series (4.3) of ω . Comparing the $v^{-\kappa}$ terms gives $c_\kappa f^\kappa = h e_\kappa$, hence $h \mid f^\kappa$. On the other hand, the winding number argument as in the proof of [BCGM18, Theorem 1.3] implies that $e_\kappa(p) \neq 0$, so that $f^\kappa \mid h$. (If B is topologically just a point, we can take any lift of the family to a polydisk, run the argument there and the conclusion persists after reduction.) Changing η by a unit in $\mathcal{O}_{B,p}$, we can assume that $h = f^\kappa$ from now on, thus verifying (ii). Coefficient comparison of the terms v^i for $i > 0$ in the equality $\omega = h\eta$ now implies that $c_{-i} = f^\kappa e_{-i}$. It also implies that $f^i c_i = f^\kappa e_i$ for $i \geq 0$. Since f^i is a non-zero function on B for those $0 \leq i < \kappa$ by the non-vanishing hypothesis of ω , this implies the remaining divisibility condition $f^{\kappa-i} \mid c_i$ for $0 \leq i < \kappa$ needed for making (4.4) equal to (4.3).

The form of ω we derived so far implies that the residue of ω is equal to $r = b_{-\kappa} f^\kappa$, which is in particular divisible by f^κ , hence proving (iii).

Finally we can multiply u by a unit and v by the inverse of the unit to make the constant term $a_0 = 1$ in g , thus completing the entire proof. \square

We write $r = r_0 f^\kappa$ from now on.

Lemma 4.5. *We may assume that B is a polydisk.*

Proof. Any (possibly reducible and non-reduced) analytic space can be embedded locally into a polydisk. We thus replace f by any of its lifts to such a polydisk. To put g into the form (4.2) we may assume that B is a polydisk with coordinates \mathbf{b} and zero is the base point. The coordinate change ϕ that puts the node in the required standard form over the polydisk then restricts to a coordinate change over B with the desired properties. \square

Proof of Theorem 4.3. We look for a solution of the form

$$(4.6) \quad \phi_{(X,Y)}(u, v) = (u e^{X(u)+Y(v)}, v e^{-X(u)-Y(v)}),$$

where $X(u) = \sum_{i=1}^{\infty} c_i(b) u^i$ is a holomorphic function of u and b with no constant term, and similarly for $Y(v)$. (In the sequel, for a holomorphic function of u, v , and b , the dependence on b will be left implicit.)

We first remark that the uniqueness of ϕ follows from the observation that any two holomorphic maps with the same pullback of a differential and that agree at a marked point in the regular locus of the differentials agree everywhere. This marked point is given by the section ζ over $f \neq 0$. Consequently, if ϕ_1 and ϕ_2 both satisfy the hypothesis of the theorem, then $\phi_1 \circ \phi_2^{-1}$ is identity on the locus in the family where $f \neq 0$, and hence $\phi_1 = \phi_2$ everywhere.

By Lemma 4.4, we may write the relative form ω as

$$\omega = u^\kappa (1 + r_0 v^\kappa + g_0(u) + h_0(v)) \frac{du}{u},$$

where g_0 is a function of u and b with no constant term, and h_0 is a function of v and b with no constant term or v^κ -term.

We first make a preliminary change of coordinate ψ so that $\psi^*\omega = \omega_0$, where

$$\omega_0 = u^\kappa(1 + r_0v^\kappa + fg(u) + fh(v))\frac{du}{u}.$$

This may be done by taking functions $\alpha(u) = ue^{A(u)}$ and $\beta(v) = ve^{-B(v)}$ such that (possibly after shrinking ϵ) on $\Delta_\epsilon \times B$,

$$\begin{aligned} \alpha^*u^\kappa(1 + g(u))\frac{du}{u} &= u^\kappa\frac{du}{u}, \quad \text{and} \\ \beta^*v^{-\kappa}(1 + r_0f^\kappa + h(v))\frac{dv}{v} &= v^{-\kappa}(1 + r_0f^\kappa)\frac{dv}{v}, \end{aligned}$$

using Strebel's normal form, Theorem 4.1. Then it is straightforward to check that $\phi_{(X,Y)}(u, v) = (ue^{X(u)+Y(v)}, ve^{-X(u)-Y(v)})$ is of the desired form.

We now wish to find functions $X(u)$ and $Y(v)$ so that $\phi_{(X,Y)}^*u^\kappa(1 + r_0v^\kappa)\frac{du}{u} = \omega_0$, and $\phi_{(X,Y)} \circ \varsigma_0 = \varsigma$. Explicitly this means that on $\Delta_\epsilon^2 \times B$, the functions X and Y satisfy the equations,

$$\begin{aligned} (e^{\kappa(X+Y)} + r_0v^\kappa) \left(1 + u\frac{\partial X}{\partial u} - v\frac{\partial Y}{\partial v} \right) - (1 + v^\kappa + fg + fh) + (uv - f)W &= 0, \\ \tau_0e^{-X(f/\tau_0)-Y(\tau_0)} - \tau &= 0, \end{aligned}$$

where $W(u, v)$ is a holomorphic function on Δ_ϵ^2 , and where the sections ς and ς_0 are written as

$$\varsigma = (f/\tau, \tau) \quad \text{and} \quad \varsigma_0 = (f/\tau_0, \tau_0)$$

for some nowhere zero functions τ and τ_0 on B . Our approach to solving these equations will be by perturbing the trivial solution $X = Y = 0$ when $g = h = 0$ and $f = 0$ via the Implicit Function Theorem. To do this, we introduce an auxiliary complex parameter s and the rescaling maps $\rho_s(b) = sb$ on B and $\tilde{\rho}_s(u, v, b) = (u, v, sb)$, so that we have the commutative diagram:

$$\begin{array}{ccc} V(f) & \xrightarrow{\varphi_{(X,Y)}} & V(f) \\ \tilde{\rho}_s \downarrow & & \tilde{\rho}_s \downarrow \\ V(f \circ \rho_s) & \xrightarrow{\varphi_{(X,Y)}} & V(f \circ \rho_s) \end{array}$$

Solving the original equations is then equivalent to solving on the polydisk $\Delta_\epsilon^2 \times B$ the equations

$$\begin{aligned} \Phi_1(W, X, Y, \tau, s) &= (e^{\kappa(X(u)+Y(v))} + (r_0 \circ \rho_s)v^\kappa) \left(1 + u\frac{\partial X}{\partial u} - v\frac{\partial Y}{\partial v} \right) \\ &\quad - (1 + v^\kappa + (fg) \circ \tilde{\rho}_s + (fh) \circ \tilde{\rho}_s) + (uv - f \circ \rho_s)W = 0, \\ \Phi_2(W, X, Y, \tau, s) &= \tau_0e^{-X(f/\tau_0)-Y(\tau_0)} - \tau = 0 \end{aligned}$$

for any nonzero s (note that only the first equation has been rescaled).

We fix some notation for the Banach spaces we need. Let $\mathcal{O}(M)_m$ denote the Banach space of holomorphic functions on M whose first m derivatives are uniformly bounded, equipped with the C^m -norm $\|F\|_m := \sum_{j=0}^m \sup_{z \in M} |F^{(j)}(z)|$. We let $U_B = \Delta_\epsilon \times B$, $V_B = \Delta_\epsilon \times B$, and $\tilde{V} = \Delta_\epsilon^2 \times B$ be polydisks with coordinates (u, b) , (v, b) , and (u, v, b) respectively. An upper index nc will refer to functions without constant term (in u resp. in v) and an upper index nr (“no residue”) will refer to functions without v^κ -term.

In this notation we can view $\Phi = (\Phi_1, \Phi_2)$ as a map

$$\Phi: \mathcal{O}(\tilde{V})_0 \oplus \mathcal{O}(U_B)_1^{\text{nc}} \oplus \mathcal{O}(V_B)_1^{\text{nc}} \oplus \mathcal{O}(B)_0 \oplus \mathbb{C} \rightarrow \mathcal{O}(\tilde{V})_0^{\text{nc,nr}} \oplus \mathcal{O}(B)_0,$$

where the domain summands parameterize W, X, Y, τ , and s respectively. In order to apply the Implicit Function Theorem, we need to show that

$$D_1\Phi: \mathcal{O}(\tilde{V})_0 \oplus \mathcal{O}(U_B)_1^{\text{nc}} \oplus \mathcal{O}(V_B)_1^{\text{nc}} \rightarrow \mathcal{O}(\tilde{V})_0^{\text{nc,nr}} \oplus \mathcal{O}(B)_0$$

is an isomorphism. Here $D_1\Phi$ refers to the derivative at $(0, 0, 0, \tau_0, 0)$ with respect to W, X , and Y . This derivative is given explicitly by

(4.7)

$$D_1\Phi_1(W, X, Y) = W \cdot uv + \left(\kappa X + u(1 + r_0(0)v^\kappa) \frac{\partial X}{\partial u} \right) + \left(\kappa Y - v(1 + r_0(0)v^\kappa) \frac{\partial Y}{\partial v} \right),$$

$$D_1\Phi_2(W, X, Y) = -\tau_0 X(f/\tau_0) - \tau_0 Y(\tau_0).$$

We will show that $D_1\Phi$ is an isomorphism by constructing an explicit inverse,

$$S: \mathcal{O}(U_B)_0^{\text{nc}} \oplus \mathcal{O}(V_B)_0^{\text{nc,nr}} \oplus \mathcal{O}(\tilde{V})_0 \oplus \mathcal{O}(B)_0 \rightarrow \mathcal{O}(\tilde{V})_0 \oplus \mathcal{O}(U_B)_1^{\text{nc}} \oplus \mathcal{O}(V_B)_1^{\text{nc}},$$

identifying $\mathcal{O}(\tilde{V})_0^{\text{nc,nr}}$ with $\mathcal{O}(U_B)_0^{\text{nc}} \oplus \mathcal{O}(V_B)_0^{\text{nc,nr}} \oplus \mathcal{O}(\tilde{V})_0$ by decomposing any holomorphic function in $\mathcal{O}(\tilde{V})_0^{\text{nc,nr}}$ uniquely as $\aleph(u) + \beth(v) + \beth(u, v)uv$.

We define bounded operators $S_X: \mathcal{O}(U_B)_0^{\text{nc}} \rightarrow \mathcal{O}(U_B)_1^{\text{nc}}$ and $S_Y: \mathcal{O}(V_B)_0^{\text{nc,nr}} \rightarrow \mathcal{O}(V_B)_1^{\text{nc}}$ to be the solutions to the differential equations

$$(4.8) \quad \kappa X + u \frac{\partial X}{\partial u} = \aleph,$$

$$(4.9) \quad \kappa Y - v(1 + r_0(0)v^\kappa) \frac{\partial Y}{\partial v} = \beth,$$

obtained from the X - and Y -components of (4.7) by deleting terms containing uv . Solving these equations explicitly using the method of integrating factors (see [Eul32]) yields

$$S_X(\aleph) = \frac{1}{u^\kappa} \int u^{\kappa-1} \aleph \, du,$$

$$S_Y(\beth) = \frac{-v^\kappa}{1 + r_0(0)v^\kappa} \int \frac{\beth}{v^{\kappa+1}} \, dv,$$

where each antiderivative is chosen to have no constant term. The second antiderivative exists because \beth was assumed to have no v^κ term. The differential operator,

$$T(X) = \kappa X + u(1 + r_0(0)v^\kappa) \frac{\partial X}{\partial u},$$

which is the X -component of (4.7), then satisfies

$$TS_X(\aleph) = r_0(0)uv^k \frac{\partial S_X(\aleph)}{\partial u} = E(\aleph),$$

where

$$E(\aleph) = -\kappa r_0(0)v^\kappa S_X(\aleph),$$

which is divisible by uv . Finally, we define S by

$$S(\aleph, \beth, \daleth, \tau) = \left(\daleth - \frac{1}{uv} E(\aleph), S_X(\aleph), S_Y(\beth) + C(\aleph, \beth, \tau)\mu(v) \right),$$

where

$$\mu(v) = \frac{v^\kappa}{1 + r_0(0)v^\kappa}$$

is the kernel of the left-hand side of (4.9), and

$$(4.10) \quad C(\aleph, \beth, \tau) = -\frac{\tau + \tau_0 S_X(\aleph)(f/\tau_0) + \tau_0 S_Y(\beth)(\tau_0)}{\tau_0 \mu(\tau_0)}$$

is chosen so that $D_1\Phi_2 \circ S(\aleph, \beth, \daleth, \tau) = \tau$. Note that since $\tau_0(0) \neq 0$, we may assume that the denominator $\tau_0 \mu(\tau_0)$ of (4.10) is nonzero by possibly shrinking B .

We then know that $D_1\Phi$ is surjective, since it has a right inverse. Injectivity of $D_1\Phi$ is easily checked, using that the solutions to (4.8) and (4.9) are unique up to the kernel of (4.9), which is of the form $C\mu(v)$, and once X and Y are fixed, there is a unique function C such that $D_1\Phi_2 = 0$.

We can now apply the Implicit Function Theorem in a neighborhood of $(s, \tau) = (0, \tau_0)$ to obtain functions $X_{s\tau}, Y_{s\tau}, W_{s\tau}$ with $\Phi(X_{s\tau}, Y_{s\tau}, W_{s\tau}, \tau, s) = 0$. Since ϕ_0 is the identity, thus mapping $V(f, \epsilon)$ into $V(f, \delta)$, this inclusion still holds for (s, τ) sufficiently small. Consequently the map ϕ we constructed maps into $V(f, \epsilon)$ as required. \square

Remark 4.6. The change of coordinates ϕ may also be represented as an explicit formal power series via the following Ansatz, as a function of the form

$$(4.11) \quad \phi(u, v) = (u(1 + Z)e^{X(u)+Y(v)}, v(1 + Z)^{-1}e^{-X(u)-Y(v)}),$$

where $X(u)$ and $Y(v)$ are holomorphic functions as before, with expressions $X(u) = \sum_{i>0} c_i u^i$ and $Y(v) = \sum_{i>0} d_i v^i$ respectively. Here the c_i, d_i , and Z are holomorphic functions on B . left-hand side Equation (4.2) is then equivalent to

$$(4.12) \quad (u^\kappa(1 + Z)^\kappa e^{\kappa(X(u)+Y(v))} + r)(1 + uX'(u) - vY'(v)) = u^\kappa g.$$

A formal solution of this differential equation can be constructed recursively. We begin with solving the equation mod f . The v^i -terms and the u^j -terms for $j \leq \kappa$ are zero mod f on both sides. The u^κ -term implies $Z = 0 \pmod{f}$. The $u^{\kappa+j}$ -term involves a linear equation for $c_j \pmod{f}$ with leading coefficient $\kappa + j$ for $j > 0$. Next we solve mod f^2 , where the $u^{\kappa-1}$ -term gives a linear equation for $b_1 \pmod{f}$. The coefficient $Z \pmod{f^2}$ is linearly determined by the u^κ -term mod f^2 and the $u^{\kappa+j}$ -term mod f^2 linearly determine $c_j \pmod{f}$. In the third round, considering terms mod f^3 , we start with the $u^{\kappa-2}$ -term, which determines $b_2 \pmod{f}$, then consider the $u^{\kappa-1}$ -term to determine $b_1 \pmod{f^2}$. The u^κ -term and higher terms to compute $Z \pmod{f^3}$ and then the $c_j \pmod{f^3}$. This clearly determines an algorithm, starting at the $u^{\kappa-n}$ -term at the

step “mod f^n ”, where the consideration of a term u^{-i} should be read as the v^i -term. The u^0 -term determines there residue, but imposes no condition on b_κ (since it appears with coefficient $\kappa - \kappa$). Making an arbitrary choice for that coefficient, the algorithm can be continued as indicated. This choice can be used to adjust the section ς .

The corresponding statement for horizontal nodes is a direct adaptation of [BHM16, Lemma 7.4]. In fact, the proof given there uses no geometry of the base, and the convergence of the given formal solution follows from straightforward estimates.

Proposition 4.7 (Normal form near horizontal nodes). *Let ω be a family of holomorphic differentials on V , whose restriction to the components X^+ and X^- of the central fiber both have a simple pole at the nodal locus $X^+ \cap X^-$.*

Then for some $\epsilon > 0$ there exists, after restricting B to a sufficiently small neighborhood of p , a change of coordinates $\phi: V(f, \epsilon) \rightarrow V(f, \delta)$ such that it is the identity on B and such that

$$(4.13) \quad \phi^* \omega = r \frac{du}{u}.$$

Moreover, given a section $\varsigma_0: B \rightarrow V$ and an (initial) section ς that both map to X^- along $f = 0$ and with ς sufficiently close to ς_0 , there is a unique change of coordinates ϕ as above that further satisfies $\phi \circ \varsigma = \varsigma_0$.

5. PRONG-MATCHED DIFFERENTIALS

In this section we construct the Teichmüller space $\Omega\mathcal{T}_\Lambda^{pm}(\mu)$ of prong-matched twisted differentials, as a topological space. Subsequently the augmented Teichmüller space of flat surfaces will be constructed as a union of quotients of such spaces $\Omega\mathcal{T}_\Lambda^{pm}(\mu)$. Along the way, we introduce the key notions of degenerations of multicurves, prong-matchings and weldings, as well as several auxiliary Teichmüller spaces.

To avoid overloading this section, we define in this section the points in the moduli spaces by specifying the objects they represent. All these objects have a natural notion of deformation that endows those spaces with a topology that we address along with the modular interpretations in Section 7.

5.1. Ordered and enhanced multicurves and their degenerations. To every multicurve $\Lambda \subset \Sigma$ in a surface Σ we can associate the dual graph $\Gamma(\Lambda)$ whose vertices correspond to connected components of $\Sigma \setminus \Lambda$, and whose edges correspond to curves in Λ . In the setting of multicurves, we will generally imitate the standard notation for level graphs from Section 2.3. We call $\bar{\Lambda} = (\Lambda, \ell)$ an *ordered multicurve* and specify the ordering relation between the components of $\Sigma \setminus \Lambda$ by \preceq . The notions *horizontal* and *vertical* are defined similarly. A multicurve is *purely vertical* (resp. *purely horizontal*) if all of its curves are vertical (resp. horizontal) edges of $\Gamma(\Lambda)$.

An *enhanced multicurve* Λ^+ is a multicurve Λ such that the associated graph $\Gamma(\Lambda)$ has been provided with the extra structure of an enhanced level graph. In order to keep the notation simple, we will mostly denote an enhanced multicurve simply by Λ . Moreover, by an abuse of notation, the enhanced level graph $\Gamma^+(\Lambda)$ associated to Λ will be denoted by Γ^+ , and often simply by Γ .

We adapt many notions for graphs to the context of multicurves. We denote by $L^\bullet(\overline{\Lambda})$ the set of all levels of the level graph associated to the multicurve, and call this set normalized if $L^\bullet(\overline{\Lambda}) = \{0, \dots, -N\}$. We denote by $L(\overline{\Lambda}) = L^\bullet(\overline{\Lambda}) \setminus \{0\}$ the set of all levels except the top one. We denote γ_e the curve of Λ corresponding to an edge e of $\Gamma(\Lambda)$, and for $i \in L^\bullet(\overline{\Lambda})$ call the union of the connected components of $\Sigma \setminus \Lambda$ at level i the *level i subsurface* $\Sigma_{(i)} \subset \Sigma$. Denote $\Sigma_v \subset \Sigma$ the subsurface corresponding to the vertex v . We write Σ_v^c and $\Sigma_{(i)}^c$ for the corresponding compact surfaces where the boundary curves have been collapsed to points.

Definition 5.1. Suppose (Λ_1, ℓ_1) and (Λ_2, ℓ_2) are ordered multicurves on a fixed topological surface. We say that (Λ_1, ℓ_1) is a *degeneration* of (Λ_2, ℓ_2) (or Λ_2 is an *undegeneration* of Λ_1), and we denote it by $\text{dg}: (\Lambda_2, \ell_2) \rightsquigarrow (\Lambda_1, \ell_1)$, if the following conditions hold:

- (i) As a set of isotopy classes of curves $\Lambda_2 \subset \Lambda_1$. Let then $\delta: \Gamma(\Lambda_1) \rightarrow \Gamma(\Lambda_2)$ be the simplicial homomorphism induced by the inclusion $\Sigma \setminus \Lambda_1 \hookrightarrow \Sigma \setminus \Lambda_2$. More concretely, the map δ is defined by collapsing every edge of $\Gamma(\Lambda_1)$ corresponding to a curve in $\Lambda_1 \setminus \Lambda_2$.
- (ii) The map δ is compatible with the orders ℓ_j in the sense that if $v_1 \preceq v_2$ then $\delta(v_1) \preceq \delta(v_2)$. It follows that if $v_1 \succ v_2$ then $\delta(v_1) \succ \delta(v_2)$, so δ induces a surjective, order non-decreasing map, still denoted by δ , on the (normalized) sets of levels $\delta: L^\bullet(\Lambda_1) \rightarrow L^\bullet(\Lambda_2)$.

The notion of degeneration of ordered multicurves extends to a notion of *degeneration of enhanced multicurves* by requiring that moreover the map dg preserves the weights κ_e of the edges e that are not contracted. △

We alert the reader that there are nontrivial degenerations that increase the number of levels without changing the underlying multicurve, see Figure 2.

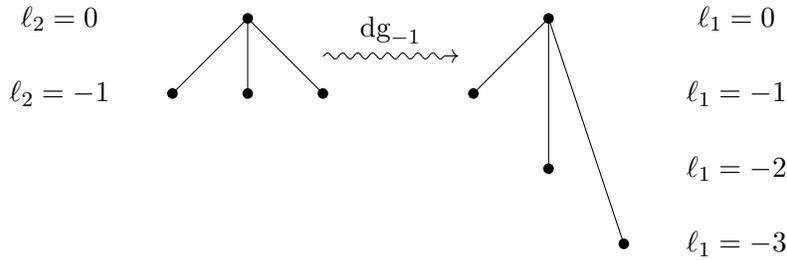


FIGURE 2. A degeneration that does not change the underlying multicurve.

There are two kinds of undegenerations of Λ_1 . First, for any subset $D^h \subseteq \Lambda_1^h$ of the set of horizontal curves we can define a *horizontal undegeneration* of Λ_1 by $\Lambda_2 = \Lambda_1 \setminus D^h$ and $\delta = \text{id}$. Geometrically this undegeneration smoothes out the horizontal nodes corresponding to D^h . Second, suppose that Λ_1 has $N + 1$ levels. Then any surjective, order non-decreasing map $\delta: \underline{N} \rightarrow \underline{M}$ defines a *vertical undegeneration* $\Lambda_2 \rightsquigarrow \Lambda_1$ of Λ_1 as follows. Let $\Lambda_2 \subseteq \Lambda_1$ be the multicurve obtained by deleting all curves that lie in the boundaries of $\Sigma_{(i)}$ and $\Sigma_{(j)}$ for $i \neq j$ such that $\delta(i) = \delta(j)$. The

level structure on Λ_2 is obtained by collapsing to a point every edge joining levels i and j such that $\delta(i) = \delta(j)$. Note that every ordered multicurve as an undegeneration of Λ_1 is obtained uniquely as the composition of a vertical undegeneration and a horizontal undegeneration. Consequently, we refer to an undegeneration by the symbol (δ, D^h) or simply by δ .

There is another way to encode vertical degenerations. Consider a decreasing sequence $J = \{j_0 = 0 > j_{-1} > \cdots > j_{-M} = -N\}$. It induces a map $\delta_J: \underline{N} \rightarrow \underline{M}$ which maps integers (i.e. levels) in each interval $[j_k, j_{k+1})$ to k . We denote the associated degeneration by $\text{dg}_J: \Lambda_J \rightsquigarrow \Lambda$. The two-level degenerations given by $J = \{i\}$, and denoted by dg_i , will be particularly useful (see Section 6.2).

5.2. The Teichmüller space of twisted differentials. For a reference surface (Σ, \mathbf{s}) let $\Omega\mathcal{T}_{(\Sigma, \mathbf{s})}(\mu)$ be the Teichmüller space of (Σ, \mathbf{s}) -marked flat surfaces of type μ and let $\mathbb{P}\Omega\mathcal{T}_{(\Sigma, \mathbf{s})}(\mu) = \Omega\mathcal{T}_{(\Sigma, \mathbf{s})}(\mu)/\mathbb{C}^*$ be its projectivization. We define the subsets $P_{\mathbf{s}}$ and $Z_{\mathbf{s}}$ of \mathbf{s} to be the marked points such that their images under f in X are respectively poles and zeros of ω . The complex structure on $\Omega\mathcal{T}_{(\Sigma, \mathbf{s})}(\mu)$ is induced by the *global period map*

$$\text{Per}: \Omega\mathcal{T}_{(\Sigma, \mathbf{s})}(\mu) \rightarrow H^1(\Sigma \setminus P_{\mathbf{s}}, Z_{\mathbf{s}}; \mathbb{C}),$$

which is locally biholomorphic (see e.g [Vee86], [HM79], [BCGGM19]).

The (classical) mapping class group $\text{Mod}_{(\Sigma, \mathbf{s})}$ of (Σ, \mathbf{s}) acts properly discontinuously on $\mathcal{T}_{(\Sigma, \mathbf{s})}$ and on the twisted Hodge bundle over it, preserving the submanifold $\Omega\mathcal{T}_{(\Sigma, \mathbf{s})}(\mu)$. The spaces $\Omega\mathcal{T}_{(\Sigma, \mathbf{s})}(\mu)$ are highly disconnected and we do not address here the question of classifying the connected components. Moreover, we do not claim that $\mathbb{P}\Omega\mathcal{T}_{(\Sigma, \mathbf{s})}(\mu)$ is simply connected.

We next define similarly strata of flat surfaces over the boundary components of the augmented Teichmüller space. We start with an auxiliary object that will play no major role further on. The upper index “no” indicates that no GRC and no matching residue condition at the horizontal nodes is imposed here. This is mainly introduced to contrast with the space defined later, where the residue conditions *are* required. Moreover, recall that we denote an enhanced multicurve Λ^+ simply by Λ .

Definition 5.2. The *Teichmüller space* $\Omega^{no}\mathcal{T}_{\Lambda}(\mu)$ of flat surfaces of type (μ, Λ) is the space of tuples (X, f, \mathbf{z}, η) where (X, f, \mathbf{z}) is a marked (in the sense of Definition 3.1) pointed stable curve with enhanced pinched multicurve Λ and where $\eta = \{\eta_v\}_{v \in V(\Lambda)}$ is a collection of not identically zero meromorphic one-forms of type μ that have order $\pm\kappa_e - 1$ at e^+ and e^- , respectively, for any edge $e \in \Gamma(\Lambda)$. \triangle

To construct $\Omega^{no}\mathcal{T}_{\Lambda}(\mu)$ as an analytic space, we take a finite unramified cover of the product of the twisted Hodge bundles over the Teichmüller spaces for the components of $\Sigma \setminus \Lambda$ that encodes the identification of the marked points that are paired to form nodes. Then the subset defined by the vanishing conditions of η along \mathbf{z} and at the nodes is the moduli space $\Omega^{no}\mathcal{T}_{\Lambda}(\mu)$.

The group $(\mathbb{C}^*)^{V(\Lambda)}$ acts on $\Omega^{no}\mathcal{T}_{\Lambda}(\mu)$ with quotient $\mathcal{T}_{\Lambda}(\mu)$, since the one-forms η_v are uniquely determined up to scale by the required vanishing conditions encoded in an enhanced multicurve.

Definition 5.3. The *Teichmüller space* $\Omega\mathcal{T}_\Lambda(\mu)$ of twisted differentials of type (μ, Λ) is the subset of $\Omega^{no}\mathcal{T}_\Lambda(\mu)$ consisting of (X, f, \mathbf{z}, η) where η is a twisted differential compatible with $\Gamma(\Lambda)$. \triangle

Said differently, $\Omega\mathcal{T}_\Lambda(\mu)$ is the subset of $\Omega^{no}\mathcal{T}_\Lambda(\mu)$ cut out by the condition of matching residues at the horizontal nodes and the global residue condition. There is an action of $(\mathbb{C}^*)^{L(\Lambda)}$ on $\Omega\mathcal{T}_\Lambda(\mu)$ preserving the fibers of the map to $\mathcal{T}_\Lambda(\mu)$, but the full group $(\mathbb{C}^*)^{V(\Lambda)}$ no longer acts on $\Omega\mathcal{T}_\Lambda(\mu)$ because it does not necessarily preserve the matching residues or the GRC.

We recall that as a consequence of Proposition 3.5, two natural topologies on $\Omega\mathcal{T}_\Lambda(\mu)$ agree. The first topology is the one used above to define the complex structure, as a subset of a finite cover of the product of the twisted Hodge bundles over a product of Teichmüller spaces. The second topology is the product of the conformal topologies on the components of $X \setminus f(\Lambda)$. By definition, this topology is the same as the conformal topology on $\Omega\mathcal{T}_\Lambda(\mu)$, where a sequence $(X_n, f_n, \mathbf{z}_n, \eta_n)$ of marked pointed twisted differentials converges to (X, f, \mathbf{z}, η) if for some exhaustion K_n of X , there is a sequence of conformal maps $g_n: K_n \rightarrow X_n$ such that $f_n \simeq g_n \circ f$ and $g_n^* \eta_n$ converges to η uniformly on compact sets.

5.3. Weldings and markings of welded surfaces. Teichmüller markings of nodal surfaces are by definition insensitive to the pre-composition by Dehn twists around the vanishing cycles. Here we introduce the concept of a welded surface to define a refined concept of markings.

Let X be a stable nodal curve with dual graph Γ , and let $\pi: X^* \rightarrow X$ be the normalization. Given a node q of X , with preimage $\pi^{-1}(q) = \{x, y\}$, a *welding of X at q* is an antilinear isomorphism $\sigma_q: T_x X^* \rightarrow T_y X^*$, modulo scaling by positive real numbers. We alternatively think of the welding as an orientation-reversing metric isomorphism $\sigma_q: S_x X^* \rightarrow S_y X^*$, where $S_p X^* = (T_p X^* \setminus \{0\})/\mathbb{R}_{>0}$ denotes the real tangent circle to X^* at p . As we will explain below, this viewpoint is natural from the perspective of real oriented blowups, which will be discussed in full generality in Section 8. The ordering of the fiber over q is not part of the structure, and we consider $\sigma_q^{-1}: T_y X^* \rightarrow T_x X^*$ to be the same welding as σ_q . The space of all weldings of a given node q is a circle S^1 .

A welding can otherwise be described in terms of a real blowup of X that we now recall, see e.g. [ACG11, Section X.9 and XV.8] and Section 8. Given the unit disk $\Delta \subset \mathbb{C}$, the real oriented blowup $p: \text{Bl}_0 \Delta \rightarrow \Delta$ is the locus

$$\text{Bl}_0 \Delta = \{(z, \tau) \in \Delta \times S^1 : z = |z|\tau\},$$

with the projection p given by $p(z, \tau) = z$. It is a real manifold with a single boundary circle $\{0\} \times S^1$. The projection p collapses the boundary circle to the origin and is otherwise a diffeomorphism.

More generally, if X is a Riemann surface and $D \subset X$ is a finite set of points, performing the above construction at each point $q \in D$ yields the *real oriented blowup* $p: \text{Bl}_D X \rightarrow X$, which is a real manifold such that its boundary maps to D , and consists of a circle over each point $q \in D$. Then p restricts to a diffeomorphism $\text{int}(\text{Bl}_D X) \rightarrow X \setminus D$, and for each $q \in D$ the boundary circle $\partial_q \text{Bl}_D X = p^{-1}(q)$ is

naturally identified with the real tangent circle $S_q X = (T_q X \setminus \{0\})/\mathbb{R}_{>0}$ of X at q . The conformal structure of X gives $\partial_q \text{Bl}_D X$ the structure of a metric circle of arc length 2π .

Given a subset $D \subset N_X$ of the set N_X of nodes of X , the *real oriented blowup* $p: \text{Bl}_D X \rightarrow X$ is the real oriented blowup of the partial normalization X^* of X at D , at the set of preimages of D on this partial normalization. In other words, for each node $q \in D$ the fiber $p^{-1}(q)$ is a pair of metric circles $S_q^+ \cup S_q^- \subset \partial \text{Bl}_D X$. In these terms, a welding of X at D is a choice for each node $q \in D$ of an orientation-reversing isometry $\sigma_q: S_q^+ \rightarrow S_q^-$.

A *(global) welding σ of X* is a choice of a welding at each node of X . If the dual graph is endowed with a level structure $\bar{\Gamma}$, then a *vertical welding σ of X* is a choice of a welding at each vertical node.

Given a vertical welding σ of a nodal curve X , we define the *associated welded surface \bar{X}_σ* to be the surface obtained by gluing the boundary components of $\text{Bl}_{N_X^v} X$ via σ . The associated welded surface has the following extra structures:

- (i) a multicurve Λ^v on \bar{X}_σ , containing for each node $q \in N_X^v$ the simple closed curve that is the image of $S_q^+ \sim S_q^-$, called the (multicurve of) *seams* of \bar{X}_σ ;
- (ii) a conformal structure on $\bar{X}_\sigma \setminus (\Lambda^v \cup N_X^h)$; and
- (iii) a metric on each component of Λ^v , of arc length 2π .

By a slight abuse of terminology, we call $\Lambda = \Lambda^v \cup N_X^h$ the *pinched multicurve* of \bar{X}_σ . Note that the surface \bar{X}_σ can have horizontal nodes and is smooth elsewhere. Such surfaces will be used throughout, motivating the following definition.

Definition 5.4. We call a surface that only has horizontal nodes an *almost smooth surface*. An *almost-diffeomorphism* $f: S \rightarrow S'$ between two almost smooth surfaces is an orientation-preserving diffeomorphism from $S \setminus (N_S^h \cup f^{-1}(N_{S'}^h))$ to its image in S' . \triangle

A *marked welded pointed stable curve* is a tuple $(X, \mathbf{z}, \sigma, f)$, where σ is a welding of the pointed stable curve (X, \mathbf{z}) , and $f: (\Sigma, \mathbf{s}) \rightarrow (\bar{X}_\sigma, \mathbf{z})$ is an oriented marked-point preserving almost-diffeomorphism. Two such surfaces $(X_j, \mathbf{z}_j, \sigma_j, f_j)$ are considered equivalent if there exists a marked-point preserving isomorphism $g: (\bar{X}_1)_{\sigma_1} \rightarrow (\bar{X}_2)_{\sigma_2}$ such that $g \circ f_1$ is isotopic to $f_2 \text{ rel } \mathbf{s}$.

We now introduce a slight generalization of the concept of a Dehn twist. Let X be a nodal curve and \bar{X}_σ be the welded surface associated to a global welding σ . Let q be a node of X and denote by $r_\theta: S_q^- \rightarrow S_q^-$ the rotation of angle θ . A *fractional Dehn twist of \bar{X}_σ of angle θ at q* is a continuous map

$$\delta_q(\theta): \bar{X}_\sigma \rightarrow \bar{X}_{r_\theta \circ \sigma}$$

which is the identity outside of a tubular neighborhood of S_q^- and is a smooth interpolation between the rotation of angle θ on S_q^- and the identity on the other boundary of the annulus.

These notions obviously extend locally to *equisingular families* $(\pi: \mathcal{X} \rightarrow B, \mathbf{z})$ of stable curves, also called *families of constant topological type*. These are families where

all the nodes are *persistent*, i.e. for each node q in each fiber of π there is a section of π passing through q and mapping to the nodal locus of \mathcal{X} . We briefly digress on these notions, aiming for the definition of the topology in Section 5.5 and the comparison in Proposition 5.14. We will return to these notions in detail in Section 8.

For an equisingular family, a *family of weldings* σ over an open set $U \subset B$ is a continuous choice of weldings for each fiber over U . Here we use the fact that π is locally trivial in the C^∞ -category to compare the tangent spaces T_q in nearby fibers. Equivalently, we can perform the real oriented blowup in families over U (see e.g. [ACG11, Section XV.9] and Section 8), and then a family of weldings is a continuous section of the S^1 -bundle at each vertical node. For each family of weldings σ the *family of welded surfaces* $\overline{\mathcal{X}}_\sigma$ is obtained by identifying the family of real oriented blowups of \mathcal{X} along the identifications provided by σ . A *marked family of welded surfaces* is defined by requiring that the fiberwise markings vary continuously.

5.4. Prongs and prong-matchings. Any point p of a meromorphic differential (X, ω) which is not a simple pole has a set of horizontal directions which we call the *prongs* of (X, ω) at p . Intuitively speaking, the prongs at p are the directions in the unit circle $S_p X = T_p X / \mathbb{R}_{>0}$ which are tangent to horizontal geodesics limiting to p under the flat structure induced by ω . In fact, the prongs can be naturally defined as vectors rather than just directions:

Definition 5.5. Suppose the meromorphic differential ω on X has order $k \neq -1$ at some point p . A *complex prong* $v \in T_p X$ of ω at p is one of the $2|k+1|$ vectors $\phi_*(\pm \frac{\partial}{\partial z})$, where ϕ is a choice of the standard coordinates of Theorem 4.1. We say that a prong is *outgoing* if it is of the form $\phi_*(\frac{\partial}{\partial z})$ and otherwise it is *incoming*.

The $2|k+1|$ vectors in $S_p X$ obtained by projectivizing the complex prongs are the *real prongs* of ω at p . \triangle

When p is a (non-simple) pole, while there are infinitely many choices of standard coordinates, there are still only $2|k+1|$ prongs, as the prongs are determined only by the first derivative of ϕ at p . Explicitly in local coordinates, if $\omega = z^k f(z) dz$ with $f(0) \neq 0$, then the prongs at 0 are the vectors $\pm \zeta \frac{\partial}{\partial z}$, where $\zeta^{k+1} = f(0)$.

Since complex and real prongs are in natural bijection, we will simply refer to them as prongs when we do not need to make the distinction.

We denote the set of incoming prongs at z by P_z^{in} and the set of outgoing prongs by P_z^{out} . Each has cardinality $\kappa_z = |1+k| = |1+\text{ord}_z \omega|$. Each set of prongs is equipped with the counterclockwise cyclic ordering when embedded in the complex plane with coordinate z .

Now suppose q is a vertical node of a twisted differential (X, η) . The matching orders condition (1) of a twisted differential equivalently says that the zero at q^+ and the pole at q^- have the same number of prongs (equal to κ_q).

Definition 5.6. A *local prong-matching* of (X, η) at q is a cyclic-order-reversing bijection $\sigma_q: P_{q^-}^{\text{in}} \rightarrow P_{q^+}^{\text{out}}$.

A (*global*) *prong-matching* σ for a twisted differential is a choice of a local prong-matching σ_q at each (vertical) node q of X . \triangle

Note that prong-matchings at horizontal nodes are void or not defined. The following equivalent definition of a local prong-matching will be useful for studying families in Section 7:

Definition 5.7. A *local prong-matching* of (X, η) at a node q is an element σ_q of $T_{q^+}^* X \otimes T_{q^-}^* X$ such that for any pair (v_+, v_-) of an outgoing and an incoming prong, the equality $\sigma_q(v_+ \otimes v_-)^{\kappa_q} = 1$ holds. \triangle

To see the equivalence of these definitions, any such σ_q corresponds to an order-preserving bijection $P_{q^-}^{\text{in}} \rightarrow P_{q^+}^{\text{out}}$ by assigning to an incoming prong v_- the unique outgoing prong v_+ such that $\sigma_q(v_- \otimes v_+) = 1$.

A prong-matching σ_q at the node q determines a welding of X at q by identifying a prong $v \in S_{q^-} X$ with the prong $\sigma_q(v) \in S_{q^+} X$, and extending this to an orientation-reversing isometry of these tangent circles. We denote by \overline{X}_σ or simply by \overline{X} the welded surface constructed using the welding defined by the prong-matching σ . Note that \overline{X}_σ is an almost smooth surface.

Definition 5.8. A *prong-matched twisted differential of type μ compatible with Λ* , or just *prong-matched twisted differential* for short, is the datum $(X, \mathbf{z}, \eta, \sigma)$ consisting of a twisted differential (X, \mathbf{z}, η) of type (μ, Λ) and a global prong-matching σ . \triangle

An isomorphism between prong-matched twisted differentials is required to commute with all the local prong-matchings. A *marked prong-matched twisted differential* contains additionally a marking of the welded surface $f: \Sigma \rightarrow \overline{X}_\sigma$.

For an equisingular family $(\pi: \mathcal{X} \rightarrow B, \mathbf{z}, \eta)$ of twisted differentials we define a *family of prong-matchings* to be a family of global weldings (in the sense of the previous subsection) that is a prong-matching in each fiber of π .

The *prong rotation group* associated with an enhanced level graph Γ is the finite group

$$(5.1) \quad P_\Gamma = \prod_{e \in \Lambda^v} \mathbb{Z} / \kappa_e \mathbb{Z}.$$

The number of prong-matchings for a given twisted differential is then equal to $|P_\Gamma|$. Moreover, for any given twisted differential (X, η) the prong rotation group acts on the set of prong-matchings as follows. An element $(j_e)_{e \in \Lambda^v} \in P_\Gamma$ acts by composing the local prong-matching at the node $q = q_e$ with the bijection $P_{q^-}^{\text{in}} \rightarrow P_{q^-}^{\text{in}}$ defined by turning counterclockwise j_e times. Here, and for other similar notions depending on graphs, we also write P_Λ or P_Γ as shorthand for $P_{\Gamma(\Lambda)}$.

5.5. The Teichmüller space of prong-matched twisted differentials. We now construct the *Teichmüller space $\Omega\mathcal{T}_\Lambda^{\text{pm}}(\mu)$ of marked prong-matched twisted differentials of type (μ, Λ)* as an analytic space. We first define it as a set.

Definition 5.9. $\Omega\mathcal{T}_\Lambda^{\text{pm}}(\mu)$ is the set of isomorphism classes of marked prong-matched twisted differentials. \triangle

Given any contractible open set $U \subset \Omega\mathcal{T}_\Lambda(\mu)$, we choose a prong-matching σ and a marking $f: \Sigma \rightarrow \overline{X}_\sigma$ for some basepoint $(X, \eta) \in U$ and then transport this to

every point in U by parallel transport. This defines a lift $U \rightarrow \Omega\mathcal{T}_\Lambda^{pm}(\mu)$. We give $\Omega\mathcal{T}_\Lambda^{pm}(\mu)$ the structure of a complex manifold such that these lifts are holomorphic local homeomorphisms. The forgetful map $\Omega\mathcal{T}_\Lambda^{pm}(\mu) \rightarrow \Omega\mathcal{T}_\Lambda(\mu)$ is then a holomorphic covering map of infinite degree.

It turns out that the space $\Omega\mathcal{T}_\Lambda^{pm}(\mu)$ is a fine moduli space for families of marked prong-matched twisted differentials. The proof of this result will be postponed to Section 10.2.

5.6. Turning numbers on welded flat surfaces. Let (X, \mathbf{z}, ω) be a pointed flat surface and let $\gamma: [a, b] \rightarrow X \setminus \mathbf{z}$ be an immersed arc. The *Gauss map* $G: [a, b] \rightarrow S^1$ of γ , defined by

$$G(t) = \frac{\omega(\gamma'(t))}{|\omega(\gamma'(t))|}$$

has a unique lift $\tilde{G}: [a, b] \rightarrow \mathbb{R}$ with $\tilde{G}(a) = 0$. The *turning number* of γ is defined to be $\tau(\gamma) = \tilde{G}(b) \in \mathbb{R}$. Turning numbers are invariant under isotopies preserving the endpoints of γ as well as the tangent vectors at these endpoints.

This notion extends naturally to *good arcs* γ on a globally welded surface \overline{X}_σ defined by a global prong-matching σ , where an arc γ is said to be good if it is transverse to the seams of the welding. In fact, under these assumptions the Gauss map is defined on the γ -preimage of the complement of the seams and admits a unique continuous extension to $[a, b]$ thanks to the welding. A good arc γ on an almost smooth surface X is defined to be an arc that admits a lift to a good arc on a global welding of X . Obviously, the vertical welding is determined by γ at all the nodes that γ passes through. (At horizontal nodes the turning number does not depend on the weldings.)

Definition 5.10. Let $h: S \rightarrow S'$ be an almost-diffeomorphism. We say that h is *turning-number-preserving*, if it preserves the turning number of all good arcs. \triangle

Let (X, η) be a differential where $X = \Delta^+ \cup \Delta^-$ consists of two closed disks joined at the origin and where η is regular except for having a zero of order $\kappa - 1 \geq 0$ at $0 \in \Delta^+$ and a pole of order $\kappa + 1 \geq 2$ at $0 \in \Delta^-$. Suppose we are given two weldings σ and σ' of X with corresponding welded surfaces $\overline{X} = \overline{X}_\sigma$ and $\overline{X}' = \overline{X}_{\sigma'}$.

Proposition 5.11. *Let $h_m: \overline{X} \rightarrow \overline{X}'$ be a sequence of diffeomorphisms that are conformal and converging to the identity on an exhaustion $\{K_m\}$ of X .*

If the maps h_m are turning-number-preserving, then the identity map on X can be lifted to a diffeomorphism $h: \overline{X} \rightarrow \overline{X}'$, i.e. the prong-matchings σ and σ' are the same.

Proof. We may assume that $\eta|_{\Delta^+} = z^\kappa dz/z$ and $\eta|_{\Delta^-} = (z^{-\kappa} + a_{-1})dz/z$ with $a_{-1} \in \mathbb{C}$ (see Theorem 4.1). Moreover we may assume that $K_m = K_m^+ \cup K_m^-$ consists of the complements of the disks $\{z \in \Delta^\pm : |z| < 1/m\}$. It suffices to show that under the assumptions of the Proposition we can isotope each h_m to the identity on $\overline{X} \setminus K_m$. Let $\pi: S \rightarrow \overline{X} \setminus K_m$ be the universal cover of the annulus $\overline{X} \setminus K_m$. We consider the form $\pi^*\eta$ on the complement in S of the preimage S_0 of the seam. We write $S \setminus S_0 = S^+ \cup S^-$ with $S^\pm = \pi^{-1}(\Delta^\pm \setminus 0)$. We can foliate S by arcs γ_{s_0} starting at points $s_0 \in S_0$ defined as follows. On S^+ , the arc γ_{s_0} is the preimage starting at s_0 of the line in Δ^+ approaching the origin in the direction corresponding to s_0 . On S^- , we choose the

preimage of the line in Δ^- approaching the origin in the direction corresponding to s_0 and orthogonal to the circle of radius $1/m$. This line is well-defined, because under the flat metric the image of the circle is convex and the line goes outside of the image circle. Note that the turning number of the preimage of every good arc γ starting in the lower boundary at a point in γ_{s_1} and ending in the upper one at a point in γ_{s_2} (and orthogonal at both boundaries) is the difference between s_1 and s_2 . Let $\pi': S' \rightarrow \overline{X'}$ be the universal cover. We define $\tilde{h}_m: S \rightarrow S'$ to be the lift of $h_m \circ \pi$ via π' , such that $\tilde{h}_m(s_0) = s_0$ for some point $s_0 \in S_0$. Since the h_m converge to the identity on the exhaustion K_m , the hypothesis on turning numbers implies that γ_s and $\tilde{h}_m(\gamma_s)$ have nearly the same start point and endpoint. Consequently, we may use the straight line homotopy to homotope \tilde{h}_m to the identity map. By equivariance, this homotopy descends to a homotopy of h_m and to a diffeomorphism $h_m^+: \overline{X} \rightarrow \overline{X'}$ such that the restrictions h_m^+ to the seam of \overline{X} agree for all m large enough. The pointwise limit of the h_m^+ is the diffeomorphism h we are looking for. \square

Remark 5.12. Note that it is straightforward to generalize Proposition 5.11 to the case where the diffeomorphisms are *nearly turning-number-preserving* in the sense that for every good arc γ the limit of $\tau(h_m(\gamma)) - \tau(\gamma)$ as $m \rightarrow \infty$ is zero.

Using these notions we can now give an alternative definition for the topology on $\Omega\mathcal{T}_\Lambda^{pm}(\mu)$, closer to what we will use later for the augmented Teichmüller space of flat surfaces.

Definition 5.13. We say that a sequence $X_m = (X_m, \mathbf{z}_m, \eta_m, \preceq, \boldsymbol{\sigma}_m, f_m)$ of marked prong-matched twisted differentials in $\Omega\mathcal{T}_\Lambda^{pm}(\mu)$ converges in the *conformal topology* to $X = (X, \mathbf{z}, \eta, \preceq, \boldsymbol{\sigma}, f)$ if and only if for any sufficiently large m there exists an almost-diffeomorphism $g_m: \overline{X} \rightarrow \overline{X}_m$ and a sequence of positive numbers ϵ_m converging to 0, such that the following conditions hold:

- (i) The function g_m is compatible with the markings in the sense that f_m is isotopic to $g_m \circ f$ rel marked points.
- (ii) The function g_m is conformal on the ϵ_m -thick part $(X, \mathbf{z})_{\epsilon_m}$.
- (iii) The differentials $g_m^* \eta_m$ converge to η uniformly on compact sets of the ϵ_m -thick part of X .
- (iv) The functions g_m are nearly turning-number-preserving, i.e. for every good arc γ in X , the turning numbers $\tau(g_m \circ \gamma)$ converge to $\tau(\gamma)$. \triangle

Note that in order to verify the last item, it suffices to choose a collection of arcs that contains, for every seam, an arc that crosses only this seam and no others, and does so exactly once. Indeed, if condition (iv) holds for these arcs, then together with (ii) this forces the convergence of the other turning numbers as well. Note moreover that the last item requires a lift of g_m to the surface where all the horizontal nodes γ passes through have been welded. However, at horizontal nodes, the turning number $\tau(g_m \circ \gamma)$ is independent of the choice of this lift.

Proposition 5.14. *The conformal topology on $\Omega\mathcal{T}_\Lambda^{pm}(\mu)$ and the topology defined in Section 5.5 as the covering space of $\Omega\mathcal{T}_\Lambda(\mu)$ agree.*

Proof. Items (ii) and (iii) in the above definition together with the composition of f_m with the map forgetting the welding show that the projection of a converging sequence in the conformal topology to $\Omega\mathcal{T}_\Lambda(\mu)$ converges by Proposition 3.5.

Open sets according to the definition of Section 5.5 are defined using a C^∞ -trivialization of the family. Such a trivialization provides maps g_m that are nearly turning-number-preserving, where “nearly” is due to the fact that g_m only nearly identifies the one-forms. Conversely, if the turning numbers converge, the prong-matchings in the trivialization have to converge, since any two possible prong-matchings differ in turning number by an integer. \square

Using a continuous trivialization is obviously impossible for a degenerating family of stable curves with varying topological types. For this reason, the conformal topology on the augmented Teichmüller space of flat surfaces defined in Section 9 will be more involved.

6. TWIST GROUPS AND LEVEL ROTATION TORI

The goal of this section is to define the twist group Tw_Λ and the level rotation torus T_Λ associated with an enhanced multicurve Λ . The twist group is generated by appropriate combinations of Dehn twists, such that the quotient of some augmented Teichmüller space by the twist group is the flat geometric counterpart of the classical Dehn space introduced in Section 3.2. This augmented Teichmüller space of flat surfaces, to be defined in Section 9, requires a level-wise projectivization of the space of prong-matched differentials, and we define here the appropriate actions of multiplicative groups, the level rotation tori, for this projectivization. We will provide various viewpoints on the level rotation torus that will be used in the definition of families of model differentials and multi-scale differentials in the later sections.

6.1. The action of $\mathbb{C}^{L^\bullet(\Lambda)}$ on the space of prong-matched differentials. Recall from Section 5.2 that $(\mathbb{C}^*)^{L^\bullet(\Lambda)}$ acts on $\Omega\mathcal{T}_\Lambda(\mu)$, by simultaneously scaling forms at the same level, and preserving the fibers of the projection to $\mathcal{T}_\Lambda(\mu)$. However, the group $(\mathbb{C}^*)^{L^\bullet(\Lambda)}$ does not act naturally on $\Omega\mathcal{T}_\Lambda^{pm}(\mu)$, since a loop around the origin in \mathbb{C}^* in general returns to the same differential with a different prong-matching and a different marking. To get a continuous action on $\Omega\mathcal{T}_\Lambda^{pm}(\mu)$, we have to pass to the universal cover $\mathbb{C}^{L^\bullet(\Lambda)}$ of $(\mathbb{C}^*)^{L^\bullet(\Lambda)}$, which acts continuously on $\Omega\mathcal{T}_\Lambda^{pm}(\mu)$ by *level rotations*, as we now describe.

- (1) On the level of forms, the tuple $\mathbf{d} = (d_i)_{i \in L^\bullet(\Lambda)} \in \mathbb{C}^{L^\bullet(\Lambda)}$ acts through the quotient $(\mathbb{C}^*)^{L^\bullet(\Lambda)}$ by multiplying the form at level i by $e(d_i)$ (recall that we denote $e(z) = \exp(2\pi\sqrt{-1}z)$).
- (2) On a prong-matching σ we act by shifting the angles by the real parts of the d_i , i.e. for a twisted differential (X, η) , a prong-matching σ , and for $\mathbf{d} = (d_i)_{i \in L^\bullet(\Lambda)}$ we define

$$(6.1) \quad \mathbf{d} \cdot (X, \eta, \sigma) = (X, \{e(d_i)\eta|_{X_{(i)}}\}_{i \in L^\bullet(\Lambda)}, \{\mathbf{d} \cdot \sigma\}),$$

where for each vertical node q we let

$$(6.2) \quad \mathbf{d} \cdot \sigma_q: P_{q^-}^{\text{in}} \rightarrow P_{q^+}^{\text{out}}$$

be the map σ_q pre-composed and post-composed with rotations by the angle $2\pi \operatorname{Re}(d_{\ell(q^-)}/\kappa_q)$ and $-2\pi \operatorname{Re}(d_{\ell(q^+)}/\kappa_q)$, so that $\mathbf{d} \cdot \sigma_q$ remains to be a prong-matching.

- (3) On the marking f , the element $\mathbf{d} \in \mathbb{C}^{L^\bullet(\Lambda)}$ acts by a fractional Dehn twist of angle $2\pi (\operatorname{Re}(d_{\ell(q^+)}) - \operatorname{Re}(d_{\ell(q^-)}))$ around the seam of the welded surface corresponding to each vertical node q .

The surfaces in the $\mathbb{C}^{L^\bullet(\Lambda)}$ -orbit are related by a useful diffeomorphism, implicitly appearing in (3) above: we define for each $\mathbf{d} \in \mathbb{C}^{L(\Lambda)}$ a *fractional Dehn twist*

$$(6.3) \quad F_{\mathbf{d}}: \overline{X}_\sigma \rightarrow \overline{X}_{\mathbf{d} \cdot \sigma}$$

to be the composition of fractional Dehn twists (as introduced in Section 5.3) of angle $2\pi (\operatorname{Re}(d_{\ell(q^+)}) - \operatorname{Re}(d_{\ell(q^-)}))$ at all vertical nodes q .

6.2. Twist groups. The restriction of the action of $\mathbb{C}^{L^\bullet(\Lambda)}$ on $\Omega\mathcal{T}_\Lambda^{pm}(\mu)$ to the subgroup $\mathbb{Z}^{L^\bullet(\Lambda)} \subset \mathbb{C}^{L^\bullet(\Lambda)}$ acts by modifying the prong-matchings and markings, while preserving the underlying differentials. The group $\mathbb{Z}^{L^\bullet(\Lambda)}$ is called the *level rotation group*. Considering only the action on prongs defines a homomorphism from the level rotation group $\mathbb{Z}^{L^\bullet(\Lambda)}$ to the prong rotation group P_Λ defined in (5.1):

$$(6.4) \quad \phi_\Lambda^\bullet: \mathbb{Z}^{L^\bullet(\Lambda)} \rightarrow P_\Lambda, \quad \mathbf{n} \mapsto (n_{\ell(e^+)} - n_{\ell(e^-)} \pmod{\kappa_e})_{e \in \Lambda^v}.$$

This map allows us to introduce an important equivalence relation.

Definition 6.1. Two prong-matchings are called *equivalent* if there exists an element of the level rotation group that transforms one into the other. \triangle

The homomorphism ϕ_Λ^\bullet fits into the following commutative diagram of group homomorphisms

$$\begin{array}{ccccc} \operatorname{Tw}_\Lambda^v & \xleftarrow{\tau_\Lambda^\bullet} & \ker(\phi_\Lambda^\bullet) & \hookrightarrow & \mathbb{Z}^{L^\bullet(\Lambda)} \\ \downarrow & & \tilde{\phi}_\Lambda^\bullet \downarrow & & \tilde{\phi}_\Lambda^\bullet \downarrow \\ \operatorname{Mod}_{(\Sigma, \mathbf{s})} & \xleftarrow{\psi} & \ker(\overline{\psi}) & \hookrightarrow & \mathbb{Z}^{\Lambda^v} \xrightarrow{\overline{\psi}} P_\Lambda \end{array}$$

that we now describe. The group \mathbb{Z}^{Λ^v} acts on the space $\Omega\mathcal{T}_\Lambda^{pm}(\mu)$ via *edge rotations* by the fractional Dehn twists, i.e. the tuple $(n_e)_{e \in \Lambda^v}$ twists the prong-matching of the edge e by $\kappa_e \{n_e/\kappa_e\}$ (i.e. the remainder of $n_e \pmod{\kappa_e}$) and pre-composes the marking by $\lfloor n_e/\kappa_e \rfloor$ left Dehn twists around the curve corresponding to e . Taking the quotient by the subgroup of full Dehn twists at such an edge e gives a map $\mathbb{Z} \rightarrow \mathbb{Z}/\kappa_e\mathbb{Z}$, and doing this for all vertical edges induces a map $\overline{\psi}: \mathbb{Z}^{\Lambda^v} \rightarrow P_\Lambda$ onto the prong rotation group. The kernel $\ker(\overline{\psi})$ is thus generated by (full) Dehn twists around Λ^v and is thus a subgroup of the mapping class group. We denote by $\psi: \ker(\overline{\psi}) \hookrightarrow \operatorname{Mod}_{(\Sigma, \mathbf{s})}$ this inclusion.

There is a natural homomorphism $\tilde{\phi}_\Lambda^\bullet: \mathbb{Z}^{L^\bullet(\Lambda)} \rightarrow \mathbb{Z}^{\Lambda^v}$ defined by

$$(6.5) \quad \tilde{\phi}_\Lambda^\bullet(\mathbf{n}) = (n_{\ell(e^+)} - n_{\ell(e^-)})_{e \in \Lambda^v}.$$

The composition of $\tilde{\phi}_\Lambda^\bullet$ followed by $\bar{\psi}$ recovers the homomorphism $\phi_\Lambda^\bullet: \mathbb{Z}^{L^\bullet(\Lambda)} \rightarrow P_\Lambda$ defined in (6.4). The kernel $\ker(\phi_\Lambda^\bullet)$ is in other words the subgroup of $\mathbb{Z}^{L^\bullet(\Lambda)}$ whose action on $\Omega\mathcal{T}_\Lambda^{pm}(\mu)$ fixes the underlying prong-matched twisted differentials, only changing the markings. This defines a homomorphism $\tau_\Lambda^\bullet = \psi \circ \tilde{\phi}_\Lambda^\bullet: \ker(\phi_\Lambda^\bullet) \rightarrow \text{Mod}_{(\Sigma, \mathbf{s})}$ sending \mathbf{n} to the product of Dehn twists,

$$(6.6) \quad \tau_\Lambda^\bullet(\mathbf{n}) = \prod_{e \in \Lambda^v} \text{tw}_{\gamma_e}^{m_e}, \quad \text{with } m_e \text{ defined by } \tilde{\phi}_\Lambda^\bullet(\mathbf{n}) = (m_e \kappa_e)_{e \in \Lambda^v},$$

where γ_e is the seam corresponding to e , and tw_{γ_e} is the Dehn twist around it. The image of τ_Λ^\bullet is called the *vertical Λ -twist group* $\text{Tw}_\Lambda^v \subset \text{Mod}_{(\Sigma, \mathbf{s})}$. Tracking the above definitions, we conclude the following.

Proposition 6.2. *The vertical Λ -twist group Tw_Λ^v is a free abelian group of rank N . Moreover, $\ker(\tau_\Lambda^\bullet) \subset \ker(\phi_\Lambda^\bullet)$ is isomorphic to \mathbb{Z} , generated by $\mathbf{1} = (1, \dots, 1)$, and $\ker(\phi_\Lambda^\bullet) = \text{Tw}_\Lambda^v \oplus \ker(\tau_\Lambda^\bullet) \cong \text{Tw}_\Lambda^v \oplus \mathbb{Z}$.*

Proof. Since P_Λ is a torsion group, the rank of $\ker(\phi_\Lambda^\bullet)$ is equal to the rank of the level rotation group $\mathbb{Z}^{L^\bullet(\Lambda)}$ which is $N + 1$. A tuple \mathbf{n} lies in $\ker(\tau_\Lambda^\bullet)$ if and only if $n_{\ell(e^+)} = n_{\ell(e^-)}$ for every vertical edge e . Since the dual graph is connected, \mathbf{n} is a multiple of $\mathbf{1}$, so $\mathbf{1}$ generates $\ker(\tau_\Lambda^\bullet)$. The vector $\mathbf{1}$ is primitive in the level rotation group, so it is also primitive in $\ker(\phi_\Lambda^\bullet)$. Hence there is a splitting of the short exact sequence

$$0 \rightarrow \mathbb{Z} \rightarrow \ker(\phi_\Lambda^\bullet) \rightarrow \text{Tw}_\Lambda^v \rightarrow 0. \quad \square$$

We define the *horizontal Λ -twist group* to be the subgroup $\text{Tw}_\Lambda^h \subset \text{Mod}_{(\Sigma, \mathbf{s})}$ generated by Dehn twists around the horizontal curves Λ^h . We then define the *Λ -twist group* to be the direct sum

$$\text{Tw}_\Lambda = \text{Tw}_\Lambda^v \oplus \text{Tw}_\Lambda^h.$$

Let $(\mathbf{f}_i)_{i=0, \dots, -N}$ be the *standard basis* of $\mathbb{C}^{L^\bullet(\Lambda)} \cong \mathbb{C}^{N+1}$, where $\mathbf{f}_i = (0^{-i}, 1, 0^{N+i})$. In order to describe the above groups in simpler terms, we will also use the *lower-triangular basis* $(\mathbf{b}_i)_{i=0, \dots, -N}$ defined as

$$\mathbf{b}_i = \sum_{k=-N}^i \mathbf{f}_k = (0^{-i}, 1^{N+i+1}).$$

Then for $v_i \in \mathbb{C}$, the element $v_i \mathbf{b}_i$ acts by simultaneously multiplying the forms on all levels $j \leq i$ by $e(v_i)$. In particular, $v_0 \mathbf{b}_0$ simultaneously scales the form on every irreducible component of X by $e(v_0)$.

Recall from Section 5.1 that for every level $i \in L(\Lambda)$ there is a two-level undegeneration $\text{dg}_i: \Lambda_i \rightsquigarrow \Lambda$ that contracts the (vertical) edges of $\Gamma(\Lambda)$ strictly above level i and the edges below or at level i . We denote by $\text{Tw}_{\Lambda, i}^{sv} = (\text{dg}_i)_*(\text{Tw}_{\Lambda_i}^v) \subset \text{Tw}_\Lambda^v$ the corresponding subgroup of the vertical Λ -twist group. Note that $\text{Tw}_{\Lambda_i}^v$ is the cyclic group generated by the element $(0, a_i)$ with $a_i = \text{lcm}_e \kappa_e$ for all edges e connecting the

graph $\Gamma_{>i}$ to $\Gamma_{\leq i}$. Moreover, $(\text{dg}_i)_*(0, a_i) = (0^{-i}, a_i^{N+i+1}) \in \mathbb{Z}^{L^\bullet(\Lambda)}$. It follows that $\text{Tw}_{\Lambda, i}^{sv}$ is the cyclic subgroup of Tw_{Λ}^v generated by

$$(6.7) \quad \tau_{\Lambda}^{\bullet}(a_i \mathbf{b}_i) = \prod_e \text{tw}_{\gamma_e}^{m_{e,i}}, \quad \text{with } a_i = \text{lcm}_e \kappa_e \quad \text{and} \quad m_{e,i} = a_i / \kappa_e,$$

where the product and the lcm are taken for the set of vertical edges e connecting $\Gamma_{>i}$ to $\Gamma_{\leq i}$. Note that $a_0 = 1$, as it is the lcm of an empty set of integers. The collection of $\text{Tw}_{\Lambda, i}^{sv}$ for all $i \in L(\Lambda)$ generates a subgroup of the twist group, which we call the *simple vertical Λ -twist group* Tw_{Λ}^{sv} .

Lemma 6.3. *The simple vertical Λ -twist group is a finite index subgroup of the vertical twist group that can be written as the direct sum*

$$(6.8) \quad \text{Tw}_{\Lambda}^{sv} = \bigoplus_{i \in L(\Lambda)} \text{Tw}_{\Lambda, i}^{sv} \subset \text{Tw}_{\Lambda}^v.$$

Proof. The vectors \mathbf{b}_i are linearly independent in the level rotation group $\mathbb{Z}^{L^\bullet(\Lambda)}$, and moreover, $\mathbf{b}_0 = (1, \dots, 1)$ generates $\ker(\tau_{\Lambda}^{\bullet})$. Combining with Proposition 6.2, it follows immediately that the N simple twists $\tau_{\Lambda}^{\bullet}(a_i \mathbf{b}_i)$ for $i < 0$ generate a rank N subgroup of the vertical twist group $\text{Tw}_{\Lambda}^v \cong \mathbb{Z}^N$. \square

In general the inclusion in (6.8) is strict; see Example 6.8. We will see that this phenomenon is responsible for the quotient singularities of our moduli space at the boundary; see Section 10.3. Finally, we denote by $\text{Tw}_{\Lambda}^s = \text{Tw}_{\Lambda}^{sv} \oplus \text{Tw}_{\Lambda}^h$ the *simple Λ -twist group*.

The $\mathbb{C}^{L^\bullet(\Lambda)}$ -action restricts to the $\mathbb{C}^{L(\Lambda)}$ -action, which acts by scaling all but the top level. All of the above objects have analogues using this restricted action, which we denote by dropping the superscript \bullet . For example, the restriction of ϕ_{Λ}^{\bullet} to $\mathbb{Z}^{L(\Lambda)}$ is denoted by ϕ_{Λ} .

The homomorphisms ϕ_{Λ}^{\bullet} and ϕ_{Λ} have the same image in the prong rotation group P_{Λ} . Similarly τ_{Λ}^{\bullet} and τ_{Λ} have the same image Tw_{Λ}^v in $\text{Mod}_{(\Sigma, s)}$. Intuitively, the actions of $\mathbb{C}^{L(\Lambda)}$ and $\mathbb{C}^{L^\bullet(\Lambda)}$ yield the same subgroups of the prong rotation group and of the mapping class group, because the top-level factor \mathbb{C} of $\mathbb{C}^{L^\bullet(\Lambda)}$ acts (in terms of the lower-triangular basis) by simultaneously scaling the differentials at all levels by the same factor, which has no effect on the markings or prong-matchings.

We next provide another characterization of the twist group as a subgroup of the full twist group $\text{Tw}_{\Lambda}^{\text{full}}$. Recall that the full twist group was defined in Section 3.2 as the group generated by Dehn twists around all curves of Λ , and is isomorphic to $\mathbb{Z}^{E(\Lambda)}$.

Proposition 6.4. *Let $(X, \eta, \sigma, f) \in \Omega \mathcal{T}_{\Lambda}^{pm}(\mu)$. The twist group Tw_{Λ} is the subgroup of $\text{Tw}_{\Lambda}^{\text{full}}$ that fixes the turning number of every good arc in \overline{X}_{σ} that starts and ends at the same level.*

Proof. Since any element of Tw_{Λ}^h fixes the turning number of every arc, we may assume that X has only vertical nodes. We denote by $\text{Tw}_{\Lambda}^{\text{rot}}$ the subgroup of $\text{Tw}_{\Lambda}^{\text{full}}$ that fixes the turning number of each good arc in \overline{X}_{σ} that starts and ends at the same level. Note that the twist group Tw_{Λ} is a subgroup of $\text{Tw}_{\Lambda}^{\text{full}}$ of rank N .

We first show that Tw_Λ is contained in $\text{Tw}_\Lambda^{\text{rot}}$. Indeed, given $\mathbf{n} \in \ker(\phi_\Lambda^\bullet) \subset \mathbb{Z}^{L^\bullet(\Lambda)}$, for any good arc in \overline{X}_σ , when it crosses a seam γ_e corresponding to a vertical edge e , the turning number changes by $n_{\ell(e^+)} - n_{\ell(e^-)}$ under the action of $\tau_\Lambda^\bullet(\mathbf{n})$. Thus for an arc that starts at level i and ends at level j , under the action $\tau_\Lambda^\bullet(\mathbf{n})$ the turning number changes in total by $n_i - n_j$. Thus for every arc that starts and ends at the same level the turning number is preserved under the action of Tw_Λ .

Next we claim that the quotients of $\text{Tw}_\Lambda^{\text{full}}$ by Tw_Λ and by $\text{Tw}_\Lambda^{\text{rot}}$ are both torsion free. For the former, note that the image of Tw_Λ in $\text{Tw}_\Lambda^{\text{full}}$ consists of all the elements of the form $((n_{\ell(e^+)} - n_{\ell(e^-)})/\kappa_e)_{e \in \Lambda}$ where $n_{\ell(e^+)} - n_{\ell(e^-)}$ is divisible by κ_e . If the gcd of the entries of such an element is $d > 1$, then after dividing it by d the resulting element is also contained in Tw_Λ , which implies that $\text{Tw}_\Lambda^{\text{full}}/\text{Tw}_\Lambda$ is torsion free. For the latter, if an element of $\text{Tw}_\Lambda^{\text{full}}$ does not lie in $\text{Tw}_\Lambda^{\text{rot}}$, then there exists an arc that starts and ends on the same level such that its turning number changes under the corresponding action by a non-zero constant c . It follows that the n -th power of this element adds nc to the turning number of the arc, and hence no higher power can lie in $\text{Tw}_\Lambda^{\text{rot}}$.

It thus suffices to show that $\text{Tw}_\Lambda^{\text{rot}}$ has rank N , equal to the rank of Tw_Λ . The corank of $\text{Tw}_\Lambda^{\text{rot}}$ in $\text{Tw}_\Lambda^{\text{full}}$ is $h_1(\Gamma) + \sum_{i \in L^\bullet(\Lambda)} (|V(\Gamma(i))| - 1) = h_1(\Gamma) + |V(\Gamma)| - N - 1$, since the conditions imposed by preserving turning numbers are given by classes in $H_1(\Gamma)$ and paths connecting a chosen vertex to the other vertices at each level (as we assumed that there are no horizontal edges). Hence the rank of $\text{Tw}_\Lambda^{\text{rot}}$ is given by $|E(\Gamma)| - h_1(\Gamma) - |V(\Gamma)| + N + 1 = N$, where we use the fact that the Euler characteristic $|V(\Gamma)| - |E(\Gamma)|$ of a connected graph Γ is equal to $h_0(\Gamma) - h_1(\Gamma) = 1 - h_1(\Gamma)$. \square

6.3. Level rotation tori. We define the *level rotation torus* T_Λ to be the quotient $T_\Lambda = \mathbb{C}^{L(\Lambda)}/\text{Tw}_\Lambda^v \cong \mathbb{C}^{L(\Lambda)}/\ker(\phi_\Lambda)$. Similarly, the *simple level rotation torus* is the quotient $T_\Lambda^s = \mathbb{C}^{L(\Lambda)}/\text{Tw}_\Lambda^{sv}$. They will play a prominent role in defining families of multi-scale differentials. The level rotation torus obviously only depends on the enhanced level graph $\Gamma(\Lambda)$ rather than on the multicurves and we will thus write T_Γ and T_Λ interchangeably.

The following is an alternative characterization of the level rotation torus. Similarly to the twist groups, the ambient $\mathbb{C}^{L(\Lambda)}$ and also $(\mathbb{C}^*)^{L(\Lambda)}$ can be parameterized using the standard and the triangular basis.

Proposition 6.5. *The level rotation torus T_Λ is the connected component containing the identity of the subgroup of*

$$(\mathbb{C}^*)^{L(\Lambda)} \times (\mathbb{C}^*)^{E(\Lambda)} = ((r_i, \rho_e))_{i \in L(\Lambda), e \in E(\Lambda)}$$

cut out by the set of equations

$$(6.9) \quad r_{\ell(e^-)} \cdots r_{\ell(e^+)-1} = \rho_e^{\kappa_e}$$

for all edges e , where the r_i are the coordinates in the triangular basis.

There is an identification $T_\Lambda^s \cong (\mathbb{C}^*)^N$ such that the quotient map $T_\Lambda^s \rightarrow T_\Lambda$ is given in coordinates by

$$(6.10) \quad (q_i) \mapsto (r_i, \rho_e) = \left(q_i^{a_i}, \prod_{i=\ell(e^-)}^{\ell(e^+)-1} q_i^{a_i/\kappa_e} \right)$$

with the numbers a_i defined in (6.7).

Proof. Consider first the projection of the subgroup of $(\mathbb{C}^*)^{L(\Lambda)} \times (\mathbb{C}^*)^{E(\Lambda)}$ cut out by Equations (6.9) (not just its identity component) onto the $(\mathbb{C}^*)^{L(\Lambda)}$ factor. Since each ρ_e is determined by the r_i 's up to roots of unity, this projection is an unramified (possibly disconnected) cover with fiber equal to the prong rotation group $P_\Gamma = \prod_e \mathbb{Z}/\kappa_e \mathbb{Z}$.

Next we determine the connected component of the identity within this subgroup. The fundamental group of $(\mathbb{C}^*)^{L(\Lambda)}$ is equal to $\mathbb{Z}^{L(\Lambda)}$, and an element $\mathbf{n} \in \mathbb{Z}^{L(\Lambda)}$ acts by multiplying each coordinate ρ_e by $e^{(n_{\ell(e^+)} - n_{\ell(e^-)})/\kappa_e}$. Recalling Equation (6.4) that defines ϕ_Λ , we see that $\ker(\phi_\Lambda)$ is precisely the set of elements $\mathbf{n} \in \mathbb{Z}^{L(\Lambda)}$ that act by trivial monodromy. Thus the connected component of the identity is an unramified cover of $(\mathbb{C}^*)^{L(\Lambda)}$ with deck transformation group being the image of monodromy, i.e. $\mathbb{Z}^{L(\Lambda)}/\ker(\phi_\Lambda)$. As by definition the level rotation torus T_Λ is a Galois cover of $(\mathbb{C}^*)^{L(\Lambda)}$ with the same Galois group, it is equal to the connected component of the identity. This shows the first statement of the claim, from which (6.10) follows, since we exhibit a map of tori of the same dimension and the right-hand side satisfies (6.9). \square

These constructions can also be regarded as covariant functors on the category of ordered enhanced multicurves on (Σ, \mathbf{s}) . More precisely, a degeneration of enhanced multicurves $\text{dg}: \Lambda_1 \rightsquigarrow \Lambda_2$ induces a monomorphism $\widehat{\text{dg}}_*: \mathbb{C}^{\Lambda_1} \rightarrow \mathbb{C}^{\Lambda_2}$. Using Proposition 6.2 to think of twist groups as kernels of the map to the prong rotation group, up to a \mathbb{Z} -summand, we obtain a monomorphism $\text{dg}_*: \text{Tw}_{\Lambda_1} \hookrightarrow \text{Tw}_{\Lambda_2}$.

Lemma 6.6. *A degeneration of enhanced multicurves $\text{dg}: \Lambda_1 \rightsquigarrow \Lambda_2$ induces an injective homomorphism $\text{dg}_*: T_{\Lambda_1} \rightarrow T_{\Lambda_2}$. In the coordinates (6.9) the image is cut out by equations $\rho_e = 1$ for every edge e of Λ_2 that is contracted in Λ_1 , and respectively $r_i = 1$ for every level $i \in L(\Lambda_2)$ such that the images of i and $i + 1$ are the same in $L(\Lambda_1)$.*

Proof. The description of the image is obvious. For injectivity we have to show that an element in Tw_{Λ_2} in the image of $\widehat{\text{dg}}_*$ already belongs to Tw_{Λ_1} . This is obvious from the description of the twist group in Proposition 6.2. \square

We will also occasionally need the rank $N + 1$ extended level rotation torus $T_\Lambda^\bullet = \mathbb{C}^{L^\bullet(\Lambda)}/(\text{Tw}_\Lambda^v \oplus \mathbb{Z})$, as well as its simple variant $T_\Lambda^{s,\bullet} = \mathbb{C}^{L^\bullet(\Lambda)}/(\text{Tw}_\Lambda^{sv} \oplus \mathbb{Z})$.

The level rotation torus T_Λ acts on a prong-matched twisted differential (X, η, σ) , where σ is a prong-matching, via

$$(6.11) \quad (r_i, \rho_e) * (X, (\eta_{(i)}), (\sigma_e)) = \left(X, (r_i \dots r_{-1} \eta_{(i)}), (\rho_e * \sigma_e) \right)$$

where $\rho_e * \sigma_e$ is the prong-matching $P_q^{\text{in}} \rightarrow P_q^{\text{out}}$ at the node q corresponding to e given by σ_e post-composed with the rotation by $\arg(\rho)$. Note that this is the exponential

version of the action described in item (2) of Section 6.1. If X is moreover marked by f we define $(\rho_e) * f$ to be the marking of $(r_i, \rho_e) * (X, (\eta_{(i)}), (\sigma_e))$ obtained by post-composing f with a fractional Dehn twist of angle $\arg(\rho_e)$ on each vertical edge e . This marking is well-defined up to an element in Tw_Λ only.

Analogously, the simple level rotation torus acts on the set of prong-matched twisted differentials. We can also assume that these differentials are marked by f , defined modulo the action of the simple twist group. Using the map $T_\Lambda^s \rightarrow T_\Lambda$ given in Proposition 6.5, the action $*$ defined in Equation (6.11) is given in the triangular basis by

$$(6.12) \quad \mathbf{q} * (X, (\eta_{(i)}), (\sigma_e), f) = (X, (q_i^{a_i} \dots q_{-1}^{a_{-1}} \eta_{(i)}), (\rho_e * \sigma_e), (\rho_e) * f),$$

where $\rho_e = \prod_{i=\ell(e^-)}^{\ell(e^+)-1} q_i^{a_i/\kappa_e}$ with the integers a_i defined in (6.7). For later use, we recast Proposition 6.5 in terms of this action.

Corollary 6.7. *Equivalence classes of prong-matched twisted differentials up to the action (6.11) of the level rotation torus are in bijection with connected components of the subgroup of $(\mathbb{C}^*)^{L(\Lambda)} \times (\mathbb{C}^*)^{E(\Lambda)}$ cut out by Equation (6.9).*

6.4. The covering viewpoint. So far we have analyzed the group $\mathbb{C}^{L(\Lambda)}$ acting on $\Omega\mathcal{T}_\Lambda^{pm}(\mu)$ and defined twist groups as cofinite subgroups of $\mathbb{Z}^{L(\Lambda)}$. In what follows we will use compactifications of quotients of $\Omega\mathcal{T}_\Lambda^{pm}(\mu)$ by twist groups. We can alternatively construct them as finite covers of $\Omega\mathcal{T}_\Lambda(\mu)$, as we explain now.

The triangular basis provides an identification of T_Λ^s with $(\mathbb{C}^*)^{L(\Lambda)}$ and we denote by $T_{\Lambda,i}^s$ the i -th factor of this torus. Recall the direct sum expression of Tw_Λ^{sv} in (6.8). We define the *level-wise ramification groups* to be $H_i = \mathbb{Z}_i / \text{Tw}_{\Lambda,i}^{sv}$, where \mathbb{Z}_i is the i -th factor of $\mathbb{Z}^{L(\Lambda)} \subset \mathbb{C}^{L(\Lambda)}$. By definition, we have the cardinality $|H_i| = a_i$ (defined in (6.7)) and the identification $H := \bigoplus_{i \in L(\Lambda)} H_i = \text{Ker}(T_\Lambda^s \rightarrow (\mathbb{C}^*)^{L(\Lambda)})$. On the other hand, we may define the *(full) ramification group* associated with an enhanced level graph Λ to be $G := \text{Ker}(T_\Lambda \rightarrow (\mathbb{C}^*)^{L(\Lambda)})$. By definition we have an exact sequence of finite abelian groups

$$(6.13) \quad 0 \rightarrow K_\Lambda = \text{Tw}_\Lambda^v / \text{Tw}_\Lambda^{sv} \rightarrow H \rightarrow G \rightarrow 0.$$

Note that the map $H_i \rightarrow G$ is injective for every $i \in L(\Lambda)$, since an element in H_i and its image in G act by the same fractional Dehn twists the seams. The situation is summarized by the following diagram:

$$\begin{array}{ccc} \Omega\mathcal{T}_\Lambda^{pm}(\mu) / \text{Tw}_\Lambda^s & & \\ \downarrow /H & \searrow /K_\Lambda & \\ \Omega\mathcal{T}_\Lambda(\mu) & & \Omega\mathcal{T}_\Lambda^{pm}(\mu) / \text{Tw}_\Lambda \\ & \swarrow /G & \end{array}$$

Of course, all the maps in the diagram are unramified covers, but they will become ramified with local ramification groups H_i at the appropriate boundary divisors, once we consider the compactifications in Section 10.3.

Example 6.8. (*The twist group quotient K_Λ can be non-trivial.*) Consider the enhanced level graphs Γ_1 and Γ_2 of our running example in Section 2.6. In both cases, the level-wise undegenerations dg_{-1} and dg_{-2} as introduced in Section 5.1 are both equal to the graph with two vertices connected by two edges, labeled by 1 and 3 respectively.

First consider the enhanced level graph Γ_1 on the left of Figure 1. Then the map

$$\phi_{\Lambda_1}^\bullet : \mathbb{Z}^{L^\bullet(\Lambda)} \cong \mathbb{Z}^3 \rightarrow P_{\Gamma_1} \cong \mathbb{Z}/3\mathbb{Z} \times \mathbb{Z}/3\mathbb{Z} \times \mathbb{Z}/\mathbb{Z}$$

is given by

$$(n_0, n_{-1}, n_{-2}) \mapsto (n_0 - n_{-1}, n_{-1} - n_{-2}, n_0 - n_{-2})$$

in the standard basis. Then $\ker(\phi_{\Lambda_1}^\bullet)$ is the subgroup of \mathbb{Z}^3 consisting of elements of the form $(m, m + 3k_1, m + 3k_2)$ for $m, k_1, k_2 \in \mathbb{Z}$, and hence $\text{Tw}_\Lambda^v = \text{Tw}_\Lambda$ is generated by the vectors $(0, 3, 0)$ and $(0, 0, 3)$. The simple Λ_1 -twist group $\text{Tw}_{\Lambda_1}^{sv}$ is the direct sum $\text{Tw}_{\Lambda_1, -1}^{sv} \oplus \text{Tw}_{\Lambda_1, -2}^{sv}$ where $\text{Tw}_{\Lambda_1, i}^{sv} = (\text{dg}_i)_*(\text{Tw}_{\Lambda_1, i}^v)$. In this case $\text{Tw}_{\Lambda_1, -1}^{sv}$ is generated by $(0, 3, 3)$ and $\text{Tw}_{\Lambda_1, -2}^{sv}$ is generated by $(0, 0, 3)$, hence $\text{Tw}_{\Lambda_1}^{sv}$ coincides with $\text{Tw}_{\Lambda_1}^v$. The level rotation torus T_{Λ_1} is an unramified cover of $(\mathbb{C}^*)^2$ of degree 9 with Galois group equal to the prong rotation group P_{Γ_1} . The action of the level rotation group by $\phi_{\Lambda_1}^\bullet$ has a unique orbit, hence the nine prong-matchings are all equivalent.

Next consider the enhanced level graph Γ_2 on the right of Figure 1. Then we have

$$\phi_{\Lambda_2}^\bullet : \mathbb{Z}^{L^\bullet(\Lambda_2)} \cong \mathbb{Z}^3 \rightarrow P_{\Gamma_2} \cong \mathbb{Z}/\mathbb{Z} \times \mathbb{Z}/\mathbb{Z} \times \mathbb{Z}/3\mathbb{Z}$$

given by

$$(n_0, n_{-1}, n_{-2}) \mapsto (n_0 - n_{-1}, n_{-1} - n_{-2}, n_0 - n_{-2})$$

in the standard basis. Then $\ker(\phi_{\Lambda_2}^\bullet)$ is the subgroup of \mathbb{Z}^3 consisting of elements of the form $(m, m + k_1, m + 3k_2)$ for $m, k_1, k_2 \in \mathbb{Z}$, and hence $\text{Tw}_{\Lambda_2}^v = \text{Tw}_{\Lambda_2}$ is generated by the vectors $(0, 1, 0)$ and $(0, 0, 3)$. The simple Λ_2 -twist group $\text{Tw}_{\Lambda_2}^{sv}$ is the direct sum $\text{Tw}_{\Lambda_2, -1}^{sv} \oplus \text{Tw}_{\Lambda_2, -2}^{sv}$ where $\text{Tw}_{\Lambda_2, -1}^{sv}$ is generated by $(0, 3, 3)$ and $\text{Tw}_{\Lambda_2, -2}^{sv}$ is generated by $(0, 0, 3)$, hence $\text{Tw}_{\Lambda_2}^{sv}$ is a subgroup of index 3 in $\text{Tw}_{\Lambda_2}^v$. The level rotation torus T_{Λ_2} is an unramified cover of $(\mathbb{C}^*)^3$ of degree 3 with Galois group equal to the prong rotation group P_{Γ_2} . The action of the level rotation group has a unique orbit, hence the three prong-matchings are all equivalent.

In both cases the local ramification groups are $H_i \cong \mathbb{Z}/3\mathbb{Z}$ for $i = -1, -2$, and the group G coincides with the prong rotation group in each case. However, in the first case $H \cong G$, while in the second case $H \twoheadrightarrow G$ has kernel $K_{\Lambda_2} = \mathbb{Z}/3$. In particular, in the second case the quotient map of a smooth space by K_{Λ_2} will produce quotient singularities in our compactification, which will be illustrated in Example 10.9.

In order to conclude this section we give a cautionary example.

Example 6.9. (*The number of non-equivalent prong-matchings may decrease under degenerations.*) We consider the degeneration of enhanced level graphs as shown in Figure 3. The first graph is a two-level graph with two edges e_1 and e_2 between two vertices. Moreover we set $\kappa_1 = \kappa_2 = 2$. We degenerate this graph to a three-level triangle with three edges e_1, e_2 and e_3 labeled as in Figure 3. The edges are labeled by $\kappa_1 = \kappa_2 = 2$ and $\kappa_3 = 1$.

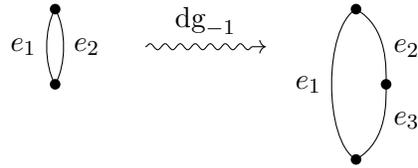


FIGURE 3. A degeneration that decreases the number of prong-matchings.

Clearly the action of the level rotation group of Equation (6.4) has two orbits in the first case. On the other hand one can check that in the second case it has only one orbit. Hence this degeneration decreases the number of non-equivalent prong-matchings from two to one.

6.5. The level rotation torus closure. The partial closures of tori we define here will give local models of the toroidal part of our compactification. Recall that the level rotation torus T_Λ is by Proposition 6.5 naturally embedded in $(\mathbb{C}^*)^{L(\Lambda)} \times (\mathbb{C}^*)^{E(\Lambda)}$, where it is the connected component of the identity of the torus cut out by Equation (6.9). We define the *level rotation torus closure* \overline{T}_Λ to be the closure of this identity component in $\mathbb{C}^{L(\Lambda)} \times \mathbb{C}^{E(\Lambda)}$.

On the other hand, the simple level rotation torus T_Λ^s is naturally identified with $(\mathbb{C}^*)^{L(\Lambda)}$, with closure $\overline{T}_\Lambda^s = \mathbb{C}^{L(\Lambda)}$. The group $K_\Lambda = \text{Tw}_\Lambda / \text{Tw}_\Lambda^s$ introduced in the previous section acts on $T_\Lambda^s = \mathbb{C}^{L(\Lambda)} / \text{Tw}_\Lambda^s$. Since each element in K_Λ acts diagonally by a tuple of roots of unity, this action extends to an action of K_Λ on \overline{T}_Λ^s . The quotient will be the local model for the toroidal part of the compactification that we will construct. Our goal here is to relate this viewpoint with the closure of the level rotation torus.

Proposition 6.10. *The projection map $p: T_\Lambda^s \rightarrow T_\Lambda$ given by Equation (6.10) extends to \overline{T}_Λ^s and descends to an isomorphism $\overline{p}: \overline{T}_\Lambda^s / K_\Lambda \rightarrow \overline{T}_\Lambda^n$ to the normalization of the level rotation torus closure.*

Proof. The map p extends to a map $p_2: \overline{T}_\Lambda^s \rightarrow \overline{T}_\Lambda$ since it is given explicitly by monomials, in the coordinates used for taking the closure. Since K_Λ acts on \overline{T}_Λ^s and since p is the quotient by K_Λ , the map p_2 factors through the quotient to give $p_3: \overline{T}_\Lambda^s / K_\Lambda \rightarrow \overline{T}_\Lambda$. Since a quotient of $\mathbb{C}^{L(\Lambda)}$ by a finite group is normal, the map p_3 factors through the normalization map \overline{p} . The map \overline{p} is finite and birational since on the open set $T_\Lambda^s / K_\Lambda = T_\Lambda$. Since the target is normal, it follows that the map \overline{p} is an isomorphism (e.g. [Sta18, Lemma 28.52.8]) □

Example 6.11. *(In general \overline{T}_Λ is not normal.)* This can be seen from the example of a graph Γ with two vertices on two levels, connected by two edges e_1 and e_2 with $\kappa_1 = 2$ and $\kappa_2 = 3$. Then \overline{T}_Λ has a cusp, locally modelled on $\mathbb{C}[f_1, f_2, s] / \langle f_1^2 - s, f_2^3 - s \rangle$. Its normalization is $\mathbb{C}[t]$, with the normalization map given by $f_1 = t^3$ and $f_2 = t^2$. This change of coordinates also describes \overline{p} here, since $\text{Tw}_\Lambda = \text{Tw}_\Lambda^s$ in this example.

7. FAMILIES OF MULTI-SCALE DIFFERENTIALS

In this section we define families of multi-scale differentials. The starting point will be a flat family of pointed stable curves $(\pi: \mathcal{X} \rightarrow B, \mathbf{z})$, for an arbitrary base B , possibly reducible and non-reduced. We will first define a germ of multi-scale differentials at a point $p \in B$. Roughly speaking, this will consist of four pieces of data: the structure of an enhanced level graph on the dual graph Γ_p of the fiber X_p , a *rescaling ensemble*, which is a germ of a morphism $R_p: B_p \rightarrow \overline{T}_{\Gamma_p}^n$ to the normalization of the level rotation torus closure, a *collection of rescaled differentials* $\omega_{(i)}$, and finally prong-matchings at all nodes of the family, such that for every non-semipersistent node (as defined below) the prong-matching is naturally induced by the family. These data satisfy some restrictions analogous to those of a single twisted differential, and there is an equivalence relation given by the action of the level rotation torus, analogous to the definition of a single multi-scale differential.

We will show in Proposition 7.9 that in favorable circumstances, for example for a family over a smooth base curve B with no persistent nodes, giving a multi-scale differential simply amounts to giving a family of stable differentials of type μ , that do not vanish identically on any fiber.

7.1. Germs of families of multi-scale differentials. We will define all the notions locally first, and thus until Section 7.2 we will assume the base $B = B_p$ to be local.

Recall that for each node q_e of X_p there is a function $f_e \in \mathcal{O}_B$ called *smoothing parameter* so that the family has the local normal form $u_e v_e = f_e$ in a neighborhood of q_e . The parameter f_e is only defined up to multiplication by a unit in \mathcal{O}_B . We will write $[f_e] \in \mathcal{O}_B / \mathcal{O}_B^*$ for the equivalence class.

Given an enhanced level graph Γ_p , suppose we have a morphism $R: B \rightarrow \overline{T}_{\Gamma_p}^n$. This morphism determines for each vertical edge e a function $f_e \in \mathcal{O}_B$ and for each level i a function $s_i \in \mathcal{O}_B$, such that if an edge e joins levels $j < i$, then

$$(7.1) \quad f_e^{\kappa_e} = s_j \dots s_{i-1}.$$

Definition 7.1. A *rescaling ensemble* is a morphism $R: B \rightarrow \overline{T}_{\Gamma_p}^n$ such that the parameters $f_e \in \mathcal{O}_B$ for each vertical edge e determined by R lie in the equivalence class $[f_e]$ determined by the family $\pi: \mathcal{X} \rightarrow B$. \triangle

The s_i and f_e will be called the *rescaling parameters* and *smoothing parameters* determined by R . The rescaling ensemble R can be thought of as a choice of these parameters that satisfies (7.1) for each edge e of Γ_p , together with the choices of appropriate roots of these which define a lift to $\overline{T}_{\Gamma_p}^n$, see Proposition 6.10 for the precise statement.

Definition 7.2. A *collection of rescaled differentials of type μ* at $p \in B$ is a collection of (germs of) sections $\omega_{(i)}$ of $\omega_{\mathcal{X}/B}$ defined on open subsets U_i of \mathcal{X} , indexed by the levels i of the enhanced level graph Γ_p . Each U_i is required to be a neighborhood of the subcurve $X_{p, \leq i}$ with the points of its intersection with $X_{p, > i} \cup \mathcal{Z}^\infty$ removed. For each level i and each edge e of Γ_p whose lower vertex is at level i or below, we define $r_{e, (i)} \in \mathcal{O}_B$ to be the period of $\omega_{(i)}$ along the vanishing cycle γ_e for the node q_e . We require the collection to satisfy the following constraints:

- (1) For any levels $j < i$ the differentials satisfy $\omega_{(i)} = s_j \cdots s_{i-1} \omega_{(j)}$ on $U_i \cap U_j$, where $s_i \in \mathcal{O}_B$ with $s_i(p) = 0$.
- (2) For any edge e joining levels $j < i$ of Γ_p , 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

$$(7.2) \quad \omega_{(i)} = (u_e^{\kappa_e} + f_e^{\kappa_e} r_{e,(i)}) \frac{du_e}{u_e} \quad \text{and} \quad \omega_{(j)} = -(v_e^{-\kappa_e} + r_{e,(j)}) \frac{dv_e}{v_e},$$

where κ_e is the enhancement of Γ_p . The irreducible components of $\mathcal{X}|_{V(f_e)}$ where $\omega_{(i)}$ is zero or ∞ are called respectively *vertical zeros* and *vertical poles*.

- (3) The $\omega_{(i)}$ have order m_k along the sections \mathcal{Z}_k that meet the level- i subcurve of X_p ; these are called *horizontal zeros and poles*. Moreover, $\omega_{(i)}$ is holomorphic and non-zero away from its horizontal and vertical zeros and poles.
- (4) (Global Residue Condition) Let Σ be the topological surface obtained by smoothing each node of X_p , and regard the vanishing cycles γ_e as oriented curves on Σ . Then each relation

$$\sum_{e: \ell(e^-) \leq i} \alpha_e \gamma_e = 0 \quad \text{in } H_1(\Sigma \setminus P_s, \mathbb{Q}) \text{ for some } \alpha_e \in \mathbb{C} \text{ implies } \sum_{e: \ell(e^-) \leq i} \alpha_e r_{e,(i)} = 0.$$

If the rescaling and smoothing parameters for the collection $\omega_{(i)}$ agree with those of the rescaling ensemble R , we call them *compatible*. We denote the collection by $\boldsymbol{\omega} = (\omega_{(i)})_{i \in L \bullet(\Gamma_p)}$ or by $\boldsymbol{\omega}_p$. \triangle

Some remarks to unravel the meaning of this definition are in order. Condition (2) is often automatic from Theorem 4.3, but this theorem does not apply in the case of semipersistent nodes defined below.

Condition (3) ensures that each $\omega_{(i)}$ is not identically zero on a neighborhood of the i -th level of X_p . Condition (1) ensures that $\omega_{(i)}$ vanishes on the components of X_p of level $j < i$. Moreover, $\omega_{(i)}$ vanishes on a neighborhood in \mathcal{X} of X_j for some level $j < i$, if some s_k with $j \leq k \leq i-1$ vanishes in a neighborhood of p .

Conditions (4) and (1) together imply the usual global residue condition. Note that $r_{e,(i)}$ agrees with $2\pi\sqrt{-1}$ times the residue of $\omega_{(i)}$ at q_e over the locus where the node q_e persists. By condition (1), given two levels $j < i$ and any edge e such that $\ell(e^-) \leq j < i$ we have

$$r_{e,(i)} = s_j \cdots s_{i-1} r_{e,(j)}.$$

In particular, if $s_j = 0$ for some $\ell(e^-) \leq j$, then $r_{e,(i)} = 0$. Consequently, the relations reflect the level filtration in an equivalent way as stated in Proposition 11.2, see also the following example.

Example 7.3. (*Definition 7.2 extends to fiberwise GRC.*) We consider the level graph given by Figure 4, where $\kappa_{e_i} = 2$ for every i . Consider a collection of rescaled differentials with $\omega_{(-1)} = t^2 \omega_{(-2)}$ (while $\omega_{(0)}$ will not matter for us) over $B = \text{Spec } \mathbb{C}[t]/(t^3)$ and let $r_{e,(i)}: B \rightarrow \mathbb{C}$ be as in Definition 7.2. The usual GRC from Section 2.4 states that the residues at the point $t = 0$ satisfy $r_{e_1,(-1)}(0) = r_{e_2,(-1)}(0) = 0$.

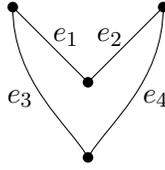


FIGURE 4. An enhanced level graph to compare different versions of the GRC.

Since the vanishing cycles corresponding to e_1 and e_3 are homologous, Definition 7.2 (4) states that

$$0 = r_{e_1,(-1)} + r_{e_3,(-1)} = r_{e_1,(-1)} + t^2 r_{e_3,(-2)}.$$

This condition reproduces the GRC when setting $t = 0$, but produces a stronger statement about the first order behavior of the residues.

In preparation for the notion of prong-matchings in families we define a subtle variant of the usual notion of persistent nodes that becomes relevant over a non-reduced base, and thus in particular for first order deformations.

Definition 7.4. Given a germ of a family $\pi: \mathcal{X} \rightarrow B_p$ at p , we say that a node e is *persistent* if $f_e = 0$. If the dual graph Γ_p has been provided with an enhanced level graph structure, we say that a node e is *semipersistent* if $f_e^{\kappa_e} = 0$. \triangle

We start with a discussion of prong-matchings in families, generalizing the definitions in Section 5.4. Suppose first that q_e is a persistent node joining levels $j < i$ of Γ_p . In local analytic coordinates, the family is of the form $u_e v_e = 0$. We write Q_e for the nodal subscheme cut out by $u_e = v_e = 0$, so that Q_e can be thought of as the image of a nodal section $B \rightarrow \mathcal{X}$. We write \mathcal{N}_e^+ and \mathcal{N}_e^- for the normal bundles to Q_e in each branch of \mathcal{X} along Q_e . These are line bundles on Q_e , because Q_e is a Cartier divisor, and by pullback via the nodal section they can be regarded as line bundles on B .

We also have the rescaled differentials $\omega_{(i)}$ and $\omega_{(j)}$ defined near Q_e in its respective branches, and choose local coordinates u_e and v_e so that these differentials are in their standard form (7.2) (with $f_e = 0$). The *prongs* of Q_e are then the κ_e sections of the dual line bundles $(\mathcal{N}_e^\pm)^*$ given by $v_+ = \theta_+ \frac{\partial}{\partial u_e}$ and $v_- = \theta_- \frac{\partial}{\partial v_e}$, where θ_\pm range over all possible κ_e -th roots of unity. A *prong-matching* at Q_e is a section σ_e of $\mathcal{P}_e = \mathcal{N}_e^+ \otimes \mathcal{N}_e^-$ such that $\sigma_e(v_+ \otimes v_-)^{\kappa_e} = 1$ for any two prongs v_+ and v_- . Intuitively each prong v_+ is matched to the unique prong v_- such that $\sigma_e(v_+ \otimes v_-) = 1$.

Prong-matchings can be defined similarly for a non-persistent node. In this case, the function f_e defines a subscheme $B_e \subset B$ over which this node persists. The entire discussion of the previous paragraph can be carried out over B_e , and one defines a prong-matching as an appropriate section of \mathcal{P}_e , which is now a line bundle over B_e .

For a non-semipersistent node e , there is a natural *induced prong-matching* σ_e over B_e which is defined 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 $\sigma_e = du_e \otimes dv_e$, where u_e and v_e are as in (7.2). Any two possible choices of u_e and v_e are of the form αu_e and $\alpha^{-1} v_e$ for some unit $\alpha \in \mathcal{O}_B^*$, so the induced prong-matching does not depend on this choice.

We can now package everything into the local version of our main notion.

Definition 7.5. Given a family of pointed stable curves $(\pi: \mathcal{X} \rightarrow B, \mathbf{z})$ and B_p a germ of B at p , the *germ of a family of multi-scale differentials* of type μ over B_p consists of the following data:

- (1) the structure of an enhanced level graph on the dual graph Γ_p of the fiber X_p ,
- (2) a rescaling ensemble $R: B \rightarrow \overline{T}_{\Gamma_p}^n$, compatible with
- (3) a collection of rescaled differentials $\omega = (\omega_{(i)})_{i \in L^\bullet(\Gamma_p)}$ of type μ , and
- (4) a collection of prong-matchings $\sigma = (\sigma_e)_{e \in E(\Gamma)^v}$, which are sections of \mathcal{P}_e over B_e . For the non-semipersistent nodes, these are required to agree with the induced prong-matchings defined above. \triangle

The level rotation torus $T_{\Gamma_p}(\mathcal{O}_B)$ acts on all of the above data, and we consider the data $(\omega_{(i)}, R, \sigma_e)$ to be *equivalent* to $(\rho \cdot \omega_{(i)}, \rho^{-1} \cdot R, \rho \cdot \sigma_e)$ for any $\rho \in T_{\Gamma_p}(\mathcal{O}_B)$. Here the torus action is defined by $\rho \cdot \omega_{(i)} = s_i \omega_{(i)}$ and $\rho \cdot \sigma_e = f_e \sigma_e$, and $\rho^{-1} \cdot (\)$ denotes post-composition with the multiplication by ρ^{-1} .

Remark 7.6. If B is a (reduced) point, a multi-scale differential is exactly a prong-matched twisted differential as defined in Definition 5.8.

The extended level rotation torus $T_{\Gamma_p}^\bullet(\mathcal{O}_B)$ also acts on this data. When we take the quotient by the extended torus, the resulting object is called a family of projectivized multi-scale differentials.

All these notions come with a simple version used later for smooth charts. We define a *simple rescaling ensemble* to be a (germ of a) morphism $R^s: B \rightarrow \overline{T}_{\Gamma_p}^s$ to the simple level rotation torus closure, such that the composition with $\overline{T}_{\Gamma_p}^s \rightarrow \overline{T}_{\Gamma_p}^n$ is a rescaling ensemble in the above sense. Concretely, the map R^s is given by functions $t_i: B \rightarrow \mathbb{C}$ from which we obtain

$$(7.3) \quad s_i = t_i^{a_i} \quad \text{and} \quad f_e = \prod_{i=\ell(e^-)}^{\ell(e^+)-1} t_i^{m_{i,e}},$$

similarly to (6.10).

7.2. Restriction of germs of multi-scale differentials to nearby points. Now we allow B to be any complex space containing a point p . Before giving the global definition of families, we need to define restrictions of germs. For this purpose consider a germ of multi-scale differentials at p given by the data $(\Gamma_p, \omega_{(i)}, R_p, \sigma_e)$. Let $U \subset B$ be a neighborhood over which R and σ_e are all defined. For every $q \in U$ we wish to define the germ of multi-scale differentials at q induced by this germ at p .

First we explain how this datum defines an undegeneration of enhanced level graphs $\text{dg}: \Gamma_q \rightsquigarrow \Gamma_p$ as in Section 5.1. There is a map of dual graphs $\delta: \Gamma_p \rightarrow \Gamma_q$ obtained by contracting every vertical edge e such that $f_e(q) \neq 0$. (Whether horizontal edges are contracted or not is determined by the fiber X_q .) If e is contracted and joins levels $j < i$, then since $f_e^{k_e} = s_j \cdots s_{i-1}$, the rescaling parameter $s_k(q) \neq 0$ for each $j \leq k < i$. We then define the order on Γ_q so that the levels of Γ_q correspond to maximal intervals $(j, i]$ in $L^\bullet(\Gamma_p)$ such that $s_i(q) = 0$ and $s_k(q) \neq 0$ for every smaller k in the interval.

Second, we define the rescaling ensemble at q . The undegeneration induces a corresponding homomorphism $\mathrm{dg}_*: T_{\Gamma_q} \rightarrow T_{\Gamma_p}$ defined in Section 6.3 and a homomorphism $\overline{\mathrm{dg}}_*: \overline{T}_{\Gamma_q}^n \rightarrow \overline{T}_{\Gamma_p}^n$ which is equivariant with respect to the action of each torus on the normalization of its closure.

Proposition 7.7. *Given a rescaling ensemble $R_p: B \rightarrow \overline{T}_{\Gamma_p}^n$ as above, there exists a neighborhood $V \subset B$ of q , a rescaling ensemble $R_q: V \rightarrow \overline{T}_{\Gamma_q}^n$, and an element $\tau \in T_{\Gamma_p}$ such that*

$$\overline{\mathrm{dg}}_* \circ R_q = \tau \cdot R_p$$

as germs of functions at q . Moreover any two such τ differ by composition with an element of T_{Γ_q} .

Proof. We consider the rescaling ensemble as $R_p: B \rightarrow \overline{T}_{\Gamma_p}^n$. To avoid identifying the connected components of the variety defined by Equations (6.9) we take the fiber product of this map with the finite quotient map $\overline{p}: \overline{T}_{\Gamma_p}^s \rightarrow \overline{T}_{\Gamma_p}^s/K_p = \overline{T}_{\Gamma_p}^n$, to get a simple rescaling ensemble $R_p^s: B^s \rightarrow \overline{T}_{\Gamma_p}^s$, defined on some ramified cover B^s of B . Let $q' \in B^s$ denote some preimage of $q \in B$. Note that levels $i, i+1 \in L^\bullet(\Gamma_p)$ have the same image in $L^\bullet(\Gamma_q) = L^\bullet(\Gamma_{q'})$ if and only if $t_i(q') \neq 0$. Similar to Lemma 6.6, the image of the monomorphism $\mathrm{dg}_*^s: \overline{T}_{\Gamma_{q'}}^s \rightarrow \overline{T}_{\Gamma_p}^s$ of simple level rotation tori is cut out by the equations $t_i = 1$ for all levels i such that the images of level i and $i+1$ are the same in $L^\bullet(\Gamma_q)$. To ensure that we land in this image we define $\tau^s \in T_{\Gamma_p}^s$ by

$$(\tau^s)_i = \begin{cases} (t_i(q'))^{-1}, & \text{if } \mathrm{dg}(i) = \mathrm{dg}(i+1) \\ 1, & \text{otherwise.} \end{cases}$$

Now $\tau^s \cdot R_p^s$ defines a simple rescaling ensemble on a neighborhood of q' .

The multiplication map τ^s is K_{Γ_p} -equivariant, since the torus $T_{\Gamma_p}^s$ is commutative. Since R_p^s is K_{Γ_p} -equivariant by construction as fiber product, the composition of the two maps is also equivariant and thus descends to the required map $\tau \cdot R_p$.

The uniqueness of τ up to the action of T_{Γ_q} follows from observing that if for some other τ' the composition $\tau' \cdot R_p$ were to also lie in the image of $\overline{\mathrm{dg}}_*$, then the values of $\tau' \cdot R_p$ and $\tau \cdot R_p$ must all be equal to 1 on all the edges of Γ_p that are contracted in Γ_q , and on all levels $i \in L^\bullet(\Gamma_p)$ such that levels i and $i+1$ have the same image in $L^\bullet(\Gamma_q)$. Thus $\tau' \cdot \tau^{-1} \in T_{\Gamma_p}$ must act trivially on all such edges and levels. But this is precisely to say that $\tau' \cdot \tau^{-1}$ lies in the image of T_{Γ_q} embedded into T_{Γ_p} . \square

From now on, we replace the germ of the multi-scale differential at p by an equivalent one, obtained by acting by τ in the preceding proposition.

Third, we define the collection of rescaled differentials at q . For each $k \in L^\bullet(\Gamma_q)$ let $(j_k, i_k]$ be its preimage in $L^\bullet(\Gamma_p)$. We have $s_\ell = 1$ for each $\ell \in (j_k, i_k)$. Moreover, for any edge e of Γ_p joining two levels in this interval, we have $f_e = 1$. By Condition (1), the differentials $\omega_{(\ell)}$ for $\ell \in (j_k, i_k]$ all agree on their overlap, so we define the differential $\omega_{(k)}$ over q by gluing all of the differentials in this interval.

The last datum to define is a prong-matching for each edge of Γ_q . An edge e of Γ_p defines a subscheme B_e of B cut out by f_e , and this edge persists in Γ_q exactly when

$q \in B_e$ (or equivalently when $f_e(q) = 0$). In that case, the prong-matching σ_e is a section of \mathcal{P}_e over B_e , and we may restrict this to a germ of a section over the neighborhood of q intersected with B_e .

7.3. The global situation. We finally obtain global objects by patching together germs using the restriction procedure of the previous subsection. Essentially, we mimic the definition of sheafification of a presheaf.

Definition 7.8. Given a family of pointed stable curves $(\pi: \mathcal{X} \rightarrow B, \mathbf{z})$, a *family of multi-scale differentials* of type μ over B is a collection of germs of multi-scale differentials of type μ for every point $p \in B$ such that if the germs at p and at p' are both defined at q , their restrictions to q are equivalent germs. \triangle

We usually refer to a multi-scale differential by $(\omega_p, \sigma)_{p \in B}$ or just by ω , suppressing Γ_p and R_p to simplify notation.

Given a family of multi-scale differentials over B and a map $\varphi: B' \rightarrow B$, we can pull back the family to a family of multi-scale differentials over B' by pulling back each germ. For this purpose we note that rescaling ensembles and prong-matchings have obvious pullbacks by pre-composing the maps with φ and the collection of rescaled differentials can be pulled back as sections of the relative dualizing sheaf. The notion of a family of multi-scale differentials can be regarded as a moduli functor $\mathbf{MS}_\mu: (\text{Analytic spaces}) \rightarrow (\text{Sets})$ that associates to an analytic space B the set of isomorphism classes of families of multi-scale differentials of type μ over B . Similarly, there is a projectivized analogue $\mathbb{P}\mathbf{MS}_\mu$. The notion of families of multi-scale differentials defines in an obvious way a *groupoid* \mathcal{MS}_μ that retains the information of isomorphisms (Section X.12 of [ACG11] provides a textbook introduction, highlighting the difference between \mathcal{MS}_μ and \mathbf{MS}_μ). In Section 14 we will see that this is a Deligne-Mumford stack.

Much of the data of multi-scale differentials is determined automatically in good circumstances. The reader should keep in mind the following situation that will be a special case of the considerations in Section 7.4.

Proposition 7.9. *If B is a smooth curve, then giving a multi-scale differential of type μ on a family $\mathcal{X} \rightarrow B$ without persistent nodes simply amounts to specifying a family ω of stable differentials of type μ in the generic fiber which is not identically vanishing in any fiber.*

Proof. Since B is smooth and one-dimensional, Proposition 7.13 below implies that the family $(\mathcal{X} \rightarrow B, \omega)$ is adjustable and hence orderly (see Definitions 7.11 and 7.14 below). The claim then follows from Proposition 7.15. \square

In contrast to this we observe:

Example 7.10. (*Lower level differentials are not determined by $\omega_{(0)}$.*) If B admits a zero divisor s , say $s \cdot y = 0$, then differentials on the lower level components of a collection of rescaled differentials with given $\omega_{(0)}$ are not uniquely determined. In fact, if $\omega_{(0)} = s\omega_{(-1)}$, then we also have $\omega_{(0)} = s(\omega_{(-1)} + y\xi)$ for any differential ξ .

7.4. Adjustable and orderly families. In this section we analyze the ingredients of multi-scale differentials and when their existence is automatic. The study here will be needed for the description of the moduli space of multi-scale differentials as a blowup of the normalization of the IVC, in Section 14.3.

For families of pointed stable differentials $(\pi: \mathcal{X} \rightarrow B, \omega, \mathbf{z})$ considered in this section, we make a *standing assumption* that ω does not vanish identically on any fiber of π .

Definition 7.11. A family of pointed stable differentials $(\pi: \mathcal{X} \rightarrow B, \omega, \mathbf{z})$ is called *adjustable of type μ* , if for every $p \in B$ and for every irreducible component X of the fiber X_p over p , there exists a non-zero regular function $h \in \mathcal{O}_{B,p} \setminus \{0\}$ and a family of differentials η defined over a neighborhood of X (minus the horizontal poles and minus the intersection with the other components of X_p), such that $\omega = h\eta$ and such that η does not vanish identically on X , and if moreover $\eta|_X$ has zero or pole order m_j prescribed by μ at every marked point $z_j \in X$ and has no other zeros or poles in the smooth locus of X .

Such a function h is called an *adjusting parameter* for (\mathcal{X}, ω) at the component X , and η is called an *adjusted differential*. \triangle

Later we will show that under some mild assumptions an adjustable family naturally yields the data of a family of multi-scale differentials (see Proposition 7.15).

The adjusting parameter h is not unique, since multiplying η by a unit in $\mathcal{O}_{B,p}$ and multiplying h by the inverse of such a unit gives another adjusting parameter. The following example shows that the existence of adjusting parameters is a non-trivial condition.

Example 7.12. (*Adjusting parameters may not exist.*) Let B be the cuspidal cubic defined by $u^2 - v^3 = 0$ in \mathbb{C}^2 , and let $p \in B$ be the origin. Consider the family of curves $xy = u$ over B and the family of differentials given by $\omega = (v + uy)dy + dx$. Over the origin $(u, v) = (0, 0)$, the form ω restricts to dx on the component $\{y = 0\}$ and vanishes identically on the component $\{x = 0\}$. Note that the family ω can be written as $(v + uy - x/y)dy$. Suppose that this family is adjustable. Then using the parametrization of B given by $u = t^3$ and $v = t^2$ for $t \in \mathbb{C}^2$, it is easy to see that the adjusting parameter h must be equal to v (times a unit at p). However, $(uy/v)dy$ has no well-defined limit on the component $\{x = 0\}$ of the central fiber, since $t = u/v$ is not a well-defined element in $\mathcal{O}_{B,p}$, which leads to a contradiction.

However if the base is sufficiently nice, adjusting parameters do exist.

Proposition 7.13. *If the base B is normal, then any family (\mathcal{X}, ω) satisfying the standing assumption is adjustable. Moreover, any two adjusting parameters for a point $p \in B$ and an irreducible component X of X_p differ by multiplication by a unit in $\mathcal{O}_{B,p}$.*

We first recall some terminology. Denote by $\mathcal{Z} = \sum_{j=1}^n m_j \mathcal{Z}_j$ with $\mathcal{Z}_j \subset \mathcal{X}$ being the image of the section of the j -th zeros or poles z_j of ω . An effective Cartier divisor $V \subset \mathcal{X}$ is called a vertical divisor if the image $\pi(V) \subset B$ is a divisor. Note that any section of \mathcal{X} is not vertical because its image is the entire base. In particular, the divisors \mathcal{Z}_i and \mathcal{Z} are not vertical. A vertical divisor is called a vertical zero divisor of ω if it is contained in the zero locus of ω (and being vertical ensures that it is not contained in \mathcal{Z}).

Proof. Suppose ω vanishes identically on an irreducible component X of the fiber X_p for some $p \in B$. Then X is contained in the vertical zero divisor of ω . More precisely, let $W \subset \mathcal{X}$ be a small neighborhood of the generic point of X away from all nodal loci of \mathcal{X} and let $U = \pi(W) \subset B$ be the corresponding neighborhood of p in B . Then $W \cong U \times \Delta$ where Δ is a disk. Let $D \subset W$ be the vertical zero divisor of ω in W , i.e. D is the zero divisor of ω regarded as a holomorphic section of the twisted dualizing line bundle $\omega_{\mathcal{X}/U}(-\mathcal{Z})$ restricted to W , so that in particular $X \subset D$ generically. Since B is normal, it is reduced, and thus \mathcal{X} is reduced. Since D is the zero locus of a holomorphic section of a line bundle, D is thus an effective Cartier divisor on W (possibly reducible and non-reduced), and we denote then by $h_D \in \mathcal{O}_W \cong \mathcal{O}_{U \times \Delta}$ the local defining equation for $D \subset W$.

Let $E = \pi(D) \subset U$ be the scheme-theoretic image of D under the projection to U . Since the only zeros and poles of ω on a generic fiber are the marked points \mathbf{z} , it follows that E does not contain a generic point of U . We claim that h_D does not depend on the second factor, i.e. $h_D = (\pi|_W)^*g_D$ is the pullback of some $g_D \in \mathcal{O}_U$. Indeed, if this were not the case, then for a generic point $b \in U$ we would be able to solve the equation $h_D(x, b) = 0$, but then D would surject onto U , which is a contradiction. We thus conclude that $E = \pi(D)$ is a Cartier divisor cut out by some function $g_D \in \mathcal{O}_U$.

Let now $F = \pi^{-1}(E) \subset \mathcal{X}|_U$ be the scheme-theoretic preimage of E . Then F is the scheme-theoretic zero locus of the function $h_F := \pi^*g_D \in \mathcal{O}_{\mathcal{X}|_U}$, i.e. of g_D viewed as a holomorphic function on $\mathcal{X}|_U$. Thus $F \subset \mathcal{X}|_U$ is an effective Cartier divisor, such that $F|_W = D$, as schemes. Let $W' \subset W$ be the smooth locus of W . Then $(h_F^{-1}\omega)|_{W'}$ is holomorphic, and its zero locus does not contain X , since locally near X the functions h_F and h_D are equal up to multiplication by a unit in the local ring. Since $U \subset B$ is normal, it implies that $W \cong U \times \Delta$ is normal and the singular locus $W \setminus W'$ has codimension two or higher. By Hartogs' theorem, $(h_F^{-1}\omega)|_{W'}$ extends to W holomorphically, and its zero locus there still does not contain X . Since the zero locus of $h_F^{-1}\omega$ must be divisorial, it follows that $(h_F^{-1}\omega)|_{X \cap W}$ is holomorphic and not identically zero. Since $X \cap W$ contains the generic point of X , it follows that $h_F^{-1}\omega$ does not vanish identically on X . Thus h_F is the desired adjusting parameter for X , by viewing the adjusted differential $\eta = h_F^{-1}\omega$ as a section of $\pi_*(\omega_{\mathcal{X}/U}(-\mathcal{Z} + F))$.

Suppose that h is another adjusting parameter for X . Note that $h^{-1}\omega = (h_F/h)\eta$. If h_F/h has zero or pole, then $h^{-1}\omega$ would have zero or pole along X , which contradicts the definition of adjusting parameter. We thus conclude that any two adjusting parameters for X differ by multiplication by a unit in $\mathcal{O}_{B,p}$. \square

Suppose for the rest of this section that (\mathcal{X}, ω) is an adjustable family of differentials over B . Given $p \in B$, let V_p be the quotient of the set $\mathcal{O}_{B,p} \setminus \{0\}$ by the multiplicative group of units $\mathcal{O}_{B,p}^\times$. The divisibility relation induces a partial order on V_p , and we write $h_2 \preceq h_1$ if $h_1 \mid h_2$. For each fiber X_p the structure of the family near p can be encoded by decorating the dual graph Γ_p . We assign to each edge corresponding to a non-persistent node the germ $f_e \in V_p$, where $uv = f_e$ is a model for the family near the node represented by e . We assign to each vertex the function $h_v \in V_p$, where h_v is an adjusting parameter for the family at the component represented by v . We emphasize

that the functions f_e and h_v are only defined up to multiplication by a unit in the local ring.

The vertices of Γ_p have the usual partial order as defined in Section 2.4. This partial order can be understood also in terms of the divisibility relation on the set of h_v . Suppose an edge e connects two vertices v and v' . Then the edge e is horizontal if $h_v \asymp h_{v'}$, and vertical otherwise, with $v \preceq v'$ if and only if $h_v \preceq h_{v'}$. In this case, we in fact have

$$(7.4) \quad h_v/h_{v'} = f_e^{\kappa_e}$$

as shown in the proof of Theorem 4.3.

In general, the divisibility relation among the h_v may not be a full order, because for two vertices v and v' that are *not* connected by an edge, it can happen that h_v and $h_{v'}$ do not divide each other (see e.g. Example 14.16). We will be especially interested in families for which it is a full order.

Definition 7.14. An adjustable family $(\pi: \mathcal{X} \rightarrow B, \omega)$ is called *orderly* if for every point $p \in B$, the divisibility relation induces a full order on the set of adjusting parameters $(h_v)_{v \in \Gamma_p}$. \triangle

After these preparations we will now show that all the ingredients in the definition of a family of multi-scale differentials can be read off from an orderly family, except possibly missing a compatible rescaling ensemble, whose existence can be further guaranteed when the base of the family is normal.

Proposition 7.15. *An orderly family $(\pi: \mathcal{X} \rightarrow B, \omega, \mathbf{z})$ over a normal base B determines an enhanced level graph, a collection of rescaled differentials of type μ , a collection of prong-matchings for every $p \in B$ and a compatible rescaling ensemble as described in Definition 7.5. Namely, $(\mathcal{X}, \omega, \mathbf{z})$ determines a family of multi-scale differentials of type μ .*

Proof. The divisibility order of the family $(\mathcal{X}, \omega, \mathbf{z})$ gives the dual graph Γ_p the structure of a level graph, which we normalize so that the level set is \underline{N} . For each level i , we denote by h_i the adjusting parameter for some arbitrarily chosen vertex of level i .

Define the germs of holomorphic functions $s_i \in \mathcal{O}_{B,p}$ by $s_0 = h_0$ and

$$(7.5) \quad s_i := h_i/h_{i+1}$$

for all $i < 0$. For each i , define the germ of a family of differentials

$$(7.6) \quad \omega_{(i)} := \omega/(s_0 \dots s_i) = \omega/h_i$$

which is generically holomorphic and non-zero on each level i component of X_p , vanishes identically on all lower level components, and has poles along each higher level component. For an edge e of Γ_p joining levels $j < i$, the pole order of $\omega_{(j)}$ (minus one) at the corresponding node determines the enhancement κ_e . Moreover, the local normal form expressions of $\omega_{(i)}$ and $\omega_{(j)}$ as in (7.2) follow from Theorem 4.3. The u_e, v_e in the normal form can also be used to define the prong-matching $\sigma_e = du_e \otimes dv_e$ at e . We thus conclude that the $\omega_{(i)}$ give a collection of rescaled differentials of type μ at p with the s_i as rescaling parameters as in Definition 7.2.

We will show the existence of a compatible rescaling ensemble $R: B \rightarrow \overline{T}_{\Gamma_p}^n$ in three steps. First, as a consequence of Theorem 4.3, a map $R': B \rightarrow \overline{T}'_{\Gamma_p}$ can always be found by using the tuples s_i and f_e as above such that they satisfy (7.1), where \overline{T}'_{Γ_p} denotes the entire torus cut out by Equation (6.9). Next, we want the image of R' to lie in the desired connected component \overline{T}_{Γ_p} of \overline{T}'_{Γ_p} , and this can be done as follows. The torus \overline{T}'_{Γ_p} has a map to $(\mathbb{C}^*)^N$ by projection, which is an isogeny by dimension considerations. Choose in each connected component of \overline{T}'_{Γ_p} an element in the kernel of the projection. Note that modifying the tuples s_i and f_e by the chosen kernel elements does not change the rescaling parameters s_i , but it changes the f_e so that the whole collection can lie in the connected component \overline{T}_{Γ_p} . Finally, if the base B is normal, the map $R': B \rightarrow \overline{T}'_{\Gamma_p}$ automatically factors through the normalization of the level rotation torus, by the universal property of normalization, and thus gives the rescaling ensemble $R: B \rightarrow \overline{T}_{\Gamma_p}^n$. \square

8. REAL ORIENTED BLOWUPS

The goal of this section is to define a canonical real oriented blowup for a family of multi-scale differentials, generalizing the single surface case in Section 5.3. This construction will be used in Section 15, where we will show that the action of $\mathrm{SL}_2(\mathbb{R})$ extends naturally to the real oriented blowup of the moduli space of multi-scale differentials along its boundary. This blowup is also used to define families of marked multi-scale differentials in Section 8.2.

8.1. The real blowup construction. We start with the local version, which only depends on the rescaling ensemble.

Proposition 8.1. *Let $\pi: \mathcal{X} \rightarrow B$ be a germ of a family of curves with a rescaling ensemble R . Then there exists the (local) level-wise real blowup, which is a map $\widehat{\pi}: \widehat{\mathcal{X}} \rightarrow \widehat{B}$ of topological spaces with the following properties:*

- (i) *There are surjective differentiable maps $\varphi_{\mathcal{X}}: \widehat{\mathcal{X}} \rightarrow \mathcal{X}$ and $\varphi_B: \widehat{B} \rightarrow B$ such that $\pi \circ \varphi_{\mathcal{X}} = \varphi_B \circ \widehat{\pi}$.*
- (ii) *All fibers of $\widehat{\pi}$ are almost smooth surfaces in the sense of Definition 5.4.*
- (iii) *The fiber of φ_B over each point $p \in B$ is a disjoint union of tori isomorphic to $(S^1)^{L(\Gamma_p)}$.*

Moreover, the level-wise real blowup is functorial under pullbacks via maps $B' \rightarrow B$ of the base.

Note that the level-wise real blowup does not modify the neighborhoods of horizontal nodes. This is the reason for the fibers of $\widehat{\pi}$ being only almost smooth.

The fibers of φ_B are connected if π has no vertical persistent nodes, but may not be connected in general. The prong-matching singles out a specific connected component in each fiber of φ_B . We perform the above construction globally.

Theorem 8.2. *A family of multi-scale differentials (ω, σ) on $\pi: \mathcal{X} \rightarrow B$ singles out a connected component \overline{B}_p of the local level-wise real blowup \widehat{B}_p for each germ B_p . We denote by $\overline{\pi}_p: \overline{\mathcal{X}}_p \rightarrow \overline{B}_p$ the restriction of the local level-wise real blowup to \overline{B}_p .*

These germs glue to a global surjective differentiable map $\overline{\pi}: \overline{\mathcal{X}} \rightarrow \overline{B}$, the (global) level-wise real blowup. Moreover, the global level-wise real blowup is functorial under pullbacks via maps $B' \rightarrow B$ of the base.

If B is a manifold, then \overline{B} is a manifold with corners.

Our construction is closely related to a number of real oriented blowup constructions that appear in the literature, e.g. the Kato-Nakayama blowup of a log structure [KN99], see also [Kat00] and [Abr+15]. The distinguishing feature here is that the blowup is determined by the level structure of multi-scale differentials.

Proof of Proposition 8.1. The rescaling ensemble R gives a collection of rescaling and smoothing parameters $(s_i, f_e)_{i \in L(\Gamma_p), e \in E(\Gamma_p)^v}$ which are germs of functions that we may suppose to be defined on B . We introduce for each of the variables an S^1 -valued partner variable, denoted by the corresponding capital letter. Concretely, we define $\widehat{B} \subset B \times (S^1)^{E(\Gamma_p)^v} \times (S^1)^{L(\Gamma_p)}$ by the equations

$$(8.1) \quad F_e |f_e| = f_e, \quad S_i |s_i| = s_i, \quad \text{and} \quad F_e^{\kappa_e} = S_j \dots S_{i-1},$$

where κ_e is the enhancement at the edge e joining levels $j < i$ of Γ_p . Note that these equations still make sense if some s_i (or f_e) is identically zero, in which case it gives an independent variable $S_i \in S^1$ (resp. F_e). The map φ_B is given by the projection onto the first factor.

Next we define the family $\widehat{\pi}: \widehat{\mathcal{X}} \rightarrow \widehat{B}$ as follows. Near a smooth point in the fiber X_p we simply pull back a neighborhood via φ_B . In the neighborhood Y of a vertical node q_e given by the equation $u_e^+ u_e^- = f_e$, we define $\widehat{Y} \subset \varphi_B^*(Y) \times (S^1)^2$ by

$$(8.2) \quad U_e^\pm |u_e^\pm| = u_e^\pm \quad \text{and} \quad U_e^+ U_e^- = F_e.$$

The fibers of $\widehat{\pi}$ are not yet (almost) smooth (as can be seen by computing the Jacobian matrix of the defining equations), but we are in the setting of [ACG11, Section X], see in particular p. 154. There it is shown that

$$(u_e^\pm, U_e^\pm) \mapsto (|u_e^+| - |u_e^-|, |u_e^+ u_e^-|, U_e^\pm) =: (r, s, U_e^\pm)$$

is a map from a real-analytic manifold to a real-analytic manifold with corners (stemming from the boundary of the base $r = 0$) that admits an inverse which is however merely continuous. The pullback of the analytic structure on the target provides the fibers of $\widehat{\pi}$ with a smooth real analytic structure away from the horizontal nodes.

The functoriality of this construction is obvious. \square

Proof of Theorem 8.2. In view of Remark 7.6, Corollary 6.7 implies the first claim.

Suppose the germs at p and at p' are both defined at q and differ there by the action of $(r_i, \rho_e) \in T_{\Gamma_q}$. Then multiplying S_i by $r_i/|r_i|$ and F_e by $\rho_e/|\rho_e|$ provides the identification of the additional parameters of the level-wise real blowup. \square

In the special case that all nodes of π are persistent, note that the base \overline{B} is isomorphic to $B \times (S^1)^{L(\Gamma_p)}$, with parameters $\mathbf{S} = (S_i)$. Denoting by $\theta(\mathbf{S})$ the argument of the S_i , the fiber of $\overline{\pi}$ over a point $(p, \mathbf{S}) \in \overline{B}$ is simply the surface \mathcal{X}_p welded according to the prong-matching $\theta(\mathbf{S}) \cdot \boldsymbol{\sigma}$, where this map is defined in (6.2). This also justifies the use of overlines for both constructions.

8.2. Families of marked multi-scale differentials. We aim to define a marked version of families of multi-scale differentials. The general strategy is that we only mark families of almost smooth surfaces and we get rid of persistent vertical nodes by welding and of non-persistent vertical nodes using the level-wise real blowup. The following construction of marking appears also in [HK14, Section 5], for curves without a differential.

Let $\mathbf{s} \subset \Sigma$ be a collection of n points on a topological surface Σ . Let $(\pi: \mathcal{Y} \rightarrow B, \mathbf{z})$ be a pointed family of almost smooth surfaces. We define the *presheaf of markings* $\text{Mark}(\mathcal{Y}/B)$ by associating with an open set $U \subseteq B$ the set of almost-diffeomorphisms $\Sigma \times U \rightarrow \pi^{-1}(U)$ respecting the marked sections \mathbf{s} and \mathbf{z} , up to isotopies over U . A *marking f of the family $\pi: \mathcal{Y} \rightarrow B$* is a global section of the sheaf associated with $\text{Mark}(\mathcal{Y}/B)$, i.e. a compatible collection of $f_{U_i} \in \text{Mark}(\mathcal{Y}/B)(U_i)$ for sets U_i that cover B .

For any fixed subgroup G of the mapping class group $\text{Mod}_{g,n}$ we similarly define the *presheaf of G -markings* $\text{Mark}(\mathcal{Y}/B_p; G)$ by enlarging the equivalence relation (from merely isotopies) to include pre-composition of the diffeomorphisms by an element in G . A *G -marking f of π* is a global section of the sheaf associated with $\text{Mark}(\mathcal{Y}/B; G)$.

We can now define the marked version of multi-scale differentials.

Definition 8.3. A *family of marked multi-scale differentials of type (μ, Λ)* over B is an equivalence class of tuples $(\pi: \mathcal{X} \rightarrow B, \mathbf{z}, \boldsymbol{\omega}, \boldsymbol{\sigma}, f)$ consisting of

- (i) a family of multi-scale differentials $(\pi, \mathbf{z}, \boldsymbol{\omega}, \boldsymbol{\sigma})$ as in Definition 7.8 such that the enhanced level graph of every fiber is an undegeneration of Λ , and
- (ii) a Tw_Λ -marking f of the (global) level-wise real oriented blowup $\bar{\pi}: \bar{\mathcal{X}} \rightarrow \bar{B}$.

Two tuples $(\mathcal{X}, \boldsymbol{\omega}, \boldsymbol{\sigma}, f)$ and $(\mathcal{X}', \boldsymbol{\omega}', \boldsymbol{\sigma}', f')$ are called isomorphic if there is an isomorphism $g: \mathcal{X} \rightarrow \mathcal{X}'$ of multi-scale differentials such that the induced map $\bar{g}: \bar{\mathcal{X}} \rightarrow \bar{\mathcal{X}'}$ commutes with the marking, up to an isotopy respecting the marked points. \triangle

If B is a (reduced) point, then \bar{B} is the arg-image of the level rotation torus, and a marked multi-scale differential is a family of markings of the family of almost smooth-welded surfaces over \bar{B} .

Given a map $\psi: B' \rightarrow B$, the functoriality of the level-wise real blowup allows to define the pullback of markings along ψ by pulling back the family as in Section 7 and by restricting the markings along the induced map $\bar{\psi}: \bar{B}' \rightarrow \bar{B}$. In this way we have thus defined a moduli functor $\mathbf{MS}_{(\mu, \Lambda)}$ of marked multi-scale differentials. This functor has its projectivized version, denoted by $\mathbb{PMS}_{(\mu, \Lambda)}$.

There is a similar definition of the auxiliary notion of a *family of simple marked multi-scale differentials of type μ* by requiring a Tw_Λ^s -marking rather than merely a Tw_Λ -marking. The resulting functors are denoted by $\mathbf{MS}_{(\mu, \Lambda)}^s$ and $\mathbb{PMS}_{(\mu, \Lambda)}^s$.

The following propositions show that the definitions of prong-matching, marking and their pullback are sensible by showing compatibility statements with the topology we will put on the augmented Teichmüller space in Section 9. The reader may skip these propositions for now and come back later when they are used in Section 12 to show the continuity of the plumbing map, when providing the augmented Teichmüller space with a complex structure.

We work locally as in Theorem 4.3 or as in (7.2) with the normal form on a plumbing fixture $V = \{uv = f\}$ with differentials $\omega = (u^\kappa + r)du/u$ and $\eta = -(v^{-\kappa} + r/f^\kappa)dv/v$, i.e. with rescaling parameter $s = f^\kappa$. Recall from (6.3) that $F_{\mathbf{d}}$ denotes the action of fractional Dehn twists.

Proposition 8.4. *The section $\sigma = du \otimes dv$ is the unique family of prong-matchings for ω on X^+ and η on X^- over the nodal locus $\{f = 0\} \subset B$, such that for any sequence of points p_n in $\{f = 0\}$ converging to p there exist diffeomorphisms $g_n: (\overline{V_p})_{(s,f)(p_n)*\sigma} \rightarrow V_{p_n}$ that are conformal on an exhaustion $\{K_n\}$ of $V_p \setminus \{(0,0)\}$, and such that for any choice of a good arc $\gamma \subset (\overline{V_p})_\sigma$ the sequence $\tau(g_n \circ F_{\log f(p_n)} \circ \gamma) - \tau(F_{\log f(p_n)} \circ \gamma)$ converges to zero.*

We remark that the action $(s, f)(p) * \sigma$ in the above is the restriction of the one defined in Equation (6.11) to each edge. Moreover, the difference of turning numbers in the proposition is independent of the choice of the logarithm.

Proof. To show uniqueness we assume the existence of two such prong-matchings σ and σ' , take a sequence of points p_n converging to p , and let g_n and g'_n be as in the statement of the proposition. Then $h_n = g_n^{-1} \circ g'_n$ is isotopic to a fractional Dehn twist along the seam of the welding. On the other hand, turning numbers of good arcs on $(\overline{V_p})_{f(p_n)*\sigma}$ and $(\overline{V_p})_{f(p_n)*\sigma'}$ differ by an integer, so h_n must be turning-number-preserving (see Definition 5.10) for n large enough. Together this implies that $\sigma = \sigma'$ by Proposition 5.11.

To show the existence, we recall that in coordinate v the differential ω can be written as $-(f^\kappa v^{-\kappa} + r)dv/v$. Thus we have to exhibit the maps g_n for a sequence of points p_n converging to p . The main observation is that if $r(p_n) \in \mathbb{R}$, then both the line $\arg(u) = 0$ and the line $\arg(v) = \arg(f(p_n))$ are straight in the ω_{p_n} -metric. Under this additional hypothesis we construct g_n as argument-preserving maps that stretch rays as follows. We denote $a_n = |f(p_n)|$, and let g_n be the identity $(u, v) \mapsto (u, v)$ on the exhaustion $\{u : a_n^{3/4} \leq |u| < 1\} \cup \{v : a_n^{3/4} \leq |v| < 1\}$. Next, take the maps $g_n^+ : \Delta_u^* \rightarrow \{u : a_n^{1/2} < |u| \leq a_n^{3/4}\}$ and $g_n^- : \Delta_v^* \rightarrow \{v : a_n^{1/2} < |v| \leq a_n^{3/4}\}$ that preserve the argument and squeeze the radius $(0, a_n^{3/4}]$ to the interval $(a_n^{1/2}, a_n^{3/4}]$ differentiably. By the definition of prong-matching the maps g_n^\pm glue to a turning-number-preserving diffeomorphism g_n of the welded surface $(\overline{V_p})_{(s,f)(p_n)*\sigma}$ to V_{p_n} .

For general residue r we note that the above lines in the u - and v -planes are nearly straight lines. More precisely, the approximation becomes better as $r \rightarrow 0$ (which it does on X^+ since $f^\kappa \mid r$) and also as $v \rightarrow 0$ (which is relevant on X^- , see [Str84, Section 7.4] for details). Consequently, there is a diffeomorphism that nearly agrees with the one defined for $r \in \mathbb{R}$ away from the seam, and that extends differentiably over the seam of the welding given by the above prong-matching. \square

The following proposition will be used in Sections 13 and 14 to prove the universal property of the Dehn space and of the moduli space of multi-scale differentials.

Proposition 8.5. *For any family of multi-scale differentials $(\pi: \mathcal{X} \rightarrow B, \mathbf{z}, \omega, \sigma)$ and any $p \in B$ there exists a neighborhood U of p such that $\pi|_U$ can be provided with a Tw_Λ -marking f where Λ is a multicurve such that $\Gamma(\Lambda)$ is degeneration of Γ_p .*

Proof. We need to provide the level-wise real blowup $\bar{\pi}|_U$ with a Tw_Λ -marking f . For this purpose we take U to be simply connected, provide some fiber of $\bar{\pi}$ with a marking and transport the marking along local smooth trivializations of $\bar{\pi}$. We only need to make sure that the monodromy in this process is contained in Tw_Λ . By the choice of U , and since by Theorem 8.2 the fibers of $\bar{U} \rightarrow U$ are (arg-images of) level rotation tori, the monodromy is generated by level rotation. From the definition of level rotation tori at the beginning of Section 6.3, it is now obvious that the monodromy is Tw_Λ . \square

9. AUGMENTED TEICHMÜLLER SPACE OF FLAT SURFACES

The aim of this section is to construct an augmented Teichmüller space of flat surfaces $\Omega\bar{\mathcal{T}}_{(\Sigma, \mathbf{s})}(\mu)$ as a topological space, in a way that closely parallels the construction of the classical case of the augmented Teichmüller space of curves without differentials, as done in Section 3. As a set, $\Omega\bar{\mathcal{T}}_{(\Sigma, \mathbf{s})}(\mu)$ is the disjoint union of the Teichmüller spaces of prong-matched twisted differentials for all enhanced multicurves Λ^+ , projectivized level by level.

9.1. Augmented Teichmüller space of flat surfaces as a set. Given an enhanced multicurve Λ^+ , we define the Λ^+ -boundary stratum as the quotient $\Omega\mathcal{B}_{\Lambda^+} = \Omega\mathcal{T}_{\Lambda^+}^{pm}(\mu)/\mathbb{C}^{L(\Lambda^+)}$ where the action of $\mathbb{C}^{L(\Lambda^+)}$ was defined in Section 6.1. We define the projectivized Λ^+ -boundary stratum to be the quotient $\mathbb{P}\Omega\mathcal{B}_{\Lambda^+} = \Omega\mathcal{T}_{\Lambda^+}^{pm}(\mu)/\mathbb{C}^{L^\bullet(\Lambda^+)}$.

As a set, the (projectivized) augmented Teichmüller space $\Omega\bar{\mathcal{T}}_{(\Sigma, \mathbf{s})}(\mu)$ of flat surfaces of type μ is defined to be the disjoint union of these (projectivized) boundary strata over all enhanced multicurves:

$$\Omega\bar{\mathcal{T}}_{(\Sigma, \mathbf{s})}(\mu) = \coprod_{\Lambda^+} \Omega\mathcal{B}_{\Lambda^+} \quad \text{and} \quad \mathbb{P}\Omega\bar{\mathcal{T}}_{(\Sigma, \mathbf{s})}(\mu) = \coprod_{\Lambda^+} \mathbb{P}\Omega\mathcal{B}_{\Lambda^+}.$$

Note that this union contains, for the trivial multicurve $\Lambda^+ = \emptyset$, the classical Teichmüller space $\Omega\mathcal{B}_\emptyset = \Omega\mathcal{T}_\emptyset^{pm}(\mu) = \Omega\mathcal{T}_{(\Sigma, \mathbf{s})}(\mu)$ that parameterizes marked flat surfaces of type μ . The mapping class group $\text{Mod}_{g, n}$ acts on $\Omega\bar{\mathcal{T}}_{(\Sigma, \mathbf{s})}(\mu)$ and $\mathbb{P}\Omega\bar{\mathcal{T}}_{(\Sigma, \mathbf{s})}(\mu)$ by pre-composition of the marking. To keep notation manageable, from now on we continue to drop the superscript $+$, writing simply Λ instead of Λ^+ , while all multicurves throughout this section are taken together with an enhancement.

The “marked” version of Remark 7.6 is:

Proposition 9.1. *A point in $\Omega\bar{\mathcal{T}}_{(\Sigma, \mathbf{s})}(\mu)$ parameterizes (the equivalence class of) a marked multi-scale differential.*

Proof. A point (X, ω, σ, f) in $\Omega\bar{\mathcal{T}}_{(\Sigma, \mathbf{s})}(\mu)$ contains the information of a prong-matched twisted differential, which defines a multi-scale differential (X, ω, σ) . We can take a representative of the marking f to mark the fiber $\bar{\pi}: \bar{X} \rightarrow T_\Lambda$ determined by the prong-matching of the level-wise real blow up and propagate it using the action (6.11) of T_Λ to a marking of each fiber. Since the fundamental group of T_Λ acts by elements in the twist group, this is indeed a Tw_Λ -marking as required in Definition 8.3. This induces a bijection since $T_\Lambda = \mathbb{C}^{L^\bullet(\Lambda^+)}/\text{Tw}_\Lambda$. \square

The following proposition follows directly from the definitions.

Proposition 9.2. *The subgroup of $\text{Mod}_{g,n}$ fixing the boundary stratum $\Omega\mathcal{B}_\Lambda$ pointwise is exactly the twist group Tw_Λ . Moreover, if Λ' is a degeneration of Λ , then the twist group Tw_Λ fixes the boundary stratum $\Omega\mathcal{B}_{\Lambda'}$ pointwise. Both statements hold as well for the projectivizations.*

Proof. For the first statement, note that the subgroup of $\text{Mod}_{g,n}$ fixing the boundary stratum $\Omega\mathcal{B}_\Lambda$ pointwise is generated by Dehn twists along the curves of Λ , that preserve the prong-matchings. Hence by definition it can be identified with Tw_Λ .

For the second statement, if Λ' is a degeneration of Λ , then $\text{Tw}_\Lambda \subset \text{Tw}_{\Lambda'}$. Hence Tw_Λ fixes $\Omega\mathcal{B}_{\Lambda'}$ pointwise. \square

9.2. Augmented Teichmüller space of flat surfaces as a topological space. We now give both augmented Teichmüller spaces $\Omega\overline{\mathcal{T}}_{(\Sigma,s)}(\mu)$ and $\mathbb{P}\Omega\overline{\mathcal{T}}_{(\Sigma,s)}(\mu)$ a topology. We give a sequential definition of the topology first (see e.g. [BJ06, Section I.8.9] for a precise discussion of defining a topology in this way). In the proofs below we also give a definition by specifying a basis of the topology. Unless stated otherwise, ω in the tuple $(X, z, \omega, \preceq, \sigma, f)$ refers to a chosen representative of the equivalence class. We also write \overline{X}_{σ_n} as shorthand for $(\overline{X}_n)_{\sigma_n}$.

Definition 9.3. A sequence $(X_n, z_n, \omega_n, \preceq_n, \sigma_n, f_n) \in \mathbb{P}\Omega\mathcal{B}_{\Lambda_n}$ converges to a point $(X, z, \omega, \preceq, \sigma, f) \in \mathbb{P}\Omega\mathcal{B}_\Lambda \subset \mathbb{P}\Omega\overline{\mathcal{T}}_{(\Sigma,s)}(\mu)$, if there exist representatives (that we denote with the same symbols) in $\Omega\mathcal{T}_{\Lambda_n}^{pm}(\mu)$ and $\Omega\mathcal{T}_\Lambda^{pm}(\mu)$, a sequence of positive numbers ϵ_n converging to 0, and a sequence of vectors $\mathbf{d}_n = \{d_{n,i}\}_{i \in L^\bullet(\Lambda)} \in \mathbb{C}^{L^\bullet(\Lambda)}$ such that the following conditions hold, where we denote $c_{n,i} = e(d_{n,i})$:

- (1) For sufficiently large n there is an undegeneration of enhanced multicurves (δ_n, D_n^h) with $\delta_n: L^\bullet(\Lambda) \rightarrow L^\bullet(\Lambda_n)$ (see Definition 5.1).
- (2) For sufficiently large n there exists an almost-diffeomorphism $g_n: \overline{X}_{\mathbf{d}_n \cdot \sigma} \rightarrow \overline{X}_{\sigma_n}$ that is compatible with the markings (in the sense that f_n is isotopic to $g_n \circ (\mathbf{d}_n \cdot f)$ rel marked points) and that is conformal on the ϵ_n -thick part $(X, z)_{\epsilon_n}$.
- (3) The restriction of $c_{n,i} g_n^*(\omega_n)$ to the ϵ_n -thick part of the level i subsurface of (X, z) converges uniformly to $\omega_{(i)}$.
- (4) For any $i, j \in L^\bullet(\Lambda)$ with $i > j$, and any subsequence along which $\delta_n(i) = \delta_n(j)$, we have

$$\lim_{n \rightarrow \infty} \frac{|c_{n,i}|}{|c_{n,j}|} = 0.$$

- (5) The almost-diffeomorphisms g_n are nearly turning-number-preserving, i.e. for every good arc γ in \overline{X}_σ , the difference of turning numbers $\tau(g_n \circ F_{\mathbf{d}_n} \circ \gamma) - \tau(F_{\mathbf{d}_n} \circ \gamma)$ converges to zero, where $F_{\mathbf{d}_n}$ is the fractional Dehn twist defined in (6.3).

For convergence in $\Omega\overline{\mathcal{T}}_{(\Sigma,s)}(\mu)$, we require moreover that $c_{n,0} \rightarrow 1$ for the rescaling parameter corresponding to the top level of Λ . \triangle

Note that the notion of convergence does not depend on the choice of representatives in $\Omega\mathcal{T}_\Lambda^{pm}(\mu)$ since if $X' = \mathbf{d}' \cdot X$ is another representative, then using $\mathbf{d}' + \mathbf{d}$ certifies convergence to X' . Note moreover, that in item (5) we could as well require the difference of winding numbers to be bounded for a fixed collection of arcs dual to the collection of seams. In fact if γ_1 and γ_2 are two homotopic good arcs, then the difference

of the turning numbers $\tau(g_n \circ \gamma_i) - \tau(\gamma_i)$ for $i = 1, 2$ is bounded by the angular distortion between $g_n^* \omega_n$ and ω on X_{ϵ_n} , which is uniformly bounded by (3).

As mentioned after the proof of Theorem 3.2, one obtains an equivalent notion of convergence if we use the ϵ -thick part of X punctured in at least one point on each irreducible component.

Theorem 9.4. *The augmented Teichmüller space of flat surfaces $\Omega\overline{\mathcal{T}}_{(\Sigma, s)}(\mu)$ and its projectivized version $\mathbb{P}\Omega\overline{\mathcal{T}}_{(\Sigma, s)}(\mu)$ are Hausdorff topological spaces.*

Proof. Instead of checking that the above definition of convergence defines a convergence class of sequences in the sense of [BJ06, Definition I.8.11], we quantify distance by some parameter ϵ and check that the resulting sets $V_\epsilon(X)$ are the open sets in a basis of this topology. More precisely, we define $V_\epsilon(X)$ to consist of $(X_0, z_0, \omega_0, \preceq_0, \sigma_0, f_0)$ such that there exists $\mathbf{d} = \{d_i\}_{i \in L \bullet(\Lambda)} \in \mathbb{C}^{L \bullet(\Lambda)}$ and

- (i) a degeneration $(\Lambda_0, \preceq_0) \rightsquigarrow (\Lambda, \preceq)$ given by a map $\delta: L(\Lambda) \rightarrow L(\Lambda_0)$ and a subset of N_X^h ,
- (ii) an almost-diffeomorphism $g: \overline{X}_{\mathbf{d}, \sigma} \rightarrow \overline{X}_{\sigma_0}$ conformal on the ϵ -thick part of X and compatible with the markings,
- (iii) letting $c_i = e(d_i)$, the bound $\|c_i g^* \omega_{0(i)} - \omega_{(i)}\|_\infty < \epsilon$ holds on the ϵ -thick part of X ,
- (iv) for any $i, j \in L(\Lambda)$ with $i > j$ and $\delta(i) = \delta(j)$, we have $|c_i/c_j| < \epsilon$,
- (v) for any good arc γ in \overline{X}_σ crossing exactly one seam once

$$|\tau(g \circ F_{\mathbf{d}} \circ \gamma) - \tau(F_{\mathbf{d}} \circ \gamma)| < \epsilon.$$

Suppose that $X_0 \in V_\epsilon(X)$. To check that the sets defined above are a basis of topology we want to find ρ such that $V_\rho(X_0) \subset V_\epsilon(X)$. Suppose that $X_1 \in V_\rho(X_0)$. Let $g_0: \overline{X}_{\mathbf{d}, \sigma} \rightarrow \overline{X}_{0, \sigma_0}$ and $g: \overline{X}_{0, \mathbf{d}_0, \sigma_0} \rightarrow \overline{X}_{1, \sigma_1}$ be the almost-diffeomorphisms given by the definitions of $V_\epsilon(X)$ and $V_\rho(X_0)$. We define $\mathbf{d}_1 = \mathbf{d} + \mathbf{d}_0$, denote the rescaling parameters by $\mathbf{c}_0 = e(\mathbf{d}_0)$ and $\mathbf{c} = e(\mathbf{d})$, and set $\mathbf{c}_1 = \mathbf{c}_0 \cdot \mathbf{c}$.

First we remark that item (i) is automatically satisfied in $V_\epsilon(X)$. For item (ii), choose ρ small enough such that the ρ -thick part of X_0 contains the g_0 -image of the ϵ -thick part of X and let $g_1 = g \circ F_{\mathbf{d}_0} \circ g_0 \circ F_{\mathbf{d}_0}^{-1}: \overline{X}_{\mathbf{d}_1, \sigma} \rightarrow \overline{X}_{1, \sigma_1}$. Then g_1 clearly satisfies (ii).

For (iii), to simplify notation and illustrate the main idea, we treat the case that X_1 is smooth and that X and X_0 have the same level graph with two levels. Moreover, we assume that all rescaling parameters on top level are equal to one and we denote the rescaling parameters on lower level by c, c_0 and $c_1 = cc_0$. The general case follows by the same idea. Under these assumptions, let ω^- and ω_0^- be differentials on the lower level of X and X_0 respectively. By assumption, the norm $\epsilon' = \|c_0 g_0^* \omega_0^- - \omega^-\|_\infty$ satisfies $\epsilon' < \epsilon$. We estimate the sup-norms on the lower level subsurface of the ϵ -thick part of X as follows:

$$\begin{aligned} \|c_1 g_1^* \omega_1 - \omega^-\|_\infty &\leq \|c_1 g_0^* g^* \omega_1 - c_0 g_0^* \omega_0^-\|_\infty + \|c_0 g_0^* \omega_0^- - \omega^-\|_\infty, \\ &< c_0 \|g_0^* (c g^* \omega_1 - \omega_0^-)\|_\infty + \epsilon' \\ &< c_0 C_{g_0} \rho + \epsilon'. \end{aligned}$$

Here C_{g_0} is a constant, depending only on g_0 , which is an upper bound for the change of the sup-norm under the pullback by g_0 on the ϵ -thick part. Note that the ϵ -thick part is compact, and thus this stretching under the pullback by g_0 is globally bounded over it. We now take ρ small enough so that $c_0 C_{g_0} \rho + \epsilon' < \epsilon$. This shows that item (iii) holds for $X_1 \in V_\epsilon(X)$.

Item (iv) follows if we moreover choose ρ such that $\rho < \epsilon c_0$. Finally, item (v) follows from the triangle inequality for the turning numbers. Consequently, the $V_\epsilon(X)$ are indeed a basis of a topology. It is obvious that in this topology a sequence converges if and only if the items (1)–(5) hold.

In order to show that the augmented Teichmüller space $\Omega\overline{\mathcal{T}}_{(\Sigma, \mathbf{s})}(\mu)$ is Hausdorff, suppose that a sequence $(X_n, \mathbf{z}_n, \omega_n, \preceq_n, \boldsymbol{\sigma}_n, f_n)$ converges to both $(X, \mathbf{z}, \omega, \preceq, \boldsymbol{\sigma}, f)$ and $(X', \mathbf{z}', \omega', \preceq', \boldsymbol{\sigma}', f')$. Forgetting all but the underlying pointed stable curves, our topology gives the conformal topology on the Deligne-Mumford compactification, which is a Hausdorff space, so the limit of the maps $(g'_n)^{-1} \circ g_n$ must define an isomorphism $h: (X, \mathbf{z}) \rightarrow (X', \mathbf{z}')$ of pointed stable curves. We next show that \preceq and \preceq' are the same (weak) full order. Suppose, for contradiction, that there exist irreducible components X_u and X_v of X such that $X_u \succ X_v$ but $X_u \preceq' X_v$. Since \preceq_n , for n sufficiently large, is an undegeneration of both \preceq and \preceq' , this is possible only if $X_u \asymp_n X_v$. We denote ℓ and ℓ' some level functions inducing the full orders \preceq and \preceq' , respectively. The specific choices of these level functions are not important, as we will only use them to match notation. Then condition (3) of convergence of sequences implies that $\|c_{n, \ell(u)} g_n^* \omega_n - \omega_u\|_\infty < \epsilon_n$ and $\|c'_{n, \ell'(u)} (g'_n)^* \omega_n - \omega'_u\|_\infty < \epsilon_n$, where ω_u is the restriction of ω to the ϵ_n -thick part of X_u . Pulling back the second inequality by h and choosing ϵ_n small enough, these conditions imply that the ratios $c_{n, \ell(u)} / c'_{n, \ell'(u)}$ are bounded away from zero and infinity. Similarly, the same holds for $c_{n, \ell(v)} / c'_{n, \ell'(v)}$. However, condition (4) of convergence implies that $|c_{n, \ell(u)} / c_{n, \ell(v)}| \rightarrow 0$, while on the other hand the hypothesis $X_u \preceq' X_v$ implies that (after possibly passing to a subsequence) $c'_{n, \ell'(u)} / c'_{n, \ell'(v)}$ is bounded away from zero. Combining these inequalities yields a contradiction.

To verify that the form ω is equal to ω' , we use that for every level i both inequalities $\|c_{n, i} g_n^*(\omega_n) - \omega_{(i)}\|_\infty < \epsilon_n$ and $\|c'_{n, i} g_n^*(\omega_n) - h^* \omega'_{(i)}\|_\infty < C \epsilon_n$ hold for some constant C that depends on the map h but not on n . We multiply the second inequality by $c_{n, i} / c'_{n, i}$, use that this quantity is bounded away from zero and infinity, and thus deduce that $\|c_{n, i} / c'_{n, i} \cdot h^* \omega'_{(i)} - \omega_{(i)}\|_\infty$ tends to zero on the ϵ_n -thick part of $X_{(i)}$. This implies the convergence of the sequence $c_{n, i} / c'_{n, i}$ for each i , and also the equivalence as projectivized differentials, which is what we aimed for.

The fact that h extends to a diffeomorphism $\bar{h}: \overline{X}_\sigma \rightarrow \overline{X}'_{\sigma'}$ is mainly the content of Proposition 5.11. We apply this proposition to the nearly turning-number-preserving map $h_n = F_{d_n}^{-1} \circ (g'_n)^{-1} \circ g_n \circ F_{d_n}$, see Remark 5.12. Consequently, the prong-matchings agree. Note that the proof of that proposition shows that \bar{h} is isotopic to h_n for n large enough. By definition there are $D_n, D'_n \in \text{Tw}_\Lambda$ such that $f \circ D_n$ is isotopic to $g_n^{-1} \circ f_n$ and such that $f' \circ D'_n$ is isotopic to $(g'_n)^{-1} \circ f'_n$. Together this shows that $h \circ f$ is isotopic

to $f' \circ D'_n \circ D_n^{-1}$, which implies that the markings agree, up to isotopy and the action of the twist group. \square

10. THE MODEL DOMAIN

In Section 12 we will provide quotients of the augmented Teichmüller space with a complex structure. For this purpose we now construct a model domain that has a complex structure by its very definition, and exhibit it as the moduli space of some auxiliary objects that we call model differentials. We will recall the universal properties and construct (in a straightforward way) the universal family over the model spaces as a starting point of the plumbing construction in Section 12.

10.1. Model differentials. In this subsection, we define the notion of model differentials of type μ closely parallel to Section 7, and then its marked version as in Section 8. Families of model differentials are constrained to be equisingular, but as a trade off they carry for each level an additional parameter t_i that is allowed to be zero, thus mimicking degenerations. Equisingularity is also the reason that we do not need to start with a germwise definition, but can start right away with the global definition.

While multi-scale differentials are based on a collection of rescaled differentials, the simpler notion of model differential is based on the simple notion of twisted differentials. We adapt the definition from Section 2.4 to families.

Definition 10.1. A *family of twisted differentials* η of type μ on an equisingular family $\pi: \mathcal{X} \rightarrow B$ of pointed stable curves compatible with Γ is a collection of families of meromorphic differentials $\eta_{(i)}$ on the subcurve $\mathcal{X}_{(i)}$ at level i , which satisfies obvious analogues the conditions in Section 2.4, interpreting the residues as regular functions on the base B . \triangle

For the construction below the “simple” version is central in order to get fine and smooth moduli spaces, so we start with that.

Definition 10.2. Let $(\pi: \mathcal{X} \rightarrow B, \mathbf{z})$ be an equisingular family of pointed stable curves. A *family of (unmarked) simple model differentials of type μ over B* consists of the following data:

- (1) the structure of an enhanced level graph on the dual graph Γ of any fiber of π ,
- (2) a simple rescaling ensemble $R^s: B \rightarrow \overline{T}_\Gamma^s$,
- (3) a collection $\eta = (\eta_{(i)})_{i \in L^\bullet(\Gamma)}$ of families of twisted differentials of type μ compatible with Γ ,
- (4) a collection $\sigma = (\sigma_e)_{e \in E(\Gamma)^v}$ of prong-matchings for η . \triangle

The simple level rotation torus $T_{\Gamma_p}^s(\mathcal{O}_B)$ acts on all of the above data, and we consider the data (η, R^s, σ) to be *equivalent* to $(\mathbf{q} * \eta, \mathbf{q}^{-1} * R^s, \mathbf{q} * \sigma)$ for any $\mathbf{q} \in T_{\Gamma_p}^s(\mathcal{O}_B)$. Here the torus action is defined in (6.12), and $\mathbf{q}^{-1} * (\)$ denotes post-composition with the multiplication by \mathbf{q}^{-1} .

Note that the prong-matchings can be defined equivalently in the form given in Section 5.4 or in Section 7.1.

Remark 10.3. As pointed out in Section 7.1, a simple rescaling ensemble is simply a collection of functions $\mathbf{t} = (t_i)_{i \in L(\Gamma)}$ in \mathcal{O}_B . In this setting, the action of $\mathbf{q} \in T_{\Gamma_p}^s$ on \mathbf{t} is simply given by t_i/q_i for each i in $L(\Gamma)$. Hence from this point on, we will denote a family of simple model differentials interchangeably by $(\boldsymbol{\eta}, R^s, \boldsymbol{\sigma})$ and by $(\boldsymbol{\eta}, \mathbf{t}, \boldsymbol{\sigma})$.

We now define the *level-wise real blowup* for families of unmarked simple model differentials over an analytic space B . The construction for unmarked model differentials is analogous to Section 8.1 and we briefly describe it. For each edge e of Γ we introduce an S^1 -valued variable F_e and for each level $i \in L(\Gamma)$ we introduce another variable $T_i \in S^1$. Then we define $\widehat{B} \subset B \times (S^1)^{E(\Gamma)^v} \times (S^1)^{L(\Gamma)}$ by

$$(10.1) \quad T_i^{a_i} |t_i^{a_i}| = t_i^{a_i} \quad \text{and} \quad F_e^{k_e} = T_j^{a_j} \dots T_{i-1}^{a_{i-1}},$$

where the edge e connects levels $j < i$. The family $\widehat{\mathcal{X}}$ is defined similarly as in Section 8.1 and the *welded family* $\overline{\pi}: \overline{\mathcal{X}}_\sigma \rightarrow \overline{B}$ is the component of $\widehat{\mathcal{X}} \rightarrow \widehat{B}$ singled out by the prong-matching $\boldsymbol{\sigma}$.

We now define families of simple model differentials by adding the datum of a marking to families of unmarked simple model differentials.

Definition 10.4. Let $(\pi: \mathcal{X} \rightarrow B, \mathbf{z})$ be an equisingular family of pointed stable curves with enhanced dual graph Γ . A *family of (marked) simple model differentials of type μ compatible with Γ* is

- (i) a tuple $(\boldsymbol{\eta}, \mathbf{t}, \boldsymbol{\sigma})$ as in Definition 10.2, and
- (ii) a Tw_Λ^v -marking f of the welded family $\overline{\pi}: \overline{\mathcal{X}}_\sigma \rightarrow \overline{B}$. \triangle

Again the simple level rotation torus $T_{\Gamma_p}^s(\mathcal{O}_B)$ acts on all of the above data, and we consider the data $(\boldsymbol{\eta}, R^s, \boldsymbol{\sigma}, f)$ to be *equivalent* to $(\mathbf{q} * \boldsymbol{\eta}, \mathbf{q}^{-1} * R^s, \mathbf{q} * \boldsymbol{\sigma}, \mathbf{q} * f)$ for any $\mathbf{q} \in T_{\Gamma_p}^s(\mathcal{O}_B)$. Here the action $\mathbf{q} * f$ is defined in (6.12). Analogously, we define the *families of simple projectivized model differentials* with the enlarged equivalence relation given by the extended simple level rotation torus $T_\Gamma^{s, \bullet}$. We denote the functor of simple model differentials by $\mathbf{MD}_{(\mu, \Lambda)}^s$ and its projectivized version by $\mathbb{P}\mathbf{MD}_{(\mu, \Lambda)}^s$.

We finally define the families of *model differentials*.

Definition 10.5. Let $(\pi: \mathcal{X} \rightarrow B, \mathbf{z})$ be an equisingular family of pointed stable curves with enhanced dual graph Γ . A *family of model differentials of type μ compatible with Γ* is a family of simple model differentials where item (2) of definition 10.2 is replaced by

- (2') a rescaling ensemble $R: B \rightarrow \overline{T}_\Gamma$. \triangle

Two families are *equivalent* if they lie in the same orbit of the action of the level rotation torus. The functor of model differentials will be denoted by $\mathbf{MD}_{(\mu, \Lambda)}$.

10.2. The universal family over $\Omega\mathcal{T}_\Lambda^{pm}(\mu)$. Our goal is to exhibit the functor that the space $\Omega\mathcal{T}_\Lambda^{pm}(\mu)$ represents, and to introduce notation for the corresponding universal family. Throughout this section, we denote by Λ an *enhanced* multicurve. The following definition extends to families the pointwise definition that appeared already in Section 5.4, using the notion of markings in families that is now at our disposal.

Definition 10.6. An equisingular family of prong-matched twisted differentials of type (μ, Λ) over an analytic space B is

- (i) a family $(\eta_v)_{v \in V(\Gamma)}$ of twisted differentials of type μ , compatible with $\Gamma(\Lambda)$ as in Definition 10.1,
- (ii) a family of prong-matchings σ , and
- (iii) a family of markings $f \in \text{Mark}(\overline{\mathcal{X}}_\sigma/\overline{B})$ of the welded family. \triangle

Note that the real oriented blowup of item (iii) is a bit different than the one constructed in Section 8. Indeed in this case, we do not have a rescaling ensemble at our disposition. But the construction is the same, putting $f_e = s_i = 0$ and $S_i = 1$ in Equation (8.1). In particular $\overline{B} \cong B$ for a family of prong-matched twisted differentials, i.e. for an equisingular family.

We can now state the universal property of the space $\Omega\mathcal{T}_\Lambda^{pm}(\mu)$.

Proposition 10.7. Let $\Lambda \subset \Sigma$ be a fixed enhanced multicurve. The Teichmüller space of prong-matched twisted differentials $\Omega\mathcal{T}_\Lambda^{pm}(\mu)$ is the fine moduli space for the functor that associates to an analytic space B the set of equisingular families of prong-matched twisted differentials of type (μ, Λ) .

The main purpose in giving the rather obvious proof is to introduce names for the corresponding universal objects. The universal property then follows immediately from the well-known universal properties of these intermediate objects.

Proof. An equisingular family of pointed stable curves defines, by normalization, a collection of families of pointed smooth curves with additional marked sections corresponding to the branches of the nodes. Conversely, such a collection of families of smooth pointed curves and a pairing of a subset of the marked sections defines an equisingular family. From this observation it is obvious that the boundary stratum \mathcal{T}_Λ of the classical augmented Teichmüller space comes with a universal family $(\pi: \mathcal{X} \rightarrow \mathcal{T}_\Lambda, \mathbf{z}, (f_v)_{v \in V(\Lambda)})$ of pointed stable curves equisingular of type $\Gamma(\Lambda)$, constructed by gluing families of smooth curves $\pi: \mathcal{X}_v \rightarrow \mathcal{T}_\Lambda$ along the nodes given by the marked sections q_e^\pm corresponding to the edges e of $\Gamma(\Lambda)$. Here $f_v \in \text{Mark}(\mathcal{X}_v/\mathcal{T}_\Lambda)$ is a Teichmüller marking by the surface Σ_v (corresponding to the component $v \in V(\Gamma)$ of $\Sigma \setminus \Lambda$, with the boundary curves contracted to points). The universal property follows from the universal properties for the Teichmüller spaces of the pieces $(\mathcal{X}_v, \mathbf{z}_v, \mathbf{q}_e^\pm, f_v)$.

Recall from Section 5.2 that there is a closed subspace $\mathcal{T}_\Lambda(\mu) \subset \mathcal{T}_\Lambda$ defined to be the quotient of $\Omega^{no}\mathcal{T}_\Lambda(\mu)$ under the action of $(\mathbb{C}^*)^{V(\Lambda)}$. The family π can be restricted to $\mathcal{T}_\Lambda(\mu)$, pulled back to $\Omega^{no}\mathcal{T}_\Lambda(\mu)$, and then restricted to $\Omega\mathcal{T}_\Lambda(\mu)$. Since the total space of a vector bundle represents the functor of sections of the bundle, $\Omega\mathcal{T}_\Lambda(\mu)$ comes with a universal family $(\pi: \mathcal{X} \rightarrow \mathcal{T}_\Lambda, \mathbf{z}, (f_v)_{v \in V(\Lambda)}, (\eta_v)_{v \in V(\Lambda)})$, where $\boldsymbol{\eta} = (\eta_v)_{v \in V(\Lambda)}$ is a twisted differential of type (μ, Λ) , and the remaining data are as above.

Now we construct the family of markings in the welded surfaces $\overline{\pi}: \overline{\mathcal{X}}_\sigma \rightarrow \overline{\Omega\mathcal{T}_\Lambda^{pm}(\mu)} \cong \Omega\mathcal{T}_\Lambda^{pm}(\mu)$. Then we mark the welded surfaces by Σ such that fiberwise after pinching Λ we obtain the collection $(f_v)_{v \in V(\Lambda)}$. The remaining data are the pullbacks of the ones we defined above. Since $\Omega\mathcal{T}_\Lambda^{pm}(\mu) \rightarrow \Omega\mathcal{T}_\Lambda(\mu)$ is an (infinite) covering map (see Section 5.5), the universal property follows from the universal properties of covering spaces. \square

10.3. The model domain and its universal property. In this section we construct the *model domain* $\overline{\mathcal{M}D}_\Lambda$ representing the functor of model differentials, together with its smooth cover $\overline{\mathcal{M}D}_\Lambda^s$, the *simple model domain*, whose superscript s recalls both that it is smooth and that all the objects are defined with respect to the simple twist group Tw_Λ^s , instead of the full twist group Tw_Λ . As a first step in the construction, we define

$$(10.2) \quad \mathcal{M}D_\Lambda^s = \mathbb{P}\Omega\mathcal{T}_\Lambda^{pm}(\mu)/\text{Tw}_\Lambda^{sv} \quad \text{and} \quad \mathcal{M}D_\Lambda = \mathbb{P}\Omega\mathcal{T}_\Lambda^{pm}(\mu)/\text{Tw}_\Lambda^v.$$

The simple level rotation torus $T_\Lambda^s = \mathbb{C}^{L(\Lambda)}/\text{Tw}_\Lambda^{sv}$ acts freely on $\mathcal{M}D_\Lambda^s$ via the action defined in (6.12). Recall from Section 9.1 that the projectivized Λ -boundary stratum $\mathbb{P}\Omega\mathcal{B}_\Lambda$ is the quotient $\mathcal{M}D_\Lambda^s/T_\Lambda^s$, so $\mathcal{M}D_\Lambda^s$ is a principal $(\mathbb{C}^*)^{L(\Lambda)}$ -bundle over $\mathbb{P}\Omega\mathcal{B}_\Lambda$. We define $\overline{\mathcal{M}D}_\Lambda^s$ as the associated $\mathbb{C}^{L(\Lambda)}$ -bundle over $\mathbb{P}\Omega\mathcal{B}_\Lambda$.

We now provide notation to describe the boundary $\partial\mathcal{M}D_\Lambda^s = \overline{\mathcal{M}D}_\Lambda^s \setminus \mathcal{M}D_\Lambda^s$ of the simple model domain. The boundary $\partial\mathcal{M}D_\Lambda^s$ is a normal crossing divisor given by $D = \cup_{i \in L(\Lambda)} D_i$ in $\overline{\mathcal{M}D}_\Lambda^s$, where D_i is fiber-wise defined by $\{t_i = 0\} \subset \mathbb{C}^{L(\Lambda)}$. Note that Λ may have horizontal nodes, but they play no role in the description of the boundary of $\overline{\mathcal{M}D}_\Lambda^s$. There is a stratification

$$(10.3) \quad \overline{\mathcal{M}D}_\Lambda^s = \coprod_{J \subset L(\Lambda)} \mathcal{M}D_\Lambda^{s, \Lambda J}$$

indexed by the vertical undegenerations of the multicurve Λ or, equivalently, by the subsets $J = \{j_1, \dots, j_m\}$ of $L(\Lambda)$ (see Section 5.1 for the correspondence), where the space $\overline{\mathcal{M}D}_\Lambda^{s, \Lambda J}$ is defined to be $D_{j_1} \cap \dots \cap D_{j_m}$, and $\mathcal{M}D_\Lambda^{s, \Lambda J}$ is simply given by

$$\mathcal{M}D_\Lambda^{s, \Lambda J} = \overline{\mathcal{M}D}_\Lambda^{s, \Lambda J} \setminus \bigcup_{J' \supsetneq J} \overline{\mathcal{M}D}_\Lambda^{s, \Lambda J'}.$$

The reader can check that the spaces $\mathcal{M}D_\Lambda^{s, \Lambda J}$ form a stratification of $\overline{\mathcal{M}D}_\Lambda^s$. Note that in these terms, $\mathcal{M}D_\Lambda^s$ corresponds to $J = \emptyset$, or equivalently to the degeneration $\text{dg}: \bullet \rightsquigarrow \Lambda$ from the trivial graph to Λ . Moreover, D_i corresponds to the subset $J = \{i\}$, or equivalently to the two-level (un)degeneration $\text{dg}_i: \Lambda_i \rightsquigarrow \Lambda$ of Λ .

Recall from Lemma 6.3 that Tw_Λ^{sv} is a finite index subgroup of Tw_Λ^v . The finite quotient group $K_\Lambda = \text{Tw}_\Lambda^v/\text{Tw}_\Lambda^{sv}$ defined in Equation (6.13) acts on $\mathcal{M}D_\Lambda^s$ with quotient $\mathcal{M}D_\Lambda$, commuting with the T_Λ^s -action. Consequently, K_Λ acts on $\overline{\mathcal{M}D}_\Lambda^s$ and we define

$$\overline{\mathcal{M}D}_\Lambda = \overline{\mathcal{M}D}_\Lambda^s/K_\Lambda.$$

The smoothness of $\mathbb{P}\Omega\mathcal{B}_\Lambda$ and this description imply immediately the following result.

Proposition 10.8. *The simple model domain $\overline{\mathcal{M}D}_\Lambda^s$ is smooth, while $\overline{\mathcal{M}D}_\Lambda$ has only finite quotient singularities.*

Example 10.9. *(A model domain with finite quotient singularities)* To see that finite quotient singularities can actually occur in this way, we analyze the second case of our running example in Section 2.6. There, $\overline{\mathcal{M}D}_{\Lambda_2}^s$ is locally the product of $\mathbb{P}\Omega\mathcal{B}_{\Lambda_2}$ with \mathbb{C}^2 , and the two boundary divisors are the coordinate axes. The generators of the level-wise ramification groups H_1 and H_2 (see Section 6.4 and Example 6.8) act

on the \mathbb{C}^2 factor by $(z_1, z_2) \mapsto (\zeta_3 z_1, z_2)$ and $(z_1, z_2) \mapsto (z_1, \zeta_3 z_2)$ respectively, where ζ_3 is a third root of unity. Consequently, the generator of $K_{\Lambda_2} = \text{Ker}(H \rightarrow G)$ acts by $(z_1, z_2) \mapsto (\zeta_3 z_1, \zeta_3^{-1} z_2)$. The ring of invariant polynomials under this action is generated by $u = z_1^3$, $v = z_2^3$ and $z = z_1 z_2$, hence the quotient has a singularity locally given by the equation $uv - z^3 = 0$.

The model domain has an obvious variant where the top level has not been projectivized. We view $\Omega\mathcal{M}D_\Lambda^s = \Omega\mathcal{T}_\Lambda^{pm}(\mu)/\text{Tw}_\Lambda^{sv}$ as a $(\mathbb{C}^*)^{L(\Lambda)}$ -bundle over $\Omega\mathcal{B}_\Lambda$, and define $\Omega\overline{\mathcal{M}D}_\Lambda^s$ as the associated $\mathbb{C}^{L(\Lambda)}$ -bundle. We let $\overline{\Omega\mathcal{M}D}_\Lambda = \Omega\overline{\mathcal{M}D}_\Lambda^s/K_\Lambda$.

Proposition 10.10. *The simple model domain $\overline{\Omega\mathcal{M}D}_\Lambda^s$ is the fine moduli space for the functor of simple model differentials $\mathbf{MD}_{(\mu,\Lambda)}^s$, and $\overline{\mathcal{M}D}_\Lambda^s$ is the fine moduli space for $\mathbb{P}\mathbf{MD}_{(\mu,\Lambda)}^s$.*

The model domain $\overline{\Omega\mathcal{M}D}_\Lambda$ is the fine moduli space for the functor of model differentials $\mathbf{MD}_{(\mu,\Lambda)}$.

Proof. For the first statement we need to construct a family of model differentials over $\overline{\Omega\mathcal{M}D}_\Lambda^s$ and prove its universal property. The family over $\Omega\mathcal{M}D_\Lambda^s$ is simply given by the $(\text{Tw}_\Lambda^{sv}$ -quotient of the) universal family $(\pi: \mathcal{X} \rightarrow \Omega\mathcal{T}_\Lambda^{pm}(\mu), \boldsymbol{\eta}, \mathbf{z}, \sigma, f)$ over $\Omega\mathcal{T}_\Lambda^{pm}(\mu)$ (see Proposition 10.7), where we consider the family of markings as a marking up to the group Tw_Λ^{sv} , together with the functions $t_i = 1$ for all $i \in L(\Lambda)$.

To extend this datum to $\overline{\Omega\mathcal{M}D}_\Lambda^s$, we cover $\mathbb{P}\Omega\mathcal{B}_\Lambda$ by charts $\Psi: U \rightarrow \mathbb{P}\Omega\mathcal{B}_\Lambda$ that are small enough to admit a lift $\tilde{\Psi}: U \rightarrow \Omega\mathcal{M}D_\Lambda^s$. Recall that $p: \Omega\mathcal{M}D_\Lambda^s \rightarrow \mathbb{P}\Omega\mathcal{B}_\Lambda$ is a principal bundle with structure group $(\mathbb{C}^*)^{L(\Lambda)} \cong T_\Lambda^s = \mathbb{C}^{L(\Lambda)}/\text{Tw}_\Lambda^{sv}$. We define the scale comparison map

$$p_\Psi = \tilde{\Psi} \circ \Psi^{-1} \circ p: p^{-1}(U) \rightarrow p^{-1}(U).$$

On the associated \overline{T}_Λ^s -bundle $\overline{p}: \overline{\Omega\mathcal{M}D}_\Lambda^s \rightarrow \mathbb{P}\Omega\mathcal{B}_\Lambda$ we can thus define locally the function

$$R_U^s: \overline{p}^{-1}(U) \rightarrow \overline{T}_\Lambda^s, \quad (\boldsymbol{\eta}, \boldsymbol{\sigma}) \mapsto (\boldsymbol{\eta}, \boldsymbol{\sigma})/p_\Psi((\boldsymbol{\eta}, \boldsymbol{\sigma})).$$

Since the equivalence class of $(\boldsymbol{\eta}, \boldsymbol{\sigma}, 1)$ agrees with $(p_\Psi(\boldsymbol{\eta}, \boldsymbol{\sigma}), R_U^s(\boldsymbol{\eta}, \boldsymbol{\sigma}))$ on $p^{-1}(U)$, the local maps R_U^s glue to the required global map $R^s: \overline{\Omega\mathcal{M}D}_\Lambda^s \rightarrow \overline{T}_\Lambda^s$. The verification of the universal property is straightforward.

Post-composing the map R^s as part of the datum of the functor $\mathbf{MD}_{(\mu,\Lambda)}^s$ with the quotient map $\overline{T}_\Lambda^s \rightarrow \overline{T}_\Lambda^s/K_\Lambda \cong \overline{T}_\Lambda^n$ and taking into account the equivalence relations, exhibits $\mathbf{MD}_{(\mu,\Lambda)}^s$ as the quotient functor of $\mathbf{MD}_{(\mu,\Lambda)}^s$ by the action of K_Λ . Since $\overline{\mathcal{M}D}_\Lambda = \overline{\mathcal{M}D}_\Lambda^s/K_\Lambda$ by definition, the second statement follows from the first. \square

Finally, we express the topology on the model domain in the language of conformal maps. The following proposition follows immediately from the definition of the quasi-conformal topology in Section 5.6 and the topology on the $\mathbb{C}^{L(\Lambda)}$ -bundle associated with a $(\mathbb{C}^*)^{L(\Lambda)}$ -bundle. Given $\mathbf{t} \in (\mathbb{C})^{L(\Lambda)}$, we define $J(\mathbf{t}) \subseteq L(\Lambda)$ to be the subset of indices i such that $t_i = 0$.

Proposition 10.11. *A sequence $(X_m, \mathbf{z}_m, \boldsymbol{\eta}_m, \mathbf{t}_m, \preceq, \sigma_m, f_m)$ of (simple) model differentials in $\overline{\Omega\mathcal{M}D}_\Lambda^s$ converges to $(X, \mathbf{z}, \boldsymbol{\eta}, \mathbf{t}, \preceq, \sigma, f)$ if and only if, taking representatives*

with $t_{m,i}, t_i \in \{0, 1\}$ for X_m and X , there exist a sequence of positive numbers ϵ_m converging to 0 and a sequence of vectors $\mathbf{d}_m = \{d_{m,j}\}_{j \in J(\mathbf{t})} \in \mathbb{C}^{J(\mathbf{t})}$ such that the following conditions hold for sufficiently large m , where we let $d_{m,0} = 0$ and denote $c_{m,j} = \mathbf{e}(d_{m,j})$:

- (i) There is an inclusion $\iota_m: J(\mathbf{t}_m) \hookrightarrow J(\mathbf{t})$.
- (ii) There exist almost-diffeomorphisms $g_m: \overline{X}_{\mathbf{d}_m \cdot \sigma} \rightarrow \overline{X}_{\sigma_m}$ that are compatible with the markings in the sense that f_m is isotopic to $g_m \circ (\mathbf{d}_m \cdot f)$ rel marked points up to an element in Tw_Λ^* . Moreover, the map g_m is conformal on the ϵ_m -thick part $(X, \mathbf{z})_{\epsilon_m}$.
- (iii) The differentials $c_{m,j} g_m^* \eta_m$ converge to η uniformly on compact sets of ϵ_m -thick part of the subsurface of X consisting of levels i with $\delta_J(i) = \delta_J(j)$.
- (iv) For any $i, j \in J(\mathbf{t})$ with $i > j$ and any subsequence along which $\delta_{J(\mathbf{t}_m)}(i) = \delta_{J(\mathbf{t}_m)}(j)$, we have

$$\lim_{n \rightarrow \infty} \frac{|c_{m,i}|}{|c_{m,j}|} = 0.$$

- (v) The almost-diffeomorphisms g_m are nearly turning-number-preserving, i.e. for every good arc γ in \overline{X}_σ , the difference of turning numbers $\tau(g_n \circ F_{\mathbf{d}_n} \circ \gamma) - \tau(F_{\mathbf{d}_n} \circ \gamma)$ converges to zero.

11. MODIFYING DIFFERENTIALS AND PERTURBED PERIOD COORDINATES

The goal of this section is first to define modifying differentials ξ as a preparation for the plumbing construction in Section 12. The second goal is to define local coordinates, which we will call perturbed period coordinates, on the simple model domain. These will be useful to describe neighborhoods of boundary points. We start with only vertical nodes and in Section 11.3 we extend the perturbed period coordinates to also include the limit of periods through horizontal nodes. The perturbed period coordinates are similar to the usual period coordinates, with the following modifications that will allow us to transition from stable curves with many nodes to curves with fewer nodes.

First, they are coordinates for the universal differential η , but rescaled by the scaling parameters \mathbf{t} , and perturbed by the modifying differentials ξ . The reason for this is that the perturbed differential lives on the universal family over the Dehn space $\Xi \mathcal{D}_\Lambda$ defined below after plumbing. Consequently, once the plumbing construction is completed, perturbed period coordinates are coordinates on $\Xi \mathcal{D}_\Lambda$.

Second, the plumbing construction cuts out the zero that used to be at the top end of any vertical node. Thus to keep track of the relative period corresponding to such a zero, we compute a period not to this zero, but to a suitably chosen nearby point. The choice of this nearby point will be made in such a way that under degeneration to the boundary of $\Omega \overline{\mathcal{M}}_\Lambda^s$ the difference between the perturbed period and the original period tends to zero.

Third, the coordinate system contains for each level one entry which measures the scale of degeneration. This is not actually a period, but rather an a_i -th root of a period of η .

The setup for this section is the following. Let $\mathcal{W} \subset \overline{\Omega\mathcal{M}D}_\Lambda^s$ be an open neighborhood of (the moduli point corresponding to) a simple marked model differential $(X, \boldsymbol{\eta}_X, \mathbf{t}, \boldsymbol{\sigma}, f)$ and let $\pi: \mathcal{X} \rightarrow \mathcal{W}$ be the restriction of the universal family to \mathcal{W} . We fix a family of differentials $\boldsymbol{\eta}$ on \mathcal{X} , not identically zero on every component of every fiber, that restricts to $\boldsymbol{\eta}_X$ on X . Recall from (6.12) the definition of the action

$$(11.1) \quad \mathbf{t} * \boldsymbol{\eta} = (t_{\lceil i \rceil} \cdot \eta_{(i)})_{i \in L(\Lambda)} = (t_{-1}^{a_{-1}} \cdot t_{-2}^{a_{-2}} \cdots t_i^{a_i} \cdot \eta_{(i)})_{i \in L(\Lambda)},$$

with the a_i as in (6.7).

11.1. Modifying differentials and the global residue condition revisited. In order to construct the plumbing map, we need modifying differentials as in [BCGGM18], but now defined on the universal family over an open subset of the model domain. The purpose of this section is to prove the existence of such families of modifying differentials.

Definition 11.1. A *family of modifying differentials* over $\mathcal{W} \subset \overline{\Omega\mathcal{M}D}_\Lambda^s$ is a family of meromorphic differentials ξ on $\pi: \mathcal{X} \rightarrow \mathcal{W}$, such that:

- (i) ξ is holomorphic, except for possible simple poles along both horizontal, vertical nodal sections and marked poles;
- (ii) ξ vanishes identically on the components of lowest level of \mathcal{X} , and $\xi_{(i)}$ is divisible by $t_{\lceil i-1 \rceil}$ for each $i \in L^\bullet(\Lambda) \setminus \{-N\}$;
- (iii) $\mathbf{t} * \boldsymbol{\eta} + \xi$ has opposite residues at every node. △

In other words, denote $\tilde{\mathcal{X}} \rightarrow \mathcal{X}$ the partial normalization at the vertical nodes and denote $\tilde{\pi}: \tilde{\mathcal{X}} \rightarrow \mathcal{W}$ its composition with π . Recall that $q_e^\pm: \mathcal{W} \rightarrow \tilde{\mathcal{X}}$ are the sections corresponding to the top and bottom preimages of the vertical node e , with images Q_e^\pm . Moreover, let \mathcal{P} be the reduced divisor associated to \mathcal{Z}^∞ . Then ξ is a holomorphic section of

$$\tilde{\pi}_* \omega_{\tilde{\mathcal{X}}/\mathcal{W}} \left(\sum_{e \in E^v} (\tilde{Q}_e^+ + \tilde{Q}_e^-) + \mathcal{P} \right),$$

which is divisible by $t_{\lceil i-1 \rceil}$ at level i and chosen so that as functions on \mathcal{W}

$$(11.2) \quad \text{Res}_{q_e^+}(\mathbf{t} * \boldsymbol{\eta} + \xi) + \text{Res}_{q_e^-}(\mathbf{t} * \boldsymbol{\eta} + \xi) = 0$$

for every vertical node $e \in E(\Gamma)^v$.

We start the construction by recalling from [BCGGM19] a topological restatement of the global residue condition. Consider the subspace $V \subseteq H_1(\Sigma \setminus P_{\mathbf{s}}; \mathbb{Q})$ spanned by the vertical curves Λ^v , where $P_{\mathbf{s}}$ is the set of marked poles. The order on Λ determines a filtration

$$(11.3) \quad 0 = V_{-N-1} \subseteq V_{-N} \subseteq \cdots \subseteq V_{-1} = V,$$

where V_i is generated by the image in V of all those vertical curves in Λ such that $\ell(e^-) \leq i$. Note that this convention differs slightly from the one of [BCGGM19]: we allow horizontal nodes, our V_i corresponds to V_{-i} there and our N corresponds to $N-1$ there.

Suppose we are given a marked differential (X, η) on a pointed stable curve that satisfies the axioms (0)-(3) of a twisted differential. Fixing an orientation of the individual

curves of Λ^v , the differential η defines a *residue assignment* $\rho: \Lambda^v \rightarrow \mathbb{C}$. With the help of these maps we give an alternative statement of the global residue condition.

Proposition 11.2 ([BCGGM19, Proposition 6.3]). *A residue assignment $\rho: \Lambda^v \rightarrow \mathbb{C}$ satisfies the global residue condition if and only if there exists period homomorphisms*

$$\rho_i: V_i/V_{i-1} \rightarrow \mathbb{C} \quad \text{for any } i \in L(\Lambda),$$

such that $\rho_i(\lambda) = \rho(\lambda)$ for all simple closed curves λ in Λ^v , where $i = \ell(\lambda^-)$.

In what follows it will be convenient for us to lift the period homomorphisms to maps $\rho_i: V_i \rightarrow \mathbb{C}$ such that $\rho_i(V_{i-1}) = 0$ for all $i \in L(\Lambda)$. We are now ready to construct the family of modifying differentials, and we will then demonstrate the constructions in the proof by an example.

Proposition 11.3. *The family $\pi: \mathcal{X} \rightarrow \mathcal{W}$ equipped with the universal differential $\mathbf{t} * \boldsymbol{\eta}$ has a family of modifying differentials.*

Proof. Choose a maximal multicurve $\Lambda_{\max} \supseteq \Lambda$ decomposing $\Sigma \setminus P_{\mathbf{s}}$ into pants. Let $V' \subset H_1(\Sigma \setminus P_{\mathbf{s}}; \mathbb{Q})$ be the subspace of homology generated by the classes of all curves in Λ_{\max} . Note that V' contains V , and projects to a Lagrangian subspace of $H_1(\Sigma; \mathbb{Q})$. The restriction of $\mathbf{t} * \boldsymbol{\eta}$ to levels i or below determines a holomorphic period map (extending ρ_i above to families)

$$\rho_i: \mathcal{W} \rightarrow \text{Hom}_{\mathbb{Q}}(V_i, \mathbb{C}),$$

such that ρ_i restricts to zero on V_{i-1} . In period coordinates, ρ_i is simply a linear projection. By (6.12), the map ρ_i is T_{Λ}^s -equivariant, i.e.

$$(11.4) \quad \rho_i(\mathbf{q} * (\mathcal{X}, \mathbf{t} * \boldsymbol{\eta})) = \prod_{j \geq i} q_j^{a_j} \cdot \rho_i(\mathcal{X}, \mathbf{t} * \boldsymbol{\eta}) \quad \text{for any } \mathbf{q} \in T_{\Lambda}^s.$$

For each $i \in L(\Lambda)$ we choose a sub-multicurve $B_i \subset \Lambda_{\max}$ whose image in V' is a basis of V'/V_i , such that for any $i \in L(\Lambda)$ the inclusion $B_i \subset B_{i-1}$ holds. We then define the extension $\tilde{\rho}_i: \mathcal{W} \rightarrow \text{Hom}_{\mathbb{Q}}(V', \mathbb{C})$ of ρ_i by the requirement $\tilde{\rho}_i(b_i) = 0$ for all $b_i \in B_i$.

Since Λ_{\max} is a maximal multicurve on $\Sigma \setminus P_{\mathbf{s}}$, a meromorphic form on \mathcal{X} , holomorphic except for at worst simple poles at the nodes and at the marked poles $P_{\mathbf{s}}$, is specified uniquely by its periods on V' . We define $\xi_{(i)}$ on $\mathcal{X}_{(i)}$ so that its V' -periods for $\gamma \in \Lambda_{\max}$ are

$$(11.5) \quad \int_{\gamma} \xi_{(i)}(u) = \sum_{j < i} \tilde{\rho}_j(u)(\gamma) \quad \text{for all } u \in \mathcal{W}.$$

By the equivariance (11.4), we see that $\tilde{\rho}_j$ is divisible by $t_{[j]}$, and hence $\xi_{(i)}$ is divisible by $t_{[i-1]}$, since the right-hand side of (11.5) has summation indices j up to $i-1$, thus verifying Definition 11.1 (ii).

Given a curve $\gamma \in \Lambda^v$ joining level i to level $j < i$, we verify the opposite residue condition (11.2) required for a modifying differential, which states that

$$\int_{\gamma} ((\mathbf{t} * \boldsymbol{\eta})_{(j)} + \xi_{(j)}) = \rho_j(\gamma) + \sum_{k < j} \tilde{\rho}_k(\gamma) = \sum_{k < i} \tilde{\rho}_k(\gamma) = \int_{\gamma} ((\mathbf{t} * \boldsymbol{\eta})_{(i)} + \xi_{(i)}).$$

In the above the first equality follows from the fact that ρ_j is the period map determined by $(\mathbf{t}*\boldsymbol{\eta})_{(j)}$ and the definition of $\xi_{(j)}$. The second equality follows from the global residue condition of $\mathbf{t}*\boldsymbol{\eta}$ as restated in Proposition 11.2 which implies that $\rho_k(\gamma) = 0$ for all $j < k \leq i$. The last equality again follows from the definition of $\xi_{(i)}$ and the fact that $\rho_i(\gamma) = 0$. \square

This proof shows in particular the following.

Corollary 11.4. *The modifying differential ξ is uniquely determined by the choice of the subspace $V' \subset H_1(\Sigma \setminus P_s; \mathbb{Q})$ and the subsets B_i . Its level-wise components $\xi_{(i)}$ depend only on t_j and η_j for $j < i$.*

Example 11.5. We illustrate the objects introduced in the proof of Proposition 11.3 in the context of a slight simplification of our running example, as pictured in Figure 5 with one pole denoted by p (so the level graph is still a triangle, but the irreducible components are simpler). We do this by analyzing the family of modifying differentials ξ along a single fiber over $\mathbf{t} = (1, \dots, 1)$, and with the $*$ -action one can similarly analyze ξ over other base points.

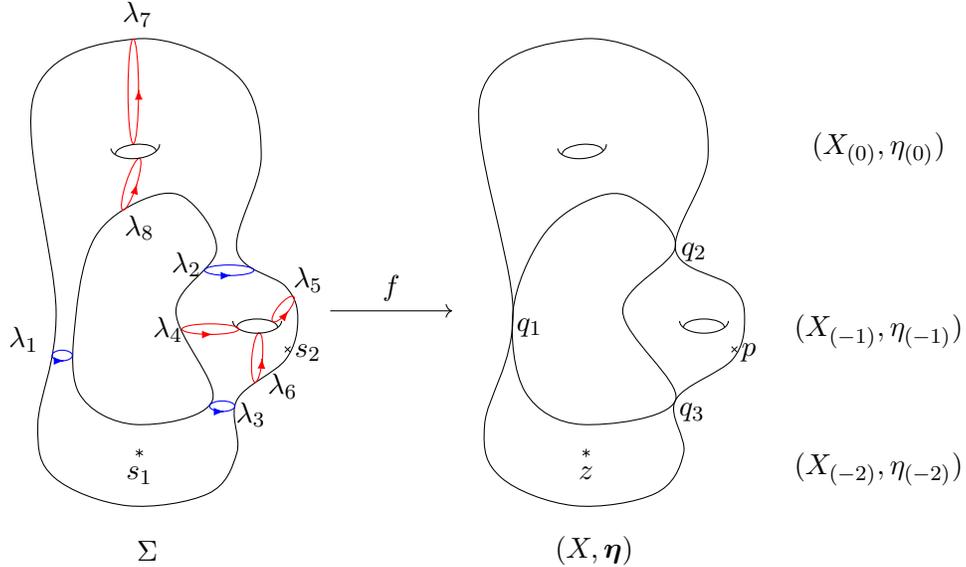


FIGURE 5. The marked surface together with the multicurves Λ and Λ_{\max} .

The vertical multicurve Λ^v (in blue in Figure 5) consists of curves λ_1, λ_2 and λ_3 , which are all homologous to each other. The filtration of the V_i induced by the multicurve Λ^v is then given by

$$0 = V_{-3} \subset V_{-2} = \langle \lambda_1 \rangle = V_{-1} = V,$$

where $\langle \cdot \rangle$ denotes the linear span. Hence the maps ρ_i are given by $\rho_{-2}(\lambda_1) = a \in \mathbb{C}$ and $\rho_{-1}(\lambda_1) = 0$. We choose the maximal multicurve $\Lambda_{\max} \supset \Lambda$ by adding the curves $\{\lambda_5, \dots, \lambda_8\}$, shown in red. Then we have in homology the equalities

$$\lambda_1 = \lambda_2 = \lambda_4 + \lambda_5 = \lambda_3 = \lambda_4 + \lambda_6 = \lambda_7 + \lambda_8,$$

and thus $V' = \langle \lambda_1, \lambda_4, \lambda_7 \rangle$. We choose the sets B_i to be $B_{-2} = \{\lambda_4, \lambda_7\} = B_{-1}$. Then the extension $\tilde{\rho}_{-2}$ of the map ρ_{-2} defined on $V_{-2} = \langle \lambda_1 \rangle$ is given by requiring $\tilde{\rho}_{-2}(B_{-2}) = 0$. That is,

$$\tilde{\rho}_{-2}(\lambda) = \begin{cases} a, & \text{if } \lambda = \lambda_1, \\ 0, & \text{if } \lambda = \lambda_4, \lambda_7. \end{cases}$$

Similarly, the extension $\tilde{\rho}_{-1}$ of ρ_{-1} from $V_{-1} = \langle \lambda_1 \rangle$ to V' is defined by requiring $\tilde{\rho}_{-1}(B_{-1}) = 0$, hence $\tilde{\rho}_{-1}$ is simply identically zero.

We can now define the modifying differentials $\xi_{(i)}$. Following the construction in the proof, we see that the differential $\xi_{(-2)} = 0$ identically. The differential $\xi_{(-1)}$ is supported on the component on the right, which is at level -1 . It has simple poles at q_2 and q_3 with residues $\pm a/2\pi\sqrt{-1}$, is holomorphic at p , has period zero over λ_4 , and has periods a over λ_5 and λ_6 . Finally, the differential $\xi_{(0)}$ has simple poles at q_1 and q_2 , with residues $\pm a/2\pi\sqrt{-1}$, has period zero over λ_7 , and period a over λ_8 . To see these, consider for example the period of $\xi_{(0)}$ at λ_8 . By definition it is given by

$$\int_{\lambda_8} \xi_{(0)} = \tilde{\rho}_{-2}(\lambda_8) + \tilde{\rho}_{-1}(\lambda_8) = \rho_{-2}(\lambda_1) + \rho_{-1}(\lambda_1) = a + 0 = a,$$

since λ_1 is homologous to $\lambda_7 + \lambda_8$ and $\tilde{\rho}_{-2}(\lambda_7) = \tilde{\rho}_{-1}(\lambda_7) = 0$ for $\lambda_7 \in B_{-2}, B_{-1}$. The other cases can be computed similarly.

11.2. Perturbed period coordinates. We will now perturb the usual notion of period coordinates, to avoid using marked points and zeros that are at the nodes, and choosing different basepoints instead. We first introduce the preparatory material in full generality, and then define the perturbed period coordinates under the simplifying assumption that there are no horizontal nodes. We extend these coordinates to the case with horizontal nodes in Section 11.3.

To define the perturbed period map we need to specify additional marked points near the vertical zeros of η and we need to recall various spaces defined by residue conditions, together with the dimension estimates from [BCGGM19].

The Teichmüller markings up to twist group of the welded surfaces in the model domain induce markings $f_i: \Sigma_{(i)}^c \rightarrow \mathcal{X}_{(i)}$ of the families of connected components of the subsurfaces at level i . Denote by $P_{s,i}$ and $Z_{s,i}$ those marked poles and zeros that lie on the compact level i subsurface $\Sigma_{(i)}^c$. Moreover, denote by $Q_{E,i}^\pm$ those zero and pole sections where one of the endpoints of the corresponding edge lies on $\Sigma_{(i)}^c$. We define the sets of points

$$(11.6) \quad P_i = P_{s,i} \cup Q_{E,i}^- \quad \text{and} \quad Z_i = Z_{s,i} \cup Q_{E,i}^+$$

for each level i .

The perturbed period coordinates are roughly modeled on the coordinates t_i from the rescaling ensemble and coordinates of the projectivization of certain subspaces $\mathcal{R}_i^{\text{grc}}$ of $H^1(\Sigma_{(i)}^c \setminus P_i, Z_i, \mathbb{C})$. Coordinates on the latter are as usual given by all but one of the periods.

To define $\mathcal{R}_i^{\text{grc}}$ we start with the map $H^1(\Sigma_{(i)}^c \setminus P_i, Z_i, \mathbb{C}) \rightarrow \mathbb{C}^{|P_i|}$ given by taking the integrals over small loops around the points P_i . Note that the image of this map

is contained in the subspace cut out by the residue theorems on each component. Let $R_i^{\text{grc}} \subset \mathbb{C}^{|P_i|}$ be the subspace cut out further by the matching residue condition and the global residue condition, as stated in Section 2.4. The *GRC space* $\mathcal{R}_i^{\text{grc}} \subseteq H^1(\Sigma_{(i)}^c \setminus P_i, Z_i, \mathbb{C})$ is then defined as the preimage of R_i^{grc} . If we denote by H the number of horizontal nodes of Λ , then [BCGGM19, Theorem 6.1] can be restated as follows.

Proposition 11.6. *The (open) simple model domain $\Omega\mathcal{M}D_\Lambda^s$ is locally modeled on the sum of the GRC spaces $\oplus_i \mathcal{R}_i^{\text{grc}}$. This space has dimension*

$$\sum_{i \in L^\bullet(\Lambda)} \dim(\mathcal{R}_i^{\text{grc}}) = \dim \Omega\mathcal{M}_{g,n}(\mu) - H.$$

For each half-edge h of $\Gamma(\Lambda)$ with non-negative m_h , i.e. for each non-polar marked point in the smooth part of \mathcal{X} , we denote by $\mathbf{z}(h)$ the corresponding section of $\mathcal{X} \rightarrow \mathcal{W}$. We choose nearby sections $\sigma_e^+ : \mathcal{W} \rightarrow \mathcal{X}$ and $\sigma_h : \mathcal{W} \rightarrow \mathcal{X}$ so that

$$(11.7) \quad \int_{q^+(e)}^{\sigma_e^+(w)} \eta_{(i)} = \text{const} \quad \text{and} \quad \int_{\mathbf{z}(h)}^{\sigma_h(w)} \eta_{(j)} = \text{const},$$

where $i = \ell(e^+)$ and $j = \ell(h)$ are the corresponding levels that contain the (short fiber-wise) integration paths respectively.

As the final preparation step, note that the form $\mathbf{t} * \boldsymbol{\eta} + \xi$ on \mathcal{X} may no longer have a zero of the prescribed order at $\mathbf{z}(h)$ because of the modifying differential ξ . In the process of plumbing in Section 12.4, we will describe a local surgery \mathcal{X}^{he} of \mathcal{X} in a neighborhood of the sections $\mathbf{z}(h)$ corresponding to the half-edges h , such that the images of the sections $\mathbf{z}(h)$ are untouched by the surgery, and the extension of $\mathbf{t} * \boldsymbol{\eta} + \xi$ to \mathcal{X}^{he} again has a zero of order $\text{ord}_{\mathbf{z}(h)} \boldsymbol{\eta}$ along a section that we still denote by $\mathbf{z}(h) : \mathcal{W} \rightarrow \mathcal{X}$.

We now define the *perturbed period map at level i* under the hypothesis that there are no horizontal nodes. We fix homology classes $\gamma_1, \dots, \gamma_{n(i)}$ such that their periods $\int_{\gamma_j} \eta_{(i)}$ form a basis of $\mathcal{R}_i^{\text{grc}}$. Stability of the curve \mathcal{X} implies that for each i at least one of the periods $\int_{\gamma_j} \eta_{(i)}$ is non-zero, say for $j = n(i)$. We thus denote by $\mathcal{R}'_i \subset \mathcal{R}_i^{\text{grc}}$ the codimension one subspace generated by the periods of $\gamma_1, \dots, \gamma_{n(i)-1}$ for all levels $i < 0$, and we let $\mathcal{R}'_0 := \mathcal{R}_0^{\text{grc}}$ as the differential on the top level is not considered up to scale. We denote by $n'(i)$ the dimension of the space \mathcal{R}'_i for all i .

The perturbed period map is then built with the help of

$$(11.8) \quad \text{PPer}_i : \begin{cases} \mathcal{W} & \rightarrow \mathcal{R}'_i, \\ [(X, \boldsymbol{\eta}, \mathbf{t})] & \mapsto \left(\int_{\gamma_j} \eta_{(i)} + t_{[i]}^{-1} \cdot \xi_{(i)} \right)_{j=1}^{n'(i)}. \end{cases}$$

Here $t_{[i]}^{-1}$ is the inverse of $t_{[i]}$, namely, $t_{[i]}^{-1} = t_{-1}^{-a-1} \cdot t_{-2}^{-a-2} \dots t_i^{-a_i}$. By Definition 11.1 (ii), $\xi_{(i)}$ is divisible by $t_{[i-1]}$, hence the integrals above are well-defined. Here the integrals are over the f_i -images of the cycles, but *we integrate from the points b_e^+ defined to be $\sigma_e^+(w)$ for cycles starting or ending at a point in $Q_{E,i}^+$, where $w = [(X, \boldsymbol{\eta}, \mathbf{t})]$, rather than from the nearby zeros of $\eta_{(i)}$.*

Proposition 11.7. *The perturbed period map*

$$(11.9) \quad \text{PPer}: \mathcal{W} \rightarrow \mathbb{C}^{L^\bullet(\Lambda)} \times \bigoplus_{i \in L^\bullet(\Lambda)} \mathcal{R}'_i, \quad [(X, \boldsymbol{\eta}, \mathbf{t})] \mapsto \left(\mathbf{t}; \bigoplus_{i \in L^\bullet(\Lambda)} \text{PPer}_i \right)$$

is open and locally injective on a neighborhood of the most degenerate stratum $\mathcal{W}_\Lambda = \bigcap_{i \in L(\Lambda)} D_i$ inside of \mathcal{W} .

Proof. We need to show that the derivative of PPer is surjective along the boundary stratum \mathcal{W}_Λ since surjectivity is an open condition and since this surjectivity implies openness. At a point of \mathcal{W}_Λ the i -th summand PPer_i consists of the usual period coordinates for $\eta_{(i)}$, shifted by a constant since we integrate from a nearby point (using the absence of horizontal nodes by our assumption). Here the integral of $t_{[i]}^{-1} \cdot \xi_{(i)}$ is zero over the most degenerate stratum \mathcal{W}_Λ , because by definition $t_{[i]}^{-1} \cdot \xi_{(i)}$ is still divisible by $t_{i-1}^{a_{i-1}}$ and \mathcal{W}_Λ is defined by $t_j = 0$ for all $j \in L(\Lambda)$. In the complementary directions, surjectivity is obvious since the t_i are coordinates on the domain and are included in the target of PPer.

Since \mathcal{W} is smooth and of the same dimension as $\bigoplus_{i \in L^\bullet(\Lambda)} \mathcal{R}'_i$ by Proposition 11.6 (under the assumption of no horizontal nodes), surjectivity of the derivative of PPer implies injectivity of the derivative map at any boundary point in \mathcal{W}_Λ and hence local injectivity in some neighborhood. \square

Example 11.8. We give the description of the perturbed period coordinates in the setting of our running example of Section 2.6. Hence the differentials that we consider are in the closure of the meromorphic stratum $\Omega\mathcal{M}_{5,4}(4, 4, 2, -2)$. More precisely, we consider the enhanced dual graph Γ_2 on the right of Figure 1. In this case the differential $\eta_{(-2)}$ is in $\Omega\mathcal{M}_{0,3}(4, -2, -4)$, $\eta_{(-1)}$ is in $\Omega\mathcal{M}_{1,4}(0, 4, -2, -2)$ and $\eta_{(0)}$ is in $\Omega\mathcal{M}_{3,3}(2, 2, 0)$. Since the global residue condition imposes precisely the condition that the residue of $\eta_{(-1)}$ at $q_{e_1}^-$ is zero, the GRC space is the product of the top and bottom H^1 with the hyperplane of the middle H^1 given by this residue condition.

We will consider the deformations over a disk $\Delta^2 = \Delta_{t_{-1}} \times \Delta_{t_{-2}}$ which parameterizes the smoothing of the levels of (X, η) . Note that the residues at the poles of the differential $\eta_{(-2)}$ are non-zero (see [BCGGM18, Lemma 3.6]). The family of modifying differentials ξ consists of $\xi_{(0)}$ on $X_{(0)}$ and $\xi_{(-1)}$ on $X_{(-1)}$, where $\xi_{(0)}$ is divisible by t_{-1}^3 , and $\xi_{(-1)}$ is divisible by t_{-2}^3 . Moreover, ξ vanishes identically on $X_{(-2)}$. In Figure 6 we show a basis of the cycles of integration before and after the plumbing construction described later in Section 12. In this basis, the map PPer_0 is given by the map which associates the integrals of $\eta_0 + \xi_{(0)}$ along the cycles belonging to $X_{(0)}$. The maps PPer_{-1} and PPer_{-2} are defined analogously.

We now describe how the perturbed period coordinates behave over the base Δ^2 . Note that since our construction is local, we can identify the circles α_j and β_j for $j = 1, \dots, 4$ with the circles α_j^0 and β_j^0 . On the subsurface $X_{(0)}$ the restriction of the differential $\eta_0 + \xi_{(0)}$ on α_j^0 and of the plumbed differential to α_j clearly coincide (where all the t_i are non-zero) under this identification. The case of the subsurface $X_{(-1)}$ is similar. Note that if the modifying differentials vanish, then the period of each cycle

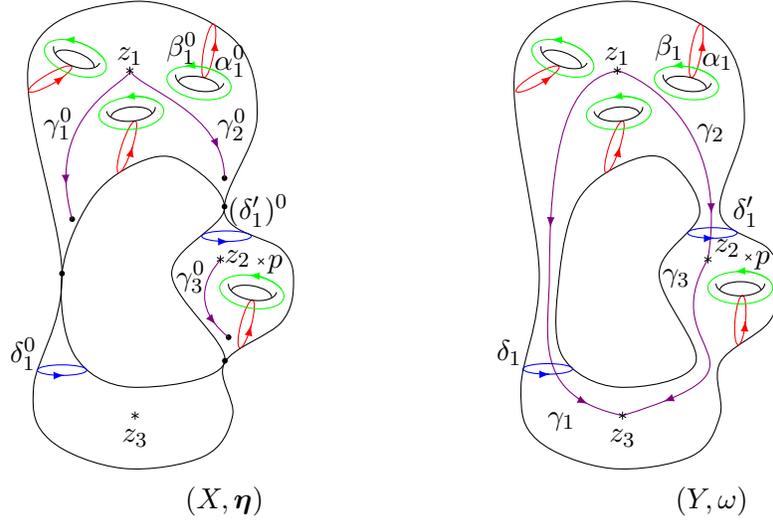


FIGURE 6. A basis of homology of our running example and of a nearby smooth differential.

on $X_{(0)}$ would be a constant and the period of each cycle on $X_{(-1)}$ would be of a constant times $(t_{-1})^3$.

We now consider the relative cycles γ_k which degenerate to the relative cycles γ_k^0 . The period for γ_1 is equal to the period for γ_1^0 plus a function of t_{-1} and t_{-2} which is zero on $\{t_{-1}t_{-2} = 0\}$. This function depends on the choice of the points near z_1 , near the node, and the way that we glue the plumbing fixture in the nodal differential. The case of the cycles γ_k for $k = 2, 3$ is similar.

Finally, note that the period of $t * \eta$ at the homotopic cycles δ_1 and δ_1' is a function $2\pi i r$ that is divisible by $(t_{-1}t_{-2})^3$, where r is the residue at the corresponding node. This is coherent with the GRC.

11.3. Horizontal extension of perturbed period coordinates. For the horizontal extension we require the full plumbing setup introduced in Section 12 below. That is, suppose that we have chosen the local gluing data to plumb each of the H horizontal nodes by a plumbing fixture $V(x)$ for $x \in \Delta_\epsilon$, see (12.5) and (12.10) below. Let $\mathcal{Y} \rightarrow U = \mathcal{W} \times \Delta_\epsilon^H$ be the family that results from plumbing the horizontal nodes of $\mathcal{X} \rightarrow \mathcal{W}$, see the ‘‘Second plumbing step’’ in Section 12.4. Our goal now is to extend the perturbed period map PPer to a local diffeomorphism whose domain is U .

Suppose the j -th horizontal node q_j lies on the level $i = i(j)$ subsurface $\Sigma_{(i)} \subset \Sigma$. Let β_j be a path that stays in $\Sigma_{(i)}$, which represents a homology class in Σ relative to Z_i (or to the points in the image of σ^+ if needed) and that crosses once the seam of q_j , and that crosses no other seams. Such a path exists, since each component of X contains at least a point in Z_i . Let α_j be the loop around q_j . We define the *perturbed period map*

$$(11.10) \quad \text{ePP} = \text{PPer} \times \text{Phor}: U \rightarrow \mathbb{C}^{L^*(\Lambda)} \times \oplus_i \mathcal{R}'_i \times \mathbb{C}^H,$$

where

$$(11.11) \quad \text{Phor: } \begin{cases} U & \rightarrow \mathbb{C}^H, \\ ([X, \boldsymbol{\eta}, \mathbf{t}], (x_j)) & \mapsto \left(e \left(\frac{\int_{\beta_j} \eta^{(i)} + t_{[i]}^{-1} \xi^{(i)}}{\int_{\alpha_j} \eta^{(i)} + t_{[i]}^{-1} \xi^{(i)}} \right)_{j=1, \dots, H} \right), \end{cases}$$

and for integration we use the f -images of the paths in the corresponding fiber of the family $\mathcal{Y} \rightarrow U$. Note that we integrate the form in the fibers of \mathcal{Y} which is the family obtained after the second plumbing map. In particular, the horizontal node q_j has been smoothed out in the fibers of $\mathcal{Y} \rightarrow U$ above the locus $x_j \neq 0$ using the plumbing fixture. The exponentiation makes this map well-defined, despite that f is only well-defined up to composition by elements in the twist group Tw_Λ . Indeed any two images of β_j differ by a power of the Dehn twist about α_j .

Proposition 11.9. *The perturbed period map ePP is a local diffeomorphism in a neighborhood of $\mathcal{W} \times \{0\}^H$.*

Proof. Using Proposition 11.7 the claim follows from the fact that the components of the map Phor are non-constant, holomorphic (since the α_j -periods tend to a non-zero residue and the imaginary part of the β_j -periods over the α_j -period goes to $+\infty$) and independent of each other by the construction of plumbing annuli disjointly and independently. \square

Example 11.10. We describe the perturbed period coordinates in the case of a curve with two irreducible components X_1 and X_2 that meet at two horizontal nodes. Take a model differential such that its restriction η_1 to the component X_1 is in $\Omega\mathcal{M}_{1,3}(2, -1, -1)$ and the restriction η_2 to X_2 is in $\Omega\mathcal{M}_{1,4}(1, 1, -1, -1)$. In this case, the GRC space is the subset of the product of the H^1 such that the sum of the residues of the η_i at each node is zero. Of course one of the two equations is redundant because of the residue theorem, and hence the GRC space is a hyperplane (i.e. of codimension one only). Moreover, since the twisted differential has only one level, there is no modifying differential, hence the perturbed period map PPer is the usual period map.

Now let us denote by (x_1, x_2) the coordinates in Δ_ϵ^2 . Moreover for $i = 1, 2$, let β_i be a good arc crossing exactly once the seam of the node q_i from the double zero of η_1 to one of the zeros of η_2 . Then the map $\text{Phor}|_{\mathbf{0} \times \Delta_\epsilon^2}$ is given by

$$(11.12) \quad (x_1, x_2) \mapsto (k_1 x_1, k_2 x_2),$$

where the k_i are non-zero constants.

To see this, we decompose the path β_i into three paths as follows. The first path β_i^1 joins the double zero of η_1 to the marked point in X_1 used to put the plumbing fixture. The second path β_i^2 is the (image in the plumbed surface of the) path in the plumbing fixture of Equation (12.8) joining the two marked points. The last path β_i^3 joins the point of X_2 used for the plumbing fixture and the endpoint of β_i . In this setting, the period of β_i is the sum of the periods of the three arcs β_i^j . Note that the periods of β_i^1 and β_i^3 are constants c_i^1 and c_i^3 above Δ^2 . An easy computation using Equation (12.9) shows that the period of β_i^2 is equal to $r_i \log(x_i)$, where r_i is the residue of η at q_i .

Hence the β_i -period is of the type $c_i^1 + r_i \log(x_i) + c_i^3$. This gives Equation (11.12), where the k_i are the exponentials of $(c_i^1 + c_i^3)/r_i$.

12. THE DEHN SPACE AND THE COMPLEX STRUCTURE

We now wish to provide the moduli space of multi-scale differentials $\Xi \overline{\mathcal{M}}_{g,n}(\mu)$ with a complex structure. In order to understand the structure of $\Xi \overline{\mathcal{M}}_{g,n}(\mu)$ at the boundary, we introduce an auxiliary space $\Xi \mathcal{D}_\Lambda^s$, the *simple Dehn space*. There is a natural open forgetful map $\Xi \mathcal{D}_\Lambda^s \rightarrow \Xi \overline{\mathcal{M}}_{g,n}(\mu)$, and $\Xi \overline{\mathcal{M}}_{g,n}(\mu)$ is covered by the images of these maps as Λ ranges over all enhanced multicurves. Our goal in this section is to give each of the topological spaces $\Xi \mathcal{D}_\Lambda^s$ the structure of a complex manifold, which will in turn give $\Xi \overline{\mathcal{M}}_{g,n}(\mu)$ its complex structure. This complex structure is induced by *plumbing maps* $\Omega \text{Pl}: U \rightarrow \Xi \mathcal{D}_\Lambda^s$, defined by a plumbing construction on the universal family of model differentials over U , where $U \subset \Omega \overline{\mathcal{M}} \overline{\mathcal{D}}_\Lambda^s$ are open sets in the model domain, which we will show give an atlas of complex coordinate charts.

Throughout this section, we fix an enhanced multicurve Λ with dual graph Γ having $N + 1 = |L^\bullet(\Lambda)|$ levels and H horizontal nodes.

12.1. The Dehn space. The *Dehn space associated with Λ* is the topological space

$$(12.1) \quad \Xi \mathcal{D}_\Lambda = \left(\coprod_{\Lambda' \rightsquigarrow \Lambda} \Omega \mathcal{B}_{\Lambda'} \right) / \text{Tw}_\Lambda^v,$$

where the disjoint union has the subspace topology induced from the topology of the augmented Teichmüller space of flat surfaces $\Omega \overline{\mathcal{T}}_{(\Sigma, s)}(\mu)$. We can write the space equivalently as

$$(12.2) \quad \Xi \mathcal{D}_\Lambda = \coprod_{\Lambda' \rightsquigarrow \Lambda} \Xi \mathcal{D}_{\Lambda'}^{\Lambda'} \quad \text{where} \quad \Xi \mathcal{D}_{\Lambda'}^{\Lambda'} = \Omega \mathcal{B}_{\Lambda'} / \text{Tw}_{\Lambda'}^v.$$

We will see that $\Xi \mathcal{D}_\Lambda$ is in fact an orbifold and has a smooth manifold cover, the *simple Dehn space*, defined by

$$(12.3) \quad \Xi \mathcal{D}_\Lambda^s = \left(\coprod_{\Lambda' \rightsquigarrow \Lambda} \Omega \mathcal{B}_{\Lambda'} \right) / \text{Tw}_\Lambda^{sv} = \coprod_{\Lambda' \rightsquigarrow \Lambda} \Xi \mathcal{D}_{\Lambda'}^{\Lambda', s},$$

where $\Xi \mathcal{D}_{\Lambda'}^{\Lambda', s} = \Omega \mathcal{B}_{\Lambda'} / \text{Tw}_{\Lambda'}^{sv}$.

We write similarly $\mathbb{P} \Xi \mathcal{D}_\Lambda$ and $\mathbb{P} \Xi \mathcal{D}_\Lambda^s$ for the corresponding spaces where the top level is projectivized, that is, $\mathbb{P} \Xi \mathcal{D}_\Lambda$ is the quotient of $\Xi \mathcal{D}_\Lambda$ under the \mathbb{C}^* -action.

We refer to a point in the (simple) Dehn space as (the moduli point of the equivalence class of) a marked multi-scale differential $(Y, \mathbf{z}, \boldsymbol{\omega}, \boldsymbol{\sigma}, f)$ where the marking is up to the action of Tw_Λ (resp. Tw_Λ^s). Pointwise this is justified by definition. The identification of families of such differentials with maps to the Dehn space is given in Section 13.

We now outline the *plumbing construction*. It starts with the universal curve $\mathcal{X} \rightarrow \Omega \overline{\mathcal{M}} \overline{\mathcal{D}}_\Lambda^s$. Taking the product with Δ^H , we pass to the family $\mathcal{X} \rightarrow \Omega \overline{\mathcal{M}} \overline{\mathcal{D}}_\Lambda^s \times \Delta^H$, and then restrict this family to a sufficiently small open set $U \subset \Omega \overline{\mathcal{M}} \overline{\mathcal{D}}_\Lambda^s \times \Delta^H$. By cutting out neighborhoods of the nodes and gluing in standard plumbing fixtures, we construct a new family of curves $\Omega \text{Pl}(\mathcal{X}) \rightarrow U$. This new family of curves carries

a Tw_Λ^s -marked family of multi-scale differentials. The plumbing construction is not canonical and depends on choices made at several points in the construction.

We remark that the classical plumbing constructions require an extra parameter for each node parameterizing the modulus of the annulus that is glued in. Our plumbing construction requires extra parameters only for the horizontal nodes, which are encoded by the Δ^H factors above.

If the universal property for $\Xi\mathcal{D}_\Lambda^s$ were available, this new family of curves $\Omega\text{Pl}(\mathcal{X}) \rightarrow U \subset \Omega\overline{\mathcal{M}\mathcal{D}}_\Lambda^s \times \Delta^H$ would give a map $U \rightarrow \Xi\mathcal{D}_\Lambda^s$. Unfortunately, the universal property is not available yet, as we wish to give $\Xi\mathcal{D}_\Lambda^s$ its complex structure, and then use it in establishing its universal property. Instead, we define a *plumbing map* $\Omega\text{Pl}: U \rightarrow \Xi\mathcal{D}_\Lambda^s$ stratum-by-stratum, using the universal property for each stratum $\Xi\mathcal{D}_\Lambda^{\Lambda',s}$, parameterizing equisingular curves with dual graph Λ' . As the plumbing construction is not canonical, neither is the plumbing map. Indeed it is a family of maps which depend on several choices.

In Section 12.4, we define the plumbing maps, and then in Section 12.7, we show that they are local homeomorphisms, which yield the following main results of this section.

Theorem 12.1. *For any point P in the deepest stratum $\Omega\mathcal{M}\mathcal{D}_\Lambda^{\Lambda',s}$ of the model domain $\Omega\overline{\mathcal{M}\mathcal{D}}_\Lambda^s$, there exists a neighborhood U of $P \times (0, \dots, 0) \in \Omega\overline{\mathcal{M}\mathcal{D}}_\Lambda^s \times \Delta^H$ and a plumbing map*

$$\Omega\text{Pl}: U \rightarrow \Xi\mathcal{D}_\Lambda^s$$

which is a local homeomorphism. This map preserves the stratifications (10.3) and (12.3), is holomorphic on each stratum and is $\text{Tw}_\Lambda/\text{Tw}_\Lambda^s$ -equivariant. Moreover, the plumbing map ΩPl is \mathbb{C}^ -equivariant and descends to a plumbing map*

$$\text{Pl}: U/\mathbb{C}^* \rightarrow \mathbb{P}\Xi\mathcal{D}_\Lambda^s$$

which is also holomorphic, stratum-preserving and $\text{Tw}_\Lambda/\text{Tw}_\Lambda^s$ -equivariant.

By doing such a plumbing construction around the deepest strata in the model domains $\Omega\overline{\mathcal{M}\mathcal{D}}_{\Lambda'}$ for all possible undegenerations Λ' of Λ , we get a collection of local homeomorphisms to various neighborhoods in the Dehn spaces $\Xi\mathcal{D}_{\Lambda'}$. We will then check that these neighborhoods cover the entire model domain and provide it with a complex structure.

Theorem 12.2. *The collection of all plumbing maps gives an atlas of charts which makes $\Xi\mathcal{D}_\Lambda^s$ a complex manifold, makes $\Xi\mathcal{D}_\Lambda$ a complex analytic space with abelian quotient singularities, and provides a family of (simple) marked multi-scale differentials over these spaces (and likewise for the projectivized versions $\mathbb{P}\Xi\mathcal{D}_\Lambda^s$ and $\mathbb{P}\Xi\mathcal{D}_\Lambda$).*

12.2. The setup and notation for plumbing. We now set up the notation for the neighborhoods in which the plumbing construction is performed, as well as for the plumbing fixtures we need. We fix as in the proof of Proposition 10.10 for the remainder of the plumbing construction a base surface $(X_0, \boldsymbol{\eta}_0) \in \Omega\mathcal{B}_\Lambda$ and a local chart $\Psi: \Delta_\epsilon^M \rightarrow \Omega\mathcal{B}_\Lambda$, where $M = \dim \Omega\mathcal{B}_\Lambda$, parameterizing a neighborhood of $\Psi(\mathbf{0}) = (X_0, \boldsymbol{\eta}_0)$. We choose a lift $\tilde{\Psi}: \Delta_\epsilon^M \rightarrow \Omega\overline{\mathcal{M}\mathcal{D}}_\Lambda^s$, which determines a trivialization

$$\Phi: \Delta_\epsilon^M \times \mathbb{C}^{L(\Lambda)} \times \mathbb{C}^* \rightarrow \Omega\overline{\mathcal{M}\mathcal{D}}_\Lambda^s.$$

We identify the universal curve $\Phi^*(\mathcal{X})$ over $\Delta_\epsilon^M \times \mathbb{C}^{L(\Lambda)} \times \mathbb{C}^*$ with the product family of curves

$$(12.4) \quad \tilde{\Psi}^*(\mathcal{X}) \times \mathbb{C}^{L(\Lambda)} \times \mathbb{C}^* \rightarrow \Delta_\epsilon^M \times \mathbb{C}^{L(\Lambda)} \times \mathbb{C}^* .$$

On $\Phi^*(\mathcal{X}) \cong \tilde{\Psi}^*(\mathcal{X}) \times \mathbb{C}^{L(\Lambda)} \times \mathbb{C}^*$, we define the universal form $\mathbf{t}*\boldsymbol{\eta}$ (using the definition of the action (11.1)), where $\boldsymbol{\eta}$ is the tautological form over $\tilde{\Psi}^*(\mathcal{X})$, with zeros and poles only at the marked points and along the nodal sections, and where $\mathbf{t} \in \mathbb{C}^{L(\Lambda)} \times \mathbb{C}^*$ is the scaling parameter of the model differential. We will subsequently work on

$$\mathcal{W}_\epsilon = \Delta_\epsilon^M \times \Delta_\epsilon^N \times \mathbb{C}^* ,$$

for $\epsilon = \epsilon(X_0, \boldsymbol{\eta}_0)$ sufficiently small and to be determined (first by Theorem 12.3, and then to be reduced a finite number of times in the course of the construction). In the remainder of the section, we will make the above identification implicit and simply write $\mathcal{X} \rightarrow \mathcal{W}_\epsilon$ for the restriction of the family (12.4) to the domain of the chart $\Phi: \mathcal{W}_\epsilon \rightarrow \Omega\overline{\mathcal{M}}\overline{\mathcal{D}}_\Lambda^s$. We will denote points in \mathcal{W}_ϵ as (\mathbf{w}, \mathbf{t}) with $\mathbf{w} \in \Delta_\epsilon^M$ or by $([X, \boldsymbol{\eta}], \mathbf{t})$.

We now introduce the notation for our standard annuli and plumbing fixtures, and families of such. We define the standard round annulus

$$A_{\delta_1, \delta_2} = \{z \in \mathbb{C} : \delta_1 < |z| < \delta_2\}$$

and use the base point $p = \sqrt{\delta_1 \delta_2} \in A_{\delta_1, \delta_2}$ unless specified differently. For $\delta = \delta(X_0, \boldsymbol{\eta}_0)$ to be determined below, and $s \in \mathbb{C}$, we define the standard *plumbing fixture*

$$(12.5) \quad V(s) = \{(u, v) \in \Delta_\delta^2 : uv = s\}$$

together with the *top plumbing annulus* and *bottom plumbing annulus*

$$A^+ = \{\delta/R < |u| < \delta\} \quad \text{and} \quad A^- = \{\delta/R < |v| < \delta\}$$

for some R still to be specified. We define $p^\pm = \delta/\sqrt{R} \in A^\pm$ respectively. For $s = 0$ the plumbing fixture is simply

$$V(0) = \Delta_\delta^+ \cup \Delta_\delta^- ,$$

i.e. two disks joined at a node, with u being the coordinate on Δ_δ^+ and v on Δ_δ^- .

We denote by \mathcal{A}^\pm the trivial families of top and bottom plumbing annuli over a base specified by the context, typically \mathcal{W}_ϵ .

For each vertical edge e of $\Gamma = \Gamma(\Lambda)$, we define the *plumbing fixture* $\mathbb{V}_e \rightarrow \mathcal{W}_\epsilon$ to be the standard model family of nodal curves over \mathcal{W}_ϵ :

$$(12.6) \quad \mathbb{V}_e = \left\{ (\mathbf{w}, \mathbf{t}, u, v) \in \mathcal{W}_\epsilon \times \Delta_\delta^2 : uv = \prod_{i=\ell(e^-)}^{\ell(e^+)-1} t_i^{m_{e,i}} \right\} ,$$

where the integers $m_{e,i}$ are defined in (6.7). Note that the fiber of $\mathbb{V}_e \rightarrow \mathcal{W}_\epsilon$ is an annulus if each t_i in the product in (12.6) is non-zero, and a pair of disks meeting at a node otherwise.

We equip \mathbb{V}_e with the relative one-form Ω_e , given in coordinates by

$$(12.7) \quad \Omega_e = (t_{[\ell(e^+)]} \cdot u^{\kappa_e} - r'_e) \frac{du}{u} \quad \text{and} \quad \Omega_e = (-t_{[\ell(e^-)]} \cdot v^{-\kappa_e} + r'_e) \frac{dv}{v} ,$$

where the notation $t_{\lceil \ell(e^+) \rceil}$ was introduced in (11.1) and where r'_e denotes the residue after adding the modifying differential (see Theorem 12.3 below). The two expressions agree, if $uv \neq 0$.

In what follows we will carefully choose the sizes of δ and ϵ for the plumbing fixtures in (12.6), so that the moduli of the annuli are sufficiently large, as required by some later parts of the plumbing construction. We start by fixing a constant $R > 1$, and denote $\delta = \delta(R)$ and $\epsilon = \epsilon(R)$ the corresponding constants provided by Theorem 12.3 below.

We define families of disjoint annuli $\mathcal{A}_e^+, \mathcal{A}_e^- \subset \mathbb{V}_e$ by

$$\begin{aligned} \mathcal{A}_e^+ &= \{(\mathbf{w}, \mathbf{t}, u, v) : |w_i|, |t_j| < \epsilon \text{ for all } i, j, \text{ and } \delta/R < |u| < \delta\} \quad \text{and} \\ \mathcal{A}_e^- &= \{(\mathbf{w}, \mathbf{t}, u, v) : |w_i|, |t_j| < \epsilon \text{ for all } i, j, \text{ and } \delta/R < |v| < \delta\}. \end{aligned}$$

We will refer to \mathcal{A}_e^+ and \mathcal{A}_e^- as the *top and bottom plumbing annuli* corresponding to the vertical edge e . For each half-edge h (in this section we use “h” to denote half-edge, as opposed to horizontal) corresponding to a marked zero of order m_h , we define a family of disks, equipped with relative one-forms

$$\mathbb{D}_h = \mathcal{W}_\epsilon \times \Delta_\delta, \quad \Omega_h = z^{m_h} dz.$$

We define $\mathcal{A}_h \subset \mathbb{D}_h$ to be the family of annuli by

$$\mathcal{A}_h = \mathcal{W}_\epsilon \times A_{\delta/R, \delta}.$$

12.3. Standard coordinates. We now apply the normal form theorems of Section 4 to the universal family $\mathcal{X} \rightarrow \mathcal{W}_\epsilon$. Recall that the normal form depends on the choice of a reference section and a root of unity ζ_j . To avoid overloading notation we keep this ambiguity of the construction for the moment and return to this topic in connection with prong-matchings when defining the plumbing map in Section 12.5.

By an application of Strebel’s original result (Theorem 4.1) in families, we know that for some $\delta_1 > 0$ and for each node there exist local coordinates

$$\phi_e^+ : \mathcal{W}_\epsilon \times \Delta_{\delta_1} \rightarrow \mathcal{X}_{\ell(e^+)} \quad \text{and} \quad \phi_e^- : \mathcal{W}_\epsilon \times \Delta_{\delta_1} \rightarrow \mathcal{X}_{\ell(e^-)}$$

(to keep the notation manageable, we write simply $\mathcal{X}_{\ell(e^+)}$ instead of $\mathcal{X}_{\ell(e^+)}$) whose restrictions to $\mathcal{W}_\epsilon \times \{0\}$ correspond to the loci Q_e^+ and Q_e^- respectively, and which put the form $\mathbf{t} * \boldsymbol{\eta}$ in the standard form. For a vertical node q_e , this standard form is

$$\begin{aligned} (\phi_e^+)^*(\mathbf{t} * \boldsymbol{\eta}) &= t_{\lceil \ell(e^+) \rceil} \cdot z^{\kappa_e} \frac{dz}{z} \quad \text{and} \\ (\phi_e^-)^*(\mathbf{t} * \boldsymbol{\eta}) &= (-t_{\lceil \ell(e^-) \rceil} \cdot z^{-\kappa_e} + r_e(\mathbf{t})) \frac{dz}{z}, \end{aligned}$$

where $r_e : \mathcal{W}_\epsilon \rightarrow \mathbb{C}$ is the residue function

$$r_e(\mathbf{t}) = \text{Res}_{q_e^-}(\mathbf{t} * \boldsymbol{\eta}).$$

For a horizontal node q_e , the standard form is

$$(\phi_e^+)^*(\mathbf{t} * \boldsymbol{\eta}) = -\frac{r_e(\mathbf{t})}{z} dz \quad \text{and} \quad (\phi_e^-)^*(\mathbf{t} * \boldsymbol{\eta}) = \frac{r_e(\mathbf{t})}{z} dz.$$

In general, the modified differential $\mathbf{t} * \boldsymbol{\eta} + \xi$ does not admit such a simple standard form in a neighborhood of a vertical node. Consider a vertical node with top section

$q_e^+ : \mathcal{W}_\epsilon \rightarrow \mathcal{X}_{\ell(e^+)}$, which is a zero of order $\kappa_e - 1$ of $\mathbf{t} * \boldsymbol{\eta}$. Then this zero breaks up into a simple pole and κ_e nearby zeros of the differential $\mathbf{t} * \boldsymbol{\eta} + \xi$. These extraneous nearby zeros should not belong to our plumbed family, so we will construct a family of disks \mathcal{U}_e^+ containing these nearby zeros, which we will then cut out of \mathcal{X} . These disks will be bounded by a family of annuli \mathcal{B}_e^+ and come with a family of *gluing maps* $\Upsilon_e^+ : \mathcal{A}_e^+ \rightarrow \mathcal{B}_e^+$ putting $\mathbf{t} * \boldsymbol{\eta} + \xi$ into a standard form on a family of annuli over \mathcal{W}_ϵ . These objects are constructed in Theorem 12.3 below. This is the basic analytic ingredient in our plumbing construction. In Section 12.4, we will use these gluing maps to glue in the standard plumbing fixture \mathbb{V}_e defined above.

Adding the modifying differential ξ creates a similar problem at the zero sections \mathbf{z} of \mathcal{X} . When the modifying differential is added, a zero of order m_h breaks into m_h nearby zeros, but we wish to construct a family where the order of the zero remains constant. The solution is similar, that is, we construct below a family of disks \mathcal{U}_h around $\mathbf{z}(h)$, gluing maps that put $\mathbf{t} * \boldsymbol{\eta}$ into the standard form on a family of annuli surrounding \mathcal{U}_h . In Section 12.4 we will then cut out these disks and glue in a standard family of disks \mathbb{D}_h .

Theorem 12.3. *For any $R > 1$, there exist constants $\epsilon, \delta > 0$ such that for each vertical edge e and for each half-edge h of Γ there are families of conformal maps of annuli*

$$\begin{aligned} v_e^+ &: \mathcal{W}_\epsilon \times A_{\delta/R, \delta} \rightarrow \mathcal{X}_{\ell(e^+)}, \\ v_e^- &: \mathcal{W}_\epsilon \times A_{\delta/R, \delta} \rightarrow \mathcal{X}_{\ell(e^-)}, \quad \text{and} \\ v_h &: \mathcal{W}_\epsilon \times A_{\delta/R, \delta} \rightarrow \mathcal{X}_{\ell(h)}. \end{aligned}$$

These families all contain the identity map on \mathcal{W}_ϵ , and have the following properties:

- (i) *The images of v_e^+ , v_e^- , v_h are families of annuli $\mathcal{B}_e^+, \mathcal{B}_e^-, \mathcal{B}_h$ that do not contain any zeros or poles of $(\mathcal{X}, \mathbf{t} * \boldsymbol{\eta})$. The families of annuli $\mathcal{B}_e^+, \mathcal{B}_e^-$, and \mathcal{B}_h bound families of disks $\mathcal{U}_e^+, \mathcal{U}_e^-$, and \mathcal{U}_h , respectively, where*

$$Q_e^+ \subset \mathcal{U}_e^+ \subset \mathcal{X}_{\ell(e^+)}, \quad Q_e^- \subset \mathcal{U}_e^- \subset \mathcal{X}_{\ell(e^-)}, \quad \text{and} \quad \mathbf{z}(h) \subset \mathcal{U}_h \subset \mathcal{X}_{\ell(h)}.$$

- (ii) *The pullback of $\mathbf{t} * \boldsymbol{\eta} + \xi$ under each of the maps v_e^+ , v_e^- , and v_h , has the standard form, that is*

$$\begin{aligned} (v_e^+)^*(\mathbf{t} * \boldsymbol{\eta} + \xi) &= (t_{\lceil \ell(e^+) \rceil} \cdot z^{\kappa_e} - r'_e(\mathbf{t})) \frac{dz}{z}, \\ (v_e^-)^*(\mathbf{t} * \boldsymbol{\eta} + \xi) &= (-t_{\lceil \ell(e^-) \rceil} \cdot z^{-\kappa_e} + r'_e(\mathbf{t})) \frac{dz}{z}, \quad \text{and} \\ v_h^*(\mathbf{t} * \boldsymbol{\eta} + \xi) &= t_{\lceil \ell(h) \rceil} \cdot z^{m_h} dz, \end{aligned}$$

where $r'_e(\mathbf{t}) = \text{Res}_{q_e^-}(\mathbf{t} * \boldsymbol{\eta} + \xi)$.

- (iii) *The maps v_e^+ , v_e^- , and v_h agree with the corresponding maps ϕ_e^+ , ϕ_e^- , and ϕ of Theorem 4.1 on the subset of $\mathcal{W}_\epsilon \times A_{\delta/R, \delta}$ where $t_{L-1} = \cdots = t_{-N} = 0$ with $L = \ell(e^\pm)$ or $L = \ell(h)$ respectively.*

Moreover, we may take ϵ, δ sufficiently small so that the maps v_e^\pm and v_h are injective and have mutually disjoint images.

The location of these annuli is illustrated in the left part of Figure 7. The images of the marked points p_e^\pm and p_h in \mathcal{X} are denoted by b_e^\pm and b_h respectively for each vertical edge or half-edge.

Proof. In the ϕ_e^+ -coordinates, the modifying differential ξ becomes

$$(\phi_e^+)^*\xi = t_{[\ell(e^+)]} \cdot \alpha_e \frac{du}{u},$$

where α_e is a holomorphic function on the product $\mathcal{W}_\epsilon \times \Delta_\epsilon^L \times \Delta_{\delta_1}$ satisfying that $t_{[\ell(e^+)]} \cdot \alpha_e(\mathbf{w}, \mathbf{t}, 0) = -r'_e(\mathbf{w}, \mathbf{t})$ and $\alpha_e(\mathbf{w}, \mathbf{0}, z) \equiv 0$. By Theorem 4.2, after possibly decreasing ϵ , there is a family of conformal maps

$$\psi_e: \mathcal{W}_\epsilon \times A_{\delta/R, \delta} \rightarrow \mathcal{W}_\epsilon \times \Delta_{\delta_1},$$

which contain the identity on \mathcal{W}_ϵ , fix the section $\mathcal{W}_\epsilon \times \{b\}$, and put $(\phi_e^+)^*(\mathbf{t} * \boldsymbol{\eta} + \xi)$ in the desired standard form as follows:

$$\begin{aligned} (\phi_e^+ \circ \psi_e)^*(\mathbf{t} * \boldsymbol{\eta} + \xi) &= \psi_e^* (t_{[\ell(e^+)]} \cdot z^{\kappa_e} + t_{[\ell(e^+)]} \cdot \alpha_e(\mathbf{w}, \mathbf{t}, z)) \frac{dz}{z}, \\ &= (t_{[\ell(e^+)]} \cdot z^{\kappa_e} - r'_e(\mathbf{t})) \frac{dz}{z}. \end{aligned}$$

Since $t_{[\ell(e^+)]}^{-1} * r'_e(\mathbf{t}) = -\alpha_e$ is holomorphic, we may in particular choose ϵ small enough so that the zeros of the rightmost form belong to the disk of radius δ/R . When all the parameters $t_{-N} = \dots = t_{L-1} = 0$, the modifying differential ξ vanishes on level L , and thus Theorem 4.2 allows us to choose ψ_e to be the identity on this locus. We then define $v_e^+ = \phi_e^+ \circ \psi_e$. The desired family of disks \mathcal{U}_e^+ is then $\phi_e^+(\mathcal{V}_e)$, where \mathcal{V}_e is the bounded component of the complement of the family of annuli $\psi_e(\mathcal{W}_\epsilon \times A_{\delta/R, \delta})$.

The construction of v_e^- is much simpler at a pole, as then we need only to apply Theorem 4.1 to construct a map v_e^- putting $\mathbf{t} * \boldsymbol{\eta} + \xi$ in its standard form. This works in a neighborhood of the node, and we may of course restrict to a family of annuli.

In the case of a half-edge, the construction of the map v_h follows from the same technique. In this case, the modifying differential ξ is holomorphic along the zero section \mathbf{z}_h , so the resulting standard form of $v_h^*(\mathbf{t} * \boldsymbol{\eta} + \xi)$ has no residue.

The sizes of ϵ and δ for which the desired injectivity and disjointness hold can be quantified in terms of the periods of saddle connections in the central fiber of the family. \square

For a horizontal edge e , choose as always the notation \pm for the ends of e arbitrarily. Then we can apply Theorem 4.1 directly to $\mathbf{t} * \boldsymbol{\eta} + \xi$ to obtain

$$v_e^+: \mathcal{W}_\epsilon \times \Delta_{\delta_1} \rightarrow \mathcal{X}_{\ell(e^+)} \quad \text{and} \quad v_e^-: \mathcal{W}_\epsilon \times \Delta_{\delta_1} \rightarrow \mathcal{X}_{\ell(e^-)}$$

such that

$$(v_e^+)^*(\mathbf{t} * \boldsymbol{\eta} + \xi) = -r'_e(\mathbf{t}) \frac{dz}{z} \quad \text{and} \quad (v_e^-)^*(\mathbf{t} * \boldsymbol{\eta} + \xi) = r'_e(\mathbf{t}) \frac{dz}{z},$$

which agree with ϕ_e^\pm when $t_{-N} = \dots = t_{L-1} = 0$, where $L = \ell(e^+) = \ell(e^-)$.

12.4. The plumbing construction. We now present the basic plumbing construction. The plumbing starts from a family $\mathcal{X} \rightarrow \mathcal{W}_\epsilon$ equipped with the universal differential $\mathbf{t} * \boldsymbol{\eta}$ (as defined in Section 12.2) together with a modifying differential ξ , and builds a family of stable differentials $(\mathcal{Y} \rightarrow \mathcal{W}_\epsilon, \omega, \mathbf{z})$ that are nowhere vanishing in the generic fiber, except for the prescribed zeros and poles $\mathbf{z}(h)$. It will be obvious from the construction that $\omega = (\omega/t_{[\bar{i}]})_{i \in L \bullet (\Lambda)}$ is a collection of rescaled differentials compatible with the simple rescaling ensemble R^s defined by \mathbf{t} , so that we actually define a family of simple multi-scale differentials over \mathcal{W}_ϵ .

Recall now the notation for annuli from Section 12.2. We define conformal isomorphisms $\Upsilon_e^\pm: \mathcal{A}_e^\pm \rightarrow \mathcal{B}_e^\pm \subset \mathcal{X}$ by

$$\Upsilon_e^+(\mathbf{w}, \mathbf{t}, u, v) = v_e^+(\mathbf{w}, \mathbf{t}, u) \quad \text{and} \quad \Upsilon_e^-(\mathbf{w}, \mathbf{t}, u, v) = v_e^-(\mathbf{w}, \mathbf{t}, v),$$

where \mathcal{B}_e^\pm and v_e^\pm are defined in Theorem 12.3. These maps identify each Ω_e with $\mathbf{t} * \boldsymbol{\eta} + \xi$ as desired. By abuse of notation, we will refer to both \mathcal{A}_e^+ and its image \mathcal{B}_e^+ as the top plumbing annuli and to both \mathcal{A}_e^- and \mathcal{B}_e^- as the bottom plumbing annuli corresponding to the edge e .

For each half-edge h , we denote the family of conformal isomorphisms provided in Theorem 12.3 by

$$\Upsilon_h = v_h: \mathcal{A}_h \rightarrow \mathcal{B}_h \subset \mathcal{X}.$$

This finally completes the preparation for the first step of the plumbing construction.

First plumbing step. We let $\mathcal{X}' \rightarrow \mathcal{W}_\epsilon$ be the family of curves obtained by removing from \mathcal{X} the families of disks \mathcal{U}_e^\pm and \mathcal{U}_h and attaching each family \mathbb{V}_e and \mathbb{D}_h by identifying the \mathcal{A} -annuli to the \mathcal{B} -annuli via the Υ -gluing maps. As the gluing maps respect the one-forms, the family \mathcal{X}' inherits a relative one-form Ω' .

It remains to plumb the horizontal nodes of \mathcal{X}' . This construction is classical and differs from plumbing the vertical nodes, as each horizontal node requires an auxiliary complex parameter. Our construction will plumb the remaining nodes of the product family

$$\mathcal{X}' \times \Delta_\epsilon^H \rightarrow \mathcal{W}_\epsilon \times \Delta_\epsilon^H,$$

to create a generically smooth family $\mathcal{Y} \rightarrow \mathcal{W}_\epsilon \times \Delta_\epsilon^H = \Delta_\epsilon^{N+M+H} \times \mathbb{C}^*$.

We enumerate the horizontal edges of Γ as e_1, \dots, e_H and define for each e_j the plumbing fixture

$$(12.8) \quad \mathbb{W}_j = \{(\mathbf{w}, \mathbf{t}, \mathbf{x}, u, v) \in \mathcal{W}_\epsilon \times \Delta_\epsilon^H \times \Delta_\delta^2 : uv = x_j\},$$

equipped with the relative holomorphic one-form

$$(12.9) \quad \Omega_j = -r'_{e_j}(\mathbf{t}) \frac{du}{u} = r'_{e_j}(\mathbf{t}) \frac{dv}{v}.$$

We define two families of conformal maps $\Upsilon_j^\pm: \mathbb{W}_j \rightarrow \mathcal{X}'$ by

$$(12.10) \quad \Upsilon_j^+(\mathbf{w}, \mathbf{t}, \mathbf{x}, u, v) = v_{e_j}^+(\mathbf{w}, \mathbf{t}, u) \quad \text{and} \quad \Upsilon_j^-(\mathbf{w}, \mathbf{t}, \mathbf{x}, u, v) = v_{e_j}^-(\mathbf{w}, \mathbf{t}, v),$$

which identify \mathbb{W}_j with two families of annuli $\mathcal{B}_j^\pm \subset \mathcal{X}'$ whose moduli tend to infinity as $x_j \rightarrow 0$. These \mathcal{B}_j^\pm bound two families of disks \mathcal{U}_j^\pm , namely the image under $v_{e_j}^\pm \times \text{id}$

of $\{(\mathbf{w}, \mathbf{t}, \mathbf{x}, z) \in \mathcal{W}_\epsilon \times \Delta_\epsilon^H \times \Delta_\delta : |z| < |x_j|/\delta\}$. We can thus proceed with the second step of the plumbing construction.

Second plumbing step. Let $\mathcal{Y} \rightarrow \mathcal{W}_\epsilon \times \Delta_\epsilon^H$ be the family of curves obtained by removing from \mathcal{X}' the families of disks $\mathcal{U}_{e_j}^\pm$ and identifying each \mathbb{W}_j with $\mathcal{B}_{e_j}^+$ and $\mathcal{B}_{e_j}^-$ by the gluing maps Υ_j^+ and Υ_j^- . As the gluing maps identify the one-forms, the family \mathcal{Y} inherits a relative one-form ω .

We denote the plumbing annuli as subsurfaces of \mathcal{Y} by \mathcal{C}_e^\pm and \mathcal{C}_h , and denote the images of the marked points p_e^\pm and p_h by c_e^\pm and c_h , respectively. These points are defined near the corresponding vertical and horizontal nodes, but for the latter the sign is an arbitrary auxiliary choice. The final result of plumbing is illustrated on the right of Figure 7.

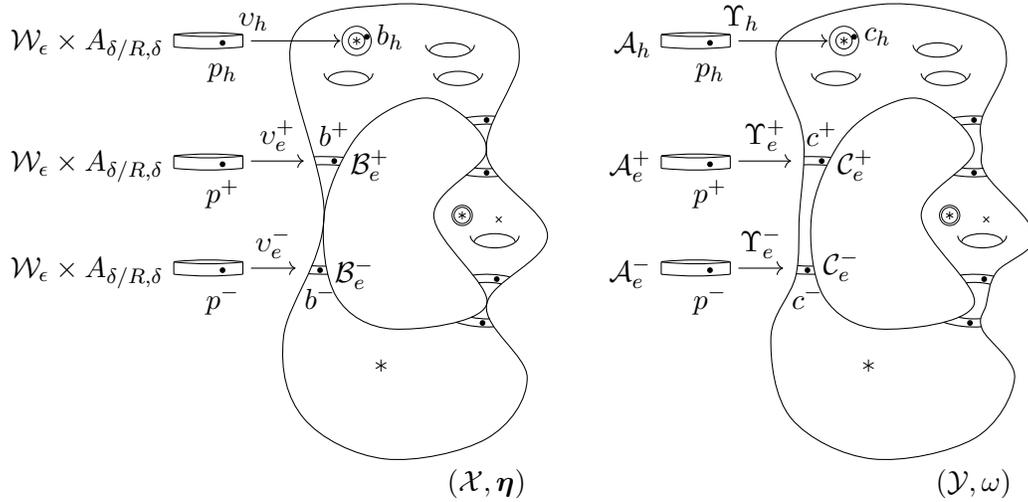


FIGURE 7. The general plumbing construction for our running example of Section 2.6.

We will later use the following consequence of the construction and the fact that the modifying differential ξ on level i depends only on the levels below i and the topological data, see Corollary 11.4.

Corollary 12.4. *For each edge e the location of the family of annuli $\mathcal{B}_e^+ \subset \mathcal{X}_{\ell(e^+)}$ depends only on the subsurfaces $(\mathcal{X}_{(i)}, \eta_{(i)})$ and on the values of t_i for $i \leq \ell(e^+)$.*

12.5. Definition of the plumbing map. In the preceding subsection we started with an arbitrary point $(X_0, \eta_0) \in \Omega\mathcal{B}_\Lambda$ and the tautological family $(\mathcal{X} \rightarrow U, \eta, \mathbf{z}, \sigma, f)$ over a neighborhood $U = \mathcal{W}_\epsilon \times \Delta_\epsilon^H$ of (X_0, η_0) in $\Omega\overline{\mathcal{M}}\mathcal{D}_\Lambda^s$ (recall the description of \mathcal{W}_ϵ in Section 12.2). We then constructed the plumbed family $(\mathcal{Y} \rightarrow U, \omega, \mathbf{z})$, leaving the location of the base points in the plumbing construction to be determined. We will now specify the location of these base points together with prong-matchings and markings on \mathcal{Y} . At the same time we will use the plumbed family together with the universal property of the strata of $\Xi\mathcal{D}_\Lambda^s$ to define a map $\Omega\text{Pl}: U \rightarrow \Xi\mathcal{D}_\Lambda^s$ that is a priori

holomorphic on each stratum individually. Later on we will establish all the desired properties of this map.

For the application of Theorem 4.1 in Theorem 12.3, for each vertical edge e we choose an arbitrary section $\varsigma_e^-: \mathcal{W}_e \rightarrow \mathcal{X}$ (with image contained in \mathcal{B}_e^-) so that the identification map of the bottom plumbing annulus given by Theorem 12.3 satisfies $v_e^-(w, p^-) = \varsigma_e^-(w)$. Similarly, for each marked point we choose an arbitrary section $\varsigma_h: \mathcal{W}_e \rightarrow \mathcal{X}$ (with image in \mathcal{B}_h) that perturbs p_h . For convenience we fix these sections by requiring their relative periods to be fixed and independent of w , as specified in (11.7). For the application of Theorem 4.2 on the upper end of the plumbing fixture in Theorem 12.3, we may choose an arbitrary section $\varsigma_e^+: \mathcal{W}_e \rightarrow \mathcal{X}$ (with image in \mathcal{B}_e^+) that perturbs $\zeta_j p_e^+$, i.e. such that $v_e^+(w, \zeta_j p_e^+) = \varsigma_e^+(w)$ if the plumbing fixture \mathbb{V}_e is nodal at $w \in \mathcal{W}_e$. The ζ_j are chosen in the unique way such that the collection of prong-matchings on each plumbing fixture given by Proposition 8.4 for the differentials $\eta_{(i)}$ is precisely the collection of prong-matchings σ given as part of the input datum, i.e. from the family of model differentials over U .

. Note that the weldings of the degenerate plumbing fixtures are naturally diffeomorphic to the corresponding plumbed annuli in the family \mathcal{Y} . Via this identification, the marking of $\mathcal{X} \rightarrow U$, i.e. the section f of $\text{Mark}(\overline{\mathcal{X}}_\sigma/\overline{U}; \text{Tw}_\Lambda)$, naturally induces a section, denoted also by f , of $\text{Mark}(\mathcal{Y}/\overline{U}; \text{Tw}_\Lambda)$, i.e. a marking of the family $\mathcal{Y} \rightarrow U$.

We now discuss the plumbing map $\Omega \text{Pl}: U^J \rightarrow \Omega \mathcal{B}_{\Lambda_J} / \text{Tw}_\Lambda^s$ for each $J \subseteq L(\Lambda)$ individually on the set $U^J = U \cap \left(\tilde{\Psi}^{-1}(\Omega \overline{\mathcal{M}D}_\Lambda^{s, \Lambda_J}) \times \Delta^H \right)$ with notation as in Section 10.3. First, for the deepest boundary stratum, i.e. for $J = L(\Lambda)$, we simply define

$$\Omega \text{Pl}(\mathcal{X} \rightarrow U^{L(\Lambda)}, \boldsymbol{\eta}, \mathbf{0}, \boldsymbol{\sigma}, f) = (\mathcal{X} \rightarrow U^{L(\Lambda)}, \boldsymbol{\eta}, \boldsymbol{\sigma}, f).$$

Next, for the open stratum $J = \emptyset$, we define

$$\Omega \text{Pl}(\mathcal{X} \rightarrow U^\emptyset, \boldsymbol{\eta}, \mathbf{t}, \boldsymbol{\sigma}, f) = (\mathcal{Y} \rightarrow U^\emptyset, \boldsymbol{\omega}, -, f)$$

since there are no prong-matchings to be recorded for the non-semipersistent nodes.

A stratum for general J combines the two extreme cases. We provide $\mathcal{Y} \rightarrow U^J$ with the collection of differential forms $\boldsymbol{\omega}^J = (\omega_{(i)}^J)_{i \in L \bullet (\Lambda_J)}$ defined as follows. We denote by $J_{\geq i}$ the subset of J of levels greater than or equal to i and define

$$(12.11) \quad \omega_{(i)}^J = \lim_{\substack{t_k \rightarrow 0 \\ \text{for all } k \in J_{\geq i}}} \left(\boldsymbol{\omega}(\mathbf{t}) \cdot \prod_{k \in J_{\geq i}} t_k^{-a_k} \right)$$

restricted to the level i component of $\mathcal{Y} \rightarrow U^J$. By construction of $\boldsymbol{\omega}$ as gluing of $\mathbf{t} * \boldsymbol{\eta}$, the collection of differential forms $\boldsymbol{\omega}^J$ is holomorphic and non-zero outside the prescribed zero and pole sections $\mathbf{z}(h)$. We then let

$$\Omega \text{Pl}(\mathcal{X} \rightarrow U^J, \boldsymbol{\eta}, \{t_j\}_{j \notin J}, \boldsymbol{\sigma}, f) = (\mathcal{Y} \rightarrow U^J, \boldsymbol{\omega}^J, \sigma_J, f),$$

where σ_J are the prong-matchings at the nodes crossing some level $j \in J$.

12.6. Continuity of the plumbing map. Suppose that $(X_n, \boldsymbol{\eta}_n, \mathbf{t}_n)$ is a sequence in $\overline{\mathcal{M}D}_\Lambda^s$ that converges to $(X, \boldsymbol{\eta}, \mathbf{t})$ as in Proposition 10.11. Our next goal is to prove the following result.

Proposition 12.5. *The sequence $(Y_n, \omega_n) = \Omega \text{Pl}(X_n, \eta_n, t_n)$ of marked multi-scale differentials converges to $(Y, \omega) = \Omega \text{Pl}(X, \eta, t)$.*

We denote by $g_n: (K_n \subset X \setminus \mathbf{z}) \rightarrow X_n$ the conformal maps exhibiting the convergence of the sequence X_n . We will show that eventually the images $g_n(K_n)$ contain the top and bottom plumbing annuli, but we will still have to extend the map g_n to the interior of the plumbing annuli. To address this problem, we first prove the following local extension statements for quasi-conformal maps.

We define the sets

$$V^+(s) = \{(u, v) \in V(s) : |u| > |s|^{1/2}\} \text{ and } V^-(s) = \{(u, v) \in V(s) : |u| < |s|^{1/2}\}$$

to be the upper and lower halves of the plumbing fixture respectively. Let α be the central curve $|x| = \delta/R^{1/2}$ of the annulus $A^\pm \subset \Delta_\delta^\pm$. We consider a sequence $\omega_n = (u^k + r_n/u)du$ of differential forms on the plumbing fixtures $V(s_n)$, with r_n and s_n converging to 0 as n goes to infinity.

The following two lemmas both state that we can extend quasi-conformally inwards in the disk, starting at the curve α . Note that since $k \geq 0$, the second lemma is metrically rather an extension “outward”, i.e. towards the pole of η .

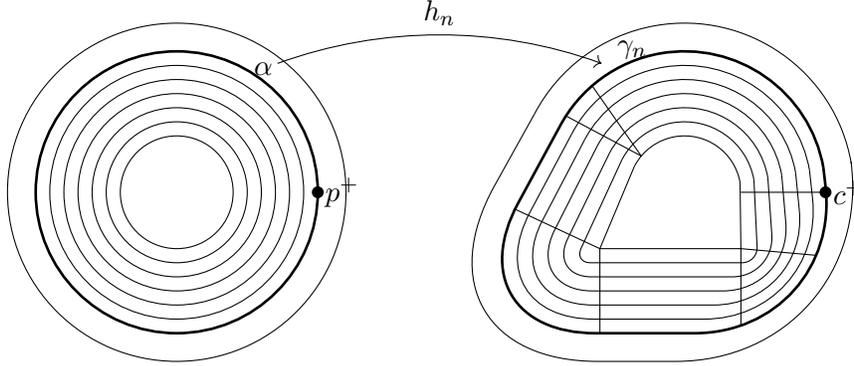


FIGURE 8. Extend from α towards the interior of the disk.

Lemma 12.6. *Let $h_n: A^+ \rightarrow V(s_n)$ be a sequence of conformal maps converging to the identity and such that $h_n^* \omega_n$ converges to $\eta = x^\kappa \frac{dx}{x}$ on A^+ . Then there exist quasi-conformal maps $\tilde{h}_n: \tilde{K}_n \rightarrow V^+(s_n)$ on an exhaustion \tilde{K}_n of Δ_δ^+ that agree with h_n on the subsurface $\{|x| > \delta/R^{1/2}\}$ exterior of α , with quasi-conformal dilatation tending to zero and with $\tilde{h}_n^* \omega_n$ converging to η in C^1 .*

In the second statement we require that r_n converges to r , with r not necessarily zero.

Lemma 12.7. *Let $h_n: A^- \rightarrow V(s_n)$ be a sequence of conformal maps converging to the identity such that $s_n^{-\kappa} h_n^* \omega_n$ converges to $\eta = (-x^{-\kappa} + r) \frac{dx}{x}$ on A^- . Then there exist quasi-conformal maps $\tilde{h}_n: \tilde{K}_n \rightarrow V^-(s_n)$ on an exhaustion \tilde{K}_n of Δ_δ^- that agree*

with h_n on the subsurface $\{|x| > \delta/R^{1/2}\}$ exterior of α , with quasi-conformal dilatation tending to zero and with $\tilde{h}_n^* \omega_n$ converging to η in C^1 .

Proof of Lemma 12.6 and Lemma 12.7. To prove Lemma 12.6, the idea is to define h_n by polar coordinates using equidistant curves to the image curve $\gamma_n = h_n(\alpha) \subset V(t_n)$ as in Figure 8. We let $D = \delta/R^{1/2}$ and define the function $I(d) = d^\kappa/\kappa$. We let γ_n^s be the equidistant curve to γ_n with ω_n -distance s . We parameterize this curve by the interval $[0, 2\pi]$ with constant ω_n -speed, starting at the ray through $h_n(p^+)$ perpendicular to the curves γ_n^s . We parameterize the subsurface interior bounded by α using polar coordinates $x = de^{i\vartheta}$ for $d \in [0, D]$.

We want to define $\tilde{h}_n(x)$ to be the point on the curve $\gamma_n^{I(D)-I(d)}$ that lies on the ray perpendicular to γ_n starting at “time” $\vartheta \in [0, 2\pi]$ in the chosen parameterizations. This procedure is well-defined and provides a diffeomorphism as long as the perpendicular rays are disjoint. We first show that this is true for d an exhaustion of the interval $[0, D]$ as $n \rightarrow \infty$. The curve α has constant η -curvature $1/I(D)$ and its total curvature is $2\pi\kappa$. Since h_n converges uniformly on compact sets to the identity by Lemma 3.4, the derivatives of h_n also converge, and hence the ω_n -curvature $\gamma_n(0)$ tends to constant curvature $1/I(D)$. By definition of curvature, the domain on which the given procedure defines a diffeomorphism converges to the punctured unit disk and we obtain the desired exhaustion. Since the map \tilde{h}_n converges to the identity uniformly on every compact set, the quasi-conformal dilatation tends to zero.

For Lemma 12.7 we view the punctured unit disk with the η -metric as the exterior of the curve α in κ slit planes, with a slit of holonomy $r \in \mathbb{C}$, glued cyclically (see e.g. the lower half of Figure 15 in [BCGGM18]). The η -curvature of α is not constant, unless $r = 0$, but the curvature has always the same sign. Since the conformal maps h_n converge uniformly on compact sets to the identity, the curvature of $\gamma_n = h_n(\alpha)$ will also have eventually the same sign. This suffices to define \tilde{h}_n by polar coordinates using equidistant curves to γ_n on the whole punctured unit disk. \square

Proof of Proposition 12.5. Let $J(\mathbf{t}) \subseteq L(\Lambda)$ as in Proposition 10.11 be the set of levels i where $t_i = 0$. Moreover, let the rescaling coefficients $\mathbf{d}_n \in \mathbb{C}^{J(\mathbf{t})}$ together with $c_{n,i} = e(d_{n,i})$, and let the maps $g_n: \overline{X}_{\mathbf{d}_n, \sigma} \rightarrow \overline{X}_n$ be as in Proposition 10.11. By property (i) and the definition of the plumbing construction the pinched multicurve on Y_n is (for n large enough) a vertical undegeneration of the pinched multicurve on Y . In fact the undegeneration is given by the map $\delta_{J(\mathbf{t}_n)}|_{J(\mathbf{t})}: J(\mathbf{t}) \rightarrow J(\mathbf{t}_n)$.

We claim that for n large enough the image of K_n under g_n contains the subsurface $X'_n = X_n \setminus (\bigcup_e \mathcal{U}_e \cup \bigcup_h \mathcal{U}_h)$. Note that $\text{Id}_{X'_n}$ identifies X'_n with a subsurface Y'_n of Y_n , but does not identify the forms because of the modifying differentials. Admitting the claim, there are two cases to discuss.

The first case concerns the levels that get identified by the undegeneration δ . For notational simplicity we restrict to the case $J(\mathbf{t}) = L(\Lambda)$ and $J(\mathbf{t}_n) = \emptyset$, i.e. to $\mathbf{t} = \mathbf{0}$ and $(Y, \omega) = (X, \eta)$. We can apply the two preceding lemmas (and a similar extension in the neighborhood of the \mathcal{U}_h) to the maps $h_n = \text{Id}_{X'_n} \circ g_n|_{B^\pm}$ because of the property (iii) of modifying differentials. Patching the maps \tilde{h}_n given by the two preceding lemmas together with $\text{Id}_{X'_n} \circ g_n$ gives quasi-conformal maps $\tilde{g}_n: \tilde{K}_n \rightarrow Y_n$ on an exhaustion \tilde{K}_n

of X . Using Proposition 3.6 we can find $\widehat{K}_n \subseteq \widetilde{K}_n$ still exhausting X and conformal maps $\widehat{g}_n: \widehat{K}_n \rightarrow Y_n$ such that the restrictions of $c_{n,i}\widehat{g}_n^*(\omega_n)$ to the level i subsurface of $(X, \mathbf{z})_{\epsilon_n}$ converge uniformly on compact sets to the restriction of $\eta_{(i)}$. This convergence statement combines Proposition 10.11 (iii), the definition of the conformal topology, and the convergence results in Lemma 12.6 and Lemma 12.7. This gives property (3) of Definition 9.3, and (4) is a restatement of (iv). The extension to \widehat{g}_n using the previous lemmas happened (without loss of generality) on an annulus and thus any twisting (compared to g_n) can be undone on an annulus even closer to the cusp. We can thus use the diffeomorphism h between \overline{X}_n and Y_n (as exhibited in the construction of the plumbing map) to extend \widehat{g}_n to a diffeomorphism $\widehat{g}_n: \overline{Y}_{d_n \cdot \sigma} = \overline{X}_{d_n \cdot \sigma} \rightarrow \overline{X}_n$. By (v) and by the choice of the plumbing construction according to Proposition 8.4, the turning-number-preserving property (5) holds. Finally, property (ii) implies (2) by the transport of markings, as in the definition of the plumbing map.

The second case concerns the levels that remain distinct in the limit. We first restrict to the case $J(\mathbf{t}) = J(\mathbf{t}_n) = L(\Lambda)$. Then the plumbed family is just a holomorphic family of curves and continuity follows trivially. Note that we can express convergence using conformal maps also in this case thanks to Proposition 3.5. The intermediate cases of arbitrary $J(\mathbf{t})$ follow by patching together the conformal maps constructed as in the two extreme cases.

Finally, to prove the claim we may neglect the \mathcal{U}_h since we can even assume g_n to be conformal near the marked points and we can assume (possibly shrinking K_n) that the boundary of K_n consists of curves $\alpha_{e,n}$ homotopic to the core curves of the seams corresponding to vertical edges e . Since the inner curves of \mathcal{B}_e^\pm have non-zero hyperbolic length on X , the length of these inner curves is also non-zero on X_n for n large enough. We thus reduce to show that the hyperbolic length of the image curve $\ell_{X_n}(g_n(\alpha_{e,n}))$ is eventually smaller than an arbitrary positive number. Since the sets K_n form an exhaustion, not only $\ell_X(\alpha_{e,n})$ but also $\ell_{K_n}(\alpha_{e,n})$ tends to zero for $n \rightarrow \infty$ as one can check by computing in the annular cover corresponding to $\alpha_{e,n}$. The Schwarz Lemma applied to g_n now proves the remaining claim. \square

12.7. Plumbing is a local homeomorphism. We first show that the plumbing map is open. The plan is to first show it for the most degenerate boundary strata with dual graph Λ in the range, then prove it for the interior points with a dual graph having only horizontal nodes, and finally combine the two approaches to obtain the result for the intermediate points.

Proposition 12.8. *Given a sequence (Y_n, ω_n) of marked multi-scale differentials converging to $(X, \boldsymbol{\eta}) = \Omega \text{Pl}(X, \boldsymbol{\eta}, \mathbf{0})$, there exists for any n large enough a marked model differential $(X_n, \boldsymbol{\eta}_n, \mathbf{t}_n)$ such that $\Omega \text{Pl}(X_n, \boldsymbol{\eta}_n, \mathbf{t}_n) = (Y_n, \omega_n)$.*

For a surface X that corresponds to a point in the subset \mathcal{W}_ϵ of the model domain we denote by $X_{(\leq i)}^-$ the subsurface consisting of the levels $\leq i$, including the bottom plumbing annulus \mathcal{B}_e^- for $\ell(e^-) \leq i$ of the plumbing fixtures to higher levels (i.e. level $\ell(e^+) > i$), but not the rest of the disks \mathcal{U}_e^- as defined in Theorem 12.3. We exclude also the disks $\mathcal{U}_{e_j}^\pm$ of the horizontal plumbing regions that are precisely at level i , but include all the horizontal plumbing regions at all lower levels. We let $X_{(\leq i)}^+$ be the

subsurface consisting of the levels $\leq i$, including all the plumbing fixtures connecting to higher levels all the way up to the top plumbing annuli \mathcal{B}_e^+ for $\ell(e^+) > i$, and also including the horizontal plumbing regions $\mathcal{U}_{e_j^\pm}$ at levels $\leq i$.

Proof. We can choose representatives of $(X, \boldsymbol{\eta})$ in $\Omega\mathcal{T}_\Lambda^{pm}(\mu)$ and $(Y_n, \boldsymbol{\omega}_n)$ in $\Omega\overline{\mathcal{T}}_{(\Sigma, \mathbf{s})}(\mu)$ so that convergence still holds. We deal only with the case that Y_n is smooth, the general case being easier (since some edges are already nodal and require no unplumbing) but notationally more involved. By the definition of convergence in $\Omega\overline{\mathcal{T}}_{(\Sigma, \mathbf{s})}(\mu)$ as given in Section 9, there is a sequence $\mathbf{d}_n = \{d_{n,i}\} \in \mathbb{C}^{L^\bullet(\Lambda)}$ and a sequence of almost-diffeomorphisms $g_n: \overline{X}_{\mathbf{d}_n, \sigma} \rightarrow Y_n$, defined up to isotopy, which are compatible with the markings, nearly turning-number-preserving and which are conformal on an exhaustion $K_{n,(i)}$ of each level $X_{(i)}$.

We start by choosing the sequence of coordinates \mathbf{t}_n defined in terms of these $d_{n,i}$ by

$$(12.12) \quad t_{n,i} = e \left(\frac{d_{n,i+1} - d_{n,i}}{a_i} \right).$$

Recall that $t_{[i]} = \prod_{j \geq i} t_j^{a_j}$ is defined in Equation (11.1). We write $t_{n,[i]}$ for the corresponding product of the $t_{n,i}$. Since we are not rescaling the top level, $d_{n,0} = 0$ and it follows that $t_{n,[i]} = e(-d_{n,i})$. By definition of convergence,

$$(12.13) \quad e(d_{n,i})g_{n,(i)}^* \omega_n \xrightarrow{n \rightarrow \infty} \eta_{(i)}, \quad \text{i.e.} \quad \frac{1}{t_{n,[i]}} g_{n,(i)}^* \omega_n \xrightarrow{n \rightarrow \infty} \eta_{(i)}.$$

We will now construct inductively the $(X_n, \boldsymbol{\eta}_n)$ such that $\Omega \text{Pl}(X_n, \boldsymbol{\eta}_n, \mathbf{t}_n) = (Y_n, \boldsymbol{\omega}_n)$. Recall that by Corollary 12.4 it makes sense to consider the effect of plumbing only to the bottom part of a surface and to write $\Omega \text{Pl}(X_{n,(\leq i)}^\pm)$, suppressing the dependence on $(\boldsymbol{\eta}_n, \mathbf{t}_n)$ for notational convenience. Note also that the set of connected components of the ϵ_n -thick part of Y_n is the disjoint union of sets of level i components $Y_{n,(i)}$, where $Y_{n,(i)}$ are those components that contain the image of the map g_n restricted to the subsurface $X_{n,(i)}$.

The base case of induction is to pick appropriately the surfaces with the correct bottom level piece $(X_{n,(-N)}^-, \eta_{n,(-N)})$ among all surfaces parameterized by \mathcal{W}_ϵ and to construct a conformal map on the bottom level $h_{n,(-N)}: \Omega \text{Pl}(X_{n,(\leq -N)}^-) \rightarrow Y_{n,(-N)}$ which identifies the two differentials. The second step of the base case is to extend this map by analytic continuation across the horizontal and vertical plumbing annuli to obtain a conformal map $h_{n,(-N)}^+: \Omega \text{Pl}(X_{n,(\leq -N)}^+) \rightarrow Y_{n,(\leq -N)}$.

The inductive step starts with the map $h_{n,(i)}^+: \Omega \text{Pl}(X_{n,(\leq i)}^+) \rightarrow Y_{n,(\leq i)}$. We choose appropriately $(X_{n,(i+1)}, \eta_{n,(i+1)})$ and also the subset of $(x_j) \in \Delta^H$ corresponding to horizontal nodes at levels $\leq i+1$ and construct a conformal map

$$h_{n,(i+1)}: \Omega \text{Pl} \left(X_{n,(\leq i+1)}^-, \eta_{n,(\leq i+1)}, \mathbf{t}_n \right) \rightarrow Y_{n,(\leq i+1)}$$

which identifies the forms and agrees with $h_{n,(i)}^+$ on its domain. We then analytically continue across the plumbing cylinders to get $h_{n,(i+1)}^+: \Omega \text{Pl}(X_{n,(\leq i+1)}^+) \rightarrow Y_{n,(\leq i+1)}$.

This procedure eventually ends at the top level when we have constructed the entire surface X_n together with a conformal isomorphism of $\Omega\text{Pl}(X_n)$ with Y_n .

We start with the details of the construction at the bottom level. The conformal map $g_{n,(-N)}$ is eventually defined on the whole subsurface $X_{(-N)}^-$ but this map only approximately identifies the differentials rescaled as explained in (12.13). We choose a sequence of surfaces $(X_{n,(-N)}, \eta_{n,(-N)})$ of the same topological type as $(X_{(-N)}, \eta_{(-N)})$ such that

$$(12.14) \quad \text{Per}(X_{n,(-N)}, \eta_{n,(-N)}) = \frac{1}{t_{n,(-N)}} \text{Per}_{(-N)}(Y_n, \omega_n)$$

and such that $(X_{n,(-N)}, \eta_{n,(-N)})$ converges to $(X_{(-N)}, \eta_{(-N)})$. In this equation $\text{Per}_{(i)} \subset H^1(\Sigma_i \setminus P_i, Z_i; \mathbb{C})$ denotes the relative periods in the level i subsurface. We may choose such a sequence as the (classical) period map is open. By the properties of the conformal topology (see Section 3.3) there exist maps $g_{n,(-N)}^X: X_{(-N)} \rightarrow X_{n,(-N)}$ that are conformal on a large subsurface. We may assume that this subsurface is $K_{n,(i)}$ by shrinking the subsurface defined above, while maintaining that the sequence forms an exhaustion of $X_{(-N)}$. We apply Theorem 3.7 to the sequences $(g_{n,(-N)}^X)^*(\mathbf{t}_n * \boldsymbol{\eta}_n)$ and $g_{n,(-N)}^* \omega_n$. The compact subsurface $K \subset X_{(-N)}$ is chosen such that $g_{n,(-N)}^X(K)$ contains $X_{n,(-N)}^-$ for n large enough. The existence of such a subsurface was isolated as a claim in the proof of Proposition 12.5. We take as in Theorem 3.7 a subsurface of $X_{(-N)}$ that is slightly larger than K . The composition of the map h_n given by the theorem with $(g_{n,(-N)}^X)^{-1}$ and $g_{n,(-N)}$ provides, for n large enough, a conformal map

$$h_{n,(-N)}: X_{n,(-N)}^- \rightarrow Y_n \quad \text{such that} \quad h_{n,(-N)}^* \omega_n = t_{n,(-N)} \eta_{n,(-N)}$$

as desired.

For the analytic continuation through the thin vertical annuli, recall that the plumbed surface is obtained by gluing for all vertical nodes e the plumbing annuli $V_{n,e} \cong V(\rho_n)$ where $\rho_n = \prod_{i=\ell(e^-)}^{\ell(e^+)-1} t_{n,i}^{m_{e,i}}$, equipped with the standard form $\Omega_{n,e}$ as in (12.7), to X_n . At this point, the gluing map on the bottom plumbing annulus is known, as we have chosen the lower level surface, and the gluing on the top annulus will be known when we have chosen the upper level surface $X_{n,(\ell(e^+))}$. By construction, the composition

$$\nu_{n,e} = h_n \circ v_{n,e}^-: A_{n,e}^- \rightarrow Y_n$$

(where $v_{n,e}^-$ was defined in Theorem 12.3) identifies the form ω_n on Y_n with the standard form $\Omega_{n,e}$ on the bottom plumbing annulus. We show in Lemma 12.9 below that $\nu_{n,e}$ can be analytically continued to a conformal map $\nu_{n,e}: V'_{n,e} \rightarrow Y_n$, where $V'_{n,e} \subset V_{n,e}$ is a round subannulus containing the basepoint $p_{n,e}^+$. *A fortiori* the analytic continuation also identifies ω_n with $\Omega_{n,e}$. The map $\nu_{n,e}$ also maps p^\pm to points c_e^\pm that have to be the points as in the plumbing construction, once we have realized Y_n as a plumbing image.

Before proceeding with the next step we have to decide the location of the point b_e^+ in our candidate surface $X_n = (X_n, \boldsymbol{\eta}_n, \mathbf{t}_n, \boldsymbol{\sigma}_n)$ for the plumbing preimage. (Recall that the possible choices of points $b_e^+ \in B_e^+$ differ by a power of a root of unity ζ_j in the

plumbing construction, and this in turn corresponds to the choice of the local prong-matching $\sigma_{n,e}$.) That is, although we have specified so far only the levels $\leq i$ of X_n , we will limit the subsequent considerations to surfaces with prong-matchings σ_n such that X_n is close to $X = (X, \boldsymbol{\eta}, \mathbf{0}, \boldsymbol{\sigma})$ in the conformal topology. This is necessary for X_n to be a point in \mathcal{W}_ϵ , and also sufficient, since in Section 12.5 the prong-matchings in the plumbing map were chosen to induce the identity on X by analytic continuation and Proposition 8.4.

To complete the bottom level argument, we apply Lemma 12.10 on the extension through the thin part of the horizontal nodes at the bottom level. By definition of $X_{n,(-N)}^-$ the map $h_{n,(-N)}$ is defined on the annuli A_j^\pm around the horizontal nodes and its restriction to them provides the maps v_j^\pm required by the lemma. If we use x_j with

$$(12.15) \quad \frac{1}{2\pi\sqrt{-1}} \left(\log(x_j) + \log(R/\delta^2) \right) = \int_{c^-}^{c^+} \omega_{n,(-N)}/r_j(\mathbf{t}),$$

then the plumbing fixture satisfies the period hypothesis of the lemma. Such a choice with $x_j \in \Delta_\epsilon$ is certainly possible for n large enough since Y_n converges to X , and thus the period ratio on the right-hand side of (12.15) will have a large positive imaginary part.

We now begin the inductive step, assuming that we have constructed conformal maps $h_{n,(i)}^+ : \Omega \text{Pl}(X_{n,(\leq i)}^-) \rightarrow Y_{n,(\leq i)}$. We now wish to construct a sequence of marked model differentials $(X_{n,(i+1)}, \eta_{n,(i+1)})$ converging to $(X_{(i+1)}, \eta_{(i+1)})$ and the conformal maps $h_{n,(i+1)}$. This is similar to the base case. The difference is that we have already constructed maps on the top plumbing annuli of the nodes connecting to level $i+1$ from below, and the new maps must agree on these annuli. To deal with this, we choose the sequence X_n so that the perturbed period coordinates satisfy

$$(12.16) \quad \text{ePP}(X_{n,(i+1)}, \eta_{n,(i+1)}) = \frac{1}{t_{n,(i+1)}} \text{ePP}(Y_{n,(i+1)}, \omega_n).$$

The ePP on the right-hand side is a shorthand to express that we compute on Y_n periods in the same way as in the definition of ePP, i.e. we consider the surface cut open at the (horizontal and vertical) polar ends and use integration at the ‘‘nearby points’’ c_e^+ determined by the induction hypothesis for all vertical cylinders whose top end is on level $(i+1)$. The choice of X_n with the required perturbed periods is possible for n large enough since the perturbed period map is open by Proposition 11.7. The surfaces of level $i+1$ of X_n^{he} with the plumbing performed at the half-edges (i.e. the marked zeros) to make $\mathbf{t} * \boldsymbol{\eta} + \xi$ have zeros of the required order also converge to X_n , and we let $g_{n,(i+1)}^X : X \rightarrow X_n^{\text{he}}$ be the maps exhibiting convergence, i.e. $g_{n,(i+1)}^X$ and $g_{n,(i+1)} : X \rightarrow Y_n$ are conformal on some exhaustion by compact sets $K_{n,(i+1)}$. The rest of the argument, by applying Theorem 3.7, now works as in the base case.

To specify X_n as a marked surface, we use the map $\varphi^{-1}|_{\Omega \text{Pl}(X_n)} \circ h_n \circ f_n : \Sigma \rightarrow X_n$, where f_n is the given marking of Y_n and where φ is the identification between the welded and the plumbed surfaces defined in Section 12.5. \square

In order to conclude the proof of the openness of the plumbing map, it remains to justify the extension of the conformal map across the thin part. We first describe this extension along vertical nodes and then describe this extension along horizontal nodes.

We summarize the situation of the previous proof for both the vertical and the horizontal cases. Let (Y, ω) in the neighborhood of $(X, \boldsymbol{\eta})$ lying in the deepest boundary stratum. Being close implies that there exists $\varepsilon > 0$ and an almost-diffeomorphism $g: \overline{X}_\sigma \rightarrow Y$ that identifies the seams of the multicurve Λ with the core curves of the ε -thin part of Y and implies that we can define \mathbf{t} as in (12.12). For a fixed level i and a component Y^0 of $Y_{(i)}$, we denote by E^0 the set of edges of $\Gamma(\Lambda)$ that connect Y^0 to $Y_{(<i)}$. For each edge $e \in E^0$ we have a conformal map $v_e^-: A_e^- \rightarrow Y_{(<i)}$ of the bottom plumbing annulus of a plumbing fixture $V(\rho_e)$ such that $(v_e^-)^*\omega = \Omega_e(\mathbf{t})$, where

$$\Omega_e(\mathbf{t}) = (t_{[i]} \cdot u^\kappa - r(\mathbf{t})) \frac{du}{u} \quad \text{and} \quad \rho_e(\mathbf{t}) = \prod_{j=\ell(e^-)}^i t_j^{m_{e,j}}$$

for a function $r(\mathbf{t})$ with $\lim_{\rho_e(\mathbf{t}) \rightarrow 0} r(\mathbf{t})/t_{[i]} = 0$.

Lemma 12.9. *Given (Y, ω) as in the preceding paragraph, for any $R > 1$, there exists a constant $C(R, \delta(R), (X, \boldsymbol{\eta}))$ such that if $|\rho_e| < C$, then the map v_e^- extends to an injective conformal map v_e whose domain contains the annulus V_e° bounded by the curve $|u| = \delta/\sqrt{R}$ through p_e^+ and such that $v_e^*(\omega) = \Omega_e$. Moreover, the images of V_e° and $V_{e'}^\circ$ are disjoint for any $e \neq e'$.*

Similarly using the notation of the paragraph preceding Lemma 12.9, for each horizontal edge e_j at level i , we have conformal maps $v_{e_j}^\pm: A_{e_j}^\pm \rightarrow Y_{e_j} = V(x_j)$ of the plumbing annuli in a plumbing fixture $V(\rho_{e_j})$ to the thin part Y_{e_j} corresponding to e_j in Y , such that $(v_{e_j}^\pm)^*\omega = \Omega_{e_j}(\mathbf{t})$, where

$$\Omega_{e_j}(\mathbf{t}) = -r_{e_j}(\mathbf{t}) \frac{du}{u} = r_{e_j}(\mathbf{t}) \frac{dv}{v}.$$

Note that the existence of the maps $v_{e_j}^\pm$ implies that the integral of ω along the core curves of Y_{e_j} is equal to $r_{e_j}(\mathbf{t})$.

Lemma 12.10. *Let (Y, ω) and e_j be given as in the preceding paragraph. If for a curve γ connecting c^- to c^+ in Y_{e_j} we have*

$$(12.17) \quad \int_{p^-}^{p^+} \Omega_{e_j}(\mathbf{t}) \equiv \int_\gamma \omega \quad \text{modulo} \quad r_{e_j}(\mathbf{t}),$$

then the maps $v_{e_j}^\pm$ extend to an injective conformal map v_{e_j} from the plumbing fixture V_{e_j} such that $v_{e_j}^(\omega) = \Omega_{e_j}$.*

Proof of Lemma 12.9. For any choice of R and δ the outer boundary $\gamma = \gamma_e$ of A_e^- given by $|v| = \delta/R$ is convex for ρ_e sufficiently small, since $r(\mathbf{t})/t_{[i]}$ tends to 0. We use orthogonal projection of equidistant curves to extend v_e^- . That is, we map the equidistant curve of distance ℓ to γ to the equidistant curve of distance ℓ to $v_e^-(\gamma)$, mapping orthogonal rays to γ into orthogonal rays to $v_e^-(\gamma)$. This procedure gives a well-defined conformal map \tilde{v}_e to the annular cover \tilde{Y} of Y with respect to the core

curve of A_e^- and identifies the forms Ω_e and ω until an equidistant curve to $v_e^-(\gamma)$ hits a zero of ω .

We now show that the domain of \tilde{v}_e contains V_e° for an appropriate choice of R , of the constant $\delta(R)$ given by Theorem 12.3 and for $|\rho_e|$ small enough. For this purpose it suffices to check that the equidistant curve to γ through any point p on $\{|u| = \delta/\sqrt{R}\}$ stays in the plumbing fixture $V(\rho_e)$, since $\delta(R)$ is chosen such that Ω_e has no zeros there. Writing $T = |t_{[i]}|$ we obtain that

$$(12.18) \quad d(p, \gamma) \leq \int_{R\rho_e/\delta}^{\delta/\sqrt{R}} (Tu^\kappa + |r(t)|) \frac{du}{u} \leq \frac{T}{\kappa} \left(\frac{\delta^\kappa}{\sqrt{R}^\kappa} - \frac{(R\rho_e)^\kappa}{\delta^\kappa} \right) + |r(t)| \log \left(\frac{\delta^2}{\rho_e R^{3/2}} \right).$$

Analogously, the distance between the curve γ and the outer circle of the plumbing fixture $\mathbb{V}(t_e)$ has the lower bound

$$(12.19) \quad d(\partial^+(A^+), \gamma) \geq \int_{R\rho_e/\delta}^{\delta} (Tu^\kappa + |r(t)|) \frac{du}{u} \geq \frac{T}{\kappa} \left(\delta^\kappa - \frac{(R\rho_e)^\kappa}{\delta^\kappa} \right) + |r(t)| \log \left(\frac{\delta^2}{\rho_e R} \right).$$

In order to ensure that $d(\partial^+(A^+), \gamma) > d(p^+, \gamma)$, combining (12.19) and (12.18), it suffices to have that

$$\frac{\rho}{\kappa} \left(\delta^\kappa - \frac{(R\rho_e)^\kappa}{\delta^\kappa} \right) + |r(t)| \log \left(\frac{1}{\rho_e R} \right) > \frac{\rho}{\kappa} \left(\frac{\delta^\kappa}{\sqrt{R}^\kappa} - \frac{\rho_e^\kappa}{\delta^\kappa} \right) + |r(t)| \log \left(\frac{\rho_e}{\sqrt{R}} \right).$$

Hence it suffices to have that

$$(12.20) \quad 1 > \frac{1}{\sqrt{R}^\kappa} + \frac{(R\rho_e)^\kappa}{\delta^{2\kappa}} + \frac{\kappa}{T\delta^\kappa} |r(t)| \log \left(\sqrt{R} \right).$$

We choose $R > 1$. Theorem 12.3 thus gives us the constant $\delta(R)$ and we can then choose C such that for $\rho_e < C$ the second term and the third term of (12.20) are both less than $(1 - 1/\sqrt{R}^\kappa)/2$, using the limit behavior of $r(t)$.

Next we want to ensure that the composition $v_e: V_e^\circ \rightarrow \tilde{Y} \rightarrow Y$ of \tilde{v}_e with the annular cover \tilde{Y} of Y is still injective. By the disjointness of plumbing annuli on X , there is no saddle connection on X starting and ending at q_e^+ , of η -length less than $2\delta^\kappa/\kappa$. Consequently, for ρ_e small enough there is no $(\omega/t_{[i]})$ -geodesic on Y starting and ending at $v_e^-(\gamma)$ and meeting Y_0 , which is of length less than $2(\delta/\sqrt{R})^\kappa/\kappa$, since the diameter of the interior of γ shrinks with ρ_e and since $r(t)/t_{[i]}$ tends to 0. If v_e were not injective, we would be able to construct such a short geodesic. The same argument ensures disjointness of the images of V_e° and $V_{e'}^\circ$ for any pair of edges. \square

Proof of Lemma 12.10. The integration of the form ω starting at c^- induces a map from Y_{e_j} to a flat cylinder contained in V_{e_j} . This map clearly coincides with the inverse of $v_{e_j}^-$ where they are both defined. Moreover, Equation (12.17) implies that this map coincides with the inverse of $v_{e_j}^+$ where they are both defined. The desired properties of this map now clearly follow from the definition. \square

We now deal with the local injectivity of the plumbing map.

Proposition 12.11. *Let (Y, ω) be a marked multi-scale differential in the neighborhood of $(X, \eta) = \Omega \text{Pl}(X, \eta, \mathbf{0})$. Then there exists a unique marked model differential $(X_1, \eta_1, \mathbf{t}_1)$ in the neighborhood of $(X, \eta_0, \mathbf{0})$ such that $\Omega \text{Pl}(X_1, \eta_1, \mathbf{t}_1) = (Y, \omega)$.*

Proof. We work with representatives (X, η) in $\Omega \mathcal{T}_\Lambda^{pm}(\mu)$ and (Y, ω) as in the proof of Proposition 12.8 and suppose moreover that Y is smooth, i.e. $(Y, \omega) \in \Omega \mathcal{T}_{(\Sigma, s)}(\mu)$. We want to show that the preimage constructed there is uniquely determined for n large enough, i.e. for Y sufficiently close to X . Since $\mathcal{MD}_\Lambda^s = \Omega \mathcal{T}_\Lambda^{pm}(\mu) / \text{Tw}_\Lambda^s$ (with horizontal twists already acting trivially on $\Omega \mathcal{T}_\Lambda^{pm}(\mu)$) and since the open stratum of $\Xi \mathcal{D}_\Lambda$ is just $\Omega \mathcal{T}_{(\Sigma, s)}(\mu) / \text{Tw}_\Lambda^s$, this proves the claim for the open stratum. The general case follows by repeating this argument in between any two levels of Y .

To show uniqueness, we revisit the previous proof. At each step the choice of $X_{n,(i)}$ we determined by the (perturbed) periods of Y_n is unique for n large enough, since the (perturbed) period map is locally injective. In the proof of openness we have seen that the prong-matching at each vertical node has to be chosen so that X_n is close to X . Since the augmented Teichmüller space is Hausdorff by Theorem 9.4, there is a unique such choice. Similarly, the marking is uniquely specified by the condition in the proof of openness for large enough n , since by the Hausdorff property at most one marking of $\overline{X}_{n, \sigma_n}$ is close to (\overline{X}_σ, f) . \square

Corollary 12.12. *The plumbing map is a local homeomorphism at every point of U .*

Proof. For points in the deepest stratum $\Omega \mathcal{MD}_\Lambda^{s, \Lambda}$ this has just been shown. Since being a local homeomorphism at a point is an open property, this implies that after possibly restricting U to a smaller neighborhood of $P \times \mathbf{0}$ the property of being a local homeomorphism holds over all of U . \square

12.8. The complex structure on the Dehn space. We can now collect the information of the preceding sections and provide the Dehn space with a complex structure.

Proof of Theorem 12.1. The desired properties of ΩPl were shown in Proposition 12.5 and Corollary 12.12. The equivariance of the map with respect to the group $K_\Lambda = \text{Tw}_\Lambda / \text{Tw}_\Lambda^s$ follows from the construction in Section 12.5, since K_Λ acts on the markings and the rescaling ensemble only, and they both have been transported from the family over the model domain to the plumbed family. \square

Proof of Theorem 12.2. We proceed inductively with respect to the partial order induced by undegenerating. The base case of the induction $\Lambda = \emptyset$ is simply $\Omega \mathcal{MD}_\emptyset^s = \Omega \mathcal{B}_\emptyset = \Omega \mathcal{T}_\emptyset^{pm}(\mu) = \Omega \mathcal{T}_{(\Sigma, s)}(\mu) = \Xi \mathcal{D}_\emptyset$.

For the induction step we consider a multicurve Λ and note that for every undegeneration $\Lambda' \rightsquigarrow \Lambda$ the complex structure on $\Xi \mathcal{D}_{\Lambda'}$ induces a complex structure on the open subset $\Xi \mathcal{D}_{\Lambda, \circ}^{\Lambda', s} = \Xi \mathcal{D}_{\Lambda'}^s / (\text{Tw}_\Lambda^s / \text{Tw}_{\Lambda'}^s)$ of $\Xi \mathcal{D}_\Lambda^s$ since $\text{Tw}_\Lambda^s / \text{Tw}_{\Lambda'}^s$ acts properly discontinuously. The complex structure on the intersections $\Xi \mathcal{D}_{\Lambda, \circ}^{\Lambda_1, s}$ and $\Xi \mathcal{D}_{\Lambda, \circ}^{\Lambda_2, s}$ agrees, since it stems from the common undegeneration of Λ_1 and Λ_2 . So far, we have obtained a complex structure on

$$\bigcup_{\Lambda' \rightsquigarrow \Lambda} \Xi \mathcal{D}_{\Lambda, \circ}^{\Lambda', s} = \Xi \mathcal{D}_\Lambda^s \setminus \Xi \mathcal{D}_\Lambda^{\Lambda, s}.$$

On the other hand, we can cover a neighborhood of the deepest boundary stratum $\Xi\mathcal{D}_\Lambda^{\Lambda,s}$ by open sets of the form $\Omega\text{Pl}(U_i)$, since ΩPl was constructed to be the identity along the deepest boundary stratum. These sets $\Omega\text{Pl}(U_i)$ inherit a complex structure from that of $U_i \subset \Omega\overline{\mathcal{M}\mathcal{D}}_\Lambda^s \times \Delta^H$. The union of all these sets cover all of $\Xi\mathcal{D}_\Lambda^s$. It remains to show that the complex structures agree, i.e. that the change of chart maps are holomorphic. Using that the change of chart maps are continuous, it suffices to show holomorphicity on the open stratum of $\Omega\overline{\mathcal{M}\mathcal{D}}_\Lambda^s$, since this is the complement of a (normal crossing) divisor. There, the change of chart maps are compositions of the moduli map for the plumbed family and of the inverse of such a map. Since the plumbed family is a holomorphic family over $\Omega\mathcal{M}\mathcal{D}_\Lambda^s$, its moduli map is holomorphic (as a map to $\Omega\mathcal{B}_\emptyset/\text{Tw}_\Lambda^s$) and this completes the proof.

The complex structure on $\Xi\mathcal{D}_\Lambda$ stems from that on $\Xi\mathcal{D}_\Lambda^s$ and the K_Λ -equivariance of the plumbing map.

Finally we need to construct the families of (simple) marked multi-scale differentials over $\Xi\mathcal{D}_\Lambda^s$ and $\Xi\mathcal{D}_\Lambda$. We start with the simple case and proceed again inductively. The base case $\Lambda = \emptyset$ is obvious. To construct the family of curves $\mathcal{Y} \rightarrow \Xi\mathcal{D}_\Lambda^s$ we use the family $\mathcal{Y}^\circ \rightarrow \Xi\mathcal{D}_\Lambda^s \setminus \Xi\mathcal{D}_\Lambda^{\Lambda,s}$ given by induction and the families $\pi_U: \mathcal{Y}_U \rightarrow \Omega\text{Pl}(U)$ resulting from plumbing for open subsets as in Theorem 12.1. To see that these families glue, we use that $\Xi\mathcal{D}_\Lambda^s$ is constructed as a union of Tw_Λ^s -quotients of $\Omega\mathcal{B}_{\Lambda'}$ for various Λ' and that the maps of strata of $\Omega\text{Pl}(U)$ to $\Omega\mathcal{B}_{\Lambda'}$ are the moduli maps for the functor that $\Omega\mathcal{B}_{\Lambda'}$ represents. The same argument allows to patch the family of stable forms ω on $\Omega\text{Pl}(U)$ to a global family of stable forms ω on $\Xi\mathcal{D}_\Lambda^s$.

To construct the family over $\Xi\mathcal{D}_\Lambda$ we use the K_Λ -equivariance of the plumbing map and the fact that on the range of the plumbing map the model domain has been provided with a (universal) family in Proposition 10.10. \square

13. THE UNIVERSAL PROPERTY OF THE DEHN SPACE

The purpose of this section is to show the following two results.

Theorem 13.1. *The Dehn space $\Xi\mathcal{D}_\Lambda$ is the fine moduli space, in the category of complex analytic spaces, for the functor $\mathbf{MS}_{(\mu,\Lambda)}$ of marked multi-scale differentials.*

As in the case of the model domain, we first prove the following related statement, and then descend by the K_Λ -action.

Proposition 13.2. *The simple Dehn space $\Xi\mathcal{D}_\Lambda^s$ is the fine moduli space, in the category of complex analytic spaces, for the functor $\mathbf{MS}_{(\mu,\Lambda)}^s$ of simple marked multi-scale differentials.*

Given a family $\pi: \mathcal{Y} \rightarrow B$ of stable curves with a family of simple marked multi-scale differentials (ω, σ, f) , we want to construct functorially a map $m: B \rightarrow \Xi\mathcal{D}_\Lambda^s$ such that the pullback of the universal family agrees with the given family. Since the complex structure on $\Xi\mathcal{D}_\Lambda^s$ stems from the model domain, the map m is constructed by using the universal property of the model domain, post-composed with the plumbing map of the previous section. To use the universal property of the model domain, we need to define an unplumbing construction that takes multi-scale differentials on \mathcal{Y} to model

differentials on an equisingular family $\mathcal{X} \rightarrow B$. Since this construction, as the plumbing construction, depends on several choices, we need to carefully arrange the unplumbing construction consistently on \mathcal{Y} and on the universal family.

We first deal with the local situation, since $\omega = (\omega_p)_{p \in B}$ is given as a collection of germs anyway, and assume that the germ ω_p is defined on all of B . We moreover assume that we are in the deepest stratum of the Dehn space, i.e., $\Gamma_p = \Gamma(\Lambda)$.

13.1. The unplumbing construction. The unplumbing construction associates with a family of multi-scale differentials a family of model differentials. The rough idea is to pinch off neighborhoods degenerating to nodes, in order to create equisingular families, and to measure the degeneracy of the nodes using the parameter t of model differentials. Technically, we cannot pinch off curves without modifying the differential due to the presence of non-trivial periods. This forces us to subtract beforehand some perturbation differentials, playing role inverse to that of the modifying differentials.

Proposition 13.3. *Given a germ of a family of Tw_Λ^s -marked multi-scale differentials with all data $(\mathcal{Y} \rightarrow B, (\omega_{(i)})_{i \in L \bullet(\Gamma_p)}, R^s, \sigma, f)$ defined over B , there is an unplumbing construction that produces a family $(\mathcal{X} \rightarrow B, (\eta_{(i)})_{i \in L \bullet(\Gamma_p)}, R^s, \sigma', f)$ of Λ -marked simple model differentials with the following properties:*

- (i) *The construction is the identity over the locus B^Λ of $q \in B$ where $\Gamma_q = \Gamma(\Lambda)$.*
- (ii) *The construction depends only on a finite number of choices of topological data and moreover a section near each vertical node.*
- (iii) *If B is an open neighborhood of p in the simple Dehn space, then the map $u: B \rightarrow \Omega\overline{\mathcal{MD}}_\Lambda^s$ induced by the unplumbing of the universal family over it is a local biholomorphism.*

Proof. For simplicity of the exposition and since the level-wise construction is similar to those in Section 12, we only treat the case where Γ has two levels, and no horizontal nodes. We may thus write $\omega_{(0)} = s \cdot \omega_{(-1)}$, and B^Λ is precisely the vanishing locus of s .

For the definition of a perturbation differential, let $\Lambda_{\max} \supseteq \Lambda$ be a maximal multi-curve. We denote by V the image of Λ in $H_1(\Sigma \setminus P_{\mathbf{s}}; \mathbb{Q})$ and, as in Proposition 11.3, we let V' be the subspace generated by Λ_{\max} and loops around points in $P_{\mathbf{s}}$. Let $\rho: B \rightarrow \text{Hom}_{\mathbb{Q}}(V, \mathbb{C})$ be the periods of $\omega_{(0)}$ along Λ and let $\tilde{\rho}$ be the extension of ρ by zero on a subset S of Λ_{\max} generating V'/V . A *perturbation differential* is a meromorphic section ξ of the relative dualizing sheaf $\pi_*\omega_{\mathcal{Y}/B}$ such that the periods of ξ are $\tilde{\rho}$. A perturbation differential exists and it is uniquely determined by the choice of the topological datum Λ_{\max} and the subset S . Since $\tilde{\rho}$ is divisible by s , the perturbation differential vanishes on the fibers over B^Λ .

Next, recall that a multi-scale differential comes with a normal form on a neighborhood of the nodes that looks like a plumbing fixture, that is, a coordinate v_e such that $\omega_{(-1)} = (-v_e^{-\kappa_e} + r_{e,(-1)}) \frac{dv_e}{v_e}$. By Theorem 4.3 the coordinates in the normal form are uniquely determined by a section near the lower end of each node. Consequently, the lower level subsurface of \mathcal{Y} with the form $\omega_{(-1)}$ can be glued together with the form $(\Delta \times B, (-v_e^{-\kappa_e} + r_{e,(-1)}) \frac{dv_e}{v_e})$ on a disk Δ times the base with a one-form $\eta_{(-1)}$.

We use Theorem 4.3 (and the same section as above to specify the coordinates uniquely) to put $\omega_{(0)} - \xi$ in standard form $\phi^*(\omega_{(0)} - \xi) = u_e^{\kappa_e} \frac{du_e}{u_e}$ on some family of

annuli in $V(f_e, \epsilon)$ near each node. Consequently, the form $\omega - \xi$ on the upper level subsurface of \mathcal{Y} and the forms $(\Delta \times B, u_e^{\kappa_e} \frac{du_e}{u_e})$ for each node glue to produce a closed surface $\mathcal{X}_{(0)}$ with a one-form $\eta_{(0)}$.

This one-form does not necessarily have the correct orders of vanishing in the smooth locus. Hence we merge the zeros. For this purpose, we specify an annulus A_{δ_1, δ_2} around each zero of ω in the upper level subsurface of \mathcal{Y} . Using Theorem 4.2 we put $\omega - \xi$ in standard form $z^m dz$ on A_{δ_1, δ_2} and we glue it with the one-form $(\Delta \times B, z^m dz)$ to obtain a differential with the correct orders. We continue to denote by $(\mathcal{X}_{(0)}, \eta_0)$ this differential. We obtain an equisingular family $\pi: \mathcal{X} \rightarrow B$ obtained by gluing the points $u = 0$ of $\mathcal{X}_{(0)}$ and $v = 0$ of $\mathcal{X}_{(-1)}$ in each plumbing fixture.

For an equisingular family, the space of prong-matchings is a covering of the base, see Section 5.4. To obtain the prong-matching σ' we simply extend the prong-matching $\sigma|_p$ in a locally constant way. The level-wise real blowup of (\mathcal{Y}, R^s) and the level-wise real blowup of \mathcal{X} as defined in Section 10 are almost-diffeomorphic. (The almost-diffeomorphism is given by the identity on the upper and lower surface, blurred near the marked zeros, and both the degenerate plumbing fixtures in \mathcal{X} and the plumbing fixtures of \mathcal{Y} are replaced by the welded fixtures as defined in (8.2).) We can thus transport the marking f via this isomorphism. The rescaling ensemble R^s is the same on both sides of the construction. Finally we verify that the equivalence relations are the same on both sides, since they stem from the T_Λ^s -action for the differentials and prong-matchings, and from Tw_Λ^s for the markings in both cases.

To prove (iii) it suffices to show that the tangent map to u is surjective at any point of B^Λ . Restricted to B^Λ , the unplumbing leaves the universal family unchanged. Using the fact that perturbed periods are coordinates on the model domain, it suffices to show that the directions corresponding to changing the parameters \mathbf{t} of the model differential are in the range of the tangent map to u . This is obvious since those parameters are encoded in R^s , which is part of the datum of the unplumbed model differential. \square

13.2. Consistent unplumbing and the proof of Proposition 13.2. In order to define the moduli map m we perform the unplumbing construction twice and consistently, for the family over B to obtain a family of model differentials on $\pi: \mathcal{X} \rightarrow B$ and to the universal family over a neighborhood W of the moduli point to which the fiber $\mathcal{Y}_p \rightarrow p$ is mapped in $\Xi\mathcal{D}_\Lambda^s$, to obtain a family of model differentials $\pi^{\text{uni}}: \mathcal{X}^{\text{uni}} \rightarrow W$. Consistent unplumbing means the following. First, we choose a maximal multicurve Λ_{max} on \mathcal{Y}_p as required for Proposition 13.3 and choose the same (which is possible since the surfaces are marked) maximal multicurve on the universal family over W . Second, we choose the normalizing sections for the unplumbing of each node (to lie in the neighborhood $V(f_e, \epsilon)$ and) of constant relative $\omega_{(\ell(e^-))}$ -period to a marked zero on the same level as the lower end of the node. The markings, which are well-defined up to Tw_Λ^{sv} -twists, allow to consistently choose the paths for those periods.

Let $m': B \rightarrow \overline{\mathcal{MD}}_\Lambda^s$ be the moduli map for the family (\mathcal{X}, η) of model differentials obtained from unplumbing π . Let u be the moduli map for the universal family as in Proposition 13.3 (iii). We claim that (after possibly shrinking B to fit domains)

$$m = u^{-1} \circ m': B \rightarrow \Xi\mathcal{D}_\Lambda^s$$

is the moduli map for B . By definition there is an isomorphism of families of model differentials $h': (m')^* \mathcal{Y}^{\text{uni}} \rightarrow \mathcal{X}$ and we need to exhibit an isomorphism of families of multi-scale differentials $h: m^* \mathcal{Y}^{\text{uni}} \rightarrow \mathcal{Y}$.

This isomorphism is constructed level by level and the idea will be clear from considering the two-level situation without horizontal nodes, as in the proof of Proposition 13.3. The lower level subsurfaces $\mathcal{X}_{(-1)}$ and $\mathcal{Y}_{(-1)}$ with their differentials are simply the same by construction, and this also holds for the universal families, which gives h on that subsurface. On the upper level surfaces a perturbation differential has been added. The consistent choice of Λ_{\max} implies that the m' -pullback of the perturbation differential on $\mathcal{Y}_{(0)}^{\text{uni}}$ agrees with that on $\mathcal{Y}_{(0)}$. This also implies that the local modifications near the marked zeros are compatible under m' -pullback. We thus obtain h on the upper level surfaces. The plumbing fixtures (i.e. the functions f_e) are compatible, since these functions can be read off from the rescaling ensemble. Last, it remains to check that the way the plumbing fixtures are glued in is compatible so that the piece-wise defined isomorphisms h glue to a global isomorphism. This follows from the consistent choice of the normalizing sections using an $\omega_{(-1)}$ -period.

Next we deal with the situation that the pinched multicurve Λ_p is a strict undegeneration of Λ . We use Proposition 8.5 to provide the family with a $\text{Tw}_{\Lambda_p}^s$ -marking and the previous argument to obtain a moduli map $B \rightarrow \Xi \mathcal{D}_{\Lambda_p}^s$. The composition with the natural maps $\Xi \mathcal{D}_{\Lambda_p}^s \rightarrow \Xi \mathcal{D}_{\Lambda}^{\Lambda_p, s} \hookrightarrow \Xi \mathcal{D}_{\Lambda}^s$ is the moduli map we need.

Finally, after having constructed $m = m_p$ locally near p we need to show that the local constructions glue. The only point that might be not clear is the prong-matching, since σ' was constructed in Proposition 13.3 by locally constant extension. However, this might make a difference only if the smoothing parameter $f_e \neq 0$, in which case the prong-matching is induced and can be retrieved as $\sigma_e = du_e \otimes dv_e$ from the other data of the family already known to agree on the overlaps.

13.3. The proof of Theorem 13.1. As for Proposition 13.2 we start with the local version and then glue the moduli maps as above. Suppose we are given a family $(\mathcal{Y} \rightarrow B, (\omega_{(i)})_{i \in L \bullet(\Gamma_p)}, R, \sigma, f)$ of marked multi-scale differentials, defined on a neighborhood of p . We proceed similarly to the proof of Proposition 7.7 and let $B^s \rightarrow B$ to be the fiber product of $R: B \rightarrow \overline{T}_{\Gamma_p}^n$ with finite quotient map $\overline{p}: \overline{T}_{\Gamma_p}^s \rightarrow \overline{T}_{\Gamma_p}^s / K_p = \overline{T}_{\Gamma_p}^n$. The pullback family $\mathcal{Y}^s \rightarrow B^s$ comes with a map $R^s: B^s \rightarrow \overline{T}_{\Gamma_p}^s$ and is thus a family of simple marked multi-scale differentials. By Proposition 13.2 we obtain a moduli map $m^s: B^s \rightarrow \Xi \mathcal{D}_{\Lambda}^s$. Composed with the quotient map $\Xi \mathcal{D}_{\Lambda}^s \rightarrow \Xi \mathcal{D}_{\Lambda}$ we get a map $B^s \rightarrow \Xi \mathcal{D}_{\Lambda}$ that is clearly K_{Γ_p} -invariant. It thus descends to the required moduli map $m: B \rightarrow \Xi \mathcal{D}_{\Lambda}$.

14. THE MODULI SPACE OF MULTI-SCALE DIFFERENTIALS

We now have all the tools that are necessary to prove the main theorems announced in the introduction.

14.1. The moduli space of multi-scale differentials as a topological space.

Theorem 14.1. *The quotient $\Xi\overline{\mathcal{M}}_{g,n}(\mu) = \Omega\overline{\mathcal{T}}_{(\Sigma,s)}(\mu)/\text{Mod}_{g,n}$ of the augmented Teichmüller space by the mapping class group provided with the quotient topology is a Hausdorff topological space. The complex structure on the Dehn space provides $\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ with the structure of a complex orbifold. The same holds for the projectivized version $\mathbb{P}\Xi\overline{\mathcal{M}}_{g,n}(\mu) = \mathbb{P}\Omega\overline{\mathcal{T}}_{(\Sigma,s)}(\mu)/\text{Mod}_{g,n}$.*

We refer to $\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ as the *moduli space of multi-scale differentials of type μ* and to $\mathbb{P}\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ as the *moduli space of projectivized multi-scale differentials of type μ* . The core of the proof of the theorem is the following lemma.

Lemma 14.2. *For each point in the stratum $\Omega\mathcal{B}_\Lambda$ of $\mathbb{P}\Omega\overline{\mathcal{T}}_{(\Sigma,s)}(\mu)$ there exists a neighborhood U such that the set of elements $\gamma \in \text{Mod}_{g,n}$ with $\gamma U \cap U \neq \emptyset$ consists of a finite number of cosets of the group Tw_Λ .*

Since Tw_Λ^s has finite index in Tw_Λ , the lemma holds also with Tw_Λ^s in place of Tw_Λ .

Proof. Consider the image of (X, ω, σ, f) under the forgetful map $\pi: \mathbb{P}\Omega\overline{\mathcal{T}}_{(\Sigma,s)}(\mu) \rightarrow \overline{\mathcal{T}}_{(\Sigma,s)}$ to the classical augmented Teichmüller space of Riemann surfaces with marked points. Then there exists a neighborhood V of $\pi(X)$ such that the set of elements $\gamma \in \text{Mod}_{g,n}$ with $\gamma V \cap V \neq \emptyset$ consists of a finite number of cosets of the group $\text{Tw}_\Lambda^{\text{full}}$. This follows from the classical fact that the action of the mapping class group acts properly discontinuously on the Teichmüller space by considering the normalization of $\pi(X)$ (see [HK14, Corollary 2.7]).

Now take one of the open neighborhoods that we denote $W = V_\epsilon(X)$ used to define the topology of the augmented Teichmüller space in Section 9 (in the proof of Theorem 9.4) with $\epsilon < 1/2$ small enough so that $\pi(W)$ is contained in the V just chosen. Suppose that $\gamma \in \text{Mod}_{g,n}$ has the property that there is some $(X', \omega', \sigma', f') \in \gamma W \cap W$. Pre-composing γ by one out of a finite number of elements, we may assume that $\gamma \in \text{Tw}_\Lambda^{\text{full}}$. It remains to prove that in fact $\gamma \in \text{Tw}_\Lambda$. By definition of the topology there are nearly turning-number-preserving almost-diffeomorphisms $g, g_1: \overline{X}_\sigma \rightarrow \overline{X}'_{\sigma'}$, conformal on the ϵ -thick part, that nearly identify the forms there, and such that there are elements $D, D' \in \text{Tw}_\Lambda$ such that $f' = g \circ f \circ D$ and $f' \circ \gamma = g_1 \circ f \circ D'$ up to isotopy. Said differently, there is an isotopy between $\tilde{f} := f \circ (D' \circ \gamma^{-1} \circ D^{-1})$ and $(g_1)^{-1} \circ g \circ f$ that we may take to be the identity on the ϵ -thick part. We compare the image under f and \tilde{f} of a path crossing once a curve in Λ , with endpoints in the ϵ -thick part. The two turning numbers differ by an integer. Since $(g_1)^{-1} \circ g$ is nearly turning-number-preserving, the two turning numbers are in fact equal. This implies that $D' \circ \gamma^{-1} \circ D^{-1}$ is trivial and thus $\gamma \in \text{Tw}_\Lambda$. \square

Proof of Theorem 14.1. The augmented Teichmüller space is covered by the open sets $\coprod_{\Lambda' \rightsquigarrow \Lambda} \Omega\mathcal{B}_{\Lambda'}$. Restricted to these open sets the quotient map to the moduli space $\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ factors through the Dehn space $\Xi\mathcal{D}_\Lambda$. Since $\Xi\mathcal{D}_\Lambda$ is Hausdorff, in fact a complex orbifold by Theorem 12.2, the claim follows from Lemma 14.2. \square

Theorem 14.3. *The moduli space $\mathbb{P}\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ of projectivized multi-scale differentials of type μ is compact.*

Proof. The fact that the moduli space of multi-scale differentials is metrizable follows from the complex orbifold structure given by Theorem 14.1. Hence it suffices to prove that this space is sequentially compact. Let $\{(X_n, z_n, \omega_n, \preccurlyeq_n, \sigma_n)\}$ be a sequence in $\mathbb{P}\Xi\overline{\mathcal{M}}_{g,n}(\mu)$. We extract from this sequence a convergent subsequence.

Since $\overline{\mathcal{M}}_{g,n}$ is compact, we can extract a subsequence such that $\{(X_n, z_n)\}$ converges to (X, z) . Moreover, since $\overline{\mathcal{M}}_{g,n}$ has a finite stratification by topological type, we can assume that the topological type of the curves X_n is constant. By definition of convergence there exists an increasing exhaustion K_n of $X' := X \setminus (z \cup N_X)$, where N_X is the set of nodes of X , and a conformal map $g_n: K_n \rightarrow X_n$. Since there is a finite number of full orders on a dual graph of a fixed topological type, we can assume that \preccurlyeq_n is constant in the following sense. We may assume that the components $X_{n,j}$ and $X_{n,k}$ of X_n satisfy $X_{n,j} \preccurlyeq_n X_{n,k}$ if and only if $g_m g_n^{-1}(X_{n,j}) \preccurlyeq_m g_m g_n^{-1}(X_{n,k})$ for all $m, n \in \mathbb{N}$. This naturally induces a full order on X that we denote by \preccurlyeq_0 .

Now we deal with the differential forms and use their order of magnitude to define the limit order \preccurlyeq on the set of components of X . (Note that in general the limit order on X will be a degeneration of \preccurlyeq_0 .) We focus on the components of X' at a fixed level for \preccurlyeq_0 and denote its components by X'_j . Let $K_{n,j}$ be the restrictions of K_n on X'_j and fix a reference point $p_j \in K_{1,j}$.

We choose $c_{n,j} \in \mathbb{C}^*$ such that $|c_{n,j}| = \lambda$ is the maximal size of the components X'_j as defined in Equation (3.6). Then Theorem 3.11 and the fact that every size differs by a bounded multiplicative constant imply that $\{g_n^* c_{n,j} \omega_n\}$ converges (up to taking a subsequence) on X'_j for each j . After extracting a further subsequence we may suppose that for each pair (j, k) either $|c_{n,j}/c_{n,k}|$ tends to zero, to infinity or that $c_{n,j}/c_{n,k}$ converges to some value in \mathbb{C}^* . We now refine the order \preccurlyeq_0 to the order \preccurlyeq by imposing that $X'_j \prec X'_k$, that $X'_j \succ X'_k$ and that $X'_j \asymp X'_k$ in these three cases.

For each level i of \preccurlyeq we pick a component $X'_{j(i)}$ at that level. Then for each component X'_k at level i the family of differentials $\{g_n^* c_{n,j(i)} \omega_n\}$ converges to some differential ω_k on X'_k . We define ω on X to be the collection of those differentials ω_k . We have to prove that ω is a twisted differential compatible with the order \preccurlyeq . The crucial conditions (matching orders, matching residues and GRC) can be verified by ω -path integrals or turning numbers (compare Section 4 in [BCGGM18]). Hence these conditions carry over from the corresponding integrals on the sequence of surfaces X_n , using the convergence of one-forms and using (for the GRC) the fact that the rescaling parameters $c_{n,j}$ depend on the levels only.

To show convergence of the prong-matchings, we first fix $d_{n,i}$ with $e(d_{n,i}) = c_{n,i}$ and let $\mathbf{d} = (d_{n,i})_{i \in L \bullet(\Lambda)}$. For each fixed n and each vertical node q choose a preliminary prong-matching $\tilde{\sigma}_q$, forming together a global prong-matching $\tilde{\sigma}$. Choose a preliminary extension $\tilde{g}_n: \overline{X}_{\mathbf{d}\tilde{\sigma}} \rightarrow \overline{X}_{\sigma_n}$ of g_n that nearly preserves turning numbers in $\mathbb{R}/2\pi\mathbb{Z}$ of arcs in $\overline{X}_{\mathbf{d}\tilde{\sigma}}$. This is possible by the convergence of differential forms established previously. We want to modify the prong-matching and \tilde{g}_n so that turning numbers are nearly preserved, without shift by $2\pi\mathbb{Z}$. For this purpose, if an arc γ crosses the seam of q and no other seam and if $\tau(\tilde{g}_n \circ F_{\mathbf{d}} \circ \gamma) - \tau(F_{\mathbf{d}} \circ \gamma)$ is approximately $2\pi m$, we rotate the prong-matching $\tilde{\sigma}_q$ by m prongs counterclockwise and post-compose \tilde{g}_n with an (m/κ_q) -th fractional Dehn twist. By performing this at all vertical nodes we

obtain a new prong-matching σ and a new extension $g_n: \overline{X}_{\mathbf{d},\sigma} \rightarrow \overline{X}_{\sigma_n}$ of the above g_n that is now nearly turning-number-preserving for all arcs. Since there are only finitely many global prong-matchings, we may assume after passing to a subsequence that the chosen σ is the same for all n . The maps g_n now exhibit the convergence in the topology of Definition 9.3 quotiented by the action of the mapping class group. \square

Proposition 14.4. *The moduli space of multi-scale differentials $\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ has a normal crossing boundary divisor. Its connected components are in bijection with the connected components of $\Omega\mathcal{M}_g(\mu)$.*

Proof. The first statement follows from the boundary structure of the model domain $\Omega\overline{\mathcal{M}}_{g,n}^s$ and the transport of structure provided by Theorem 12.1. The second statement follows since every stratum of $\mathbb{P}\Xi\mathcal{D}_\Lambda$ is adjacent to the stratum $\Omega\mathcal{B}_\emptyset$ of abelian differentials on smooth curves. This implies that $\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ has at most as many components as $\Omega\mathcal{M}_g(\mu)$. If the closures of two connected components of $\Omega\mathcal{M}_g(\mu)$ intersected in $\Xi\overline{\mathcal{M}}_{g,n}(\mu)$, this would contradict the fact that $\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ has only finite quotient singularities by Theorem 14.1. \square

14.2. The universal property. Recall that for a complex orbifold with local orbifold charts (U, G) there is an underlying complex space with charts being the (in general singular) complex spaces U/G .

Theorem 14.5. *The complex space associated with the moduli space of multi-scale differentials $\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ is the coarse moduli space for the functor \mathbf{MS}_μ of multi-scale differentials of type μ .*

Proof. In order to construct the moduli map $m: B \rightarrow \Xi\overline{\mathcal{M}}_{g,n}(\mu)$ for a family of multi-scale differentials $(\pi: \mathcal{X} \rightarrow B, \omega) \in \mathbf{MS}_\mu(B)$, we want to provide the family locally near any point p with a marking and define m as the composition of the moduli map from Theorem 13.1 and the natural quotient map.

For this purpose, we choose for any point $p \in B$ an enhanced multicurve Λ_p on Σ with $\Gamma(\Lambda_p) = \Gamma_p$. For a sufficiently small neighborhood U_p of p we apply Proposition 8.5 to provide the family with a marking. The moduli maps composed with the projection $U_p \rightarrow \Xi\mathcal{D}_\Lambda \rightarrow \Xi\overline{\mathcal{M}}_{g,n}(\mu)$ glue, since any two choices of marking differ by the action of an element in the mapping class group. This argument, together with the universal property of $\Xi\mathcal{D}_\Lambda$, also implies the bijection on complex points and the maximality required as properties of a coarse moduli space. \square

We also take the first steps towards the proof of Theorem 1.3. See [ACG11, Chapter XII] for a general introduction to (algebraic) stacks and [Toë99] for analytic stacks and analytification.

Lemma 14.6. *The groupoid \mathcal{MS}_μ is an analytic Deligne-Mumford stack.*

Proof. The effectivity of descent data in the complex analytic (or equivalently in the étale topology, compare [PY16, Section 3.2]) follows from the definition of multi-scale differentials by gluing germs, see Section 7.3. The Isom-functor for multi-scale differentials is represented by a subspace of the Isom-space representing the Isom-functor

for curves. As in [ACGH85, Section XII.8] this implies that Isom is a sheaf, separated and quasicompact.

We claim that the covering of \mathcal{MS}_μ by the union of Dehn spaces $\Xi\mathcal{D}_\Lambda$ for all Λ is indeed an étale covering. To see this we revisit the proof of Lemma 14.2 and check that the elements with $\gamma U \cap U \neq \emptyset$ actually come from automorphisms of curves and are thus recorded in the stack structure. This is the content of [HK14, Proposition 2.6]. \square

14.3. A blowup description. The incidence variety compactification in general can have bad singularities. For instance, it can fail to be normal, as it can have multiple local irreducible components along the locus of pointed stable differentials that admit more than one compatible enhanced structure on the dual graph (see e.g. [BCGGM18, Example 3.2]), and its normalization may still be quite singular, e.g. not being \mathbb{Q} -factorial as shown in the following example.

Example 14.7. (*The IVC may be not \mathbb{Q} -factorial.*) Consider a level graph with three levels such that the top level has one vertex X_0 , the level -1 has two vertices X_1 and X'_1 , and the bottom level has one vertex X_2 , where X_0 is connected to each of X_1 and X'_1 by one edge and X_2 is connected to each of X_1 and X'_1 by one edge. In other words, the graph looks like a *rhombus*. Since it has three levels and no horizontal edges, the corresponding stratum has codimension two in the moduli space of multi-scale differentials $\mathbb{P}\Xi\overline{\mathcal{M}}_{g,n}(\mu)$. On the other hand, since X_1 and X'_1 are disjoint, when considering the incidence variety compactification $\mathbb{P}\Omega\overline{\mathcal{M}}_{g,n}^{\text{inc}}(\mu)$ we lose the information of relative sizes of rescaled differentials $\lambda\eta_1$ and $\lambda'\eta'_1$ on X_1 and X'_1 , where $\lambda, \lambda' \in \mathbb{C}^*$, and hence the corresponding locus has codimension three in $\mathbb{P}\Omega\overline{\mathcal{M}}_{g,n}^{\text{inc}}(\mu)$. Namely, the map $\mathbb{P}\Xi\overline{\mathcal{M}}_{g,n}(\mu) \rightarrow \mathbb{P}\Omega\overline{\mathcal{M}}_{g,n}^{\text{inc}}(\mu)$ locally around these loci looks like a \mathbb{P}^1 -fibration, where $\mathbb{P}^1 = [\lambda, \lambda']$ (in the degenerate case $\lambda = 0$, X_1 goes lower than X'_1 and the graph has four levels). One can check that locally outside of these loci the map does not have positive dimensional fibers. We thus obtain locally a *small contraction* (which means no divisors get contracted), and consequently the target space $\mathbb{P}\Omega\overline{\mathcal{M}}_{g,n}^{\text{inc}}(\mu)$ (as well as its normalization) is not \mathbb{Q} -factorial (see e.g. [KM98, Corollary 2.63]).

Denote by $\mathbb{P}\Omega\overline{\mathcal{M}}_{g,n}^{\text{nic}}(\mu)$ the normalization of the incidence variety compactification, considered as a substack of $\mathbb{P}\Omega\overline{\mathcal{M}}_{g,n}$. In this section we will show that the stack of multi-scale differentials can be obtained from $\mathbb{P}\Omega\overline{\mathcal{M}}_{g,n}^{\text{nic}}(\mu)$ as the normalization of a certain explicit (complex algebraic, not real oriented) blowup. We will then be able to conclude the proof of the main theorem about the moduli space of multi-scale differentials, in particular proving algebraicity of $\Xi\overline{\mathcal{M}}_{g,n}(\mu)$.

Given an adjustable but not necessarily orderly family $(\mathcal{X} \rightarrow B, \omega)$ (see Section 7.4), we first describe a canonical way to blow up the base B so that the pullback family under this base change becomes orderly. Let X_v and $X_{v'}$ be two irreducible components of the fiber X_p over some $p \in B$. The family fails to be orderly if neither of the adjusting parameters h and h' for X_v and $X_{v'}$, respectively, divides the other one, as elements in $\mathcal{O}_{B,p}$. Therefore, we perform the following blowup construction.

Let $U \subset B$ be a (sufficiently small) neighborhood of p such that there exist adjusting parameters $\{h_1, \dots, h_n\}$ for the family $\mathcal{X}|_U$. The *disorderly ideal* $\mathcal{D}_U \subset \mathcal{O}_{U,p}$ for $\mathcal{X}|_U$ at p is the product of all ideals of the form $(h_{i_1}, \dots, h_{i_k})$, where $\{i_1, \dots, i_k\}$ ranges over all subsets of components of X_p on which ω vanishes identically.

We denote by \widetilde{U} the blowup of U along \mathcal{D}_U , and call it the *orderly blowup*. If U' is an open subset of U such that $\mathcal{X}|_{U'}$ becomes less degenerate, namely, some h_i becomes a unit in U' or the ratio of some h_i and h_j becomes a unit, then $\mathcal{D}_U|_{U'}$ possibly differs from $\mathcal{D}_{U'}$ by some repeated factors of ideals. Note that blowing up the principal ideal of a non-zero-divisor (i.e. the underlying subscheme is an effective Cartier divisor) is simply the identity map, and moreover, blowing up a product of ideals is the same as successively blowing up (the total transform of) each ideal (see e.g. [Sta18, Tag 01OF]). It implies that for two open subsets U_1 and U_2 , we can glue \widetilde{U}_1 and \widetilde{U}_2 along their common restrictions $\widetilde{U}_1 \cap \widetilde{U}_2$. In other words, this local blowup construction chart by chart leads to a well-defined global space, which we denote by \widetilde{B} , and there exists a blowdown morphism $\widetilde{B} \rightarrow B$ locally given by $\widetilde{U} \rightarrow U$.

Example 14.8. We illustrate the behavior of disorderly ideals by the following example. Suppose the special fiber X_p consists of four irreducible components X_0, X_1, X_2, X'_2 such that X_0 is on top level which connects to X_1 on level -1 , and X_1 connects to X_2 and X'_2 on lower levels which we cannot order. Let h_1, h_2, h'_2 be the adjusting parameters for X_1, X_2, X'_2 respectively, and assume that they are not zero divisors. Then the partial order implies that h_1 divides both h_2 and h'_2 , and hence

$$\mathcal{D}_U = (h_1)(h_2)(h'_2)(h_1, h_2)(h_1, h'_2)(h_2, h'_2)(h_1, h_2, h'_2) = (h_1)^4(h_2)(h'_2)(h_2, h'_2)$$

for a sufficiently small neighborhood U of p . Suppose $q \in U$ is a nearby point such that the fiber X_q is less degenerate in the sense that the nodes connecting X_2, X'_2 to X_1 are smoothed, i.e. suppose X_q has only one lower level component with adjusting parameter h_1 and both h_2, h'_2 become h_1 multiplied by some units in a neighborhood $U' \subset U$ of q . Then $\mathcal{D}_{U'} = (h_1)$, which differs from $\mathcal{D}_U|_{U'} = (h_1)^7$ by a power of (h_1) . In particular, the ideals (h_1) and $(h_1)^7$ define different subschemes in U' . However, since both ideals are principal, blowup along each of them is thus the identity map, so the resulting spaces are isomorphic to each other.

We need the following lemmas about the properties of disorderly ideals.

Lemma 14.9. *Let R be a local ring and $I, J \subset R$ be two ideals such that the product ideal IJ is a principal ideal generated by a non-zero-divisor. Then both I and J are principal ideals generated by non-zero-divisors.*

Proof. Suppose $IJ = (a)$ for some non-zero-divisor a . Then there exist $b_i \in I$ and $c_i \in J$ such that $b_1c_1 + \dots + b_nc_n = a$, which implies that $b_1(c_1/a) + \dots + b_n(c_n/a) = 1$ as a relation in the ring of fractions. Since the (unique) maximal ideal of R consists exactly of all non-unit elements, it follows that some $b_i(c_i/a)$ must be a unit, hence $I = (b_i)$. \square

Lemma 14.10. *Let R be a local ring and let $h_1, \dots, h_n \in R$ be some elements that are non-zero-divisors. Let $D = \prod (h_{i_1}, \dots, h_{i_k})$ be the product of ideals where $\{i_1, \dots, i_k\}$*

ranges over all subsets of $\{1, \dots, n\}$. Then D is a principal ideal (h) with h being a non-zero-divisor if and only if h_1, \dots, h_n are fully ordered by the divisibility relation.

Proof. If h_1, \dots, h_n are fully ordered by divisibility, it is clear that $D = (h)$ where h is given by certain products of powers of the h_i , and by assumption each h_i is a non-zero-divisor. Conversely if $D = (h)$ is principal with h being a non-zero-divisor, then the same holds for each factor $(h_{i_1}, \dots, h_{i_k})$ by Lemma 14.9. Suppose $(h_1, \dots, h_n) = (b)$ such that $h_i = bt_i$ for t_i in R and b being a non-zero-divisor. Then there exist u_i in R such that $u_1t_1 + \dots + u_nt_n = 1$. If all of t_1, \dots, t_n are not units, then the ideal (t_1, \dots, t_n) is contained in the (unique) maximal ideal of the local ring R , which is absurd because it also contains 1. Hence we may assume that t_1 is a unit in R , which implies that h_1 divides h_2, \dots, h_n . Carrying out the same analysis for the ideal (h_2, \dots, h_n) and repeating the process thus implies the desired claim. \square

The orderly blowup construction possesses some functorial property.

Proposition 14.11. *Given an adjustable family of differentials $(\pi: \mathcal{X} \rightarrow B, \omega, \mathbf{z})$, the pullback family $\tilde{\pi}: \tilde{\mathcal{X}} \rightarrow \tilde{B}$ is orderly. Moreover, any dominant map $\pi: B' \rightarrow B$, such that the pullback family $\mathcal{X}' \rightarrow B'$ is orderly, factors through \tilde{B} .*

Proof. It suffices to check the claim locally over each U with the disorderly ideal \mathcal{D}_U in the preceding setup. The first statement then follows from Lemma 14.10. More precisely, after blowing up the pullback of \mathcal{D}_U becomes a principal ideal, hence at every point of \tilde{U} the pullback family of differentials has adjusting parameters (given by the pullback of the functions h_i) fully ordered by divisibility, which implies that the family is orderly over \tilde{U} .

The second statement follows from the universal property of blowup (see e.g. [Sta18, Tag 01OF]). Let $U' = \pi^{-1}(U)$. Since π is dominant, the pullback of any adjusting parameter π^*h_i is a non-zero-divisor, and moreover $\pi^*\eta_{(i)} = \pi^*\omega/\pi^*h_i$ holds for the adjusted differential η on any irreducible component X_i of any fiber X_p over a point $p \in U$. Hence these π^*h_i can be used as adjusting parameters for the pullback family over U' . Since the pullback family is orderly, the adjusting parameters π^*h_i in U' are fully ordered by divisibility with respect to the full order, and consequently the corresponding disorderly ideal $\pi^*\mathcal{D}_U$ in U' is principal (generated by a non-zero-divisor). Since the blowup of \mathcal{D}_U is the final object that turns \mathcal{D}_U into a principal ideal (generated by a non-zero-divisor), it implies that $\pi: U' \rightarrow U$ factors through \tilde{U} . \square

We remark that there is some flexibility in choosing the local disorderly ideals. For instance, we can alternatively take $D = \prod(h_{i_1}, \dots, h_{i_k})$ to be the product of ideals ranging over all subsets of cardinality at least two. Then it differs from the original setting by a product of principal ideals, and hence the blowup gives the same space. We can also take the product $D = \prod(h_i, h_j)$ over all pairs of h_i and h_j that do not satisfy the divisibility relation. Then after blowing up the adjusting parameters are pairwise orderly, hence are orderly altogether.

We also warn the reader that the orderly blowup of a normal base may fail to be normal, as illustrated by the following example.

Example 14.12. (*A non-normal orderly blowup*) Let x and y be the standard coordinates of $B = \mathbb{C}^2$. Then x^2 and y^3 do not divide each other in the local ring of the origin. The orderly blowup \widetilde{B} for the ideal (x^2, y^3) can be described by

$$(14.1) \quad \{(x, y) \times [u, v] \in \mathbb{C}^2 \times \mathbb{P}^1 : x^2v - y^3u = 0\}.$$

Then we see that \widetilde{B} is singular along the entire exceptional curve over $x = y = 0$. It implies that \widetilde{B} is not normal, since a normal algebraic surface can have isolated singularities only.

Now we apply the previous considerations to the IVC.

Lemma 14.13. *The incidence variety compactification $\mathbb{P}\Omega\overline{\mathcal{M}}_{g,n}^{\text{inc}}(\mu)$ can be considered as a closed substack of $\mathbb{P}\Omega\overline{\mathcal{M}}_{g,n}$.*

Proof. We can restrict to the neighborhood of a stable curve with dual graph Γ . It suffices to realize that conditions of the existence of a twisted differential compatible with an enhanced level structure are closed conditions that can be read off from a family of pointed stable curves. This is clear both for the existence of differentials (i.e. sections of a line bundle determined by the family and the marked points) and the global residue condition (vanishing conditions of residues associated with these differentials). \square

Theorem 14.14. *The moduli stack of multi-scale differentials \mathcal{MS}_μ is equivalent to the normalization $\widetilde{\Omega\mathcal{M}}_{g,n}^{\text{n}}(\mu)$ of the orderly blowup of the normalization $\Omega\overline{\mathcal{M}}_{g,n}^{\text{inc}}(\mu)$ of the incidence variety compactification. In particular, \mathcal{MS}_μ is the analytification of an algebraic stack.*

Proof. First, we make sure that the operations of normalization (e.g. [ACG11, Example 8.3]) and orderly blowup make sense in this context, considering $\Omega\overline{\mathcal{M}}_{g,n}^{\text{inc}}(\mu)$ both as an algebraic and as an analytic stack. Proposition 7.13 ensures that over $\Omega\overline{\mathcal{M}}_{g,n}^{\text{inc}}(\mu)$ the family of one-forms that are the top level forms of the twisted differential is adjustable. From the proof of this proposition we see that the adjusting parameters are defined not only locally in the analytic topology, but also locally in the Zariski topology. In a local quotient groupoid presentation $[U/G]$ we see that G -pullbacks of adjusting parameters are again adjusting parameters. Consequently, the disorderly ideal is G -invariant and the blowup is well-defined. In particular we can consider $\widetilde{\Omega\mathcal{M}}_{g,n}^{\text{n}}(\mu)$ both as an algebraic and as an analytic stack.

Proposition 7.15 then ensures that the resulting orderly family over $\widetilde{\Omega\mathcal{M}}_{g,n}^{\text{n}}(\mu)$ gives a family of multi-scale differentials of type μ . This family induces a map of stacks $\widetilde{\Omega\mathcal{M}}_{g,n}^{\text{n}}(\mu) \rightarrow \mathcal{MS}_\mu$.

Conversely, a family in the stack \mathcal{MS}_μ is orderly by definition, hence by Proposition 14.11 we obtain a map of stacks $\mathcal{MS}_\mu \rightarrow \widetilde{\Omega\mathcal{M}}_{g,n}^{\text{n}}(\mu)$. Since \mathcal{MS}_μ is normal, this map factors through $\widetilde{\Omega\mathcal{M}}_{g,n}^{\text{n}}(\mu)$, which gives the desired inverse map. \square

It is well-known that the blowup of a projective scheme along a globally defined ideal sheaf (or equivalently a globally defined subscheme) remains to be projective. Nevertheless, we remark that in general gluing local blowups can lead to a non-projective

global space (recall the famous Hironaka's example, see e.g. [Har77, Appendix B.3, Example 3.4.1]).

Proof of Theorem 1.2 and Theorem 1.3 completed. The density (1) and the description of the boundary divisor (2) have been taken care of in Proposition 14.4. Compactness (3) is the content of Theorem 14.3 and the coarse moduli space (5) has been addressed in Theorem 14.5. The algebraicity (4) is a consequence of Theorem 14.14 and the remark that the coarse moduli spaces are glued from the quotients U/G (as schemes). The forgetful map (6) is obvious, e.g. it follows from Theorem 14.14 and its proof. We have thus completed the proof of Theorem 1.2.

Finally, the property of being a proper Deligne-Mumford stack carries over from $\overline{\mathcal{M}}_{g,n}$ all the way up through $\mathbb{P}\Omega\overline{\mathcal{M}}_{g,n}$, $\mathbb{P}\Omega\overline{\mathcal{M}}_{g,n}^{\text{inc}}(\mu)$, and $\widetilde{\mathbb{P}\Omega\mathcal{M}}_{g,n}^n(\mu)$. The isomorphism in the statement of Theorem 1.3 is obvious since our compactification does not alter the interior $\mathbb{P}\Omega\mathcal{M}_{g,n}(\mu)$ and on a smooth curve a multi-scale differential is simply an abelian differential of type μ . \square

Remark 14.15. We remark that the forgetful map in Theorem 1.3 is not an isomorphism at the boundary strata where K_Γ is non-trivial. Consider such a boundary point (X, ω, σ) , e.g. in our running example, and suppose for simplicity (X, ω, σ) has no non-trivial automorphisms. Then in a neighborhood of that point the stack \mathcal{MS}_μ is represented by the singular space U/K_Γ where U is an open set in \mathbb{C}^2 times a product of moduli spaces. Unwinding definitions, one checks that there is no map from this singular space to the orbifold $[U/K_\Gamma]$ that provides an inverse to the natural quotient map $[U/K_\Gamma] \rightarrow U/K_\Gamma$.

The proof and the previous considerations show that ideally $\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ should be considered in a hybrid category between orbifolds and analytic spaces, where local orbifold charts of the form $[U/G]$ with $U \subset \mathbb{C}^N$ open (and thus smooth) and singular analytic spaces (in fact with abelian quotient singularities) are both permitted. Since such a category is not common in the literature and since in practice computations (such as intersection numbers, see e.g. [CMZ19b]) are performed in covering charts in both cases, i.e. on the orbifold $\Xi\overline{\mathcal{M}}_{g,n}(\mu)$, we decided not to introduce such a category.

14.4. Some moduli spaces in genus zero and cherry divisors. To illustrate the necessity of both orderly blowup and subsequent normalization in the passage from IVC to $\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ we consider the following class of divisors.

A *cherry divisor* is a boundary divisor of $\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ such that the generic multi-scale differential has one top-level component and two components at the second level, each connected to the top level by one node. Note that the forgetful map from the moduli space of multi-scale differentials to the incidence variety compactification (and hence to the Deligne-Mumford compactification) contracts any cherry divisor, as we saw in Example 14.7.

Example 14.16. (*The cherry requires orderly blowup and normalization*) We consider the incidence variety compactification of $\mathbb{P}\Omega\mathcal{M}_{0,5}(2, 1, 0, 0, -5)$, with marked points. Note that in this case the IVC is simply $\overline{\mathcal{M}}_{0,5}$, and in particular it is smooth. On the right in Figure 9 we schematically depict the local structure of $\mathbb{P}\Omega\overline{\mathcal{M}}_{0,5}^{\text{inc}}(2, 1, 0, 0, -5)$

near the point that is the image of the cherry divisor in the moduli space of multi-scale differentials. We will study the cherry where the marked points meet the zeros of orders 1 and 2, respectively. This point is the intersection of two boundary divisors of the IVC, the first one parameterizing the differentials where the zero of order 1 meets a marked point and the second parameterizing the differentials where the zero of order 2 meets the other marked point. We introduce local coordinates x, y on the IVC near this cherry point, such that the first divisor is the locus $\{x = 0\}$ and the second one is $\{y = 0\}$. Note that the number of prongs is respectively equal to $\kappa_1 = 2$ and $\kappa_2 = 3$ along these two divisors.

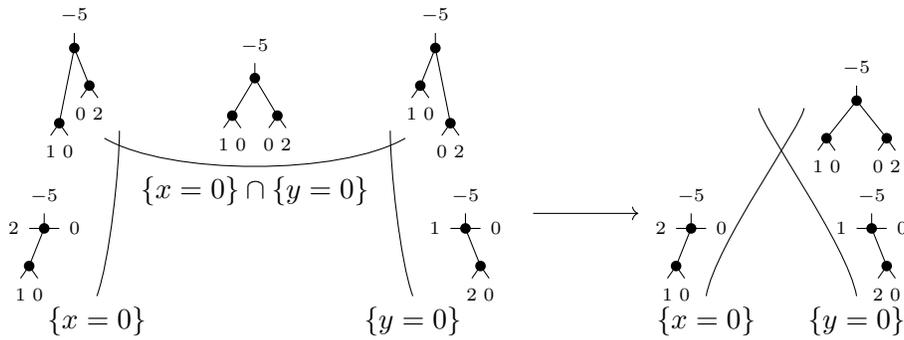


FIGURE 9. The orderly blowup of the incidence variety compactification of $\mathbb{P}\Omega\mathcal{M}_{0,5}(2, 1, 0, 0, -5)$ at a cherry point.

Let us perform the orderly blowup in the neighborhood of the cherry point. We have to blow up the ideal (x^2, y^3) discussed in Example 14.12 (see (14.1) for the description). We recall that this space is not normal, and that the exceptional locus of this orderly blowup is a \mathbb{P}^1 which is parameterized by the ratio of the differentials on the two lower level components. This exceptional locus meets the strict transforms of the two divisors $\{x = 0\}$ and $\{y = 0\}$ in two distinct points. The complete picture of this orderly blowup is represented in Figure 9. Hence in this case the moduli space of multi-scale differentials is obtained by normalizing the orderly blowup of the IVC, and this normalization is not the identity map. We note moreover that in this case all prong-matchings are equivalent, and thus this difficulty is not due to the choice of a prong-matching.

We now illustrate the fact that the orderly blowup does not see the prong-matchings in general. Consider the stratum $\mathbb{P}\Omega\mathcal{M}_{0,5}(1, 1, 0, 0, -4)$. We will study the cherry where the marked points meet respectively the simple zeros, so that the number of prongs is $\kappa_1 = \kappa_2 = 2$, and there are two non-equivalent prong-matchings on the generic cherry curve.

The orderly blowup is given by the equation

$$\{(x, y) \times [u, v] \in \mathbb{C}^2 \times \mathbb{P}^1 : x^2v - y^2u = 0\} .$$

Note that this space has two locally irreducible branches meeting along the exceptional divisor. In the moduli space of multi-scale differentials, the limits from these

two branches will give non-equivalent prong-matchings for the limiting twisted differential. But in the orderly blowup, both branches converge to the same limit. Hence it is not possible to distinguish the prong-matchings from the orderly blowup. However, the normalization of the orderly blowup precisely separates these two branches corresponding to the two non-equivalent prong-matchings.

15. EXTENDING THE $\mathrm{SL}_2(\mathbb{R})$ -ACTION TO THE BOUNDARY

The goal of this section is to modify the boundary of the moduli space of multi-scale differentials in such a way that the $\mathrm{SL}_2(\mathbb{R})$ -action on the open stratum extends to this boundary, and such that the quotient of this compactification by rescaling by positive real numbers is compact. The reason we need to consider rescaling by $\mathbb{R}_{>0}$ instead of by \mathbb{C}^* is essentially due to the fact that $\mathrm{SL}_2(\mathbb{R})$ does not act meaningfully on $\Omega\mathcal{M}_{g,n}(\mu)/\mathbb{C}^*$ but it does act on $\Omega\mathcal{M}_{g,n}(\mu)/\mathbb{R}_{>0}$, as $\mathrm{SO}_2(\mathbb{R})$ is not a normal subgroup of $\mathrm{SL}_2(\mathbb{R})$ but $\mathbb{R}_{>0}$ is. The concept of level-wise real blowup provides the setup for this purpose. A related bordification, also a manifold with corners, is also studied in an ongoing project of Smillie and Wu with the goal of understanding the $\mathrm{SL}_2(\mathbb{R})$ -action near the boundary. While the constructions have certain similarities, they differ e.g. in the treatment of horizontal nodes.

Theorem 15.1. *The $\mathrm{SL}_2(\mathbb{R})$ -action on the moduli space $\Omega\mathcal{M}_{g,n}(\mu)$ extends to a continuous $\mathrm{SL}_2(\mathbb{R})$ -action on the level-wise real blowup $\Xi\widehat{\mathcal{M}}_{g,n}(\mu)$ of the moduli space of multi-scale differentials $\Xi\overline{\mathcal{M}}_{g,n}(\mu)$.*

In comparison to Section 8, note that $\Xi\widehat{\mathcal{M}}_{g,n}(\mu)$ agrees with $\overline{\Xi\overline{\mathcal{M}}_{g,n}(\mu)}$ (where the long upper bar refers to the level-wise real blowup) because the generic fiber is smooth, and there are no persistent nodes.

The basic objects parameterized by $\Xi\widehat{\mathcal{M}}_{g,n}(\mu)$ are *real* multi-scale differentials, replacing multi-scale differentials. The definition is very similar to Definition 1.1, simply replacing the equivalence relation.

Definition 15.2. *A real multi-scale differential of type μ on a stable curve X is*

- (i) a full order \preceq on the dual graph Γ of X ,
- (ii) a differential $\omega_{(i)}$ on each level $X_{(i)}$, such that the collection of these differentials satisfies the properties of a twisted differential of type μ compatible with \preceq , and
- (iii) a prong-matching $\sigma = (\sigma_e)$ where e runs through all vertical edges of Γ .

Two real multi-scale differentials are considered equivalent if they differ by rescaling at each level (but the top level) by multiplication by an element in $\mathbb{R}_{>0}$. \triangle

To properly state families of such differentials, we have to also leave the category of complex spaces. Recall that *manifolds with corners* are topological spaces locally modelled on $[0, \infty)^k \times \mathbb{R}^{n-k}$. These spaces form a category (with a notion of smooth maps, see [Joy12] for a recent account with definitions and caveats, but we will not detail here). Since $\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ already has non-trivial orbifold structures, we in fact work with *orbifolds with corners*, where the local orbifold charts are manifolds with corners and where the local group actions are smooth maps preserving the boundary.

Theorem 15.3. *The level-wise real blowup $\Xi\widehat{\mathcal{M}}_{g,n}(\mu)$ is an orbifold with corners. Its points correspond bijectively to isomorphism classes of real multi-scale differentials.*

The reason for orbifold structures is due to automorphisms of flat surfaces, as for $\Omega\mathcal{M}_{g,n}$, and also the boundary points where $\text{Tw}_\Lambda/\text{Tw}_\Lambda^s$ introduces quotient singularities.

Given Theorem 15.3, we define the action of $A \in \text{SL}_2(\mathbb{R})$ on $\Xi\widehat{\mathcal{M}}_{g,n}(\mu)$ by

$$A \cdot (X = (X_{(i)}), \omega = (\omega_{(i)}), \sigma) = ((A \cdot (X_{(i)}, \omega_{(i)})), A \cdot \sigma) ,$$

where $i \in L^\bullet(\Gamma)$. The first argument is the usual $\text{SL}_2(\mathbb{R})$ -action on the components of the stable curve. For the second argument we use the action of A on the set of directions and note that a matching of horizontal directions (for ω) gives a matching of directions of slope $A \cdot (\frac{1}{0})$ (for $A \cdot \omega$) that can be reconverted into a matching of horizontal directions.

The notion of families of real multi-scale differentials (over bases being orbifolds with corners) can now be phrased as in Section 7, with the equivalence relation changed from $T_\Gamma(\mathcal{O}_{B,p})$ -action to level-wise $\mathbb{R}_{>0}$ -multiplication, as above. We will now essentially construct a universal object.

All the statements in the above theorems are local, since continuity of the $\text{SL}_2(\mathbb{R})$ -action can also be probed by a neighborhood of the identity element. We thus pick a point $p \in \Xi\overline{\mathcal{M}}_{g,n}(\mu)$ and work in a neighborhood U that will be shrunk for convenience, e.g. to apply Proposition 8.5 and to find an enhanced multicurve Λ with $\Gamma(\Lambda) = \Gamma_p$ and provide the restriction of the universal family over (the orbifold chart of) U with a Λ -marking. We may thus view $U \subset \Xi\mathcal{D}_\Lambda$.

We provide the pullback to the level-wise real blowup \widehat{U} of the universal family over U with real multi-scale differentials. (Note that here as in Definition 15.2 above, real multi-scale differentials live on stable complex curves.) The smooth (differentiable) family $\widehat{X} \rightarrow \widehat{U}$ constructed in Theorem 8.2 is tacitly used for the marking, but we do not treat the issue whether differentials can be pulled back there.) Let t_i be the rescaling parameters of the multi-scale differential $\omega = (\omega_{(i)})_{i \in L^\bullet(\Lambda)}$ on U and let T_i and F_e be the S^1 -valued functions used in the blowup construction (Section 8.1).

Proof of Theorem 15.3. The second statement is an immediate consequence of Theorem 14.5 about the points in $\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ and of Theorem 8.2 that describes the fibers of the level-wise real blowup as the argument-images of the level rotation tori.

For the first statement we may use charts of the level-wise real blowup of $\Xi\mathcal{D}_\Lambda^s$ as orbifold charts. There, the boundary is a normal crossing divisor with one component $D_i = \{t_i = 0\}$ for each level (but the top level). The real blowup of a normal crossing divisor is then known to be a manifold with corners (see e.g. [ACG11], Section X.9, in particular page 150). \square

Proof of Theorem 15.1. We need to justify continuity. Consider a sequence $\{\widehat{p}_n\}$ converging to \widehat{p} in $\widehat{U} \subset \Xi\widehat{\mathcal{D}}_\Lambda$. By definition of the topology, this is equivalent (with notations as in Section 8) to the convergence of the image points $\varphi_U(\widehat{p}_n)$ to $\varphi_U(\widehat{p})$ in $U \subset \Xi\mathcal{D}_\Lambda$ and to the convergence of $F_e(\widehat{p}_n)$ to $F_e(\widehat{p})$ and $T_i(\widehat{p}_n)$ to $T_i(\widehat{p})$. In turn, the convergence in $\Xi\mathcal{D}_\Lambda$ is manifested by diffeomorphism g_n satisfying Definition 9.3 with

the compatibility with markings relaxed up to elements in Tw_Λ . We aim to justify the convergence of the sequence of image points $\varphi_U(A \cdot \widehat{p}_n)$ to $\varphi_U(A \cdot \widehat{p}) \in U \subset \overline{\Xi \mathcal{M}}_{g,n}(\mu)$. Since $T_i(A \cdot \widehat{p}_n)$ converges to $T_i(A \cdot \widehat{p})$, and similarly for F_e , by continuity of the $\text{SL}_2(\mathbb{R})$ -action on S^1 , this will conclude the proof.

For our aim, we use the maps $A \cdot g_n \cdot A^{-1}$, where $A \cdot (\cdot)$ denotes the induced $\text{SL}_2(\mathbb{R})$ -action on pointed flat surfaces. This map is well-defined away from the seams and we observe that it can be extended to a differentiable map across each seam using the action of $\text{SL}_2(\mathbb{R})$ on the seam identified with S^1 . \square

INDEX OF NOTATION

In this section we summarize the notations thematically. In each theme we mainly give the notations in chronological order, omitting the introduction.

Surfaces.

(Σ, \mathbf{s})	“Base” compact n -pointed oriented differentiable surface	14
(X, \mathbf{z})	Pointed stable curve of genus g	12
X_v	Irreducible component of X	12
N_X	Set of nodes of X	11
$X^s = X \setminus N_X$	The smooth part of X	15
N_X^v	Set of vertical nodes of X	11
N_X^h	Set of horizontal nodes of X	11
$f: \Sigma \rightarrow X$	Marking	14
P_s, Z_s	Subset of \mathbf{s} mapped respectively to the poles and zeros of ω	36
$f: S \rightarrow S'$	An almost-diffeomorphism between two almost smooth surfaces	38
$\delta_q(\theta)$	Fractional Dehn twist of \overline{X}_σ of angle θ at q	38
\overline{X}_σ	Welded surface associated to the prong-matching σ	40
$X_{(i)}$	Components of X at level i	11
$\eta_{(i)}$	Restriction of the twisted differential η on $X_{(i)}$	12
$X_{>i}$	Components of X at level $> i$	12
$X_\epsilon, (X, \mathbf{z})_\epsilon$	ϵ -thick part of X , resp. $X \setminus \mathbf{z}$	15

Graphs and Levels.

$\overline{\Gamma} = (\Gamma, \succ), \Gamma$	Level graph with full order \succ	11
$V(\Gamma)$	Vertices of Γ	11
$E(\Gamma)$	Edges of Γ	11
$E(\Gamma)^v, E(\Gamma)^h$	Set of vertical, resp. horizontal, edges of Γ	11
$\text{val}(v)$	Valence of the vertex v	11
$L^\bullet(\overline{\Gamma})$	Set of levels of the level graph $\overline{\Gamma}$	11
$L(\overline{\Gamma})$	Set of all but the top level of the level graph $\overline{\Gamma}$	11
N	Number of levels strictly below 0	11
$\ell: \Gamma \rightarrow \underline{N}$	Normalized level function	11
$\overline{\Gamma}_{(i)}, \overline{\Gamma}_{>i}$	Subgraph at (resp. above) level i of $\overline{\Gamma}$	11

$\ell(q^\pm), \ell(e^\pm)$	Bottom and top levels of the ends of a node	12
Γ^+, Γ	Enhanced level graph	13
Λ	Multicurve in Σ	34
Λ^+, Λ	Enhanced multicurve	34
$\Gamma^+(\Lambda^+), \Lambda$	Enhanced graph associated to the enhanced multicurve Λ	34
$\text{dg}: \Lambda_2 \rightsquigarrow \Lambda_1$	Degeneration of the ordered multicurve Λ_2	35
$\delta: \underline{N} \rightarrow \underline{M}$	Map defining a vertical undegeneration	35
$D^h \subseteq \Lambda_1^h$	Subset of horizontal curves inducing a horizontal undegeneration	35
$(\delta, D^h), \delta$	Undegeneration of an enhanced multicurve	36
dg_J, δ_J	(Un)degenerations associated with the subset J	36

Teichmüller and Moduli Spaces. Most of the following spaces have a projectivized variant which is indicated with the symbol \mathbb{P} .

$\mathcal{T}_{g,n} = \mathcal{T}_{(\Sigma,s)}$	Teichmüller space	14
$\text{Mod}_{g,n}$	Classical mapping class group	15
$\overline{\mathcal{T}}_{g,n} = \overline{\mathcal{T}}_{(\Sigma,s)}$	Augmented Teichmüller space	15
\mathcal{D}_Λ	Classical Dehn space	19
$\Omega\mathcal{T}_{(\Sigma,s)}(\mu)$	Teichmüller space of marked flat surfaces of type μ	36
$\Omega^{no}\mathcal{T}_\Lambda(\mu)$	Teichmüller space of flat surfaces of type (μ, Λ) without GRC	36
$\Omega\mathcal{T}_\Lambda(\mu)$	Teichmüller space of twisted differentials of type (μ, Λ)	37
$\Omega\mathcal{T}_\Lambda^{pm}(\mu)$	Teichmüller space of prong-matched twisted differentials	40
$\Omega\mathcal{B}_\Lambda$	The Λ^+ -boundary stratum $\Omega\mathcal{T}_\Lambda^{pm}(\mu)/\mathbb{C}^{L(\Lambda^+)}$	65
$\Omega\overline{\mathcal{T}}_{(\Sigma,s)}(\mu)$	Augmented Teichmüller space of marked flat surfaces of type μ	65
$\Omega\mathcal{D}_\Lambda$	Hodge bundle over the Dehn space	20
$\overline{\mathcal{M}\mathcal{D}}_\Lambda^s, \overline{\mathcal{M}\mathcal{D}}_\Lambda$	(Smooth) model domain	72
$\Xi\mathcal{D}_\Lambda$	Dehn space associated with Λ	83
$\Xi\mathcal{D}_\Lambda^s$	Simple Dehn space associated with Λ	83
$\Xi\overline{\mathcal{M}}_{g,n}(\mu)$	Moduli space of multi-scale differentials of type μ	105
$\Xi\widehat{\mathcal{M}}_{g,n}(\mu)$	Level-wise real blowup of $\Xi\overline{\mathcal{M}}_{g,n}(\mu)$	114
\mathbf{MS}_μ	Functor of multi-scale differentials	57
\mathcal{MS}_μ	Grupoid of multi-scale differentials	57
$\mathbf{MS}_{(\mu,\Lambda)}$	Functor of marked multi-scale differentials	63
$\mathbf{MD}_{(\mu,\Lambda)}$	Functor of model differentials	70

Families.

\mathcal{Z}_j	Image of the section z_j	10
\mathcal{Z}^0	Horizontal zero divisor	10
\mathcal{Z}^∞	Horizontal polar divisor	10
$uv = f$	Local equation of a nodal family	29
$u_e^+ u_e^- = f_e$	Local equation of a family near the node q_e	62
$\widehat{\pi}: \widehat{\mathcal{X}} \rightarrow \widehat{\mathcal{B}}$	Level-wise real blowup of a family of multi-scale differentials	61

$\bar{\pi}: \bar{\mathcal{X}} \rightarrow \bar{B}$	Global level-wise real blowup	62
$\tilde{\pi}: \tilde{\mathcal{X}} \rightarrow \tilde{B}$	Orderly blowup of an adjustable family of differentials	110
h, h_v, h_i	Adjusting parameter	58
$R_p: B_p \rightarrow \bar{T}_{\Gamma_p}^n$	Rescaling ensemble	52
s_i	Rescaling parameters	52
f_e	Smoothing parameters	52
(x_j)	Smoothing parameters for the horizontal nodes	82
$R^s: B \rightarrow \bar{T}_{\Gamma_p}^s$	Simple rescaling ensemble	55
t_i	Roots of rescaling parameters given by a simple rescaling ensemble	55

Plumbing Constructions.

$(X_0, \boldsymbol{\eta}_0)$	Base surface in the Λ -boundary stratum	84
ϕ_e^\pm	Local coordinates at a node e	86
$r_e(\mathbf{t})$	Residue of $\mathbf{t} * \boldsymbol{\eta}$ at the node e	86
ξ	Family of modifying differentials	75
ρ_i	Period homomorphism	76
A_{δ_1, δ_2}	Annulus of inner radius δ_1 and outer radius δ_2	85
A^\pm	Top and bottom plumbing annuli	85
R, δ	Defining constants of the plumbing annuli	85
$V(s)$	Standard plumbing fixture	85
p^\pm	Top and bottom marked points	85
U	Open in $\Omega \overline{\mathcal{M}D}_\Lambda^s \times \Delta^H$ on which the plumbing map is defined	83
\mathcal{W}_e	Base of the plumbing construction	85
Ψ	Local chart center at the base surface	84
$\tilde{\Psi}$	Lift of Ψ to $\Omega \overline{\mathcal{M}D}_\Lambda^s$	84
$\mathcal{A}^-, \mathcal{A}^+$	Bottom and top plumbing annuli in the plumbing fixture	85
$\mathcal{B}^-, \mathcal{B}^+$	Bottom and top plumbing annuli in \mathcal{X}	89
$\mathcal{C}^-, \mathcal{C}^+$	Bottom and top plumbing annuli in \mathcal{Y}	90
v_e^+, v_e^-, v_h	Conformal maps on annuli putting $\mathbf{t} * \boldsymbol{\eta} + \xi$ in standard form	87
b_e^\pm, b_h	Image of p^\pm in \mathcal{B}^\pm	88
c_e^\pm	Image of the points p^\pm in \mathcal{C}^\pm	90

Prong-matchings and Rotation Groups. The groups below usually have “extended” analogues which we denote by a superscript \bullet .

κ_q	Number of prongs, equal to $\text{ord}_{q^+} \eta + 1$	13
$\boldsymbol{\sigma}$	Global prong-matching for X	39
$\mathbf{d} = (d_i)_{i \in L^\bullet(\Lambda)}$	Tuple in $\mathbb{C}^{L(\Lambda)}$ acting on prong-matched differentials	43
\cdot	Action of $\mathbb{C}^{L(\Lambda)}$ on prong-matched twisted differentials	43
$*$	Action of T_Λ on prong-matched twisted differentials	48
$F_{\mathbf{d}}$	Fractional Dehn twist	44
$a_i, m_{e,i}$	Defined by $a_i = \text{lcm}_e \kappa_e$ and $m_{e,i} = a_i / \kappa_e$	46

$t_{[i]}$	Product of the $t_j^{a_j}$ for $j \geq i$	75
P_Γ	Prong rotation group	40
$\text{Tw}_\Lambda^{\text{full}}$	Classical Λ -twist group	19
$\mathbb{Z}^{L^\bullet(\Lambda)}$	Level rotation group	44
ϕ_Λ^\bullet	Map from level rotation group to prong rotation group	44
Tw_Λ^v	Vertical twist group	44
Tw_Λ^h	Horizontal twist group	45
Tw_Λ	Twist group	45
$\text{Tw}_{\Lambda,i}^{sv}$	Simple vertical twist group of level i	46
Tw_Λ^{sv}	Simple vertical twist group	46
Tw_Λ^s	Simple twist group	46
T_Λ^s, T_Λ	(Simple) level rotation torus	47
H, G	Ramifications groups	49
K_Λ	Finite group defined by $\text{Tw}_\Lambda^v/\text{Tw}_\Lambda^{sv}$	49

Other Notations.

$\Delta_r = \{z \in \mathbb{C} : z < r\}$	18	$\bar{n} = \{1, \dots, n\}$	10
$e(z) = \exp(2\pi\sqrt{-1}z)$	28	$\underline{N} = \{0, -1, \dots, -N\}$	11

REFERENCES

[ABF16] F. Acquistapace, F. Broglia, and J. F. Fernando. *On the Nullstellensätze for Stein spaces and C -analytic sets*. Trans. Amer. Math. Soc. 368.6 (2016), pp. 3899–3929.

[Abi77] W. Abikoff. *Degenerating families of Riemann surfaces*. Ann. of Math. (2) 105.1 (1977), pp. 29–44.

[Abr+15] D. Abramovich et al. *Logarithmic geometry and moduli*. In: *Handbook of moduli. Volume I*. Somerville, MA: International Press, 2015, pp. 1–61.

[ACG11] E. Arbarello, M. Cornalba, and P. Griffiths. *Geometry of algebraic curves. Volume II*. Vol. 268. Grundlehren der Mathematischen Wissenschaften. With a contribution by J. Harris. Heidelberg: Springer, 2011, pp. xxx+963.

[ACGH85] E. Arbarello, M. Cornalba, P. A. Griffiths, and J. Harris. *Geometry of algebraic curves. Vol. I*. Vol. 267. Grundlehren der Mathematischen Wissenschaften. New York: Springer-Verlag, 1985, pp. xvi+386.

[Bar18] I. Barros. *Uniruledness of strata of holomorphic differentials in small genus*. Adv. Math. 333 (2018), pp. 670–693.

[BCGGM18] M. Bainbridge, D. Chen, Q. Gendron, S. Grushevsky, and M. Möller. *Compactification of strata of abelian differentials*. Duke Math. J. 167.12 (2018), pp. 2347–2416.

[BCGGM19] M. Bainbridge, D. Chen, Q. Gendron, S. Grushevsky, and M. Möller. *Strata of k -differentials*. Algebr. Geom. 6.2 (2019), pp. 196–233.

- [Ber74] L. Bers. *Spaces of degenerating Riemann surfaces*. In: *Discontinuous groups and Riemann surfaces (Proc. Conf., Univ. Maryland, College Park, Md., 1973)*. Princeton, N.J.: Princeton Univ. Press, 1974, 43–55. *Ann. of Math. Studies*, No. 79.
- [BF82] C. Bănică and O. Forster. *Complete intersections in Stein manifolds*. *Manuscripta Math.* 37.3 (1982), pp. 343–356.
- [BHM16] M. Bainbridge, P. Habegger, and M. Möller. *Teichmüller curves in genus three and just likely intersections in $\mathbf{G}_m^n \times \mathbf{G}_a^n$* . *Publ. Math. Inst. Hautes Études Sci.* 124 (2016), pp. 1–98.
- [BJ06] A. Borel and L. Ji. *Compactifications of symmetric and locally symmetric spaces*. Mathematics: Theory & Applications. Birkhäuser Boston, Inc., 2006, pp. xvi+479.
- [CMSZ19] D. Chen, M. Möller, A. Sauvaget, and D. Zagier. *Masur-Veech volumes and intersection theory on moduli spaces of abelian differentials*. (2019). arXiv: 1901.01785.
- [CMZ19a] M. Costantini, M. Möller, and J. Zachhuber. *The area is a good metric*. (2019). arXiv: 19???.?????
- [CMZ19b] M. Costantini, M. Möller, and J. Zachhuber. *The Euler characteristic of strata*. (2019). arXiv: 19???.?????
- [CW19] D. Chen and A. Wright. *The WYSIWYG compactification*. (2019). arXiv: 1908.07436.
- [DM69] P. Deligne and D. Mumford. *The irreducibility of the space of curves of given genus*. *Inst. Hautes Études Sci. Publ. Math.* 36 (1969), pp. 75–109.
- [EFW18] A. Eskin, S. Filip, and A. Wright. *The algebraic hull of the Kontsevich-Zorich cocycle*. *Ann. Math. (2)* 188.1 (2018), pp. 281–313.
- [EM12] C. Earle and A. Marden. *Holomorphic plumbing coordinates*. In: *Quasi-conformal mappings, Riemann surfaces, and Teichmüller spaces*. Vol. 575. *Contemp. Math.* Amer. Math. Soc., Providence, RI, 2012, pp. 41–52.
- [EM18] A. Eskin and M. Mirzakhani. *Invariant and stationary measures for the $SL(2, \mathbb{R})$ action on moduli space*. *Publ. Math. Inst. Hautes Études Sci.* 127 (2018), pp. 95–324.
- [EMM15] A. Eskin, M. Mirzakhani, and A. Mohammadi. *Isolation, equidistribution, and orbit closures for the $SL(2, \mathbb{R})$ action on moduli space*. *Ann. Math. (2)* 182.2 (2015), pp. 673–721.
- [EMMW18] A. Eskin, C. McMullen, R. Mukamel, and A. Wright. *Billiards, quadrilaterals and moduli spaces* (2018). preprint.
- [Eul32] L. Euler. *Nova methodus innumerabiles aequationes differentiales secundi gradus reducendi ad aequationes differentiales primi gradus*. *Commentarii academiae scientiarum Petropolitanae* 3 (1732), pp. 124–137.
- [Fil16] S. Filip. *Splitting mixed Hodge structures over affine invariant manifolds*. *Ann. Math. (2)* 183.2 (2016), pp. 681–713.
- [FP18] G. Farkas and R. Pandharipande. *The moduli space of twisted canonical divisors*. *J. Inst. Math. Jussieu* 17.3 (2018), pp. 615–672.

- [FV14] G. Farkas and A. Verra. *The geometry of the moduli space of odd spin curves*. Ann. of Math. (2) 180.3 (2014), pp. 927–970.
- [Gen18] Q. Gendron. *The Deligne-Mumford and the incidence variety compactifications of the strata of ΩM_g* . Ann. Inst. Fourier (Grenoble) 68.3 (2018), pp. 1169–1240.
- [GT17] Q. Gendron and G. Tahar. *Différentielles à singularités prescrites*. (2017). arXiv: 1705.03240.
- [Har77] R. Hartshorne. *Algebraic geometry*. Graduate Texts in Mathematics, No. 52. New York: Springer-Verlag, 1977, pp. xvi+496.
- [HK14] J. Hubbard and S. Koch. *An analytic construction of the Deligne-Mumford compactification of the moduli space of curves*. J. Differential Geom. 98.2 (2014), pp. 261–313.
- [HKP18] D. Holmes, J. Kass, and N. Pagani. *Extending the double ramification cycle using Jacobians*. Eur. J. Math. 4.3 (2018), pp. 1087–1099.
- [HM79] J. Hubbard and H. Masur. *Quadratic differentials and foliations*. Acta Math. 142.3-4 (1979), pp. 221–274.
- [HS19] D. Holmes and J. Schmitt. *Infinitesimal structure of the pluricanonical double ramification locus*. (2019). arXiv: 1909.11981.
- [Hub06] J. Hubbard. *Teichmüller theory and applications to geometry, topology, and dynamics. Vol. 1*. Matrix Editions, Ithaca, NY, 2006, pp. xx+459.
- [Joy12] D. Joyce. *On manifolds with corners*. In: *Advances in geometric analysis*. Vol. 21. Adv. Lect. Math. Int. Press, Somerville, 2012, pp. 225–258.
- [Kat00] F. Kato. *Log smooth deformation and moduli of log smooth curves*. Internat. J. Math. 11.2 (2000), pp. 215–232.
- [KM98] J. Kollár and S. Mori. *Birational geometry of algebraic varieties*. Vol. 134. Cambridge Tracts in Mathematics. With the collaboration of C. H. Clemens and A. Corti, Translated from the 1998 Japanese original. Cambridge University Press, Cambridge, 1998, pp. viii+254.
- [KN99] K. Kato and C. Nakayama. *Log Betti cohomology, log étale cohomology, and log de Rham cohomology of log schemes over \mathbf{C}* . Kodai Math. J. 22.2 (1999), pp. 161–186.
- [LV73] O. Lehto and K. Virtanen. *Quasiconformal mappings in the plane. Translated from the German by K. W. Lucas. 2nd ed.* Vol. 126. Grundlehren der Mathematischen Wissenschaften. Springer, Berlin, 1973.
- [Mar87] A. Marden. *Geometric complex coordinates for Teichmüller space*. In: *Mathematical aspects of string theory (San Diego, Ca., 1986)*. Vol. 1. Adv. Ser. Math. Phys. World Sci. Publishing, Singapore, 1987, pp. 341–354.
- [McM94] C. McMullen. *Complex dynamics and renormalization*. Vol. 135. Annals of Mathematics Studies. Princeton University Press, Princeton, NJ, 1994, pp. x+214.
- [Mon09] G. Mondello. *A criterion of convergence in the augmented Teichmüller space*. Bull. Lond. Math. Soc. 41.4 (2009), pp. 733–746.

- [Muj86] J. Mujica. *Complex analysis in Banach spaces. Holomorphic functions and domains of holomorphy in finite and infinite dimensions*. Vol. 120. North-Holland Math. Stud. Elsevier, Amsterdam, 1986.
- [MW17] M. Mirzakhani and A. Wright. *The boundary of an affine invariant submanifold*. *Invent. Math.* 209.3 (2017), pp. 927–984.
- [Nac69] L. Nachbin. *Topology on spaces of holomorphic mappings*. *Ergebnisse der Mathematik und ihrer Grenzgebiete, Band 47*. Springer-Verlag New York Inc., New York, 1969, pp. v+66.
- [Pom92] C. Pommerenke. *Boundary behaviour of conformal maps*. Vol. 299. *Grundlehren der Mathematischen Wissenschaften*. Springer-Verlag, Berlin, 1992, pp. x+300.
- [PY16] M. Porta and T.Y. Yu. *Higher analytic stacks and GAGA theorems*. *Adv. Math.* 302 (2016), pp. 351–409.
- [Raf07] K. Rafi. *Thick-thin decomposition for quadratic differentials*. *Math. Res. Lett.* 14.2 (2007), pp. 333–341.
- [Sau18] A. Sauvaget. *Volumes and Siegel–Veech constants of $\mathcal{H}(2g - 2)$ and Hodge integrals*. *Geom. Funct. Anal.* 28.6 (2018), pp. 1756–1779.
- [Str84] K. Strebel. *Quadratic differentials*. Vol. 5. *Ergebnisse der Mathematik und ihrer Grenzgebiete (3)*. Springer-Verlag, Berlin, 1984, pp. xii+184.
- [Szp79] L. Szpiro. *Lectures on equations defining space curves*. Notes by N. Mohan Kumar. Tata Institute of Fundamental Research, Bombay; by Springer-Verlag, Berlin-New York, 1979, pp. v+81.
- [Toë99] B. Toën. *K-theory and cohomology of algebraic stacks: Riemann-Roch theorems, D-modules and GAGA theorems*. (1999). arXiv:math/9908097.
- [Vee86] W. Veech. *The Teichmüller geodesic flow*. *Ann. Math. (2)* 124 (1986), pp. 441–530.
- [Wol90] S. Wolpert. *The hyperbolic metric and the geometry of the universal curve*. *J. Differential Geom.* 31.2 (1990), pp. 417–472.
- [Sta18] The Stacks project authors. *The Stacks project*. 2018.

DEPARTMENT OF MATHEMATICS, INDIANA UNIVERSITY, BLOOMINGTON, IN 47405, USA
Email address: `mabainbr@indiana.edu`

DEPARTMENT OF MATHEMATICS, BOSTON COLLEGE, CHESTNUT HILL, MA 02467, USA
Email address: `dawei.chen@bc.edu`

CENTRO DE CIENCIAS MATEMÁTICAS-UNAM, ANTIGUA CAR. A PÁTZCUARO 8701, COL. EX HACIENDA SAN JOSÉ DE LA HUERTA, MORELIA, MICH., MÉXICO
Email address: `gendron@matmor.unam.mx`

MATHEMATICS DEPARTMENT, STONY BROOK UNIVERSITY, STONY BROOK, NY 11794-3651, USA
Email address: `sam@math.stonybrook.edu`

INSTITUT FÜR MATHEMATIK, GOETHE-UNIVERSITÄT FRANKFURT, ROBERT-MAYER-STR. 6-8, 60325 FRANKFURT AM MAIN, GERMANY
Email address: `moeller@math.uni-frankfurt.de`