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 smooth 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 $\mathrm{GL}_2^+(\mathbb{R})$-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: March 1, 2022.

Research of the first author is supported in part by the Simons Foundation grant #713192.

Research of the second author was supported in part by the National Science Foundation grant DMS-20-01040, the CAREER grant DMS-13-50396, the Simons Foundation grant #635235, and a von Neumann Fellowship at the Institute for Advanced Study in Spring 2019.

Research of the third author was supported in part by the project CONACyT A1-S-9029 "Moduli de curvas y Curvatura en A_g" of Abel Castorena.

Research of the fourth author was supported in part by the National Science Foundation under the grant DMS-18-02116 and DMS-21-01631.

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

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 $\mu = (m_1, \ldots, m_n)$ 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 14.1 and Example 14.1), 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 $\ell\colon V(\Gamma) \to \mathbb{R}$, and we may assume that its image, called the set of levels $L^\bullet(\Gamma)$, is the set $\{0, -1, \ldots, -N\}$ 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 pointed curve $(X, \boldsymbol{z})$ consists of

(i) an enhanced level structure on the dual graph Γ of $(X, \boldsymbol{z})$,
(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. △

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 $\mathbf{MS}_\mu$ on the category of complex spaces, and the groupoid $\mathcal{MS}_\mu$ will be given in Section 11. They come with projectivized versions, denoted by $\mathbb{P}\mathbf{MS}_\mu$ 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) *The complex space underlying* $\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ *is a coarse moduli space for* $\mathbf{MS}_\mu$.
(5) *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. For the next theorem note that we may view an orbifold such as $\mathbb{P}\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ as a smooth stack glued from quotient 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 morphism of proper algebraic Deligne-Mumford stacks* $\mathbb{P}\Xi\overline{\mathcal{M}}_{g,n}(\mu) \to \mathbb{P}\mathcal{MS}_\mu$, *which is an isomorphism over the open substack* $\mathbb{P}\Omega\mathcal{M}_{g,n}(\mu)$.

The groupoid $\mathbb{P}\mathcal{MS}_\mu$ is thus a hybrid object, smooth non-trivial isomorphism groups (due to automorphisms) at some places, and with finite quotient singularities at other places. The description as orderly blowup $\mathbb{P}\mathcal{MS}_\mu$ in Theorem 1.6 makes this functorial viewpoint even more natural. In fact the map in Theorem 1.3 is an isomorphism over the the substack where the local groups K_Λ introduced in Section 6.4 are trivial. 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 [CMZ19] for details.

Other compactifications. We briefly mention the relation with other compactifications in the literature. The space constructed in [FP18] can have extra components, hence in general it is a reducible space that contains the IVC only 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 [HS21], 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 [CW20].

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 9.2) near the boundary. The first application in [CMZ19] shows that the area form is a good enough metric on the tautological bundle. This is required in [Sau18; CMS19; Sau20] for direct computations of Masur-Veech volumes, and in [CMSZ20] to justify the volume formula for the spin components.

A second application in [CMZ20] is the construction of an analogue of the Euler sequence for projective spaces 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 $\mathbb{P}\Omega\mathcal{M}_{g,n}(\mu)$. Previously the Kodaira dimension was known only for some series of special cases, see [FV14; Gen18; Bar18].

A third application is towards the classification of connected components of the strata of k-differentials in [CG20]. Crucially, the viewpoint in that paper relies on the smoothness of a k-differential version of $\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ developed later in [CMZ19]. In particular, the concept of prong matchings helps to construct certain multi-scale k-differentials in the boundary that can provide information for generalized spin and hyperelliptic structures after smoothing into the interior of the strata.

A fourth large circle of applications concerns the dynamics of the action of $\mathrm{GL}_2^+(\mathbb{R})$ 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 defined by linear equations in period coordinates. The classification of these orbit closures is an important goal towards which significant progress has been made recently, see e.g. constraints found by Eskin-Filip-Wright [EFW18] and the constructions of special orbit closures by Eskin-McMullen-Mukamel-Wright in [EMMW20]. Using our moduli space $\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ as a black box, in [CW20] the Mirzakhani-Wright formula for the tangent space to the boundary of an orbit closure in the WYSIWYG space is generalized to the case of multi-component surfaces. Using the details of the construction of $\Xi\overline{\mathcal{M}}_{g,n}(\mu)$, it is further shown by Benirschke [Ben20a] that the boundary of any orbit closure in $\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ is again given by linear equations in generalized period coordinates of the boundary. Besides $\mathrm{SL}_2(\mathbb{R})$-orbits, taking the closures of other moduli spaces in $\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ such as certain Hurwitz spaces (e.g., the double ramification loci) can provide nice boundary structures, see [BDG20; Ben20b].

Moreover, it is also important yet challenging to explore dynamical invariants associated to orbit closures, such as saddle connections and related counting problems. The space $\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ is also used as a key tool by Dozier [Doz20] to prove regularity for the $\mathrm{SL}_2(\mathbb{R})$-invariant measure of the set of translation surfaces with multiple short saddle connections in the strata; his results for example immediately imply the finiteness of volume of strata of k-differentials first proven by Nguyen [Ngu19].

As a further evidence of potential applications towards the study of orbit closures, we show that our compactification provides a natural bordification of $\Omega\mathcal{M}_{g,n}(\mu)$ on which the action of $\mathrm{GL}_2^+(\mathbb{R})$ extends.

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 $\Xi\widehat{\mathcal{M}}_{g,n}(\mu) \to \Xi\overline{\mathcal{M}}_{g,n}(\mu)$ 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 $\mathrm{GL}_2^+(\mathbb{R})$ 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 12.

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 (6.12) 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 $\mathrm{Tw}_\Gamma^s \subset \mathrm{Tw}_\Gamma$ 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 $K_\Gamma = \mathrm{Tw}_\Gamma/\mathrm{Tw}_\Gamma^s$ 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 the 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 [Ber74a] 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 $g_n\colon X \to X_n$ that are conformal on compact subsets that exhaust the complement of nodes and punctures. In the *quasiconformal topology* one relaxes form conformal to quasiconformal, but requires that the quasiconformal dilatation tends to zero. 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 conformal and quasiconformal topologies 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* $\Omega\overline{\mathcal{T}}_{(\Sigma,\boldsymbol{s})}(\mu)$ 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 $\Omega\overline{\mathcal{T}}_{(\Sigma,\boldsymbol{s})}(\mu)$ with a topology that makes it a Hausdorff topological space in Theorem 7.7. For each multicurve Λ, the subspace of $\Omega\overline{\mathcal{T}}_{(\Sigma,\boldsymbol{s})}(\mu)$ 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 10. As a topological space, $\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ is the quotient of the augmented Teichmüller space $\Omega\overline{\mathcal{T}}_{(\Sigma,\boldsymbol{s})}(\mu)$ 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 asymptotically turning-number-preserving maps. In order to provide the augmented Teichmüller space $\Omega\overline{\mathcal{T}}_{(\Sigma,\boldsymbol{s})}(\mu)$ 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 sequence of diffeomorphisms g_n exhibiting convergence to preserve turning numbers in the limit $n \to \infty$. This needs to be done with care, to consider only a finite collection of good arcs; the 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 $\widehat{B} \to B$ and define markings using the pullback of the family to $\widehat{B}$, see Section 12. 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 $\mathcal{M}D_\Lambda$ 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 $\overline{\mathcal{M}D}_\Lambda$ of $\mathcal{M}D_\Lambda$ roughly by allowing the rescalings to attain the value zero. The actual definition given in Section 8 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, $\overline{\mathcal{M}D}_\Lambda$ 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 9.

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 (9.7) for the precise amount of perturbation.

Finally, in Section 10 we complete this setup and define the plumbing map in full generality and prove in Theorem 10.1 that plumbing is a local diffeomorphism. This is used in Theorem 10.2 to show that $\Xi\mathcal{D}_\Lambda$ is a complex orbifold.

Families and the universal properties. The functorial viewpoint and the proof of Theorem 1.2 (4) rely on showing first in Theorem 13.1 that the Dehn space $\Xi\mathcal{D}_\Lambda$ is a fine moduli space for a functor of marked multi-scale differentials. In order to provide families of multi-scale differentials with a Teichmüller marking, we need to deal with the problem that the equivalence relation in Definition 1.1 will twist the marking around the vertical nodes. To counterbalance this, the marking will not be defined on the original family, but on the family of welded surfaces over a real oriented blowup of the base, where the blowup structure is triggered by the level structure. The existence of the appropriate *level-wise real blowup* is proven in Theorem 12.2.

The proof of the universal property of $\Xi\mathcal{D}_\Lambda$ uses the (obvious) universal property of the model domain. To make use of it, we introduce an unplumbing construction that is roughly an inverse of the plumbing construction. This unplumbing construction is based on the normal form Theorem 4.3 for families of differentials on plumbing fixtures to pinch off the node and get an equisingular family, after subtracting differentials that play the inverse role of the modification differentials above. Finally, Theorem 14.10 completes the program by taking a family of multi-scale differentials, providing it locally with a marking, and gluing the moduli maps that universal property of $\Xi\mathcal{D}_\Lambda$ gives.

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 14.1 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 11.1. It ultimately reflects the normality of the toroidal compactifications by Δ^N/K_Γ used above. This procedure, culminating in Theorem 14.8, is summarized as follows.

Theorem 1.6. *The moduli stack of projectivized multi-scale differentials* $\mathbb{P}\mathcal{MS}_\mu$ *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 Sebastian Casalaina-Martin, Qile Chen, Matteo Costantini, David Holmes, John Smillie, Jakob Stix, Michael Temkin, Ilya Tyomkin, Scott Wolpert, Alex Wright, Jonathan Zachhuber, and Anton Zorich for inspiring discussions on differentials, log structures, moduli, algebraicity, and stacks.

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 $m_1 \geq \cdots \geq m_r > m_{r+1} = \cdots = m_{r+s} = 0 > m_{r+s+1} \geq \cdots \geq m_n$. Note that $m_j = 0$ is allowed, representing an ordinary marked point. We use the abbreviation $\overline{n} = \{1, \dots, n\}$.

A *(pointed) flat surface* or equivalently a *(pointed) abelian differential* is a triple $(X, \boldsymbol{z}, \omega)$, where X is a (smooth and connected) compact genus g Riemann surface, ω is a non-zero meromorphic one-form on X, and $\boldsymbol{z} \colon \overline{n} \hookrightarrow X$ is an injective function such that $\operatorname{ord}_{\boldsymbol{z}(j)} \omega = m_j$ for each j, and moreover every zero or pole of ω is marked by some $\boldsymbol{z}(j)$. We also denote by z_j the marked point $\boldsymbol{z}(j)$.

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

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

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

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

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 $\Omega\mathcal{M}_{g,n}(\mu) \subseteq K\overline{\mathcal{M}}_{g,n}(\widetilde{\mu})$, 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 $\mathbb{P}\Omega\overline{\mathcal{M}}^{\mathrm{inc}}_{g,n}(\mu)$ inside $\mathbb{P}K\overline{\mathcal{M}}_{g,n}(\widetilde{\mu})$ of the (projectivized) moduli space of abelian differentials of type μ. A point $(X, \omega, z_1, \dots, z_n) \in \mathbb{P}\Omega\overline{\mathcal{M}}_{g,n}(\widetilde{\mu})$ 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 labeling by the** $\mathrm{Sym}(\mu)$**-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 labeled. We let $\mathrm{Sym}(\mu) \subset S_n$ 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 $\mathrm{Sym}(\mu)$ acts everywhere and in particular on $\Xi\overline{\mathcal{M}}_{g,n}(\mu)$. The quotient $\Xi\overline{\mathcal{M}}_{g,n}(\mu)/\mathrm{Sym}(\mu)$ 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, which is allowed to have loops as well as half edges, and connected unless explicitly stated otherwise. We denote the set of vertices by $V(\Gamma)$, the set of edges by $E(\Gamma)$, and the set of half-edges by $H(\Gamma)$. . We denote by $\mathrm{val}(v)$ the *valence* of a vertex $v \in V$, the number of ordinary edges incident to 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* $\ell\colon V(\Gamma) \to \mathbb{Z}_{\leq 0}$ 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*

$$\ell\colon \Gamma \to \underline{N} \;=\; \{0, -1, \ldots, -N\}\,, \tag{2.2}$$

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 $n_v = \mathrm{val}(v) + |\mu_v|$ 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 $(X, \boldsymbol{z})$ 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 $(X, \boldsymbol{z})$ 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 $\operatorname{ord}_{z_j} \eta_v = m_j$.

(1) **(Matching orders)** For any node of X that identifies $q_1 \in X_{v_1}$ with $q_2 \in X_{v_2}$,
$$\operatorname{ord}_{q_1} \eta_{v_1} + \operatorname{ord}_{q_2} \eta_{v_2} \;=\; -2\,.$$

(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 $\operatorname{ord}_{q_1} \eta_{v_1} = \operatorname{ord}_{q_2} \eta_{v_2} = -1$ holds, then $\operatorname{Res}_{q_1} \eta_{v_1} + \operatorname{Res}_{q_2} \eta_{v_2} = 0$.

Let $\overline{\Gamma} = (\Gamma_X, \succcurlyeq)$ 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* $\overline{\Gamma}$ 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 $v_1 \succcurlyeq v_2$ if and only if $\operatorname{ord}_{q_1} \eta_{v_1} \geq -1$. Moreover, $v_1 \asymp v_2$ if and only if $\operatorname{ord}_{q_1} \eta_{v_1} = -1$.

We remark that this condition only uses the partial order induced by $\overline{\Gamma}$ 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 $q_1, \ldots, q_b$ denote the set of all nodes where Y intersects $X_{(i)}$. Then
$$\sum_{j=1}^{b} \operatorname{Res}_{q_j^-} \eta \;=\; 0\,,$$
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 $\preccurlyeq$ by $(X, \boldsymbol{z}, \eta, \preccurlyeq)$. 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* $(X, \omega, \boldsymbol{z})$ *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* $\succcurlyeq$ *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* $\overline{\Gamma} = (\Gamma_X, \succcurlyeq)$*.*
(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 $\overline{\Gamma}$ together with a numbering of the half-edges by $\overline{n}$ and with an assignment of a positive number $\kappa_e \in \mathbb{N}$ for each vertical edge $e \in E(\overline{\Gamma})^v$. The *degree* of a vertex v in Γ^+ is defined to be

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

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 $\{e \in E(\overline{\Gamma}) : e^+ = v\}$. 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 [CMSZ20]. The main example is the enhanced level graph Γ_X^+ of a twisted differential $(X, \boldsymbol{z}, \eta)$, obtained by assigning to each vertical node q the weight

$$\kappa_q = \mathrm{ord}_{q^+} \eta + 1\,. \tag{2.3}$$

In these terms, the above stability condition is equivalent to stability of $(X, \boldsymbol{z})$. 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 $\overline{\Gamma}$ 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 $\overline{\Gamma}$ 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.

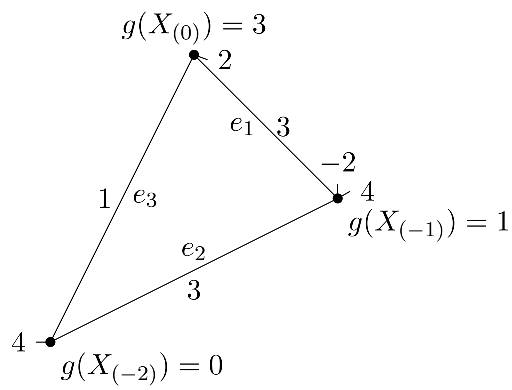

FIGURE 1. Two different enhanced orders Γ_1 and Γ_2 on Γ.

We denote the twisted differentials compatible with the level graphs Γ_i by $(X, \boldsymbol{z}, \eta^i)$. The enhanced structure Γ_1 tells us that the differential $\eta^1_{(-2)}$ is in $\Omega\mathcal{M}_{0,3}(4, -2, -4)$, the differential $\eta^1_{(-1)}$ is in $\Omega\mathcal{M}_{1,4}(4, 2, -2, -4)$ and $\eta^1_{(0)}$ is in $\Omega\mathcal{M}_{3,3}(2, 2, 0)$. Similarly $\eta^2_{(-2)}$ is in $\Omega\mathcal{M}_{0,3}(4, -2, -4)$, $\eta^2_{(-1)}$ is in $\Omega\mathcal{M}_{1,4}(4, 0, -2, -2)$ and $\eta^2_{(0)}$ is in $\Omega\mathcal{M}_{3,3}(2, 2, 0)$. The global residue condition in both cases says that the differential $\eta^i_{(-1)}$ has no residue at its pole at the point $q^-_{e_1}$. Note that it follows from [GT21, Theorem 1.1] 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 $(\Sigma, \boldsymbol{s})$ of genus g. We regard $\boldsymbol{s}$ as a set of $n \geq 0$ distinct labeled points $\{s_1, \ldots, s_n\} \subset \Sigma$, or alternatively as an injective function $\boldsymbol{s} \colon \overline{n} \hookrightarrow \Sigma$. Let $\mathcal{T}_{g,n} = \mathcal{T}_{(\Sigma,\boldsymbol{s})}$ be the Teichmüller space of $(\Sigma, \boldsymbol{s})$. Next, recall that a *multicurve* Λ on $\Sigma \setminus \boldsymbol{s}$ is a collection of disjoint simple closed curves, such that no two curves are isotopic on $\Sigma \setminus \boldsymbol{s}$, 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* $(X, \boldsymbol{z}, f)$ is a pointed stable curve $(X, \boldsymbol{z})$ together with a marking $f \colon (\Sigma, \boldsymbol{s}) \to (X, \boldsymbol{z})$, where a *marking* of a pointed stable curve is a continuous map $f \colon \Sigma \to X$ such that

(i) the inverse image of every node $q \in X$ is a simple closed curve on $\Sigma \setminus \boldsymbol{s}$,
(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 $\Sigma \setminus \Lambda \to X \setminus N_X$,
(iii) the map f preserves the marked points, that is, $f \circ \boldsymbol{s} = \boldsymbol{z}$.

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

The *augmented Teichmüller space* $\overline{\mathcal{T}}_{g,n} = \overline{\mathcal{T}}_{(\Sigma,\boldsymbol{s})}$ is the set of all equivalence classes of pointed stable curves marked by $(\Sigma, \boldsymbol{s})$. 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* $\mathrm{Mod}_{g,n}$ 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 $\Lambda \subset \Sigma \setminus \boldsymbol{s}$, 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 $\mathcal{T}_\Lambda$ is itself a finite unramified cover of the product of the Teichmüller spaces of the components of $(\Sigma, \boldsymbol{s}) \setminus \Lambda$ that takes into account the identification of the branches of the nodes. In particular each $\mathcal{T}_\Lambda$ 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 $X \setminus \boldsymbol{z}$, denoted by $(X, \boldsymbol{z})_{\epsilon_m}$, form an exhaustion of $X \setminus \boldsymbol{z}$. 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 $(X, \boldsymbol{z}, f)$ be a marked pointed stable curve in $\overline{\mathcal{T}}_{(\Sigma,\boldsymbol{s})}$, 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 $(X_m, \boldsymbol{z}_m, f_m)$ in $\overline{\mathcal{T}}_{g,n}$ *converges quasiconformally* to $(X, \boldsymbol{z}, f)$ if for some exhaustion $\{K_m\}$ of X^s, there exists a sequence of quasiconformal maps $g_m \colon K_m \to X_m$ 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. $g_m \circ \boldsymbol{z} = \boldsymbol{z}_m$), and the quasiconformal dilatations $\|\bar\partial g_m / \partial g_m\|_\infty$ tend to 0 as $m \to \infty$. The sequence *converges conformally* if the K_m are instead an exhaustion of $X^s \setminus \boldsymbol{z}$ and the g_m can be taken to be conformal. We call the topologies on $\overline{\mathcal{T}}_{(\Sigma,\boldsymbol{s})}$ 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 $(X_m, \boldsymbol{z}_m, f_m)$ converges *strongly conformally* to $(X, \boldsymbol{z}, f)$ 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. $g_m \circ \boldsymbol{z} = \boldsymbol{z}_m$).

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 $(X, \boldsymbol{z})$ 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 $f_\nu \colon X \to X$, preserving the marked points, such that the Beltrami differential of f_ν restricted to $X \setminus K$ agrees with ν, and $f(K) \subset U$.*

Moreover, the collection of such maps f_ν may be regarded as a holomorphic map $\mathcal{M}^k(X \setminus K) \to \mathrm{QC}^0(X)$, 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

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

Consider the derivative operator defined as

$$D = D_1\Phi_{(0,0)} \colon \mathcal{M}(K) \to T_{(X,\boldsymbol{z})}\mathcal{T}_{g,n}$$

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 $D^* \colon T^*_{(X,\boldsymbol{z})}\mathcal{T}_{g,n} \to \mathcal{M}(K)^*$

is injective. Under the usual identification of the cotangent space to Teichmüller space at $(X, \boldsymbol{z})$ with $Q(X, \boldsymbol{z})$, the space of quadratic differentials q on X with at worst simple poles contained in $\boldsymbol{z}$, 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 $T_{(X,\boldsymbol{z})}\mathcal{T}_{g,n}$ 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 $\psi\colon \mathcal{M}^k(X \setminus K) \to \mathcal{M}^1(K)$ such that $\Phi(\psi(\nu), \nu) = (X, \boldsymbol{z})$. In other words for each $\nu \in \mathcal{M}^k(X \setminus K)$, there is a quasiconformal map $f_\nu\colon X \to X$ with Beltrami differential given by $\psi(\nu) + \nu$.

The map $\nu \mapsto f_\nu$ can be regarded as a map $\Psi\colon \mathcal{M}^k(X \setminus K) \to \mathrm{QC}^0(X)$. 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 $f_\nu(K) \subset U$. ∎

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 $(X_m, \boldsymbol{z}_m, f_m)$ converges to $(X, \boldsymbol{z}, f)$ in the quasiconformal topology, so that there is an exhaustion of X^s by compact sets K_m and quasiconformal maps $g_m\colon K_m \to X_m$ 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 $h_m\colon X \setminus U \to X_m$ 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 $k_m\colon X \to X$ 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 $X^s \setminus \boldsymbol{z}$ so that the inclusion $K_m \hookrightarrow X \setminus \boldsymbol{z}$ is a homotopy equivalence, and let $g_m\colon K_m \to X_m$ 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 $\mathbb{H}/\Gamma$ 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 $\mathbb{H}/\Gamma_m$ for some Fuchsian group Γ_m. Let $\widetilde{J}_m^f \subset \mathbb{H}$ and $\widetilde{L}_m^f \subset \mathbb{H}$ be the unique connected subsurfaces, invariant under the Fuchsian groups, with $\widetilde{J}_m^f/\Gamma = J_m^f$ and $\widetilde{L}_m^f/\Gamma_m = L_m^f$. The conformal map g_m then lifts to a conformal map $\widetilde{g}_m\colon \widetilde{J}_m^f \to \widetilde{L}_m^f$ which is equivariant in the sense that

$$\rho_m(\gamma)\cdot \widetilde{g}_m(z) \;=\; \widetilde{g}_m(\gamma\cdot z) \tag{3.1}$$

for some isomorphism $\rho_m\colon \Gamma \to \Gamma_m$ 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 $\widetilde{g}_m$ to this limit set fixes these three points.

We claim now that $\widetilde{g}_m$ 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 $\widetilde{g}_m$ has a further subsequence which converges uniformly on compact sets to some $G\colon \mathbb{H} \to \overline{\mathbb{H}}$. Since each $\widetilde{g}_m$ is conformal and fixes three points on the boundary of $\mathbb{H}$, in fact G must be the identity map. Since every subsequence of $\widetilde{g}_m$ 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 $p\colon \Delta \to \widetilde{J}_m^f$ 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 $\widetilde{J}_m$. The composition $\widetilde{g}_m \circ p$ sends the boundary circle to a smooth curve $\gamma_m \subset \widetilde{L}_m^f$ 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 $p_m\colon \Delta \to \widetilde{L}_m^f$ be the Riemann mapping of Δ onto U_m which is normalized so that $p_m(w) = \widetilde{g}_m \circ p(w)$ for the points $w = a_1, a_2$, and 0. Since $\widetilde{g}_m$ 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 $\alpha_m\colon \Delta \to \Delta$ be the Douady-Earle extension of $p_m^{-1} \circ \widetilde{g}_m \circ p|_{\partial\Delta}$. The boundary map is uniformly close to the identity, so $\alpha_m(0)$ is close to 0, and we may construct a quasiconformal map $\beta_m\colon \Delta \to \Delta$ 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 $z_{1,m}$. ∎

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 $(X, \boldsymbol{z}, p, f)$ such that p is not a node, we say that a sequence $(X_m, \boldsymbol{z}_m, p_m, f_m)$ converges to $(X, \boldsymbol{z}, p, f)$ if $(X_m, \boldsymbol{z}_m, f_m) \to (X, \boldsymbol{z}, f)$ as pointed stable curves; and moreover, if $g_m \colon K_m \to X_m$ 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 $\mathcal{D}_\Lambda$, 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 $\Lambda \subset \Sigma \setminus \boldsymbol{s}$, the *full* Λ*-twist group* $\mathrm{Tw}_\Lambda^{\mathrm{full}} \subset \mathrm{Mod}_{g,n}$ is the free abelian subgroup generated by Dehn twists around the curves of Λ. The *Dehn space* $\mathcal{D}_\Lambda$ 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 $\mathrm{Tw}_\Lambda^{\mathrm{full}}$. That is,

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

(Bers [Ber74a] called this space the "deformation space".) Each $\mathcal{D}_\Lambda$ 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 $\pi \colon \mathcal{X}_\Lambda \to \mathcal{D}_\Lambda$ is the quotient

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

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 $\mathcal{D}_\Lambda$, 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 $\overline{\mathcal{T}}_{g,n}/\operatorname{Mod}_{g,n}$. For each multicurve Λ, the natural map $\mathcal{D}_\Lambda \to \overline{\mathcal{M}}_{g,n}$ is a local homeomorphism. The image of $\mathcal{D}_\Lambda$ 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 $\overline{\mathcal{M}}_{g,n}^{\mathrm{alg}}$ (see [DM69] or [ACG11]). Hubbard-Koch [HK14] showed that $\overline{\mathcal{M}}_{g,n}^{\mathrm{alg}} \cong \overline{\mathcal{M}}_{g,n}$ as complex orbifolds, so the natural topology of the algebraic variety $\overline{\mathcal{M}}_{g,n}^{\mathrm{alg}}$ 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 sometimes more convenient than the quasiconformal topology, as pullbacks of holomorphic one-forms by quasiconformal maps are in general only locally L^2 as quasiconformal maps have locally L^2 derivatives, but we will use these topologies both. 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 $\pi\colon \mathcal{X}_\Lambda \to \mathcal{D}_\Lambda$, 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* $\Omega\mathcal{D}_\Lambda \to \mathcal{D}_\Lambda$, 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 $\mathcal{D}_\Lambda$ gives a second natural topology on $\Omega\mathcal{D}_\Lambda$. A sequence $(X_m, \boldsymbol{z}_m, \omega_m, f_m)$ of marked pointed stable forms converges to $(X, \boldsymbol{z}, \omega, f)$ *in the conformal topology* if for some exhaustion K_m of $X^s \setminus \boldsymbol{z}$, there is a sequence of conformal maps $g_m\colon K_m \to X_m$ 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 will occasionally want to allow the maps f_m to be only quasiconformal. We say a sequence $(X_m, \boldsymbol{z}_m, \omega_m, f_m)$ converges to $(X, \boldsymbol{z}, \omega, f)$ *in the quasiconformal topology* if for some exhaustion K_m of X^s, there is a sequence of L_m-conformal maps $g_m\colon K_m \to X_m$ which exhibit convergence in the quasiconformal topology on $\mathcal{D}_\Lambda$ and such that $g_m^*\omega_m$ converges to ω in the topology of weak locally L^2 convergence.

We show below that these topologies agree.

Lemma 3.4. *Suppose $3g-3+n > 0$. Let $(X_m, \boldsymbol{z}_m, f_m)$ be a sequence in $\mathcal{D}_\Lambda$ converging to $(X, \boldsymbol{z}, f)$ in the quasiconformal topology, and let $g_m\colon K_m \to X_m$ be a sequence of L_m-quasiconformal maps exhibiting this convergence, where K_m is an exhaustion of X. Then the maps g_m (regarded as maps into the universal curve $\mathcal{X}_\Lambda$) 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 $K = (X, \boldsymbol{z})_{\epsilon_0}$, and convergence on compact sets follows from the usual diagonal trick. We fix also $K' = (X, \boldsymbol{z})_{\epsilon_1}$ for some $\epsilon_1 < \epsilon_0$ and assume m is large enough that g_m is defined on K'.

Choose a Riemannian metric ρ' on $\mathcal{X}_\Lambda^s$, the complement of the nodes and marked points in $\mathcal{X}_\Lambda$, whose restriction to the fibers is the vertical hyperbolic metric ρ. The functions g_m have a uniform modulus of continuity on K by [LV73, § 3.3.3] and so are uniformly equicontinuous.

To apply Arzelá-Ascoli, we just need that the g_m are contained in a compact subset of $\mathcal{X}^s_\Lambda$. By Mumford's compactness criterion, the ϵ-thick part in the vertical hyperbolic metric $(\mathcal{X}^s_\Lambda)_\epsilon$ is compact, so it suffices to show that g_m maps K into $(X_m, \boldsymbol{z}_m)_\epsilon$ for some uniform ϵ. Again, since the g_m have a uniform modulus of continuity, there are small constants $\epsilon < \epsilon'$ so that if $g_m(K)$ intersects the ϵ-thin part, it must be contained in the ϵ'-thin part 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 $X \setminus \boldsymbol{z}$ is finite type and hyperbolic, so $g_m(K)$ cannot be contained in the ϵ'-thin part. It follows immediately that $g_m(K)$ is contained in $(X_m, \boldsymbol{z}_m)_\epsilon$ for some uniform ϵ.

We now 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 L_m-quasiconformal, with $L_m \to 1$, any subsequential limit is a conformal automorphism of $X \setminus \boldsymbol{z}$. As the g_m are compatible with the markings, this map is homotopic to the identity, so must in fact be the identity, since $X \setminus \boldsymbol{z}$ 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. ◻

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

Proof. On the base surface $(\Sigma, \boldsymbol{s})$, choose g disjoint homologically independent "α-curves" $\alpha_1, \ldots, \alpha_g$, 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 $\mathrm{Tw}^{\mathrm{full}}_\Lambda$, they are dual to a basis of relative forms. Let α_i^* be smooth, compactly supported 1-forms on $\Sigma \setminus \boldsymbol{s}$ which are dual to the α_i. Then there are relative one-forms $\eta_1, \ldots, \eta_g$ on the universal curve $\mathcal{X}_\Lambda$ over $\mathcal{D}_\Lambda$ such that in each fiber,

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

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

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

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 $g_m^*\omega_m \wedge \alpha_i^*$, this follows from the weak convergence of $g_m^*\omega_m$ to ω.

Now 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 $(X_m, \boldsymbol{z}_m, f_m)$ converge to $(X, \boldsymbol{z}, f)$ as marked pointed surfaces, and by Theorem 3.2 there is a sequence of maps $g_m \colon K_m \to X_m$, 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

$g_m^*\eta_i \to \eta_i|_X$ uniformly on compact sets, so $g_m^*\omega_m \to \omega$ as well, and so the sequence converges in the conformal topology.

Since conformal convergence obviously implies quasiconformal convergence, it follows that the three topologies are equivalent. ∎

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 $(X_m, \boldsymbol{z}_m, \omega_n)$ converges to $(X, \boldsymbol{z}, \omega)$ on an irreducible component $Y \subset X$ if there are conformal maps $g_m \colon K_m \to X_m$ 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.

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 3.6. *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 $\operatorname{ord}_z \nu_m$, $\operatorname{ord}_z \eta_m$ and $\operatorname{ord}_z \omega$ 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 $h_n \colon U \to W$ 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 $m \to \infty$.

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 $U \to X$, 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 $V \to E$, 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 3.7. *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 $g\colon U \to \mathbb{C}^3$ which fix S pointwise, let $\mathcal{H}$ be the locus of those functions sending U into Y, and let $\mathcal{B}_\epsilon$ 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 $F_1, F_2\colon \mathbb{C}^3 \to \mathbb{C}$ so that Y is defined by the equations $F_1 = F_2 = 0$ and the derivative $DF\colon \mathbb{C}^3 \to \mathbb{C}^2$ (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 $\Phi\colon \mathcal{O}_U(\mathbb{C}^3)_S \to \mathcal{O}_U(\mathbb{C}^2)_S$ by $\Phi(g) = F \circ (\mathrm{id} + g) - \mathrm{id}$, a holomorphic map with derivative $D\Phi_{\mathrm{id}}(g) = DF \cdot g$. The space $\mathcal{H}$ is then the fiber of Φ over 0. If we could show that $D\Phi_{\mathrm{id}}$ is a split surjection by constructing a right-inverse to DF, it would then follow that $\mathcal{H}_\epsilon$ 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 ijth entry is the entire function $\frac{\partial F_i}{\partial z_j}$. Let M_i be the 3×2 matrix obtained by replacing the ith row of DF^T by zeros, and let μ_i be the ith minor of DF, so that

$$DF \cdot M_i \;=\; \mu_i I. \tag{3.3}$$

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. ∎

Remark 3.8. 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 3.6. 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 $U \to Y$ which fix the subset S and are sufficiently close to the identity may be identified with a δ-ball $\mathcal{V}(U)^S_\delta$.

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 $m \to \infty$.

Let $\mathcal{B}_\epsilon$ 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 $\Psi\colon \mathcal{B}_\epsilon \to \mathcal{O}_U^0$ defined by

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

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 $\mathcal{B}_\epsilon$. We have a universal holomorphic map $\Phi\colon \mathcal{U} \to \mathcal{X}$ whose fiber over a point in $\mathcal{V}(U)^S$ is the map which that point represents. Similarly, there are universal bounded holomorphic functions $F, G\colon \mathcal{W} \to \mathbb{C}$ associated to the factors $\mathcal{O}_W$ and $\mathcal{O}_W^0$ of $\mathcal{B}_\epsilon$. The form ω can be regarded as a relative one-form Ω on $\mathcal{W}$. The function

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

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 3.9, 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 $D_1\Psi_{(id,0,0)}\colon \mathcal{V}(U)^S \to \mathcal{O}_U^0$ 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

$$\Upsilon\colon \mathcal{O}_U^0 \to \mathcal{V}(U)^S\,, \qquad \Upsilon(f) \;=\; \frac{1}{\omega}\int_{z_0} f\omega$$

and argue now that this is well-defined. The integral is over any path starting at z_o 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 $V\colon \mathcal{O}_U^0 \to \mathcal{V}(U)^S$ 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. ◻

To complete the proof of Theorem 3.6 it remains to verify the following universal property.

Lemma 3.9. *Let E and F be complex Banach spaces containing open sets U and V respectively, and let $f\colon U \times V \to \mathbb{C}$ a bounded holomorphic function. Then the map $F\colon U \to \mathcal{O}_V$ defined by $F(z)(w) = f(z,w)$ is a holomorphic function with*

$$DF_z(w) \;=\; D_1 f_{(z,w)}\,. \tag{3.4}$$

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

$$\|D_1^p f_{(z,w)}\| \;\le\; p!\,\frac{M}{R^p}\,, \tag{3.5}$$

where M is a uniform bound for $|f|$ on $U \times V$. Given $z \in U$, let $D_z\colon E \to \mathcal{O}_V$ be the bounded operator $D_z(h)(w) = D_1 f_{(z,w)}(h)$. 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

$$\begin{aligned} |F(z+h) - F(z) - D_z(h)| \;&=\; \sup_{w\in V} |f(z+h,w) - f(z,w) - D_1 f_{(z,w)}(h)| \\ &\le\; \frac{M}{(R-|h|)^2}|h|^2, \end{aligned}$$

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

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 7.4 to obtain compactness of the moduli space of multi-scale differentials. For an alternative approach to these compactness results, see the appendix to [McM89].

Given a pointed stable curve $(X, \boldsymbol{z})$ we denote the punctured surface $X^s \setminus \boldsymbol{z}$ 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 $(X_m, \boldsymbol{z}_m, \omega_m)$ in $\Omega\mathcal{M}_{g,n}(\mu)$ such that the underlying pointed curves converge to some pointed stable curve $(X, \boldsymbol{z})$. 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

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

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

Theorem 3.10. *Suppose $(X_m, \boldsymbol{z}_m, \omega_m)$ is a sequence of meromorphic differentials in $\Omega\mathcal{M}_{g,n}(\mu)$ such that $(X_m, \boldsymbol{z}_m)$ converges to some $(X, \boldsymbol{z})$ 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 $(X, \boldsymbol{z}, \omega)$ as the function $\daleth\colon X' \to \mathbb{R}$ defined by

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

This function measures how quickly the flat metric $|\omega|$ varies with respect to the hyperbolic metric ρ. Note that $\daleth$ 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 3.11. *There is a constant C depending only on μ and ϵ so that for any $(X, \boldsymbol{z}, \omega) \in \Omega\mathcal{M}_{g,n}(\mu)$, the distortion function $\daleth$ 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 $\daleth$ 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 $\pi\colon \widetilde{\mathcal{X}} \to \mathbb{P}\Omega\overline{\mathcal{M}}_{g,n}^{\text{ninc}}(\mu)$. 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 $D \subset \widetilde{\mathcal{X}}$ is an irreducible vertical component of the zero divisor of ω lying over D'. Since the base is normal, by Proposition 11.13 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 the pullback of a regular function which does not vanish at p. The family of one-forms

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

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 $\widetilde{\mathcal{X}}'$ by [Wol90]. Since we may extend β on a neighborhood of any such vertical zero divisor, this gives a continuous extension of $\widetilde{\beta}$ over all of $\mathcal{X}'$. The function $\widetilde{\daleth}(p) = |\widetilde{\beta}|_p$ is then the desired continuous extension of $\daleth$ to $\widetilde{\mathcal{X}}'$. Since the ϵ-thick part of $\widetilde{\mathcal{X}}$ is compact, we see that $\daleth$ is bounded on the ϵ-thick part. ∎

Corollary 3.12. *There exists a constant L that depends only on μ and ϵ, such that for any pointed meromorphic differential $(X, \boldsymbol{z}, \omega) \in \Omega\mathcal{M}_{g,n}(\mu)$, for any points p and q in the ϵ-thick part of X, we have*

$$|\omega|_q \;\leq\; L|\omega|_q \,.$$

Proof. By Lemma 3.11, $\log|\omega/\rho|$ 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 $L = Ce^E$. ∎

Proof of Theorem 3.10. Let $f_m\colon K_m \to X_m$ be conformal maps on an exhaustion $\{K_m\}$ of Y' that exhibit the convergence of the $(X_m, \boldsymbol{z}_m)$. By Corollary 3.12, 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. ∎

4. Normal forms for differentials

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, Section 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 $k = \operatorname{ord}_0 \omega$ and its residue $r = \operatorname{Res}_0 \omega$. 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 $k = \operatorname{ord}_0 \omega$ and $r = \operatorname{Res}_0 \omega$. Then for some $\epsilon > 0$, there exists a conformal map $\phi\colon (\Delta_\epsilon, 0) \to (\Delta_\delta, 0)$ such that*

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

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* $\widetilde{q} \in U$ *there exists a map* $\widetilde{\phi}$ *satisfying* (4.1) *and with* $\widetilde{\phi}(p) = \widetilde{q}$.

This statement also holds for families of differentials $\omega_{\boldsymbol{t}}$ on families of disks, as long as the order $\mathrm{ord}_0\,\omega_{\boldsymbol{t}} = k$ is the same for all $\boldsymbol{t}$. For families of differentials such that the order $\mathrm{ord}_0\,\omega_{\boldsymbol{t}}$ is not constant, the situation is more complicated. This has essentially been dealt with in [BCGGM18, Section 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 $\zeta_j := \boldsymbol{e}(j/(k+1))$ be a $(k+1)$-st root of unity (where we denote $\boldsymbol{e}(z) = \exp(2\pi\sqrt{-1}z)$).

Theorem 4.2 (**Normal form of a deformation on an annulus**). *Let* $\omega_{\boldsymbol{t}}$ *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* $\varsigma\colon U \to A_{\delta_1,\delta_2}$ *such that* $\varsigma(Y) = \zeta_j p$.

Then there exists a neighborhood $U_{\boldsymbol{0}} \subset U$ *of* Y, *together with* $\delta_1 < \epsilon_1 < \epsilon_2 < \delta_2$ *and a holomorphic map* $\phi\colon U_{\boldsymbol{0}} \times A_{\epsilon_1,\epsilon_2} \to A_{\delta_1,\delta_2}$ *such that* $\phi_{\boldsymbol{t}}^*(\omega_{\boldsymbol{t}})$ *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* $\phi_{\boldsymbol{t}}(p) = \varsigma(\boldsymbol{t})$ *for all* $\boldsymbol{t} \in U_{\boldsymbol{0}}$.

We now pass to families where the topology of the underlying Riemann surfaces changes. We establish below a normal form in a neighborhood of a node, which will be used in the unplumbing construction of Proposition 13.6. 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 $\widetilde{V} = \widetilde{V}_\delta = \Delta_\delta^2 \times B$, 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 (see [ACG11, Proposition X.2.2.1]). 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* $\phi\colon V(f,\epsilon) \to V(f,\delta)$, *which lifts the identity map of* B *to itself, such that*

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

Moreover, given a section $\varsigma_0 \colon B \to V$ *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 11.11. 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

$$\omega \;=\; \Big(\sum_{i\geq 0} c_i u^i + \sum_{i>0} c_{-i} v^i\Big)\frac{du}{u} \tag{4.3}$$

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

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

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

$$\eta \;=\; \Big(\sum_{i\geq 0} e_i v^{-i} + \sum_{i>0} e_{-i} v^i\Big)\frac{du}{u}\,. \tag{4.5}$$

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 [BCGGM18, 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. ∎

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 $\boldsymbol{b}$ 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. ◻

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

$$\phi_{(X,Y)}(u,v) \;=\; (ue^{X(u)+Y(v)}, ve^{-X(u)-Y(v)})\,, \tag{4.6}$$

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 ς 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

$$\omega_0 \;=\; u^{\kappa}(1 + r_0 v^{\kappa} + fg(u) + fh(v))\frac{du}{u}\,.$$

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 $\Delta_\epsilon \times B$,

$$\begin{aligned} \alpha^* u^{\kappa}(1 + g(u))\frac{du}{u} &\;=\; u^{\kappa}\frac{du}{u}\,, \quad \text{and} \\ \beta^* v^{-\kappa}(1 + r_0 f^{\kappa} + h(v))\frac{dv}{v} &\;=\; v^{-\kappa}(1 + r_0 f^{\kappa})\frac{dv}{v}\,, \end{aligned}$$

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 $\phi^*_{(X,Y)} u^{\kappa}(1 + r_0 v^{\kappa})\frac{du}{u} = \omega_0$, and $\phi_{(X,Y)} \circ \varsigma_0 = \varsigma$. Explicitly this means that on $\Delta^2_\epsilon \times B$, the functions X and Y satisfy the equations,

$$\begin{aligned} (e^{\kappa(X+Y)} + r_0 v^{\kappa})\left(1 + u\frac{\partial X}{\partial u} - v\frac{\partial Y}{\partial v}\right) - (1 + v^{\kappa} + fg + fh) + (uv - f)W &\;=\; 0\,, \\ \tau_0 e^{-X(f/\tau_0) - Y(\tau_0)} - \tau &\;=\; 0\,, \end{aligned}$$

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 $\widetilde{\rho}_s(u,v,b) = (u,v,sb)$, so that we have the commutative diagram:

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

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

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

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 $\widetilde{V} = \Delta_\epsilon^2 \times B$ 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

$$\Phi\colon \mathcal{O}(\widetilde{V})_0 \oplus \mathcal{O}(U_B)_1^{\mathrm{nc}} \oplus \mathcal{O}(V_B)_1^{\mathrm{nc}} \oplus \mathcal{O}(B)_0 \oplus \mathbb{C} \to \mathcal{O}(\widetilde{V})_0^{\mathrm{nc,nr}} \oplus \mathcal{O}(B)_0\,,$$

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

$$D_1\Phi\colon \mathcal{O}(\widetilde{V})_0 \oplus \mathcal{O}(U_B)_1^{\mathrm{nc}} \oplus \mathcal{O}(V_B)_1^{\mathrm{nc}} \to \mathcal{O}(\widetilde{V})_0^{\mathrm{nc,nr}} \oplus \mathcal{O}(B)_0$$

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)

$$\begin{aligned} D_1\Phi_1(W,X,Y) &= W\cdot uv + \left(\kappa X + u(1+r_0(0)v^\kappa)\frac{\partial X}{\partial u}\right) + \left(\kappa Y - v(1+r_0(0)v^\kappa)\frac{\partial Y}{\partial v}\right), \\ D_1\Phi_2(W,X,Y) &= -\tau_0 X(f/\tau_0) - \tau_0 Y(\tau_0)\,. \end{aligned}$$

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

$$S\colon \mathcal{O}(U_B)_0^{\mathrm{nc}} \oplus \mathcal{O}(V_B)_0^{\mathrm{nc,nr}} \oplus \mathcal{O}(\widetilde{V})_0 \oplus \mathcal{O}(B)_0 \to \mathcal{O}(\widetilde{V})_0 \oplus \mathcal{O}(U_B)_1^{\mathrm{nc}} \oplus \mathcal{O}(V_B)_1^{\mathrm{nc}}\,,$$

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

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

$$\kappa X + u\frac{\partial X}{\partial u} \;=\; \aleph\,, \tag{4.8}$$

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

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

$$\begin{aligned} S_X(\aleph) \;&=\; \frac{1}{u^\kappa}\int u^{\kappa-1}\aleph\,du\,,\\ S_Y(\beth) \;&=\; \frac{-v^\kappa}{1+r_0(0)v^\kappa}\int \frac{\beth}{v^{\kappa+1}}\,dv\,, \end{aligned}$$

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

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

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. ◻

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

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

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

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

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 $Z = 0 \mod f$. The $u^{\kappa+j}$-term involves a linear equation for $c_j \mod f$ 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 $b_1 \mod f$. The coefficient $Z \mod f^2$ is linearly determined by the u^κ-term mod f^2 and the $u^{\kappa+j}$-term mod f^2 linearly determine $c_j \mod f$. In the third round, considering terms mod f^3, we start with the $u^{\kappa-2}$-term, which determines $b_2 \mod f$, then consider the $u^{\kappa-1}$-term to determine $b_1 \mod f^2$. The u^κ-term and higher terms to compute $Z \mod f^3$ and then the $c_j \mod f^3$. 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 $\phi\colon V(f,\epsilon) \to V(f,\delta)$ such that it is the identity on B and such that

$$\phi^*\omega \;=\; r\frac{du}{u}\,. \tag{4.13}$$

Moreover, given a section $\varsigma_0\colon B \to V$ 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 $\Omega\mathcal{T}^{pm}_\Lambda(\mu)$ 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 $\Omega\mathcal{T}^{pm}_{\Lambda}(\mu)$. 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 in Section 7, as well as modular interpretations that we address in Section 11.

5.1. **Ordered and enhanced multicurves and their degenerations.** We continue to fix an n-pointed topological surface $(\Sigma, \boldsymbol{s})$, as in Section 3. To every multicurve $\Lambda \subset \Sigma \setminus \boldsymbol{s}$, we can associate the dual graph $\Gamma(\Lambda)$ whose vertices correspond to connected components of $\Sigma \setminus \Lambda$, whose edges correspond to curves in Λ, and whose half-edges correspond to the marked points. In the setting of multicurves, we will generally imitate the standard notation for level graphs from Section 2.3. We call $\overline{\Lambda} = (\Lambda, \ell)$ an *ordered multicurve* and specify the ordering relation between the components of $\Sigma \setminus \Lambda$ by $\preccurlyeq$. 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^c_v and $\Sigma^c_{(i)}$ 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 pointed 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 $\mathrm{dg}\colon (\Lambda_2, \ell_2) \rightsquigarrow (\Lambda_1, \ell_1)$, if the following conditions hold:

- As a set of isotopy classes of curves $\Lambda_2 \subset \Lambda_1$. Let then $\delta\colon \Gamma(\Lambda_1) \to \Gamma(\Lambda_2)$ 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 $\Lambda_1 \setminus \Lambda_2$.
- The map δ is compatible with the orders ℓ_j in the sense that if $v_1 \preccurlyeq v_2$ then $\delta(v_1) \preccurlyeq \delta(v_2)$. It follows that if $v_1 \asymp v_2$ then $\delta(v_1) \asymp \delta(v_2)$, so δ induces a surjective, order non-decreasing map, still denoted by δ, on the (normalized) sets of levels $\delta\colon L^\bullet(\Lambda_1) \twoheadrightarrow L^\bullet(\Lambda_2)$.
- The map δ respects the labeling of the half-edges.

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. △

As these constraints are all phrased only in terms of the dual graphs of these multicurves, there is an analogous notion of a degeneration of enhanced level graphs.

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

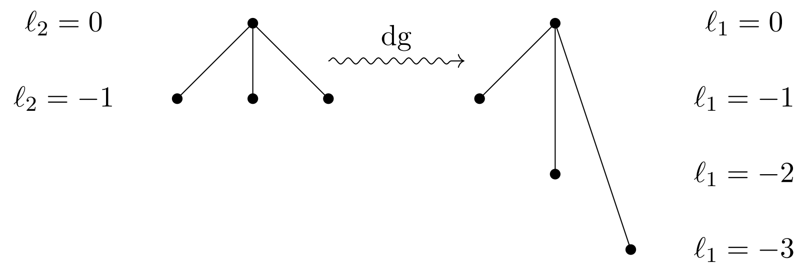

FIGURE 2. A degeneration that does not change the underlying multicurve.

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 $\delta = \mathrm{id}$. 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 $\delta\colon \underline{N} \twoheadrightarrow \underline{M}$ 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 $J = \{0 > j_{-1} > \cdots > j_{-M} \geq -N\}$. We define $j_0 = 0$ and $j_{-M-1} = -N-1$ (though they are not part of J). The subset J induces a map $\delta_J\colon \underline{N} \to \underline{M}$ which maps integers (i.e. levels) in each interval $(j_{k-1}, j_k]$ to k. We denote the associated degeneration by $\mathrm{dg}_J\colon \Lambda_J \rightsquigarrow \Lambda$. The two-level degenerations given by $J = \{i\}$, and denoted by dg_i, will be particularly useful (see Section 6.2). In the example of Figure 2, we have $J = \{-1\}$. The level $(j_{-1}, j_0] = (-1, 0]$ is mapped to 0 and the levels $(j_{-2}, j_{-1}] = (-4, -1]$ are mapped to -1.

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

$$\mathrm{Per}\colon \Omega\mathcal{T}_{(\Sigma,\boldsymbol{s})}(\mu) \to H^1(\Sigma \setminus P_{\boldsymbol{s}}, Z_{\boldsymbol{s}}; \mathbb{C}),$$

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

The (classical) mapping class group $\mathrm{Mod}_{(\Sigma,\boldsymbol{s})}$ of $(\Sigma, \boldsymbol{s})$ acts properly discontinuously on $\mathcal{T}_{(\Sigma,\boldsymbol{s})}$ and on the twisted Hodge bundle over it, preserving the submanifold $\Omega\mathcal{T}_{(\Sigma,\boldsymbol{s})}(\mu)$. The spaces $\Omega\mathcal{T}_{(\Sigma,\boldsymbol{s})}(\mu)$ are highly disconnected and we do not address here the question of classifying the connected components. Moreover, we do not claim that $\mathbb{P}\Omega\mathcal{T}_{(\Sigma,\boldsymbol{s})}(\mu)$ 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 $(X, f, \boldsymbol{z}, \eta)$ where $(X, f, \boldsymbol{z})$ 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 $e \in \Gamma(\Lambda)$. △

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 $\boldsymbol{z}$ 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 $(X, f, \boldsymbol{z}, \eta)$ where η is a twisted differential compatible with $\Gamma(\Lambda)$. △

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 $(X_n, f_n, \boldsymbol{z}_n, \eta_n)$ of marked pointed twisted differentials converges to $(X, f, \boldsymbol{z}, \eta)$ if for some exhaustion K_n of X, there is a sequence of conformal maps $g_n\colon K_n \to X_n$ such that $f_n \simeq g_n \circ f$ and $g_n^*\eta_n$ converges to η uniformly on compact sets.

5.3. **Welded surfaces.** Teichmüller markings of nodal surfaces are by definition insensitive to the precomposition 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 $\pi\colon X^* \to X$ 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 $\sigma_q\colon T_xX^* \to T_yX^*$, modulo scaling by positive real numbers. We alternatively think of the welding as an orientation-reversing metric isomorphism $\sigma_q\colon S_xX^* \to S_yX^*$, where $S_pX^* = (T_pX^* \setminus \{0\})/\mathbb{R}_{>0}$ 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 12. The ordering of the fiber over q is not part of the structure, and we consider $\sigma_q^{-1}\colon T_yX^* \to T_xX^*$ 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 12. Given the unit disk $\Delta \subset \mathbb{C}$, the real oriented blowup $p\colon \operatorname{Bl}_0 \Delta \to \Delta$ is the locus

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

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* $p\colon \operatorname{Bl}_D X \to X$, 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 $\operatorname{int}(\operatorname{Bl}_D X) \to X \setminus D$, and for each $q \in D$ the boundary circle $\partial_q \operatorname{Bl}_D X = p^{-1}(q)$ is naturally identified with the real tangent circle $S_qX = (T_qX \setminus \{0\})/\mathbb{R}_{>0}$ of X at q. The conformal structure of X gives $\partial_q \operatorname{Bl}_D X$ 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* $p\colon \operatorname{Bl}_D X \to X$ 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 $S_q^+ \cup S_q^- \subset \partial \operatorname{Bl}_D X$. In these terms, a welding of X at D is a choice for each node $q \in D$ of an orientation-reversing isometry $\sigma_q\colon S_q^+ \to S_q^-$.

A *global welding $\boldsymbol{\sigma}$ of X* is a choice of a welding at each node of X. If the dual graph is endowed with a level structure $\overline{\Gamma}$, then a *vertical welding $\boldsymbol{\sigma}$ of X* is a choice of a welding at each vertical node (horizontal nodes will never be welded in this paper).

Given a vertical welding $\boldsymbol{\sigma}$ of X, we define the *associated welded surface* $\overline{X}_{\boldsymbol{\sigma}}$ to be the surface obtained by gluing the boundary components of $\operatorname{Bl}_{N_X^v} X$ via $\boldsymbol{\sigma}$. The associated welded surface has the following extra structures:

(1) a multicurve Λ^v on $\overline{X}_{\boldsymbol{\sigma}}$, 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 $\overline{X}_{\boldsymbol{\sigma}}$;
(2) a conformal structure on $\overline{X}_{\boldsymbol{\sigma}} \setminus \big(\Lambda^v \cup N_X^h\big)$; and
(3) 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 $\overline{X}_{\boldsymbol{\sigma}}$. Note that the surface $\overline{X}_{\boldsymbol{\sigma}}$ can have horizontal nodes and is smooth elsewhere.

These notions obviously extend locally to *equisingular families* $(\pi\colon \mathcal{X} \to B, \boldsymbol{z})$ 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 though 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 5.15. We will return to these notions in detail in Section 12.

For an equisingular family, a *family of weldings* $\boldsymbol{\sigma}$ 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 12), and then a family of weldings is a continuous section of the S^1-bundle at each vertical node. For each family of weldings $\boldsymbol{\sigma}$ the *family of welded surfaces* $\overline{\mathcal{X}}_{\boldsymbol{\sigma}}$ is obtained by identifying the family of real oriented blowups of $\mathcal{X}$ along the identifications provided by $\boldsymbol{\sigma}$. 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 $S_pX = T_pX/\mathbb{R}_{>0}$ 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 5.4. Suppose the meromorphic differential ω on X has order $k \neq -1$ at some point p. A *complex prong* $v \in T_pX$ 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_pX obtained by projectivizing the complex prongs are the *real prongs* of ω at p. △

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 $\kappa_z = |1+k| = |1+\operatorname{ord}_z \omega|$. 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 5.5. A *local prong-matching* of (X, η) at q is a cyclic-order-reversing bijection $\sigma_q \colon P^{\mathrm{in}}_{q^-} \to P^{\mathrm{out}}_{q^+}$.

A *(global) prong-matching* $\boldsymbol{\sigma}$ for a twisted differential is a choice of a local prong-matching σ_q at each vertical node q of X. △

Note that prong-matchings at horizontal nodes are not defined. The following equivalent definition of a local prong-matching will be useful for studying families in Section 11:

Definition 5.6. 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. △

To see the equivalence of these definitions, any such σ_q corresponds to an order-preserving bijection $P^{\mathrm{in}}_{q^-} \to P^{\mathrm{out}}_{q^+}$ 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 $\overline{X}_{\boldsymbol{\sigma}}$ or simply by $\overline{X}$ the associated welded surface constructed using the welding defined by the prong-matching $\boldsymbol{\sigma}$.

Definition 5.7. A *prong-matched twisted differential of type* μ *compatible with* Γ, or just *prong-matched twisted differential* for short, is the datum $(X, \boldsymbol{z}, \eta, \boldsymbol{\sigma})$ consisting of a twisted differential $(X, \boldsymbol{z}, \eta)$ of type μ, compatible with Γ and a global prong-matching $\boldsymbol{\sigma}$. △

An isomorphism between prong-matched twisted differentials is an isomorphism of stable curves which identifies the forms on each component and is additionally required to commute with all of the local prong-matchings.

We will occasionally need to consider non-holomorphic maps between prong-matched twisted differentials, which we define as follows.

Definition 5.8. Given two prong-matched twisted differentials X_1 and X_2 with associated welded surfaces $\overline{X}_1$ and $\overline{X}_2$, an *almost-diffeomorphism* $f \colon \overline{X}_1 \to \overline{X}_2$ is a continuous map which satisfies:

(1) The preimage of each horizontal node is either a horizontal node or a simple closed curve disjoint from the nodes of $\overline{X}_1$, and the restriction of f to each component of $\overline{X}_1 \setminus f^{-1}(N^h_{X_2})$ is a diffeomorphism onto a component of $X_2 \setminus N^h_{X_2}$.
(2) The map $\Gamma(X_1) \to \Gamma(X_2)$ induces a degeneration of enhanced level graphs.

If f contracts no simple closed curves, we call it a *diffeomorphism.* △

For an equisingular family $(\pi \colon \mathcal{X} \to B, \boldsymbol{\mathcal{Z}}, \eta)$ of twisted differentials we define *a family of prong-matchings* to be a family of global weldings 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

$$P_\Gamma \;=\; \prod_{e\in\Lambda^v} \mathbb{Z}/\kappa_e\mathbb{Z}\,. \tag{5.1}$$

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 $P^{\rm in}_{q^-} \to P^{\rm in}_{q^-}$ 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 define the notion of a marking of a prong-matched twisted differential and construct the Teichmüller space $\Omega\mathcal{T}^{pm}_\Lambda(\mu)$ of marked prong-matched twisted differentials of type (μ,Λ) as a complex manifold. The notion of a marking is modeled on the definition of a marked stable curve, except the target of the marking is the associated welded surface.

Definition 5.9. A *marking* of a prong-matched twisted differential $(X,\boldsymbol{z},\eta,\boldsymbol{\sigma})$ is a continuous map $f\colon(\Sigma,\boldsymbol{s}) \to (\overline{X}_{\boldsymbol{\sigma}},\boldsymbol{z})$ which satisfies:

1. The preimage of every horizontal node is a simple closed curve on Σ.
2. If we denote by $\Lambda^h \subset \Sigma$ the "horizontal" multicurve consisting of the preimage of the set of horizontal nodes N^h_X of X, then the restriction of f to $\Sigma\setminus\Lambda^h$ is an orientation-preserving diffeomorphism $\Sigma\setminus\Lambda^h \to \overline{X}_{\boldsymbol{\sigma}}\setminus N^h_X$.
3. The map f preserves the marked points, that is, $f\circ\boldsymbol{s} = \boldsymbol{z}$.

We say that a *marked prong-matched twisted differential* $(X,f,\boldsymbol{z},\eta,\boldsymbol{\sigma})$ is *of type* Λ if Λ is the enhanced multicurve obtained by pulling back the seams of $\overline{X}_{\boldsymbol{\sigma}}$.

Two marked prong-matched twisted differentials are equivalent if there is an isomorphism of prong-matched twisted differentials that identifies the marking of their associated welded surfaces up to isotopy rel $\boldsymbol{s}$. △

Definition 5.10. The *Teichmüller space* $\Omega\mathcal{T}^{pm}_\Lambda(\mu)$ *of marked prong-matched twisted differentials of type* (μ,Λ) is the set of isomorphism classes of marked prong-matched twisted differentials of type μ with marking of type Λ. △

Given any contractible open set $U \subset \Omega\mathcal{T}_\Lambda(\mu)$ together with a prong-matching $\boldsymbol{\sigma}$ and a marking $f\colon\Sigma\to\overline{X}_{\boldsymbol{\sigma}}$ for some basepoint $(X,\eta)\in U$, we may uniquely extend $\boldsymbol{\sigma}$ to a continuous family of prong-matchings over U. The corresponding family of welded surfaces is then topologically trivial over U, so the marking f may be extended uniquely (up to isotopy) to a continuous family of markings over the base U. This defines a lift $U\to\Omega\mathcal{T}^{pm}_\Lambda(\mu)$. We give $\Omega\mathcal{T}^{pm}_\Lambda(\mu)$ the structure of a complex manifold such that these lifts are holomorphic local homeomorphisms. The forgetful map $\Omega\mathcal{T}^{pm}_\Lambda(\mu)\to\Omega\mathcal{T}_\Lambda(\mu)$ is then a holomorphic covering map of infinite degree.

Once we define families of marked prong-matched twisted differentials, in Proposition 13.4 we will observe that $\Omega\mathcal{T}^{pm}_\Lambda(\mu)$ is the fine moduli space for families of marked prong-matched twisted differentials, i.e. that every family of marked prong-matched twisted differentials can be obtained as a pullback from the family $\mathcal{X} \to \Omega\mathcal{T}_\Lambda^{pm}(\mu)$ of prong-matched twisted differentials over $\Omega\mathcal{T}_\Lambda^{pm}(\mu)$.

5.6. **Turning numbers on prong-matched twisted differentials.** Given a prong-matched twisted differential $(X, \boldsymbol{z}, \omega, \boldsymbol{\sigma})$, we call a *good arc* on $\overline{X}_{\boldsymbol{\sigma}}$ an arc γ which is transverse to the seams, disjoint from the horizontal nodes, and whose endpoints are disjoint from the seams. More generally, given an extension $\boldsymbol{\sigma}'$ of the vertical welding $\boldsymbol{\sigma}$ to a global welding, the image of an arc γ (satisfying the above constraints) on $\overline{X}_{\boldsymbol{\sigma}'}$ under the projection to $\overline{X}_{\boldsymbol{\sigma}}$ will be called a *generalized good arc*.

The *Gauss map* $G\colon S(X \setminus (\boldsymbol{z} \cup N_X)) \to S^1$, where SX denotes the tangent circle bundle $TX/\mathbb{R}_{>0}$, is defined by

$$G(v) \;=\; \frac{\omega(v)}{|\omega(v)|}.$$

Given any welding $\boldsymbol{\sigma}'$ of X whose restriction to the vertical nodes is a prong-matching, the Gauss map naturally extends to the tangent circle bundle of the blowup of X at the welded nodes. This extension is compatible with the weldings and descends to a Gauss map G from the tangent circle bundle of $\overline{X}_{\boldsymbol{\sigma}'}$ to S^1.

We define the *turning number* $\tau(\gamma)$ of any generalized good arc as $\tau(\gamma) = g(b) - g(a)$, where $g\colon [a, b] \to \mathbb{R}$ is a continuous lift of $G \circ \gamma$, that is $e^{2\pi i g} = G \circ \gamma$.

A good arc can alternatively thought of as a chain of arcs which may begin or end vertical nodes, subject to the constraint that at any vertical node, incoming and outgoing tangent vectors are identified by the prong matching. In these terms, the turning number of a good arc is simply the sum of the turning numbers of its pieces.

Turning numbers are invariant under regular isotopies (meaning isotopies through immersed curves) preserving the endpoints of γ as well as the tangent vectors at these endpoints.

Definition 5.11. Consider a prong-matched twisted differential X, a sequence of prong-matched twisted differentials $\{X_m\}_{m\in\mathbb{N}}$, and a sequence of almost-diffeomorphisms $h_m\colon X_m \to X$. We say that the sequence $\{h_m\}$ is *asymptotically turning number preserving* if for any good arc γ on $\overline{X}_\sigma$, we have $\tau(h_m^{-1}(\gamma)) \to \tau(\gamma)$ as $m \to \infty$. △

Proposition 5.12. *Consider a twisted differential X with a prong-matching $\boldsymbol{\sigma}$ and a sequence of prong-matchings $\boldsymbol{\sigma}_n$. Let h_n be an asymptotically turning number preserving sequence of diffeomorphisms of prong-matched twisted differentials $h_n\colon \overline{X}_{\boldsymbol{\sigma}_n} \to \overline{X}_{\boldsymbol{\sigma}}$. Suppose that the h_n^{-1} converge C^1-uniformly on compact sets of $X^s \setminus \boldsymbol{z}$ to the identity map. Then for n sufficiently large, $\boldsymbol{\sigma} = \boldsymbol{\sigma}_n$, and moreover h_n is isotopic to the identity map on $\overline{X}_{\boldsymbol{\sigma}}$.*

Proof. Fix a compact subsurface K which is a deformation retract of $X^s \setminus \boldsymbol{z}$, and let γ be a curve joining two boundary components and crossing a single seam corresponding to a node p. For n large, h_n^{-1} is C^1-close to the identity on K, so the endpoints of γ and $h_n^{-1}(\gamma)$ are close, as well as their corresponding tangent vectors. Since the turning number of a curve is determined (mod $\mathbb{Z}$) by its endpoints and tangent vectors, we have $\tau(h_n^{-1}(\gamma)) - \tau(\gamma) \sim \theta_n \pmod{\mathbb{Z}}$, where the prong-matchings at p are related by $e^{2\pi i\theta_n}\sigma = \sigma_n$. Since the h_n are asymptotically turning number preserving, we have $\theta_n \to 0$, so $\sigma_n = \sigma$ eventually, as the set of prong-matchings is discrete.

Now take n large enough that $\boldsymbol{\sigma} = \boldsymbol{\sigma}_n$ and moreover, h_n^{-1} moves each point q of K by a distance smaller than the injectivity radius of $X^s \setminus \boldsymbol{z}$ at q. We may then on K take the nearest-point isotopy from h_m^{-1} to the identity, and extend it to an isotopy on $\overline{X}_{\boldsymbol{\sigma}}$ from h_m^{-1} to a map k which is the identity on K. Each seam is contained in an annular component A of $\overline{X}_{\boldsymbol{\sigma}} \setminus K$, and to show that k is isotopic to the identity on A (rel ∂A) it suffices to show that $k(\gamma)$ is isotopic to γ (rel ∂A), with γ as before joining the two boundary components. But the isotopy class of γ (through immersed curves with the endpoints and vectors fixed) is determined by the turning number of γ, since the core curve of A has nonzero turning number. As $\tau(h_,^{-1}(\gamma)) - \tau(\gamma) \to 0$, and this difference of turning numbers is integral, the difference is eventually zero, so these curves are isotopic. ∎

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 $X_m = (X_m, \boldsymbol{z}_m, \eta_m, \preccurlyeq, \boldsymbol{\sigma}_m, f_m)$ of marked prong-matched twisted differentials in $\Omega\mathcal{T}_\Lambda^{pm}(\mu)$ converges in the *conformal topology* to $X = (X, \boldsymbol{z}, \eta, \preccurlyeq, \boldsymbol{\sigma}, f)$ if and only if for any sufficiently large m there exists a diffeomorphism $g_m \colon \overline{X}_m \to \overline{X}$ 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 is isotopic to $g_m \circ f_m$ rel marked points.
(ii) The function g_m^{-1} is conformal on the ϵ_m-thick part $(X, \boldsymbol{z})_{\epsilon_m}$.
(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 asymptotically turning-number-preserving. △

Here (and in the sequel) we use the pushforward notation $g_* = (g^{-1})^*$ for the action on differentials.

Remark 5.14. In order to verify that the g_m are asymptotically turning-number-preserving, 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 turning numbers converge for these arcs, then together with (ii) this forces the convergence of the other turning numbers as well.

Proposition 5.15. *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. Consider a sequence $X_m \to X$ as in the above definition. Let $U \subset \Omega\mathcal{T}_\Lambda(\mu)$ be a contractible neighborhood of the point corresponding to X, and let $\mathcal{X} \to U$ be the restriction of the universal curve over $\Omega\mathcal{T}_\Lambda(\mu)$ to U. The prong-matching $\boldsymbol{\sigma}$ of X extends uniquely to a continuous family of prong-matchings over U. Choose a smooth (in the sense of Definition 5.8) trivialization $h\colon X \times U \to \mathcal{X}$ of this family whose restriction to the fiber over X is the identity. This defines a lift $U \to \Omega\mathcal{T}_\Lambda^{pm}(\mu)$, and our goal is to show that X_m is eventually in U.

Restricting the marking maps to the complement of the seams defines a projection $\pi\colon \Omega\mathcal{T}^{pm}_\Lambda(\mu) \to \Omega\mathcal{T}_\Lambda(\mu)$. Items (i)–(iii) above, together with Proposition 3.5 ensures that this projection is continuous, so that $\pi(X_m) \to \pi(X)$ in $\Omega\mathcal{T}_\Lambda(\mu)$. The trivialization h then defines smooth maps $\overline{h}_m\colon \overline{X}_{\boldsymbol{\sigma}} \to \overline{(X_m)}_{\boldsymbol{\sigma}_m}$. The maps $\overline{h}_m^{-1} \circ g_m$ then satisfy the hypotheses of Proposition 5.12, which ensures g_m eventually lies in the lift of U.

Conversely, consider a sequence X_m so that $\pi(X_m) \to \pi(X)$ in $\Omega\mathcal{T}_\Lambda(\mu)$, with the points corresponding to X_m eventually in the lift of U constructed above. The desired maps g_m may be constructed by modifying (as in the proof of Theorem 3.2) the maps h_m constructed above so that they are conformal on an exhaustion. ◻

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 7 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 7, 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 $\Omega\mathcal{T}^{pm}_\Lambda(\mu)$, 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 $\Omega\mathcal{T}^{pm}_\Lambda(\mu)$, we have to pass to the universal cover $\mathbb{C}^{L^\bullet(\Lambda)}$ of $(\mathbb{C}^*)^{L^\bullet(\Lambda)}$, which acts continuously on $\Omega\mathcal{T}^{pm}_\Lambda(\mu)$ by *level rotations*, as we now describe.

(1) On the level of forms, the tuple $\boldsymbol{d} = (d_i)_{i\in L^\bullet(\Lambda)} \in \mathbb{C}^{L^\bullet(\Lambda)}$ acts through the quotient $(\mathbb{C}^*)^{L^\bullet(\Lambda)}$ by multiplying the form at level i by $\boldsymbol{e}(d_i)$ (recall that we denote $\boldsymbol{e}(z) = \exp(2\pi\sqrt{-1}z)$).

(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 $\boldsymbol{d} = (d_i)_{i\in L^\bullet(\Lambda)}$ we define

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

where for each vertical node q we let

$$\boldsymbol{d}\cdot\sigma_q\colon P^{\mathrm{in}}_{q^-} \to P^{\mathrm{out}}_{q^+} \tag{6.2}$$

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

Alternatively, following Definition 5.6, σ_q can be regarded as an element of $T^*_{q^+}X \otimes T^*_{q^-}X$ in which case $\boldsymbol{d}\cdot\sigma_q = \boldsymbol{e}((d_{\ell(q^+)} - d_{\ell(q^-)})/\kappa_q)\sigma_q$.

(3) On the marking f, the element $\boldsymbol{d} \in \mathbb{C}^{L^\bullet(\Lambda)}$ acts by composition with a *fractional Dehn twist*

$$F_{\boldsymbol{d}}\colon \overline{X}_\sigma \to \overline{X}_{\boldsymbol{d}\cdot\sigma} \tag{6.3}$$

which is the identity map outside a union of annular neighborhoods A_q of the corresponding seams.

6.2. **Twist groups.** The restriction of the action of $\mathbb{C}^{L^\bullet(\Lambda)}$ on $\Omega\mathcal{T}^{pm}_\Lambda(\mu)$ 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):

$$\phi^\bullet_\Lambda\colon \mathbb{Z}^{L^\bullet(\Lambda)} \to P_\Lambda\,, \qquad \boldsymbol{n} \mapsto \big(n_{\ell(e^+)} - n_{\ell(e^-)} \mod \kappa_e\big)_{e\in\Lambda^v}\,. \tag{6.4}$$

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. △

The homomorphism $\phi^\bullet_\Lambda$ fits into the following commutative diagram of group homomorphisms

$$\begin{array}{ccccccc}
\mathrm{Tw}^v_\Lambda & \xleftarrow{\tau^\bullet_\Lambda} & \ker(\phi^\bullet_\Lambda) & \hookrightarrow & \mathbb{Z}^{L^\bullet(\Lambda)} & & \\
\big\downarrow & & \big\downarrow{\scriptstyle \widetilde{\phi}^\bullet_\Lambda} & & \big\downarrow{\scriptstyle \widetilde{\phi}^\bullet_\Lambda} & \searrow^{\phi^\bullet_\Lambda} & \\
\mathrm{Mod}_{(\Sigma,\boldsymbol{s})} & \xleftarrow[\psi]{} & \ker(\overline{\psi}) & \hookrightarrow & \mathbb{Z}^{\Lambda^v} & \xrightarrow[\overline{\psi}]{} & P_\Lambda
\end{array}$$

that we now describe. The group $\mathbb{Z}^{\Lambda^v}$ acts on the space $\Omega\mathcal{T}^{pm}_\Lambda(\mu)$ 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 $\kappa_e\{n_e/\kappa_e\}$ (the remainder of n_e mod κ_e) and precomposes 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 $\mathbb{Z} \to \mathbb{Z}/\kappa_e\mathbb{Z}$, and doing this for all vertical edges induces a map $\overline{\psi}\colon \mathbb{Z}^{\Lambda^v} \to P_\Lambda$ 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 $\psi\colon\ \ker(\overline{\psi}) \hookrightarrow \mathrm{Mod}_{(\Sigma,\boldsymbol{s})}$ this inclusion.

There is a natural homomorphism $\widetilde{\phi}^\bullet_\Lambda\colon \mathbb{Z}^{L^\bullet(\Lambda)} \to \mathbb{Z}^{\Lambda^v}$ defined by

$$\widetilde{\phi}^\bullet_\Lambda(\boldsymbol{n}) \;=\; \big(n_{\ell(e^+)} - n_{\ell(e^-)}\big)_{e\in\Lambda^v}\,. \tag{6.5}$$

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

$$\tau^\bullet_\Lambda(\boldsymbol{n}) = \prod_{e\in\Lambda^v} \mathrm{tw}^{m_e}_{\gamma_e}, \quad \text{with } m_e \text{ defined by} \quad \widetilde{\phi}^\bullet_\Lambda(\boldsymbol{n}) = (m_e\kappa_e)_{e\in\Lambda^v}\,, \tag{6.6}$$

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

Proposition 6.2. *The vertical* Λ*-twist group* Tw^v_Λ *is a free abelian group of rank* N*. Moreover,* $\ker(\tau^\bullet_\Lambda) \subset \ker(\phi^\bullet_\Lambda)$ *is isomorphic to* $\mathbb{Z}$*, generated by* $\mathbf{1} = (1,\ldots,1)$*, and* $\ker(\phi^\bullet_\Lambda) = \mathrm{Tw}^v_\Lambda \oplus \ker(\tau^\bullet_\Lambda) \cong \mathrm{Tw}^v_\Lambda \oplus \mathbb{Z}$*.*

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

$$0 \to \mathbb{Z} \to \ker(\phi^\bullet_\Lambda) \to \mathrm{Tw}^v_\Lambda \to 0. \qquad \square$$

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

$$\mathrm{Tw}_\Lambda \;=\; \mathrm{Tw}^v_\Lambda \oplus \mathrm{Tw}^h_\Lambda\,.$$

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

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

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

Recall from Section 5.1 that for every level $i \in L(\Lambda)$ there is a two-level undegeneration $\mathrm{dg}_i \colon \Lambda_i \rightsquigarrow \Lambda$ that contracts the (vertical) edges of $\Gamma(\Lambda)$ strictly above level i and the edges below or at level i. We denote by $\mathrm{Tw}^{sv}_{\Lambda,i} = (\mathrm{dg}_i)_*(\mathrm{Tw}^v_{\Lambda_i}) \subset \mathrm{Tw}^v_\Lambda$ the corresponding subgroup of the vertical Λ-twist group. Note that $\mathrm{Tw}^v_{\Lambda_i}$ is the cyclic group generated by the element $(0,a_i)$ with $a_i = \mathrm{lcm}_e\,\kappa_e$ over all edges e connecting the graph $\Gamma_{>i}$ to $\Gamma_{\le i}$. Moreover, $(\mathrm{dg}_i)_*(0,a_i) = (0^{-i},a_i^{N+i+1}) \in \mathbb{Z}^{L^\bullet(\Lambda)}$. It follows that $\mathrm{Tw}^{sv}_{\Lambda,i}$ is the cyclic subgroup of Tw^v_Λ generated by

$$\tau^\bullet_\Lambda(a_i\boldsymbol{b}_i) \;=\; \prod_e \mathrm{tw}^{m_{e,i}}_{\gamma_e}, \quad \text{with} \quad a_i \;=\; \operatorname*{lcm}_e \kappa_e \quad \text{and} \quad m_{e,i} \;=\; a_i/\kappa_e\,, \tag{6.7}$$

where the product and the lcm are taken for the set of vertical edges e connecting $\Gamma_{>i}$ to $\Gamma_{\le i}$. The collection of $\mathrm{Tw}^{sv}_{\Lambda,i}$ for all $i \in L(\Lambda)$ generates a subgroup of the twist group, which we call the *simple vertical* Λ*-twist group* $\mathrm{Tw}^{sv}_{\Lambda}$.

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*

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

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

We denote by K_Λ the finite quotient $K_\Lambda = \mathrm{Tw}^v_\Lambda/\mathrm{Tw}^{sv}_\Lambda$. 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 also Section 8. Finally, we denote by $\mathrm{Tw}^s_\Lambda = \mathrm{Tw}^{sv}_\Lambda \oplus \mathrm{Tw}^h_\Lambda$ 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 $\phi^\bullet_\Lambda$ to $\mathbb{Z}^{L(\Lambda)}$ is denoted by ϕ_Λ.

The homomorphisms $\phi^\bullet_\Lambda$ and ϕ_Λ have the same image in the prong rotation group P_Λ. Similarly $\tau^\bullet_\Lambda$ and τ_Λ have the same image Tw^v_Λ in $\mathrm{Mod}_{(\Sigma,\boldsymbol{s})}$. Intuitively, the actions both 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 remark that there is another characterization of the twist group as a subgroup of the full twist group $\mathrm{Tw}^{\mathrm{full}}_\Lambda$. 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 $\mathbb{Z}^{E(\Lambda)}$. To prove the following proposition one checks that the quotient of $\mathrm{Tw}^{\mathrm{full}}_\Lambda$ by Tw_Λ and by the group $\mathrm{Tw}^{\mathrm{rot}}_\Lambda$ defined implicitly in the proposition are torsion free, and that Tw_Λ and $\mathrm{Tw}^{\mathrm{rot}}_\Lambda$ have the same rank.

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

6.3. **Level rotation tori.** We define the *level rotation torus* T_Λ to be the quotient $T_\Lambda = \mathbb{C}^{L(\Lambda)}/\mathrm{Tw}^v_\Lambda \cong \mathbb{C}^{L(\Lambda)}/\ker(\phi_\Lambda)$. Similarly, the *simple level rotation torus* is the quotient $T^s_\Lambda = \mathbb{C}^{L(\Lambda)}/\mathrm{Tw}^{sv}_\Lambda$. They will play a prominent role in defining families of multi-scale differentials. The level rotation tori depend obviously only 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*

$$(\mathbb{C}^*)^{L(\Lambda)} \times (\mathbb{C}^*)^{E(\Lambda)} \;=\; ((r_i, \rho_e))_{i \in L(\Lambda), e \in E(\Lambda)} \tag{6.9}$$

cut out by the set of equations

$$r_{\ell(e^-)} \cdots r_{\ell(e^+)-1} \;=\; \rho_e^{\kappa_e} \tag{6.10}$$

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 $T_\Lambda^s \to T_\Lambda$ is given in coordinates by

$$(q_i) \;\mapsto\; (r_i, \rho_e) \;=\; \left(q_i^{a_i}, \prod_{i=\ell(e^-)}^{\ell(e^+)-1} q_i^{a_i/\kappa_e} \right) \tag{6.11}$$

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 (6.10) (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 $\boldsymbol{n} \in \mathbb{Z}^{L(\Lambda)}$ acts by multiplying each coordinate ρ_e by $\boldsymbol{e}((n_{\ell(e^+)} - n_{\ell(e^-)})/\kappa_e)$. Recalling Equation (6.4) that defines ϕ_Λ, we see that $\ker(\phi_\Lambda)$ is precisely the set of elements $\boldsymbol{n} \in \mathbb{Z}^{L(\Lambda)}$ 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 (6.11) follows, since we exhibit a map of tori of the same dimension and the right-hand side satisfies (6.10). ∎

These constructions can also be regarded as covariant functors on the category of ordered enhanced multicurves on $(\Sigma, \boldsymbol{s})$. More precisely, a degeneration of enhanced multicurves $\mathrm{dg}\colon \Lambda_1 \rightsquigarrow \Lambda_2$ induces a monomorphism $\widehat{\mathrm{dg}_*}\colon \mathbb{C}^{\Lambda_1} \to \mathbb{C}^{\Lambda_2}$. 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 $\mathrm{dg}_*\colon \mathrm{Tw}_{\Lambda_1} \hookrightarrow \mathrm{Tw}_{\Lambda_2}$.

Lemma 6.6. *A degeneration of enhanced multicurves $\mathrm{dg}\colon \Lambda_1 \rightsquigarrow \Lambda_2$ induces an injective homomorphism $\mathrm{dg}_*\colon T_{\Lambda_1} \to T_{\Lambda_2}$. 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 $\widehat{\mathrm{dg}_*}$ already belongs to Tw_{Λ_1}. This is obvious from the description of the twist group in Proposition 6.2. ∎

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

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

$$(r_i, \rho_e) * \big(X, (\eta_{(i)}), (\sigma_e)\big) \;=\; \Big(X, \big(r_i \dots r_{-1}\eta_{(i)}\big), (\rho_e * \sigma_e)\Big) \tag{6.12}$$

where $\rho_e * \sigma_e$ is the prong-matching $P^{\mathrm{in}}_{q^-} \to P^{\mathrm{out}}_{q^+}$ at the node q corresponding to e given by σ_e post-composed with the rotation by $\arg(\rho_e)$. 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 $(r_i, \rho_e) * \big(X, (\eta_{(i)}), (\sigma_e)\big)$ 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 $T_\Lambda^s \to T_\Lambda$ given in Proposition 6.5, the action $*$ defined in Equation (6.12) is given in the triangular basis by

$$\boldsymbol{t} * (X, (\eta_{(i)}), (\sigma_e), f) \;=\; (X, (t_i^{a_i} \dots t_{-1}^{a_{-1}} \eta_{(i)}), (f_e * \sigma_e), (f_e) * f)\,, \tag{6.13}$$

where $f_e = \prod_{i=\ell(e^-)}^{\ell(e^+)-1} t_i^{a_i/\kappa_e}$ 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* (6.12) *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* (6.10).

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 finite index 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 $H_i = \mathbb{Z}_i/\mathrm{Tw}_{\Lambda,i}^{sv}$, 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 $H := \oplus_{i\in L(\Lambda)} H_i = \mathrm{Ker}(T_\Lambda^s \to (\mathbb{C}^*)^{L(\Lambda)})$. On the other hand, we may define the *(full) ramification group* associated with an enhanced level graph Λ to be $G := \mathrm{Ker}(T_\Lambda \to (\mathbb{C}^*)^{L(\Lambda)})$. By definition we have an exact sequence of finite abelian groups

$$0 \to K_\Lambda = \mathrm{Tw}_\Lambda^v/\mathrm{Tw}_\Lambda^{sv} \to H \to G \to 0\,. \tag{6.14}$$

Note that the map $H_i \to G$ 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:

$$\begin{array}{ccc} \Omega\mathcal{T}^{pm}_{\Lambda}(\mu)/\mathrm{Tw}^{s}_{\Lambda} & \xrightarrow{/K_\Lambda} & \\ \Big\downarrow /H & & \Omega\mathcal{T}^{pm}_{\Lambda}(\mu)/\mathrm{Tw}_{\Lambda} \\ \Omega\mathcal{T}_{\Lambda}(\mu) & \xleftarrow{/G} & \end{array}$$

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.

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

$$\phi^{\bullet}_{\Lambda_1}\colon \mathbb{Z}^{L^\bullet(\Lambda)} \cong \mathbb{Z}^3 \to P_{\Gamma_1} \cong \mathbb{Z}/3\mathbb{Z} \times \mathbb{Z}/3\mathbb{Z} \times \mathbb{Z}/\mathbb{Z}$$

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 $\ker(\phi^{\bullet}_{\Lambda_1})$ 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 $\mathrm{Tw}^{v}_{\Lambda} = \mathrm{Tw}_{\Lambda}$ is generated by the vectors $(0,3,0)$ and $(0,0,3)$. The simple Λ_1-twist group $\mathrm{Tw}^{sv}_{\Lambda_1}$ is the direct sum $\mathrm{Tw}^{sv}_{\Lambda_1,-1} \oplus \mathrm{Tw}^{sv}_{\Lambda_1,-2}$ where $\mathrm{Tw}^{sv}_{\Lambda_1,i} = (\mathrm{dg}_i)_*(\mathrm{Tw}^{v}_{\Lambda_{1,i}})$. In this case $\mathrm{Tw}^{sv}_{\Lambda_1,-1}$ is generated by $(0,3,3)$ and $\mathrm{Tw}^{sv}_{\Lambda_1,-2}$ is generated by $(0,0,3)$, hence $\mathrm{Tw}^{sv}_{\Lambda_1}$ coincides with $\mathrm{Tw}^{v}_{\Lambda_1}$. 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 $\phi^{\bullet}_{\Lambda_1}$ 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

$$\phi^{\bullet}_{\Lambda_2}\colon \mathbb{Z}^{L^\bullet(\Lambda_2)} \cong \mathbb{Z}^3 \to P_{\Gamma_2} \cong \mathbb{Z}/\mathbb{Z} \times \mathbb{Z}/\mathbb{Z} \times \mathbb{Z}/3\mathbb{Z}$$

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 $\ker(\phi^{\bullet}_{\Lambda_2})$ 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 $\mathrm{Tw}^{v}_{\Lambda_2} = \mathrm{Tw}_{\Lambda_2}$ is generated by the vectors $(0,1,0)$ and $(0,0,3)$. The simple Λ_2-twist group $\mathrm{Tw}^{sv}_{\Lambda_2}$ is the direct sum $\mathrm{Tw}^{sv}_{\Lambda_2,-1} \oplus \mathrm{Tw}^{sv}_{\Lambda_2,-2}$ where $\mathrm{Tw}^{sv}_{\Lambda_2,-1}$ is generated by $(0,3,3)$ and $\mathrm{Tw}^{sv}_{\Lambda_2,-2}$ is generated by $(0,0,3)$, hence $\mathrm{Tw}^{sv}_{\Lambda_2}$ is a subgroup of index 3 in $\mathrm{Tw}^{v}_{\Lambda_2}$. 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 8.2.

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 $\kappa_3 = 1$.

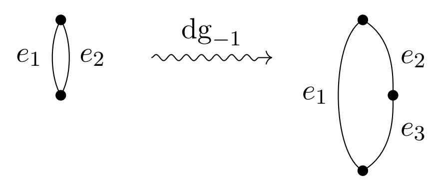

FIGURE 3. A degeneration that decreases the number of prong-matchings.

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 (6.10). We define the *level rotation torus closure* $\overline{T}_\Lambda$ 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^s_Λ is naturally identified with $(\mathbb{C}^*)^{L(\Lambda)}$, with closure $\overline{T}^s_\Lambda = \mathbb{C}^{L(\Lambda)}$. The group $K_\Lambda = \mathrm{Tw}_\Lambda/\mathrm{Tw}^s_\Lambda$ introduced in the previous section acts on $T^s_\Lambda = \mathbb{C}^{L(\Lambda)}/\mathrm{Tw}^s_\Lambda$. 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 $\overline{T}^s_\Lambda$. 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* $p\colon T^s_\Lambda \to T_\Lambda$ *given by Equation* (6.11) *extends to* $\overline{T}^s_\Lambda$ *and descends to an isomorphism* $\overline{p}\colon \overline{T}^s_\Lambda/K_\Lambda \to \overline{T}^n_\Lambda$ *to the normalization of the level rotation torus closure.*

Proof. The map p extends to a map $p_2\colon \overline{T}^s_\Lambda \to \overline{T}_\Lambda$ since it is given explicitly by monomials, in the coordinates used for taking the closure. Since K_Λ acts on $\overline{T}^s_\Lambda$ and since p is the quotient by K_Λ, the map p_2 factors through the quotient to give $p_3\colon \overline{T}^s_\Lambda/K_\Lambda \to \overline{T}_\Lambda$. 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 $T^s_\Lambda/K_\Lambda = T_\Lambda$. Since the target is normal, it follows that the map $\overline{p}$ is an isomorphism (see [Sta18, Lemma 28.52.8]). ∎

Example 6.11. (*In general, $\overline{T}_\Lambda$ 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 $\overline{T}_\Lambda$ has a cusp, locally modeled 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 $\mathrm{Tw}_\Lambda = \mathrm{Tw}^s_\Lambda$ in this example.

7. Augmented Teichmüller space of marked multi-scale differentials

In this section, we formally introduce the notion of a multi-scale differential, as well as markings. We introduce the augmented Teichmüller space, parameterizing equivalence classes of marked multi-scale differentials, define its topology, and establish basic topological properties. The main results are the Hausdorff property of quotients of augmented Teichmüller space by any subgroup of the mapping class group in Theorem 7.7 and the compactness of the moduli space in Theorem 7.12.

7.1. Multi-scale differentials and markings.

The notation we developed so far enables us to make the Definition 1.1 of multi-scale differentials precise. The notion of general families of multi-scale differentials is not needed for the construction of augmented Teichmüller space and Dehn space and will be given in Section 11.

Definition 7.1. Given an enhanced level graph Γ, a *multi-scale differential of type* (μ, Γ) consists of the following data:

(1) A pointed stable curve $(X, \boldsymbol{z})$,
(2) an identification of the dual graph Γ_X with Γ,
(3) a collection of differentials $\boldsymbol{\omega} = (\omega_{(i)})_{i \in L^\bullet(\Gamma)}$ that give $(X, \boldsymbol{z})$ the structure of a twisted differential of type (μ, Γ) compatible with Γ, and
(4) a global prong matching $\boldsymbol{\sigma} = (\sigma_e)_{e \in E(\Gamma)}$,

where $(\boldsymbol{\omega}, \boldsymbol{\sigma})$ is defined to be equivalent to $(\boldsymbol{\omega}', \boldsymbol{\sigma}')$ if and only if there exists an element of the level rotation torus T_Γ that sends $(\boldsymbol{\omega}, \boldsymbol{\sigma})$ to $(\boldsymbol{\omega}', \boldsymbol{\sigma}')$ under the action (6.12). △

Given an enhanced multicurve $\Lambda \subset \Sigma \setminus \boldsymbol{s}$, a *marked multi-scale differential of type* (μ, Λ) is the data of $(\boldsymbol{\omega}, \boldsymbol{\sigma}, f)$, where f is a marking (in the sense of Definition 5.9) of the prong-matched twisted differential defined by $(\boldsymbol{\omega}, \boldsymbol{\sigma})$, considered up to the action of $\mathbb{C}^{L(\Lambda)}$, as in Section 6.1.

Note that, since the vertical twist group $\mathrm{Tw}^v_\Lambda \subset \mathbb{C}^{L(\Lambda)}$ fixes $(\boldsymbol{\omega}, \boldsymbol{\sigma})$, a marked multi-scale differential may alternatively be defined as a T_Γ-equivalence class of the data $(\boldsymbol{\omega}, \boldsymbol{\sigma}, [f])$, where $[f]$ is the equivalence class of markings up to the action of Tw_Λ. More generally, we will call this a G-marked multi-scale differential if the marking is taken up to the action of a subgroup G of the mapping class group.

Projectivized multi-scale differentials and their marked analogues are defined similarly, replacing T_Γ with $T_\Gamma^\bullet$.

7.2. **Augmented Teichmüller space of multi-scale differentials as a set.** As a first step towards defining the augmented Teichmüller space of multi-scale differentials of type μ, we introduce it as a set.

Definition 7.2. The *(projectivized) augmented Teichmüller space of multi-scale differentials of type* μ is the set of equivalence classes of (projectivized)marked multi-scale differentials of type μ, which we denote by $\Omega\overline{\mathcal{T}}_{(\Sigma,\boldsymbol{s})}(\mu)$ or $\mathbb{P}\Omega\overline{\mathcal{T}}_{(\Sigma,\boldsymbol{s})}(\mu)$ respectively. △

Given an enhanced multicurve $\Lambda \subset \Sigma \setminus \boldsymbol{s}$, we define the *(projectivized)* Λ*-boundary stratum* to be the set $\Omega\mathcal{B}_\Lambda \subset \Omega\overline{\mathcal{T}}_{(\Sigma,\boldsymbol{s})}(\mu)$ or $\mathbb{P}\Omega\mathcal{B}_\Lambda \subset \mathbb{P}\Omega\overline{\mathcal{T}}_{(\Sigma,\boldsymbol{s})}(\mu)$ of (projectivized) marked multi-scale differentials of type (μ,Λ). Note that the special case $\Omega\mathcal{B}_\emptyset = \Omega\mathcal{T}_\emptyset^{pm}(\mu) = \Omega\mathcal{T}_{(\Sigma,\boldsymbol{s})}(\mu)$ is the Teichmüller space of marked flat surfaces of type μ. These strata have the structure of complex manifolds, as they are naturally identified with the quotients $\Omega\mathcal{T}_\Lambda^{pm}(\mu)/\mathbb{C}^{L(\Lambda)}$ or $\Omega\mathcal{T}_\Lambda^{pm}(\mu)/\mathbb{C}^{L^\bullet(\Lambda)}$ of the Teichmüller spaces of prong-matched twisted differentials of the corresponding type. The augmented Teichmüller spaces are the disjoint union of these strata over the set of enhanced multicurves $\Lambda \subset \Sigma \setminus \boldsymbol{s}$:

$$\Omega\overline{\mathcal{T}}_{(\Sigma,\boldsymbol{s})}(\mu) \;=\; \coprod_\Lambda \Omega\mathcal{B}_\Lambda \quad \text{and} \quad \mathbb{P}\Omega\overline{\mathcal{T}}_{(\Sigma,\boldsymbol{s})}(\mu) \;=\; \coprod_\Lambda \mathbb{P}\Omega\mathcal{B}_\Lambda\,. \tag{7.1}$$

The mapping class group $\mathrm{Mod}_{g,n}$ acts on $\Omega\overline{\mathcal{T}}_{(\Sigma,\boldsymbol{s})}(\mu)$ and $\mathbb{P}\Omega\overline{\mathcal{T}}_{(\Sigma,\boldsymbol{s})}(\mu)$ by pre-composition of the marking, that is, given $g \in \mathrm{Mod}_{g,n}$, an equivalence class of markings $f\colon \Sigma \to \overline{X}_{\boldsymbol{\sigma}}$ is replaced with $[f \circ g^{-1}]$

Proposition 7.3. *The subgroup of* $\mathrm{Mod}_{g,n}$ *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. The first statement follows directly from the definition of the equivalence relation on markings. For the second statement, if Λ' is a degeneration of Λ, then $\mathrm{Tw}_\Lambda \subset \mathrm{Tw}_{\Lambda'}$. Hence Tw_Λ fixes $\Omega\mathcal{B}_{\Lambda'}$ pointwise. ∎

7.3. **Augmented Teichmüller space as a topological space.** We now give both augmented Teichmüller spaces $\Omega\overline{\mathcal{T}}_{(\Sigma,\boldsymbol{s})}(\mu)$ and $\mathbb{P}\Omega\overline{\mathcal{T}}_{(\Sigma,\boldsymbol{s})}(\mu)$ 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, $\boldsymbol{\omega}$ in the tuple $(X,\boldsymbol{z},\boldsymbol{\omega},\preccurlyeq,\boldsymbol{\sigma},f)$ refers to a chosen representative of the equivalence class. We also write $\overline{X}_{\boldsymbol{\sigma}_n}$ as shorthand for $(\overline{X}_n)_{\boldsymbol{\sigma}_n}$.

Definition 7.4. A sequence $(X_n,\boldsymbol{z}_n,\boldsymbol{\omega}_n,\preccurlyeq_n,\boldsymbol{\sigma}_n,f_n) \in \mathbb{P}\Omega\mathcal{B}_{\Lambda_n}$ converges to a point $(X,\boldsymbol{z},\boldsymbol{\omega},\preccurlyeq,\boldsymbol{\sigma},f) \in \mathbb{P}\Omega\mathcal{B}_\Lambda \subset \mathbb{P}\Omega\overline{\mathcal{T}}_{(\Sigma,\boldsymbol{s})}(\mu)$, 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 $\boldsymbol{d}_n = \{d_{n,i}\}_{i\in L^\bullet(\Lambda)} \in \mathbb{C}^{L^\bullet(\Lambda)}$ such that the following conditions hold, where we denote $c_{n,i} = \boldsymbol{e}(d_{n,i})$:

(1) For sufficiently large n there is an undegeneration of enhanced multicurves (δ_n, D_n^h) with $\delta_n\colon L^\bullet(\Lambda) \to L^\bullet(\Lambda_n)$ (see Definition 5.1).
(2) For sufficiently large n there exists an almost-diffeomorphism $g_n\colon \overline{X}_{\boldsymbol{d}_n\cdot\boldsymbol{\sigma}_n} \to \overline{X}_{\boldsymbol{\sigma}}$ that is compatible with the markings (in the sense that $g_n \circ (\boldsymbol{d}_n \cdot f_n)$ is isotopic to f rel marked points) and such that g_n^{-1} is conformal on the ϵ_n-thick part $(X, \boldsymbol{z})_{\epsilon_n}$.
(3) The restriction of $c_{n,i}(g_n)_*(\boldsymbol{\omega}_n)$ to the ϵ_n-thick part of the level i subsurface of $(X, \boldsymbol{z})$ converges uniformly on compact sets 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
$$\lim_{n\to\infty} \frac{|c_{n,i}|}{|c_{n,j}|} = 0\,.$$
(5) The almost-diffeomorphisms g_n are asymptotically turning number preserving.

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

Note that the notion of convergence does not depend on the choice of representative of X in $\Omega\mathcal{T}_\Lambda^{pm}(\mu)$ since if $X' = \boldsymbol{d}' \cdot X$ is another representative, then using $\boldsymbol{d}' + \boldsymbol{d}$ certifies convergence to X'. Note moreover, that in item (5) we could as well require the difference of turning numbers to tend to 0 for a fixed collection of arcs dual to the collection of seams, since for any fixed arc γ disjoint from the seams, $\tau(g_n^{-1}(\gamma)) \to \tau(\gamma)$, as g_n converges uniformly C^1 to the identity on γ.

This topology can be given equivalently by a basis of open sets which we now describe. For a given marked multi-scale differential X, let $V_\epsilon(X)$ be the set of marked multi-scale differentials $(X_0, \boldsymbol{z}_0, \boldsymbol{\omega}_0, \preccurlyeq_0, \boldsymbol{\sigma}_0, f_0)$ such that there exists $\boldsymbol{d} = \{d_i\}_{i\in L^\bullet(\Lambda)} \in \mathbb{C}^{L^\bullet(\Lambda)}$ and

(i) a degeneration $(\Lambda_0, \preccurlyeq_0) \rightsquigarrow (\Lambda, \preccurlyeq)$ given by a map $\delta\colon L(\Lambda) \to L(\Lambda_0)$ and a subset of N_X^h,
(ii) an almost-diffeomorphism $g\colon \overline{X}_{\boldsymbol{d}\cdot\boldsymbol{\sigma}_0} \to \overline{X}_{\boldsymbol{\sigma}}$ with g^{-1} conformal on the ϵ-thick part of X and compatible with the markings,
(iii) letting $c_i = \boldsymbol{e}(d_i)$, 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 each good arc γ on $\overline{X}_{\boldsymbol{\sigma}}$,
$$|\tau(g^{-1} \circ \gamma) - \tau(\gamma)| < \epsilon.$$

Proposition 7.5. *The sets $V_e(X)$ are a basis for a topology on $\Omega\overline{\mathcal{T}}_{(\Sigma,\boldsymbol{s})}(\mu)$, and the convergent sequences in this topology agree with the notion of convergence of Definition 7.4.*

Note that the topology defined by this basis is apparently stronger than that defined by Definition 7.4 because (v) requires uniform control on turning numbers of all good arcs. In the next lemma we show that one can uniformly control the change in turning numbers of all good arcs by restricting to a finite collection of arcs dual to the seams. For this, fix a collection $\{\gamma_1, \dots, \gamma_k\}$ of good arcs on a multi-scale differential X which are dual to the seams of X, meaning that to each seam corresponds exactly one γ_i which crosses that seam once and is disjoint from the others. We define $V'_\epsilon(X)$ to be the set of marked multi-scale differentials which satisfy conditions (i)-(iv) above as well as

(vi) for each arc γ_i, we have $|\tau(g^{-1} \circ \gamma_i) - \tau(\gamma_i)| < \epsilon$.

Lemma 7.6. *Consider $X \in \Omega\overline{\mathcal{T}}_{(\Sigma,\boldsymbol{s})}(\mu)$ with a collection of arcs dual to its vertical seams as above. Then for any $\epsilon > 0$, there is a $\delta > 0$ such that $V'_\delta(X) \subset V_\epsilon(X)$.*

Proof. We first claim that there is a $\delta_1 > 0$ such that for any $g \in V_{\delta_1}(X)$ for closed immersed loop γ in $(X, \boldsymbol{z})_\epsilon$ (where ϵ is small enough that $(X, \boldsymbol{z})_\epsilon$ is the complement of neighborhoods of the punctures and nodes), g preserves the turning number of γ. To see this, choose a finite set of immersed curves $\{\alpha_i\}$ which form a basis of $H_1(T^1(X, \boldsymbol{z})_\epsilon; \mathbb{Z})$. Choose δ_1 small enough that g is C^1-close-enough to the identity that it changes the turning number of each α_i by at most $1/2$. Since closed curves have integral turning numbers, they must be fixed. Since the α_i are a basis of homology, it follows that all closed curves must be fixed.

Now let γ be any good arc which crosses exactly one seam, and let γ_i be the chosen arc which crosses the same seam. We may take arcs β_1, β_2 in $(X, \boldsymbol{z})_\epsilon$ which have bounded length (in terms of the genus of X) and join the endpoints of γ to those of γ_i, so that $\gamma' = \beta_1 + \gamma + \beta_2$ is a good arc which has the same endpoints as γ_i. They then differ by closed curves in the thick part, so have the same turning number. Take δ_2 small enough that g changes the turning number of the β_i by at most $\epsilon/2$. Taking $\delta < \max(\epsilon/2, \delta_1, \delta_2)$ then ensures that g changes the turning number of γ by at most ϵ. ∎

Proof of Proposition 7.5. 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 $g_0 \colon \overline{X}_{0,\boldsymbol{d}_0\cdot\boldsymbol{\sigma}_0} \to \overline{X}_{\boldsymbol{\sigma}}$ and $g_1 \colon \overline{X}_{1,\boldsymbol{d}_1\cdot\boldsymbol{\sigma}_1} \to \overline{X}_{0,\boldsymbol{\sigma}_0}$ be the almost-diffeomorphisms given by the definitions of $V_\epsilon(X)$ and $V_\rho(X_0)$. We define $\boldsymbol{d} = \boldsymbol{d}_1 + \boldsymbol{d}_0$, denote the rescaling function by $\boldsymbol{c}_0 = \boldsymbol{e}(\boldsymbol{d}_0)$ and $\boldsymbol{c}_1 = \boldsymbol{e}(\boldsymbol{d}_1)$, and set $\boldsymbol{c} = \boldsymbol{e}(\boldsymbol{d}_0 + \boldsymbol{d}_1)$.

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 $g = g_0 \circ F_{\boldsymbol{d}_0} \circ g_1 \circ F^{-1}_{\boldsymbol{d}_0} \colon \overline{X}_{1,\boldsymbol{d}_1\cdot\boldsymbol{\sigma}_1} \to \overline{X}_{\boldsymbol{\sigma}}$. Then g 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 function on top level are equal to one and we denote the rescaling function on lower level by c_0, c_1 and $c = c_0 c_1$. 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:

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

Here C_{g_0} is the supremum of the norm of the derivative Dg_0, with respect to the hyperbolic metrics, on the ϵ-thick part. 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. Consequently, the $V_\epsilon(X)$ are indeed a basis of a topology.

It is obvious that a sequence which converges in the topology defined by this basis also converges in the sense of Definition 7.4. Conversely, given any $V_\epsilon(X)$ choose curves $\{\gamma_i\}$ and δ such that $V'_\delta(X) \subset V_\epsilon(X)$. Any sequence converging to X in the sense of Definition 7.4 is then eventually in $V'_\delta(X) \subset V_\epsilon(X)$. ◻

Theorem 7.7. *For any subgroup $G < \mathrm{Mod}_{g,n}$, the quotient $\Omega\overline{\mathcal{T}}_{(\Sigma,\boldsymbol{s})}(\mu)/G$ and its projectivized version $\mathbb{P}\Omega\overline{\mathcal{T}}_{(\Sigma,\boldsymbol{s})}(\mu)/G$ are Hausdorff topological spaces.*

Proof. Suppose that $(X_n, \boldsymbol{z}_n, \boldsymbol{\omega}_n, \preccurlyeq_n, \boldsymbol{\sigma}_n, f_n)$ is a sequence of marked multi-scale differentials which converges to $(X, \boldsymbol{z}, \boldsymbol{\omega}, \preccurlyeq, \boldsymbol{\sigma}, f)$, and suppose moreover we have a sequence $\{\gamma_n\}$ in G so that $(X_n, \boldsymbol{z}_n, \boldsymbol{\omega}_n, \preccurlyeq_n, \boldsymbol{\sigma}_n, f_n\gamma_n^{-1})$ converges to $(X', \boldsymbol{z}', \boldsymbol{\omega}', \preccurlyeq', \boldsymbol{\sigma}', f')$. Let g_n and g'_n be the respective sequences of maps exhibiting this convergence, and let $h_n = g'_n \circ g_n^{-1} \colon \overline{X}_{\boldsymbol{\sigma}} \to \overline{X}'_{\boldsymbol{\sigma}'}$ (strictly speaking, h_n may be only defined on an exhaustion of the complement of the horizontal nodes).

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 maps h_n must converge uniformly on compact sets to an isomorphism $h\colon (X, \boldsymbol{z}) \to (X', \boldsymbol{z}')$ of pointed stable curves. We next show that $\preccurlyeq$ and $\preccurlyeq'$ 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 $X_u \preccurlyeq' X_v$. Since $\preccurlyeq_n$, for n sufficiently large, is an undegeneration of both $\preccurlyeq$ and $\preccurlyeq'$, this is possible only if $X_u \asymp_n X_v$. We denote ℓ and ℓ' some level functions inducing the full orders $\preccurlyeq$ and $\preccurlyeq'$, 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 $||c_{n,\ell(u)}(g_n)_*\boldsymbol{\omega}_n - \omega_u||_\infty < \epsilon_n$ and $||c'_{n,\ell'(u)}(g'_n)_*\boldsymbol{\omega}_n - \omega'_u||_\infty < \epsilon_n$, where ω_u is the restriction of $\boldsymbol{\omega}$ 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 $|c_{n,\ell(u)}/c_{n,\ell(v)}| \to 0$, while on the other hand the hypothesis $X_u \preccurlyeq' X_v$ 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 $\boldsymbol{\omega}$ is equal to $h^*\boldsymbol{\omega}'$, we use that for every level i both inequalities $||c_{n,i}(g_n)_*(\boldsymbol{\omega}_n) - \omega_{(i)}||_\infty < \epsilon_n$ and $||c'_{n,i}(g_n)_*(\boldsymbol{\omega}_n) - h^*\omega'_{(i)}||_\infty < C\epsilon_n$ 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, as desired.

Finally, the maps h_n are asymptotically turning number preserving, so by Proposition 5.12, $h^*\boldsymbol{\sigma}' = \boldsymbol{\sigma}$, and moreover the induced map $\overline{h}\colon \overline{X}_{\boldsymbol{\sigma}} \to \overline{X}'_{\boldsymbol{\sigma}'}$ is eventually isotopic to h_n. We have $h_n f \simeq f'\gamma_n$ for each n, so eventually $hf \simeq f'\gamma_n$, so h exhibits a G-equivalence of X and X', as desired. ◻

Our next goal is to show that $\Omega\overline{\mathcal{T}}_{(\Sigma,\boldsymbol{s})}(\mu)$ is a second countable topological space with countable basis $\{V_\epsilon(X_n)\}$, where $\epsilon \in \mathbb{Q}$ and $\{X_n\}$ is a dense sequence.

Lemma 7.8. *Let $X_n \to X$ be a convergent sequence in a single stratum of $\Omega\overline{\mathcal{T}}_{(\Sigma,\boldsymbol{s})}(\mu)$. For any $\epsilon > 0$, the neighborhoods $V_\epsilon(X_n)$ eventually contain X, and $V_\epsilon(X_n) \subset V_{4\epsilon}(X)$.*

Proof. Let $g_n\colon \overline{X}_n \to \overline{X}_{\boldsymbol{\sigma}}$ be a sequence of maps which exhibit convergence of this sequence as in Definition 7.4. We wish to show that eventually g_n^{-1} exhibits $X \in V_\epsilon(X_n)$. Since X and X_n lie in the same stratum, it suffices to check items (ii), (iii), and (v) in the definition of $V_\epsilon(X)$.

Let N be large enough so that $\epsilon_N < \epsilon$. By Lemma 3.4, the g_n^{-1} converge uniformly to the identity as maps to the universal curve. Since the vertical hyperbolic metric is continuous, the image of $(X,\boldsymbol{z})_{\epsilon_N}$ under g_n^{-1} of eventually contains $(X_n,\boldsymbol{z}_n)_\epsilon$, and moreover

$$\frac{1}{C_n} \le \|Dg_n^{-1}\|_{\epsilon_N,\infty} \le C_n$$

(meaning the sup-norm on the ϵ_N-thick part) for $C_n \searrow 1$, where the norm is defined via the Poincaré metrics.

As a consequence the map g_n is eventually defined on $(X_n,\boldsymbol{z}_n)_\epsilon$, and moreover on $(X_n,\boldsymbol{z}_n)_\epsilon$

$$\|\boldsymbol{\omega}_n - g_n^*\boldsymbol{\omega}\|_{\epsilon,\infty} \le C_n\|(g_n)_*\omega_n - \omega\|_{\epsilon_N,\infty} \to 0,$$

so $\|\boldsymbol{\omega}_n - g_n^*\boldsymbol{\omega}\|_\infty < \epsilon$ eventually.

Finally, g_n changes turning numbers of good arcs as much as g_n^{-1} does, so eventually g_n changes turning numbers of good arcs by at most ϵ, so $X \in V_\epsilon(X_n)$.

Now, suppose $X' \in V_\epsilon(X_n)$ is exhibited by $g_n'\colon \overline{X}'_{\boldsymbol{\sigma}'} \to \overline{X}_n$. We wish to show that $g_n \circ g_n'$ eventually exhibits $X' \in V_{4\epsilon}(X)$. The composition $(g_n \circ g_n')^{-1}$ is eventually conformal on $(X,\boldsymbol{z})_{\epsilon_N}$, and moreover

$$\begin{aligned}\|(g_n \circ g_n')_*\boldsymbol{\omega}' - \boldsymbol{\omega}\|_{\epsilon_N,\infty} &\le C_n\|(g_n')_*\boldsymbol{\omega}' - g_n^*\boldsymbol{\omega}\|_{\epsilon,\infty}\\ &\le C_n\|(g_n')_*\boldsymbol{\omega}' - \boldsymbol{\omega}_n\|_{\epsilon,\infty} + C_n\|\boldsymbol{\omega}_n - g_n^*\boldsymbol{\omega}\|_{\epsilon,\infty}\\ &\le 2C_n\epsilon < 4\epsilon.\end{aligned}$$

◻

Proposition 7.9. *The augmented Teichmüller space $\Omega\overline{\mathcal{T}}_{(\Sigma,\boldsymbol{s})}(\mu)$ is second countable.*

Proof. Each stratum $\Omega\mathcal{B}_\Lambda$ is a complex manifold and thus separable, so $\Omega\overline{\mathcal{T}}_{(\Sigma,\boldsymbol{s})}(\mu)$ is separable as well as it is a countable union of these strata. Let $\{X_n\}$ be a sequence whose intersection with each stratum is dense. We claim that the family $\mathcal{F} = \{V_\epsilon(X_n) : \epsilon \in \mathbb{Q}\}$, is a basis of the topology on $\Omega\overline{\mathcal{T}}_{(\Sigma,\boldsymbol{s})}(\mu)$.

Consider any $X \in \Omega\mathcal{B}_\Lambda$ and $\epsilon > 0$. Take any subsequence $X_{n_k} \to X$ within $\Omega\mathcal{B}_\Lambda$ and rational $\epsilon' < \epsilon/4$. By the previous Lemma, eventually $X \in V_{\epsilon'}(X_{n_k}) \subset V_\epsilon(X)$, so $\mathcal{F}$ is a countable basis, as desired. ◻

We finish this section with the following Proposition which allows us to relax some of the conditions of Definition 7.4.

Proposition 7.10. *For convergence in Definition 7.4, it suffices that the maps g_m satisfy all of the conditions listed, with the following modifications:*

(2') The g_n^{-1} are conformal on an open set U which intersects each irreducible component of X, but may be only K_n-quasiconformal on the rest of the ϵ_n-thick part, where $K_n \to 1$.
(3') These forms converge in the weak locally L^2 topology.
(5') Convergence of turning numbers is only required for simple arcs whose endpoints lie in U.

Proof. Apply Lemma 3.3 to produce a sequence $k_n\colon X \to X$ of K_n-quasiconformal maps, converging uniformly to the identity, such that $\widetilde{g}_n = g_n \circ k_n^{-1}$ is conformal and the restriction of k_n to U is holomorphic. We claim that these maps $\widetilde{g}_n$ satisfy all of the requirements of Definition 7.4.

The uniform convergence of forms required by (3) follows from Proposition 3.5. For (6), note that Lemma 7.6 allows us to consider only turning numbers of simple arcs whose endpoints are contained in U. Since the maps k_n converge uniformly, their derivatives converge as well on U. The turning number of a simple arc can be computed using only its homotopy class (rel endpoints) and its tangent vectors at the endpoints. It follows that (5') implies that the $\widetilde{g}_n$ are asymptotically turning number preserving. □

7.4. **The moduli space of multi-scale differentials as a topological space.** We are now in a position to define our central moduli space as a topological space and establish its main topological properties.

Definition 7.11. The moduli space of multi-scale differentials is the quotient space $\Xi\overline{\mathcal{M}}_{g,n}(\mu) = \Omega\overline{\mathcal{T}}_{(\Sigma,\boldsymbol{s})}(\mu)/\operatorname{Mod}_{g,n}$, and its projectivization is the quotient $\mathbb{P}\Xi\overline{\mathcal{M}}_{g,n}(\mu) = \mathbb{P}\Omega\overline{\mathcal{T}}_{(\Sigma,\boldsymbol{s})}(\mu)/\operatorname{Mod}_{g,n} = \Xi\overline{\mathcal{M}}_{g,n}(\mu)/\mathbb{C}^*$. △

It follows immediately from Theorem 7.7 and Proposition 7.9 that these spaces are Hausdorff and second countable.

Theorem 7.12. *The moduli space $\mathbb{P}\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ of projectivized multi-scale differentials of type μ is compact.*

Proof. In a second countable space, compactness is equivalent to sequential compactness (see [BB02, Proposition 1.6.23]), so it suffices to establish sequential compactness.

Let $\{(X_n, \boldsymbol{z}_n, \boldsymbol{\omega}_n, \preccurlyeq_n, \sigma_n, f_n)\}$ be a sequence in $\mathbb{P}\Omega\overline{\mathcal{T}}_{(\Sigma,\boldsymbol{s})}(\mu)$. We wish to exhibit a convergent subsequence, after pre-composing the markings f_n by a sequence in $\operatorname{Mod}_{g,n}$.

Since $\overline{\mathcal{M}}_{g,n}$ is compact, after pre-composing the markings, we can pass to a subsequence so that $\{(X_n, \boldsymbol{z}_n, f_n)\}$ converges in $\overline{\mathcal{T}}_{g,n}$ to a marked surface $(X, \boldsymbol{z}, f)$. Since there are finitely many enhanced multicurves up to the action of $\operatorname{Mod}_{g,n}$, we may pass to a subsequence so that these surfaces lie in a single stratum $\Omega\mathcal{B}_\Lambda$.

By definition of convergence in $\overline{\mathcal{T}}_{g,n}$, there exists an exhaustion K_n of $X^s \setminus \boldsymbol{z}$, and conformal maps $h_n\colon K_n \to X_n$ compatible with the markings. Transporting the orders on X_n to X by h_n induces a full order on X that we denote by $\preccurlyeq_0$.

Now we use the sizes of the forms $\boldsymbol{\omega}_n$ to refine the order $\preccurlyeq_0$. Choose ϵ small enough that for each irreducible component Y of X, the ϵ-thick part of Y minus the marked points is connected. Let $\lambda_n(Y)$ be the size of the corresponding component of X_n, in the sense of Equation (3.6). Passing to a subsequence, we may assume for any pair of components, the ratios $\lambda_n(Y)/\lambda_n(Y')$ converge in the extended real line. We then define an order $\preccurlyeq$ on the components of X so that $Y \preccurlyeq Y'$ when $Y \preccurlyeq_0 Y'$, and moreover if $Y \asymp_0 Y'$, we define $Y \preccurlyeq Y'$ if $\lambda_n(Y)/\lambda_n(Y') \not\to \infty$. Theorem 3.10 then implies that we may pass to a subsequence so that on each component Y of X the rescaled forms $\{h_n^* \boldsymbol{\omega}_n/\lambda_n(Y)\}$ converges to some form $\boldsymbol{\omega}$.

For each level i of $\preccurlyeq$ we pick a component Y_i at that level and define $c_{n,i} = \lambda(Y_i)^{-1}$. Then the sequence of differentials $\{c_{n,i} h_n^* \boldsymbol{\omega}_n\}$ converges to some differential $\omega_{(i)}$ on the i-th level $X_{(i)}$ of X with respect to $\preccurlyeq$. We define $\boldsymbol{\omega}$ on X to be the collection of those differentials $\omega_{(i)}$. We have to prove that $\boldsymbol{\omega}$ is a twisted differential compatible with the order $\preccurlyeq$. The crucial conditions (matching orders, matching residues and GRC) can be verified by $\boldsymbol{\omega}$-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 functions $c_{n,j}$ depend on the levels only.

For each vertical node q choose a preliminary prong-matching $\widetilde{\sigma_q}$, forming together a global prong-matching $\widetilde{\sigma}$, and choose a preliminary almost-diffeomorphism $\widetilde{g_n}\colon \overline{X}_{\boldsymbol{d}\cdot\widetilde{\sigma}} \to \overline{X}_{\sigma_n}$ which is isotopic to h_n^{-1} on the complement of the seams. Conditions (1)–(4) of Definition 7.4 are clearly satisfied for these g_n and $d_{n,i} = \frac{1}{2\pi i} c_{n,i}$, and it remains to show that the prong-matchings and g may be modified so that the g_n are asymptotically turning number preserving.

Now choose a collection of good arcs γ_i, dual to the seams of $\overline{X}_{\widetilde{\sigma}}$ as in the Remark 5.14. By convergence of the forms, $\tau(g_n^{-1}(\gamma_i)) \to \tau(\gamma_i) \pmod{\mathbb{Z}}$. We first modify each prong-matching so that these turning numbers converge mod κ_i, and then modify g by an appropriate twist around each seam so that the turning numbers converge. By Remark 5.14 convergence of the turning numbers of these arcs ensure the g_n are asymptotically turning number preserving. ◻

8. The model domain

In this section, we construct the *model domain* $\overline{\mathcal{MD}}_\Lambda$, an orbifold which will serve as a local model for the boundary of the moduli space $\Xi\overline{\mathcal{M}}_{g,n}(\mu)$. The model domain is constructed as the finite quotient orbifold of the *simple model domain* $\overline{\mathcal{MD}}_\Lambda^s$, a complex manifold which is in turn constructed as a bundle over a product of Teichmüller spaces. We moreover construct a family of differentials over the model domain we will call the *universal family of model differentials.* These objects will be used in Section 10, where we provide $\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ with an atlas of "plumbing maps", which are defined on open subsets of the model domain by a plumbing construction on the universal family. The families of model differentials over arbitrary bases will be defined by combining the definition of families of multi-scale differentials in Section 11 and the families of markings in Section 12. We establish that the family over the model domain is indeed universal in Section 13.1.

8.1. **Construction of the model domain.** As a first step in the construction, we define

$$\mathcal{MD}^s_\Lambda \;=\; \mathbb{P}\Omega\mathcal{T}^{pm}_\Lambda(\mu)/\mathrm{Tw}^{sv}_\Lambda \quad\text{and}\quad \mathcal{MD}_\Lambda \;=\; \mathbb{P}\Omega\mathcal{T}^{pm}_\Lambda(\mu)/\mathrm{Tw}^{v}_\Lambda\,, \tag{8.1}$$

whose points represent prong-matched twisted differentials defined up to rescaling all components simultaneously, together with a marking defined up to the action of Tw^{sv}_Λ or Tw^{v}_Λ, respectively. Recall from Lemma 6.3 that Tw^{sv}_Λ is a finite index subgroup of Tw^{v}_Λ. The finite quotient group $K_\Lambda = \mathrm{Tw}^{v}_\Lambda/\mathrm{Tw}^{sv}_\Lambda$ defined in Equation (6.14) thus acts on $\mathcal{MD}^s_\Lambda$ with quotient $\mathcal{MD}_\Lambda$.

The model domains $\overline{\mathcal{MD}}_\Lambda$ and $\overline{\mathcal{MD}}^s_\Lambda$ are constructed informally by adding a boundary consisting of differentials which are allowed to be identically zero on some set of levels below the top level.

More precisely, the simple level rotation torus $T^s_\Lambda = \mathbb{C}^{L(\Lambda)}/\mathrm{Tw}^{sv}_\Lambda \cong (\mathbb{C}^*)^{L(\Lambda)}$ acts freely on $\mathcal{MD}^s_\Lambda$ via the action defined in (6.13). Recall from Section 7.2 that the projectivized Λ-boundary stratum $\mathbb{P}\Omega\mathcal{B}_\Lambda$ is the quotient $\mathcal{MD}^s_\Lambda/T^s_\Lambda$, so $\mathcal{MD}^s_\Lambda$ is a principal $(\mathbb{C}^*)^{L(\Lambda)}$-bundle over $\mathbb{P}\Omega\mathcal{B}_\Lambda$. We define $\overline{\mathcal{MD}}^s_\Lambda$ as the associated $\mathbb{C}^{L(\Lambda)}$-bundle over $\mathbb{P}\Omega\mathcal{B}_\Lambda$, where $(\mathbb{C}^*)^{L(\Lambda)}$ acts on $\mathbb{C}^{L(\Lambda)}$ by coordinate-wise multiplication in the usual way.

As the K_Λ-action on $\mathcal{MD}^s_\Lambda$ commutes with the T^s_Λ-action, K_Λ acts on $\overline{\mathcal{MD}}^s_\Lambda$ and we define (as a complex orbifold)

$$\overline{\mathcal{MD}}_\Lambda \;=\; \overline{\mathcal{MD}}^s_\Lambda/K_\Lambda\,.$$

We now provide notation to describe the boundary $\partial\mathcal{MD}^s_\Lambda = \overline{\mathcal{MD}}^s_\Lambda \setminus \mathcal{MD}^s_\Lambda$ of the simple model domain. The boundary $\partial\mathcal{MD}^s_\Lambda$ is a normal crossing divisor given by $D = \cup_{i\in L(\Lambda)} D_i$ in $\overline{\mathcal{MD}}^s_\Lambda$, where D_i is fiber-wise defined by $\{t_i = 0\} \subset \mathbb{C}^{L(\Lambda)}$. There is a stratification

$$\overline{\mathcal{MD}}^s_\Lambda \;=\; \coprod_{J\subset L(\Lambda)} \mathcal{MD}^{s,\Lambda_J}_\Lambda \tag{8.2}$$

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 we define $D_J = D_{j_1}\cap\cdots\cap D_{j_m}$, and $\mathcal{MD}^{s,\Lambda_J}_\Lambda$ is defined by

$$\mathcal{MD}^{s,\Lambda_J}_\Lambda \;=\; D_J \setminus \bigcup_{J'\supsetneq J} D_{J'}\,.$$

In these terms, the space $\mathcal{MD}^s_\Lambda$ corresponds to the subset $J=\emptyset$, or equivalently to the degeneration $\mathrm{dg}\colon \bullet \rightsquigarrow \Lambda$ from the trivial graph to Λ. Moreover, D_i corresponds to the subset $J=\{i\}$, or equivalently to the two-level (un)degeneration $\mathrm{dg}_i\colon \Lambda_i \rightsquigarrow \Lambda$ of Λ.

The model domain has an obvious non-projectivized variant. The quotient $\Omega\mathcal{MD}^s_\Lambda = \Omega\mathcal{T}^{pm}_\Lambda(\mu)/\mathrm{Tw}^{sv}_\Lambda$ is a $(\mathbb{C}^*)^{L(\Lambda)}$-bundle over $\Omega\mathcal{B}_\Lambda$. We define $\Omega\overline{\mathcal{MD}}^s_\Lambda$ as the associated $\mathbb{C}^{L(\Lambda)}$-bundle, and let $\Omega\overline{\mathcal{MD}}_\Lambda = \Omega\overline{\mathcal{MD}}^s_\Lambda/K_\Lambda$.

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

Proposition 8.1. *The simple model domain $\overline{\mathcal{MD}}^s_\Lambda$ is smooth, while $\overline{\mathcal{MD}}_\Lambda$ has only finite quotient singularities.*

Example 8.2. (*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, $\overline{\mathcal{M}D}^s_{\Lambda_2}$ 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 $K_{\Lambda_2} = \mathrm{Ker}(H \to G)$ 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 $uv - z^3 = 0$.

8.2. **The universal family.** We now construct the universal family of model differentials over the model domain. Once we formally define the notion of families of model differentials over arbitrary bases, we will see in Proposition 13.5 that this family is in fact the universal family of model differentials.

Over the open part $\Omega\mathcal{M}D^s_\Lambda$, this universal family of model differentials is simply given by the Tw^{sv}_Λ-quotient of the equisingular family $(\pi\colon \mathcal{X} \to \Omega\mathcal{T}^{pm}_\Lambda(\mu), \boldsymbol{\eta}, \boldsymbol{z}, \boldsymbol{\sigma}, f)$ over $\Omega\mathcal{T}^{pm}_\Lambda(\mu)$, where $\boldsymbol{\omega}$ is a universal relative one-form, $\boldsymbol{\sigma}$ is a family of prong-matchings, $\boldsymbol{z}$ are sections marking the zeros and poles, and f is a family of markings (to be defined precisely in Section 12) up to the group Tw^{sv}_Λ.

As the action of T^s_Λ on $\Omega\mathcal{T}^{pm}_\Lambda(\mu)$ is trivial on the level of underlying curves, the universal curve over $\Omega\mathcal{M}D^s_\Lambda$ is the pullback of the universal curve over the quotient $\mathbb{P}\Omega\mathcal{B}_\Lambda$. It follows that the universal curve over $\Omega\mathcal{M}D^s_\Lambda$ extends to a universal curve $\mathcal{X} \to \Omega\overline{\mathcal{M}D}^s_\Lambda$ as the pullback of this bundle.

Consider an open set $\mathcal{V} \subset \Omega\mathcal{B}_\Lambda$ together with a section $\mathscr{S}\colon \mathcal{V} \to \Omega\mathcal{M}D^s_\Lambda$ of the T^s_Λ-bundle. Let $\mathcal{W} \subset \Omega\overline{\mathcal{M}D}^s_\Lambda$ be the preimage of $\mathcal{V}$ and $\mathcal{X}_\mathcal{W} \to \mathcal{W}$ its universal curve. Informally, a point in $\mathcal{V}$ represents a T^s_Λ-orbit of forms and compatible prong-matchings, and $\mathscr{S}$ represents a holomorphic choice of representative "rescaled forms" $\boldsymbol{\eta}$ and compatible prong-matchings $\boldsymbol{\sigma}$. The section $\mathscr{S}$ determines a trivialization $\mathcal{W} \to \mathcal{V} \times \mathbb{C}^{L(\Lambda)} \times \mathbb{C}^*$, and composing with the projection to $\mathbb{C}^{L(\Lambda)} \times \mathbb{C}^*$ determines a tuple of holomorphic functions $\boldsymbol{t}\colon \mathcal{W} \to \mathbb{C}^{L(\Lambda)} \times \mathbb{C}^*$ which we call *simple rescaling parameters*.

The rescaled differentials $\boldsymbol{\eta}$ can be regarded as a tuple of relative one-forms on $\mathcal{X}_\mathcal{W}$ which do not vanish on any vertical component of the universal curve, and which satisfy $\boldsymbol{t} * \boldsymbol{\eta} = \boldsymbol{\omega}$, where $\boldsymbol{\omega}$ is the universal relative one-form over $\Omega\overline{\mathcal{M}D}^s_\Lambda$, and we recall from (6.13) the definition of the action

$$\boldsymbol{t} * \boldsymbol{\eta} \;=\; \left(\boldsymbol{t}^{\boldsymbol{a}}_{\lceil i \rceil} \cdot \eta_{(i)}\right)_{i \in L(\Lambda)} \;=\; \left(t_{-1}^{a_{-1}} \cdot t_{-2}^{a_{-2}} \dots t_i^{a_i} \cdot \eta_{(i)}\right)_{i \in L(\Lambda)} \;=\; \left(s_{-1} \cdot \dots \cdot s_i \cdot \eta_{(i)}\right)_{i \in L(\Lambda)}, \tag{8.3}$$

where we define the *rescaling parameters* $\boldsymbol{s}$ to be the powers $(s_i) = (t_i^{a_i})$. The prong-matchings $\boldsymbol{\sigma}$ can be regarded as a continuously varying family of prong-matchings for $\boldsymbol{\eta}$. Together $\boldsymbol{\eta}$, $\boldsymbol{\sigma}$, and the markings give each fiber of $\mathcal{X}_\mathcal{W}$ the structure of a marked multi-scale differential.

To summarize, we have defined locally, depending on a choice of section $\mathscr{S}\colon \mathcal{V} \to \Omega\mathcal{M}D^s_\Lambda$, which we will in the sequel call a *local trivialization* of $\Omega\mathcal{M}D^s_\Lambda$ over $\mathcal{V}$, a collection of rescaled differentials $\boldsymbol{\eta}$, compatible prong-matchings $\boldsymbol{\sigma}$, holomorphic functions $\boldsymbol{t}$, and rescaling parameters $\boldsymbol{s}$ such that the product $\boldsymbol{t}*\boldsymbol{\eta}$ agrees with the universal one-form $\boldsymbol{\omega}$ on the universal curve over $\Omega\overline{\mathcal{M}D}^s_\Lambda$.

The product of $\boldsymbol{t}$ with the projection to $\Omega\mathcal{B}_\Lambda$ determines an isomorphism $\mathcal{W} \to \mathcal{V} \times \mathbb{C}^{L(\Lambda)} \times \mathbb{C}^*$. In working with model differentials, we will often implicitly assume we have chosen such a local trivialization. In the sequel, the functions $\boldsymbol{t}$ together with local coordinates on $\mathcal{V}$ will give a convenient system of local coordinates on $\Omega\mathcal{M}D_\Lambda$. In Section 11, the equivalence class of $(\boldsymbol{\eta}, \boldsymbol{\sigma}, \boldsymbol{t})$ will be part of the data describing a general family of model differentials.

8.3. **Topology of $\Omega\mathcal{M}D_\Lambda$.** The topology on the model domain may also be expressed in the language of conformal maps. The following proposition follows immediately from the definition of the 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 $\boldsymbol{t} \in (\mathbb{C})^{L(\Lambda)}$, we define $J(\boldsymbol{t}) \subseteq L(\Lambda)$ to be the subset of indices i such that $t_i = 0$.

Proposition 8.3. *A sequence $(X_m, \boldsymbol{z}_m, \boldsymbol{\eta}_m, \boldsymbol{t}_m, \preccurlyeq, \sigma_m, f_m)$ of (simple) model differentials in $\Omega\overline{\mathcal{M}D}^s_\Lambda$ converges to $(X, \boldsymbol{z}, \boldsymbol{\eta}, \boldsymbol{t}, \preccurlyeq, \sigma, f)$ 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 $\boldsymbol{d}_m = \{d_{m,j}\}_{j\in J(\boldsymbol{t})} \in \mathbb{C}^{J(\boldsymbol{t})}$ such that the following conditions hold for sufficiently large m, where we let $d_{m,0} = 0$ and denote $c_{m,j} = \boldsymbol{e}(d_{m,j})$:*

- *(i) There is an inclusion $\iota_m \colon J(\boldsymbol{t}_m) \hookrightarrow J(\boldsymbol{t})$.*
- *(ii) For sufficiently large m there exists an almost-diffeomorphism $g_m \colon \overline{X}_{\boldsymbol{d}_m\cdot\boldsymbol{\sigma}_m} \to \overline{X}_{\boldsymbol{\sigma}}$ that is compatible with the markings (in the sense that $g_m \circ \boldsymbol{d}_m \cdot f_m$ is isotopic to f rel marked points) and such that g_m^{-1} is conformal on the ϵ_m-thick part $(X, \boldsymbol{z})_{\epsilon_m}$.*
- *(iii) The restriction of $(g_m)_*(c_{m,i}\boldsymbol{\omega}_m)$ to the ϵ_m-thick part of the level i subsurface of $(X, \boldsymbol{z})$ converges uniformly to $\omega_{(i)}$.*
- *(iv) For any $i, j \in J(\boldsymbol{t})$ with $i > j$ and any subsequence along which $[j,i] \cap \operatorname{im}(\iota_m) = \emptyset$, we have*
$$\lim_{m\to\infty} \frac{|c_{m,i}|}{|c_{m,j}|} \;=\; 0\,.$$
- *(v) The almost-diffeomorphisms g_m are asymptotically turning number preserving.*

9. Modifying differentials and perturbed period coordinates

The first goal of this section is to define *modifying differentials* $\boldsymbol{\xi}$ as a preparation for the plumbing construction in Section 10, which will enable us to give complex charts on $\Xi\overline{\mathcal{M}}_{g,n}(\mu)$. The second goal is to define local coordinates, which we will call *perturbed period coordinates*, on the simple model domain. Once we define the plumbing construction and define families of multi-scale differentials, the universal property of the family of model differentials over the model domain will allow us to prove the universal property of $\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ in Section 13.

For defining perturbed period coordinates, in this section we restrict to the case with only vertical nodes. In Section 10.10 we will define extended perturbed period coordinates, to also account for the periods through horizontal nodes. This extension will require the plumbing setup introduced in Section 10. 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 $\boldsymbol{\eta}$, but perturbed by the modifying differentials $\boldsymbol{\xi}$, and rescaled by $\boldsymbol{t}$ as defined in (8.3). The reason for this is that the perturbed differential lives on the universal family over the Dehn space $\Xi\mathcal{D}_\Lambda$, which will be defined after plumbing. Consequently, once the plumbing construction is completed, perturbed period coordinates will turn out to be 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 $\Omega\overline{\mathcal{M}D}^s_\Lambda$ the difference between the perturbed period and the original period tends to zero.

Third, the perturbed period 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 $\boldsymbol{\eta}$.

In the whole of this section we work in the preimage $\mathcal{W} \subset \Omega\overline{\mathcal{M}D}^s_\Lambda$ of an open set $\mathcal{V} \subset \Omega\mathcal{B}_\Lambda$ where we have chosen a local trivialization as in Section 8.2.

9.1. **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. In this section, we prove the existence of such families of modifying differentials, for families that may have both horizontal and vertical nodes.

Definition 9.1. A *family of modifying differentials* over $\mathcal{W} \subset \Omega\overline{\mathcal{M}D}^s_\Lambda$ is a family of meromorphic differentials $\boldsymbol{\xi}$ on $\pi\colon \mathcal{X} \to \mathcal{W}$, such that:

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

In other words, denote $\widetilde{\mathcal{X}} \to \mathcal{X}$ the partial normalization at the vertical nodes and denote $\widetilde{\pi}\colon \widetilde{\mathcal{X}} \to \mathcal{W}$ its composition with π. Recall that $q_e^\pm\colon \mathcal{W} \to \widetilde{\mathcal{X}}$ denote 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 $\mathcal{Z}^\infty$. Then $\boldsymbol{t} * \boldsymbol{\xi}$ is a holomorphic section of

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

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

$$\mathrm{Res}_{q_e^+}\,\boldsymbol{t} * (\boldsymbol{\eta} + \boldsymbol{\xi}) + \mathrm{Res}_{q_e^-}\,\boldsymbol{t} * (\boldsymbol{\eta} + \boldsymbol{\xi}) \;=\; 0 \tag{9.1}$$

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

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

$$0 \;=\; V_{-N-1} \,\subseteq\, V_{-N} \,\subseteq\, \ldots \,\subseteq\, V_{-1} \;=\; V\,, \tag{9.2}$$

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* $\rho\colon \Lambda^v \to \mathbb{C}$. With the help of these maps we give an alternative statement of the global residue condition.

Proposition 9.2 ([BCGGM19, Proposition 6.3])**.** *A residue assignment $\rho\colon \Lambda^v \to \mathbb{C}$ satisfies the global residue condition if and only if there exist* period homomorphisms

$$\rho_i\colon V_i/V_{i-1} \to \mathbb{C} \quad \textit{for any } i \in L(\Lambda)\,,$$

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 $\rho_i\colon V_i \to \mathbb{C}$ 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 9.3. *The family $\pi\colon \mathcal{X} \to \mathcal{W}$ equipped with the universal differential $\boldsymbol{t} * \boldsymbol{\eta}$ has a family of modifying differentials $\boldsymbol{\xi}$.*

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

$$\rho_i\colon \mathcal{W} \to \operatorname{Hom}_{\mathbb{Q}}(V_i, \mathbb{C})\,,$$

such that ρ_i restricts to zero on V_{i-1}. In period coordinates, ρ_i is simply a linear projection. By (6.13), the map ρ_i is T^s_Λ-equivariant, i.e.

$$\rho_i(\boldsymbol{q} * (\mathcal{X}, \boldsymbol{t} * \boldsymbol{\eta})) \;=\; \prod_{j\geq i} q_j^{a_j} \,\cdot\, \rho_i(\mathcal{X}, \boldsymbol{t} * \boldsymbol{\eta}) \quad \text{for any} \quad \boldsymbol{q} \in T^s_\Lambda\,. \tag{9.3}$$

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 $\widetilde{\rho_i}\colon \mathcal{W} \to \operatorname{Hom}_{\mathbb{Q}}(V', \mathbb{C})$ of ρ_i by the requirement $\widetilde{\rho_i}(b_i) = 0$ for all $b_i \in B_i$.

Since $\Lambda_{\max}$ is a maximal multicurve on $\Sigma \setminus P_{\boldsymbol{s}}$, a meromorphic form on $\mathcal{X}$, holomorphic except for at worst simple poles at the nodes and at the marked poles $P_{\boldsymbol{s}}$, is specified uniquely by its periods on V'. We define $\boldsymbol{t} * \boldsymbol{\xi}$ on $\mathcal{X}$ so that its V'-periods for $\gamma \in \Lambda_{\max}$ are

$$\int_\gamma \xi_{(i)}(u) \;=\; \sum_{j<i} \widetilde{\rho}_j(u)(\gamma) \quad \text{for all} \quad u \in \mathcal{W}\,. \tag{9.4}$$

By the equivariance (9.3), we see that $\widetilde{\rho}_j$ is divisible by $\boldsymbol{t}^{\boldsymbol{a}}_{\lceil i-1 \rceil}$, and hence $\boldsymbol{t} * \boldsymbol{\xi}$ is also divisible by $\boldsymbol{t}^{\boldsymbol{a}}_{\lceil i-1 \rceil}$.

Given a curve $\gamma \in \Lambda^v$ joining level i to level $j < i$, we verify the opposite residue condition (iii) for a modifying differential, which is given by (9.1), and states that

$$\int_\gamma ((\boldsymbol{t} * (\boldsymbol{\eta} + \boldsymbol{\xi}))_{(j)} \;=\; \rho_j(\gamma) + \sum_{k<j} \widetilde{\rho}_k(\gamma) \;=\; \sum_{k<i} \widetilde{\rho}_k(\gamma) \;=\; \int_\gamma ((\boldsymbol{t} * (\boldsymbol{\eta} + \boldsymbol{\xi}))_{(i)}\,.$$

In the above the first equality follows from the fact that ρ_j is the period map determined by $(\boldsymbol{t} * \boldsymbol{\eta})_{(j)}$, and from the definition of $\xi_{(j)}$. The second equality follows from the global residue condition of $\boldsymbol{t} * \boldsymbol{\eta}$ as restated in Proposition 9.2, which implies that $\rho_k(\gamma) = 0$ for all $j < k \le i$. The last equality again follows from the definition of $\xi_{(i)}$ and the fact that $\rho_i(\gamma) = 0$. ∎

This proof shows in particular the following.

Corollary 9.4. *The modifying differential $\boldsymbol{\xi}$ is uniquely determined by the choice of the subspace $V' \subset H_1(\Sigma \setminus P_{\boldsymbol{s}}; \mathbb{Q})$ and the multicurves B_i. Its level-wise components $\xi_{(i)}$ depend only on t_j and η_j for $j < i$.*

Example 9.5. We illustrate the objects introduced in the proof of Proposition 9.3 in the context of a slight simplification of our running example, as pictured in Figure 4, with one pole denoted by p (so the level graph is still a triangle, but the irreducible components are simpler). The family of modifying differentials $\boldsymbol{\xi}$ depends on the parameters $\boldsymbol{t} = (t_0, t_{-1}, t_{-2})$.

The vertical multicurve Λ^v (in blue in Figure 4) 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 $\rho_{-2}(\boldsymbol{t})(\lambda_1) = t_{-2}^{a_{-2}} \cdot a \in \mathbb{C}$ and $\rho_{-1}(\boldsymbol{t})(\lambda_1) = 0$, where a is $2\pi\sqrt{-1}$ times the residue of η at the corresponding node. We choose the maximal multicurve $\Lambda_{\max} \supset \Lambda$ by adding the curves $\{\lambda_4, \dots, \lambda_8\}$, 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 $\widetilde{\rho}_{-2}$ of the map ρ_{-2} defined on $V_{-2} = \langle \lambda_1 \rangle$ is given by requiring $\widetilde{\rho}_{-2}(\boldsymbol{t})(B_{-2}) = 0$. That is,

$$\widetilde{\rho}_{-2}(\boldsymbol{t})(\lambda) = \begin{cases} t_{-2}^{a_{-2}} \cdot a\,, & \text{if } \lambda \;=\; \lambda_1\,, \\ 0\,, & \text{if } \lambda \;=\; \lambda_4, \lambda_7\,. \end{cases}$$

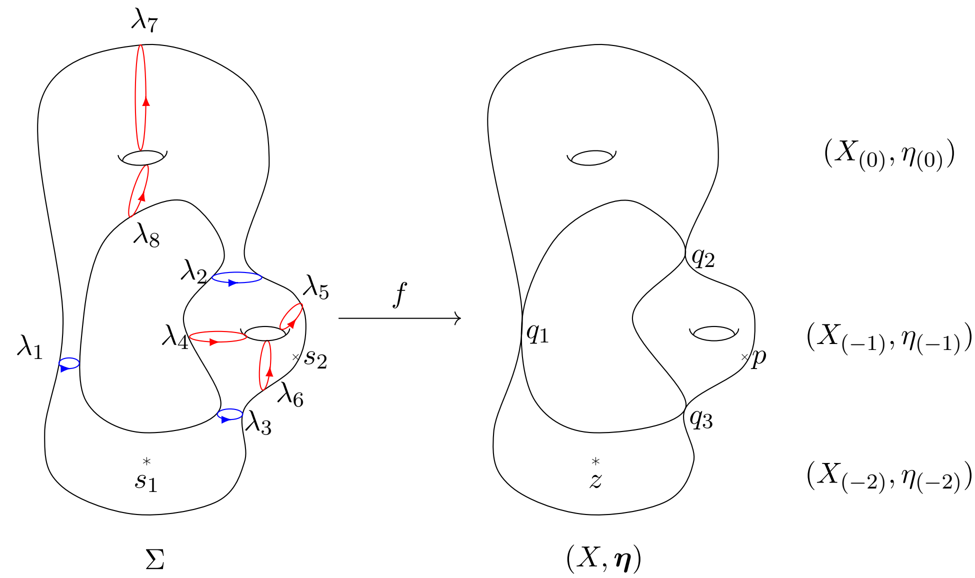

FIGURE 4. The marked surface together with the multicurves Λ and $\Lambda_{\max}$.

Similarly, the extension $\widetilde{\rho}_{-1}$ of ρ_{-1} from $V_{-1} = \langle\lambda_1\rangle$ to V' is defined by requiring $\widetilde{\rho}_{-1}(\boldsymbol{t})(B_{-1}) = 0$, hence $\widetilde{\rho}_{-1}$ 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 $\pm t_{-2}^{a_{-2}} \cdot a/2\pi\sqrt{-1}$, is holomorphic at p, has period zero over λ_4, and has periods $t_{-2}^{a_{-2}} \cdot a$ over λ_5 and λ_6. Finally, the differential $\xi_{(0)}$ has simple poles at q_1 and q_2, with residues $\pm t_{-1}^{a_{-1}} t_{-2}^{a_{-2}} \cdot a/2\pi\sqrt{-1}$, has period zero over λ_7, and period $t_{-1}^{a_{-1}} t_{-2}^{a_{-2}} \cdot a$ over λ_8. To see these, consider for example the period of $\xi_{(0)}$ at λ_8. By definition it is given by

$$\int_{\lambda_8} \xi_{(0)} \;=\; \widetilde{\rho}_{-2}(\lambda_8) + \widetilde{\rho}_{-1}(\lambda_8) \;=\; \rho_{-2}(\lambda_1) + \rho_{-1}(\lambda_1) \;=\; t_{-2}^{a_{-2}} \cdot a + 0 \;=\; t_{-2}^{a_{-2}} \cdot a\,,$$

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

9.2. **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 10.10.

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].

Recall that $\Sigma^c_{(i)}$ was defined in the paragraph preceding Definition 5.1 as the subsurface of Σ at level i where the boundary curves have been collapsed to points. The Teichmüller markings up to twist group of the welded surfaces in the model domain induce markings $f_i\colon \Sigma^c_{(i)} \to \mathcal{X}_{(i)}$ of the families of connected components of the subsurfaces at level i. Denote by $P_{\boldsymbol{s},i}$ and $Z_{\boldsymbol{s},i}$ those marked poles and zeros that lie on the compact level i subsurface $\Sigma^c_{(i)}$. We denote by $Q^{\pm}_{E,i}$ the set of those zero and pole sections on $\Sigma^c_{(i)}$ mapping to the preimages of the nodes E. We define the sets of points

$$P_i = P_{\boldsymbol{s},i} \cup Q^-_{E,i} \quad \text{and} \quad Z_i = Z_{\boldsymbol{s},i} \cup Q^+_{E,i} \tag{9.5}$$

for each level i.

The perturbed period coordinates are roughly the product of the coordinates t_i and coordinates of the projectivization of certain subspaces $\mathcal{R}^{\mathrm{grc}}_i$ of $H^1(\Sigma^c_{(i)} \setminus P_i, Z_i; \mathbb{C})$. Coordinates on the latter are as usual given by all but one of the periods.

To define $\mathcal{R}^{\mathrm{grc}}_i$ we start with the map $H^1(\Sigma^c_{(i)} \setminus P_i, Z_i; \mathbb{C}) \to \mathbb{C}^{|P_i|}$ 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 $R^{\mathrm{grc}}_i \subset \mathbb{C}^{|P_i|}$ be the subspace cut out further by the matching residue condition at the horizontal nodes, and the global residue condition, as stated in Section 2.4. The *GRC space* $\mathcal{R}^{\mathrm{grc}}_i \subseteq H^1(\Sigma^c_{(i)} \setminus P_i, Z_i; \mathbb{C})$ is then defined as the preimage of R^{grc}_i. If we denote by H the number of horizontal nodes of Λ, then [BCGGM19, Theorem 6.1] can be restated as follows.

Proposition 9.6. *The (open) simple model domain $\Omega\mathcal{M}D^s_\Lambda$ is locally modeled on the sum of the GRC spaces $\oplus_i \mathcal{R}^{\mathrm{grc}}_i$. This space has dimension*

$$\sum_{i \in L^\bullet(\Lambda)} \dim(\mathcal{R}^{\mathrm{grc}}_i) \;=\; \dim \Omega\mathcal{M}_{g,n}(\mu) - H\,.$$

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 $\boldsymbol{z}(h)$ the corresponding section of $\mathcal{X} \to \mathcal{W}$. We choose nearby sections $\sigma^+_e\colon \mathcal{W} \to \mathcal{X}$ and $\sigma_h\colon \mathcal{W} \to \mathcal{X}$ so that

$$\int_{q^+(e)}^{\sigma^+_e(w)} \eta_{(i)} \;=\; \text{const} \quad \text{and} \quad \int_{\boldsymbol{z}(h)}^{\sigma_h(w)} \eta_{(j)} \;=\; \text{const}\,, \tag{9.6}$$

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 $\boldsymbol{t} * (\boldsymbol{\eta} + \boldsymbol{\xi})$ on $\mathcal{X}$ may no longer have a zero of the prescribed order at $\boldsymbol{z}(h)$ because of the modifying differential $\boldsymbol{\xi}$. In the process of plumbing in Section 10.4, we will describe a local surgery of $\mathcal{X}$ in a neighborhood of the sections $\boldsymbol{z}(h)$ corresponding to the half-edges h, such that the images of the sections $\boldsymbol{z}(h)$ are untouched by the surgery, and the extension of $\boldsymbol{t}*(\boldsymbol{\eta}+\boldsymbol{\xi})$ to the resulting family again has a zero of order $\mathrm{ord}_{\boldsymbol{z}(h)} \boldsymbol{\eta}$ along a section that we still denote by $\boldsymbol{z}(h)\colon \mathcal{W} \to \mathcal{X}$.

Finally, we can now define the *perturbed period map at level* i under the hypothesis that there are no horizontal nodes. We fix homology classes $\gamma_1, \ldots, \gamma_{n(i)}$ such that their periods $\int_{\gamma_j} \eta_{(i)}$ form a basis of $\mathcal{R}_i^{\mathrm{grc}}$. 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 $\mathcal{R}_i' \subset \mathcal{R}_i^{\mathrm{grc}}$ the codimension one subspace generated by the periods of $\gamma_1, \dots, \gamma_{n(i)-1}$ for all levels $i < 0$, and we let $\mathcal{R}_0' = \mathcal{R}_0^{\mathrm{grc}}$, 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

$$\mathrm{PPer}_i \colon \begin{cases} \mathcal{W} & \to \ \mathcal{R}_i' \,, \\ [(X, \boldsymbol{\eta}, \boldsymbol{t})] & \mapsto \ \left(\int_{\gamma_j} \eta_{(i)} + \xi_{(i)} \right)_{j=1}^{n'(i)} \,. \end{cases} \tag{9.7}$$

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 $w = [(X, \boldsymbol{\eta}, \boldsymbol{t})]$, rather than from the nearby zeros of $\eta_{(i)}$.

Proposition 9.7. *The perturbed period map*

$$\mathrm{PPer} \colon \mathcal{W} \to \mathbb{C}^{L^\bullet(\Lambda)} \times \bigoplus_{i \in L^\bullet(\Lambda)} \mathcal{R}_i' \,, \qquad [(X, \boldsymbol{\eta}, \boldsymbol{t})] \mapsto \Big(\boldsymbol{t}\,; \oplus_{i \in L^\bullet(\Lambda)} \mathrm{PPer}_i \Big) \tag{9.8}$$

is open and locally injective on a neighborhood of the most degenerate stratum $\mathcal{W}_\Lambda = \cap_{i \in L(\Lambda)} D_i$ inside of $\mathcal{W}$.

Proof. We need to show that the derivative of PPer is surjective along the boundary stratum $\mathcal{W}_\Lambda$, since surjectivity is an open condition, and since this surjectivity implies openness. At a point of $\mathcal{W}_\Lambda$ 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 along the most degenerate stratum $\mathcal{W}_\Lambda$, the integral of $\xi_{(i)}$ is identically zero, because by definition $\xi_{(i)}$ is divisible by $t_{i-1}^{a_{i-1}}$, and $\mathcal{W}_\Lambda$ is defined by the equations $t_j = 0$ 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 $\oplus_{i \in L^\bullet(\Lambda)} \mathcal{R}_i'$ by Proposition 9.6 (under the assumption of no horizontal nodes), surjectivity of the derivative of PPer implies injectivity of the derivative map at any boundary point in $\mathcal{W}_\Lambda$, and hence local injectivity in some neighborhood. ∎

Example 9.8. 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 disc $\Delta^2 = \Delta_{t_{-1}} \times \Delta_{t_{-2}}$ 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

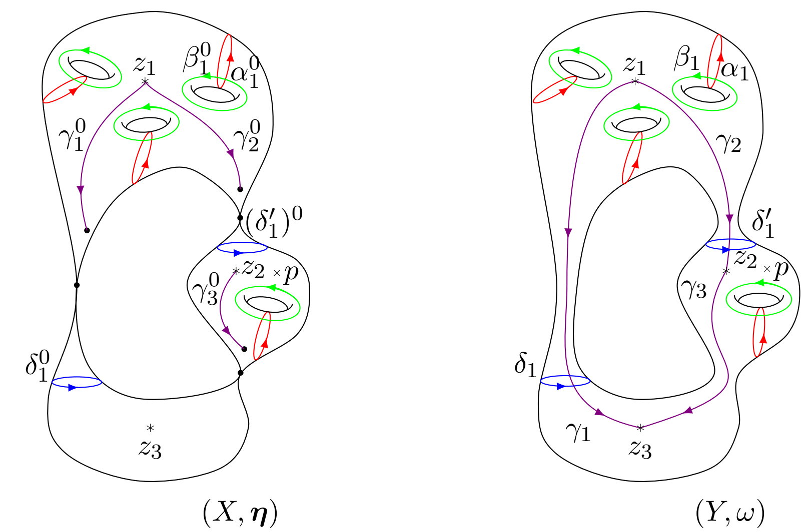

FIGURE 5. A basis of homology of our running example and of a nearby smooth differential.

differentials $\boldsymbol{\xi}$ 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, $\boldsymbol{\xi}$ vanishes identically on $X_{(-2)}$. In Figure 5 we show a basis of the cycles of integration before and after the plumbing construction described later in Section 10. 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 $j = 1, \ldots, 4$ 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 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 $\boldsymbol{t} * \boldsymbol{\eta}$ 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 consistent with the GRC.

10. The Dehn space and the complex structure

In order to understand the structure of $\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ at the boundary, we introduce an auxiliary space $\Xi\mathcal{D}^s_\Lambda$, the *simple Dehn space*, which is a direct analogue of the classical Dehn space. The goal of this section is to give, for each Λ, the topological space $\Xi\mathcal{D}^s_\Lambda$ the structure of a complex manifold, which will then be used to give $\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ its complex structure. This complex structure is induced by *plumbing maps* $\Omega\,\mathrm{Pl}\colon U \to \Xi\mathcal{D}^s_\Lambda$, defined by a local plumbing construction on the universal family of model differentials, which we will show give an atlas of complex coordinate charts on the simple Dehn space.

There is a natural open forgetful map $\Xi\mathcal{D}^s_\Lambda \to \Xi\overline{\mathcal{M}}_{g,n}(\mu)$, and $\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ is covered by the images of these maps as Λ ranges over all enhanced multicurves. The conclusion of this section is summarized in Theorem 10.3, where we use the complex structures on the $\Xi\mathcal{D}^s_\Lambda$ to give $\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ the structure of a smooth complex orbifold.

Throughout this section, we fix an enhanced multicurve Λ with dual graph Γ having $N+1 = |L^\bullet(\Lambda)|$ levels and H horizontal nodes.

10.1. **The Dehn space.** Informally, $\Xi\mathcal{D}_\Lambda$ is the moduli space of Tw_Λ-marked multi-scale differentials of type (μ,Λ'), where Λ' is any undegeneration of Λ. The simple Dehn space $\Xi\mathcal{D}^s_\Lambda$ is the analogous space of Tw^s_Λ-marked differentials.

More formally, the *Dehn space associated with* Λ is the topological space

$$\Xi\mathcal{D}_\Lambda \;=\; \Big(\coprod_{\Lambda'\rightsquigarrow\Lambda} \Omega\mathcal{B}_{\Lambda'}\Big)/\mathrm{Tw}^v_\Lambda\,, \tag{10.1}$$

where we endow this disjoint union with the subspace topology induced from the topology of the augmented Teichmüller space of flat surfaces $\Omega\overline{\mathcal{T}}_{(\Sigma,\boldsymbol{s})}(\mu)$ and recall from (7.1) that $\Omega\mathcal{B}_{\Lambda'}$ are boundary strata in $\Omega\overline{\mathcal{T}}_{(\Sigma,\boldsymbol{s})}(\mu)$). We can write the space equivalently as

$$\Xi\mathcal{D}_\Lambda \;=\; \coprod_{\Lambda'\rightsquigarrow\Lambda} \Xi\mathcal{D}^{\Lambda'}_\Lambda \quad \text{where} \quad \Xi\mathcal{D}^{\Lambda'}_\Lambda \;=\; \Omega\mathcal{B}_{\Lambda'}/\mathrm{Tw}^v_\Lambda\,. \tag{10.2}$$

The *simple Dehn space* is defined by

$$\Xi\mathcal{D}^s_\Lambda \;=\; \Big(\coprod_{\Lambda'\rightsquigarrow\Lambda} \Omega\mathcal{B}_{\Lambda'}\Big)/\mathrm{Tw}^{sv}_\Lambda \;=\; \coprod_{\Lambda'\rightsquigarrow\Lambda} \Xi\mathcal{D}^{\Lambda',s}_\Lambda\,, \tag{10.3}$$

where $\Xi\mathcal{D}^{\Lambda',s}_\Lambda \;=\; \Omega\mathcal{B}_{\Lambda'}/\mathrm{Tw}^{sv}_\Lambda$. The *simple vertical Dehn space* $\Xi\mathcal{D}^{sv}_\Lambda \subset \Xi\mathcal{D}^s_\Lambda$ is the locus consisting of surfaces where every horizontal edge of Λ corresponds to a horizontal node. In other words,

$$\Xi\mathcal{D}^{sv}_\Lambda = \coprod_{\Lambda'\rightsquigarrow\Lambda} \Xi\mathcal{D}^{\Lambda',s}_\Lambda\,,$$

where the union is over all *vertical* undegenerations $\Lambda' \rightsquigarrow \Lambda$.

The finite group $K_\Lambda = \mathrm{Tw}^v_\Lambda/\mathrm{Tw}^{sv}_\Lambda$ acts on $\Xi\mathcal{D}^s_\Lambda$ with topological quotient $\Xi\mathcal{D}_\Lambda$. We will see that $\Xi\mathcal{D}^s_\Lambda$ is in fact a smooth manifold with quotient orbifold $\Xi\mathcal{D}_\Lambda$.

We write similarly $\mathbb{P}\Xi\mathcal{D}_\Lambda$ and $\mathbb{P}\Xi\mathcal{D}^s_\Lambda$ 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 Dehn space (resp. simple Dehn space) as (the moduli point of the equivalence class of) a marked multi-scale differential $(Y,\boldsymbol{z},\boldsymbol{\omega},\boldsymbol{\sigma},f)$ where the marking is up to the action of Tw_Λ (resp. Tw_Λ^s). Pointwise this is justified by definition.

We now outline the *plumbing construction*. We divide the construction into two steps, *vertical plumbing* and *horizontal plumbing*. The vertical plumbing construction starts with the universal curve $\mathcal{X} \to \Omega\overline{\mathcal{M}D}_\Lambda^s$, which we restrict to a neighborhood $\mathcal{W}_\epsilon$ of a point P in the deepest boundary stratum. By cutting out neighborhoods of the vertical nodes and gluing in standard plumbing fixtures, we construct a new family of curves $\mathcal{Y}^v \to \mathcal{W}_\epsilon$ whose generic fiber has only horizontal nodes.

The horizontal plumbing construction requires an extra complex parameter for each horizontal node, parameterizing the modulus and a twist parameter for the annulus that is glued in. We consider $\mathcal{Y}$ as a family over the product $\mathcal{W}_\epsilon \times \Delta^H$ (which does not depend on the second factor). By cutting out a neighborhood of each horizontal node and gluing in a standard plumbing fixture, with parameter given by the second factor, we then construct a new generically smooth family of curves $\mathcal{Y} \to \mathcal{W}_\epsilon \times \Delta^H$.

We equip our standard plumbing fixtures with families of one-forms and choose our gluing maps to identify these forms with those on the target, so that the family $\mathcal{Y}$ comes with a degenerating family of one-forms ω. In fact, with more care we give the fibers of $\mathcal{Y}$ the structure of Tw_Λ^s-marked multi-scale differentials.

This horizontal plumbing construction is essentially the standard plumbing construction used to construct coordinates near the boundary of $\overline{\mathcal{M}}_{g,n}$, dating back at least to Bers [Ber74b]). We emphasize that the vertical construction differs from the usual plumbing in that it does not require extra parameters to describe opening up the nodes—this is rather prescribed by the relative size of the differentials, so that they would glue on the plumbed surface. These plumbing constructions are not canonical and depend on choices made at several points in the construction.

If the universal property for $\Xi\mathcal{D}_\Lambda^s$ were available, the plumbed family $\mathcal{Y} \to \mathcal{W}_\epsilon \times \Delta^H$ would give a holomorphic map $\mathcal{W}_\epsilon \times \Delta^H \to \Xi\mathcal{D}_\Lambda^s$. 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 the universal property. Instead, we define a *plumbing map* $\Omega\,\mathrm{Pl}\colon \mathcal{W}_\epsilon \times \Delta^H \to \Xi\mathcal{D}_\Lambda^s$ stratum-by-stratum, using the universal property for the boundary strata $\Omega\mathcal{B}_{\Lambda'}$ parameterizing equisingular loci in the augmented Teichmüller space. Similarly, the family $\mathcal{Y}^v \to \mathcal{W}_\epsilon$ will give rise to a *vertical plumbing map* $\Omega\,\mathrm{Pl}^v\colon \mathcal{W}_\epsilon \to \Xi\mathcal{D}_\Lambda^{sv}$. As the plumbing constructions are not canonical, neither are these plumbing maps, as they depend on several choices.

In Section 10.3 below, we use the normal forms from Section 4 to construct the gluing maps used in the vertical plumbing construction. In Section 10.4, we define a vertical plumbing construction and a vertical plumbing map. In Sections 10.6 and 10.7, we show that this map is a local homeomorphism. In Section 10.8, we introduce the horizontal plumbing construction and the full plumbing map, and show that it is a local homeomorphism, which yields the following main results of this section.

Theorem 10.1. *For any point P in the deepest stratum $\Omega\mathcal{M}\mathcal{D}_\Lambda^{\Lambda,s}$ of the model domain $\Omega\overline{\mathcal{M}D}_\Lambda^s$, there exists a neighborhood $\mathcal{W}_\epsilon \times \Delta^H$ of $P \times \mathbf{0} \in \Omega\overline{\mathcal{M}D}_\Lambda^s \times \Delta^H$, and a plumbing map*

$$\Omega\,\mathrm{Pl}\colon \mathcal{W}_\epsilon \times \Delta^H \to \Xi\mathcal{D}_\Lambda^s\,,$$

which is a local homeomorphisms. This map preserves the stratifications (8.2) *and* (10.3), *is holomorphic on each stratum, and is K_Λ-equivariant. Moreover, the plumbing map* $\Omega\,\mathrm{Pl}$ *can be chosen to be $\mathbb{C}^*$-equivariant, and thus to descend to a plumbing map*

$$\Omega\,\mathrm{Pl}\colon (\mathcal{W}_\epsilon \times \Delta^H)/\mathbb{C}^* \to \mathbb{P}\Xi\mathcal{D}^s_\Lambda$$

which is also holomorphic, stratum-preserving, and K_Λ-equivariant.

Note that the deepest stratum of the model domain $\Omega\overline{\mathcal{M}D}^s_\Lambda$ can be canonically identified with the deepest stratum of the corresponding Dehn space (and we implicitly do so throughout this section).

Theorem 10.2. *The collection of all plumbing maps gives an atlas of charts which makes $\Xi\mathcal{D}^s_\Lambda$ and the projectivized version $\mathbb{P}\Xi\mathcal{D}^s_\Lambda$ a complex manifold. Moreover, for each point of these spaces the plumbing construction provides a corresponding multi-scale differential.*

The spaces $\Xi\mathcal{D}_\Lambda$ and $\mathbb{P}\Xi\mathcal{D}_\Lambda$ have the structure of complex analytic spaces (or orbifolds) with at worst abelian quotient singularities.

We will discuss the universal properties of the Dehn spaces in Section 13.

We now collect some of the properties already proved, which proves the first of our main results, Theorem 1.2 except for item (4). Recall that a divisor in an orbifold is said to be normal crossing if it is the image of a normal crossing divisor in orbifold chart.

Theorem 10.3. *The moduli space of multi-scale differentials is a complex orbifold $\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ containing $\Omega\mathcal{M}_g$ as an open dense suborbifold with complement a normal crossing boundary divisor. The quotient $\mathbb{P}\Xi\overline{\mathcal{M}}_{g,n}(\mu) = \Xi\overline{\mathcal{M}}_{g,n}(\mu)/\mathbb{C}^*$ is compact.*

The connected components of $\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ are in bijection with the connected components of $\Omega\mathcal{M}_g(\mu)$.

Proof. The statement combines Theorem 7.12 and Theorem 10.2. The orbifold structure and the normal crossing boundary carry over from the model domain, as defined in Proposition 8.1 and along with (8.2). The statement about components follows from the smoothness of the orbifold charts. ∎

10.2. **The setup and notation for vertical plumbing.** We now set up the notation for the neighborhoods in which the vertical plumbing construction is performed, as well as for the plumbing fixtures we need. We fix for the remainder of the plumbing construction a base surface $(X_0, \boldsymbol{\eta}_0) \in \Omega\mathcal{B}_\Lambda$ and a local chart $\Psi\colon \Delta^M_\epsilon \to \Omega\mathcal{B}_\Lambda$, where $M = \dim \Omega\mathcal{B}_\Lambda$, parameterizing a neighborhood $\mathcal{V}$ of $\Psi(\mathbf{0}) = (X_0, \boldsymbol{\eta}_0)$. As in Section 8.2, we let $\mathcal{W} \subset \Omega\overline{\mathcal{M}D}^s_\Lambda$ be the neighborhood of $(X_0, \boldsymbol{\eta}_0)$ that is the preimage of $\mathcal{V}$. We fix for the rest of this section a local trivialization of the model domain over $\mathcal{W}$, which determines holomorphic functions $\boldsymbol{t}\colon \mathcal{W} \to \mathbb{C}^N \times \mathbb{C}^*$, rescaled forms $\boldsymbol{\eta}$, and prong-matchings $\boldsymbol{\sigma}$ so that the tautological one-form is $\boldsymbol{t} * \boldsymbol{\eta}$.

The product $\Psi \times \boldsymbol{t}$ identifies $\mathcal{W}$ with $\Delta^M_\epsilon \times \mathbb{C}^N \times \mathbb{C}^*$. We will subsequently work on

$$\mathcal{W}_\epsilon \;=\; \Delta^M_\epsilon \times \Delta^N_\epsilon \times \mathbb{C}^* \subset \mathcal{W}\,,$$

for $\epsilon = \epsilon(X_0, \boldsymbol{\eta}_0)$ sufficiently small and to be determined (first by Theorem 10.4, 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 $\mathcal{X} \to \mathcal{W}_\epsilon$ for the restriction of the universal curve to the domain of the chart. We will denote points in $\mathcal{W}_\epsilon$ as $(\boldsymbol{w}, \boldsymbol{t})$ with $\boldsymbol{w} \in \Delta_\epsilon^M$, or by $([X, \boldsymbol{\eta}], \boldsymbol{t})$. The boundary stratification of the model domain induces a stratification of $\mathcal{W}_\epsilon$. Given a subset $J \subset L(\Lambda)$, we define $\mathcal{W}_\epsilon^J = \mathcal{W}_\epsilon \cap \mathcal{MD}_\Lambda^{s,\Lambda_J}$. In other words, $\mathcal{W}_\epsilon^J$ is the locus where $t_i = 0$ if and only if $i \in J$.

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*

$$V(s) \;=\; \{(u, v) \in \Delta_\delta^2 : uv = s\} \tag{10.4}$$

together with the *top plumbing annulus* and *bottom plumbing annulus*

$$A^+ \;=\; \{\delta/R < |u| < \delta\} \quad \text{and} \quad A^- \;=\; \{\delta/R < |v| < \delta\} \tag{10.5}$$

for some R still to be specified. Unless specified otherwise, we will use the basepoints

$$p^\pm \;=\; \delta/\sqrt{R} \in A^\pm \tag{10.6}$$

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^-.

For each vertical edge e of $\Gamma = \Gamma(\Lambda)$, we define the *plumbing fixture* $\mathbb{V}_e \to \mathcal{W}_\epsilon$ to be the standard model family of nodal curves over $\mathcal{W}_\epsilon$:

$$\mathbb{V}_e \;=\; \left\{(\boldsymbol{w}, \boldsymbol{t}, u, v) \in \mathcal{W}_\epsilon \times \Delta_\delta^2 :\; uv = \prod_{i=\ell(e^-)}^{\ell(e^+)-1} t_i^{m_{e,i}}\right\}, \tag{10.7}$$

where the integers $m_{e,i}$ are defined in (6.7). Note that the fiber of $\mathbb{V}_e \to \mathcal{W}_\epsilon$ is an annulus if each t_i in the product in (10.7) is non-zero, and a pair of disks meeting at a node otherwise.

We denote by $r_e, r'_e \colon \mathcal{W}_\epsilon \to \mathbb{C}$ the residue functions

$$r_e \;=\; \operatorname{Res}_{q_e^-} \boldsymbol{t} * \boldsymbol{\eta} \quad \text{and} \quad r'_e \;=\; \operatorname{Res}_{q_e^-} \boldsymbol{t} * (\boldsymbol{\eta} + \boldsymbol{\xi}),$$

where $\boldsymbol{\xi}$ denotes the modifying differential, satisfying conditions of Definition 9.1, constructed in Proposition 9.3.

We equip $\mathbb{V}_e$ with the relative one-form Ω_e, given in coordinates by

$$\Omega_e \;=\; \left(\boldsymbol{t}^{\boldsymbol{a}}_{\lceil \ell(e^+)\rceil} \cdot u^{\kappa_e} - r'_e\right)\frac{du}{u} \quad \text{and} \quad \Omega_e \;=\; -\left(\boldsymbol{t}^{\boldsymbol{a}}_{\lceil \ell(e^-)\rceil} \cdot v^{-\kappa_e} - r'_e\right)\frac{dv}{v}, \tag{10.8}$$

where the notation $\boldsymbol{t}^{\boldsymbol{a}}_{\lceil \ell(e^+)\rceil}$ was introduced in (8.3). 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 (10.7), 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 that will be provided by Theorem 10.4 below.

We define families of disjoint annuli $\mathcal{A}_e^+, \mathcal{A}_e^- \subset \mathbb{V}_e$ by

$$\begin{aligned} \mathcal{A}_e^+ &= \{(\boldsymbol{w}, \boldsymbol{t}, u, v) : |w_i|, |t_j| < \epsilon \text{ for all } i, j, \text{ and } \delta/R < |u| < \delta\} \quad \text{and} \\ \mathcal{A}_e^- &= \{(\boldsymbol{w}, \boldsymbol{t}, u, v) : |w_i|, |t_j| < \epsilon \text{ for all } i, j, \text{ and } \delta/R < |v| < \delta\}. \end{aligned}$$

We will refer to $\mathcal{A}_e^+$ and $\mathcal{A}_e^-$ as the *top and bottom plumbing annuli* corresponding to the vertical edge e.

10.3. **Standard coordinates.** We now apply the normal form theorems of Section 4 to the family $\mathcal{X} \to \mathcal{W}_\epsilon$. Several of these normal forms are not unique, in which case we simply make an arbitrary choice.

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

$$\phi_e^+ \colon \mathcal{W}_\epsilon \times \Delta_{\delta_1} \to \mathcal{X}_{\ell(e^+)} \quad \text{and} \quad \phi_e^- \colon \mathcal{W}_\epsilon \times \Delta_{\delta_1} \to \mathcal{X}_{\ell(e^-)}$$

(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 $\boldsymbol{t} * \boldsymbol{\eta}$ in the standard form. For a vertical node q_e, this standard form is

$$\begin{aligned} (\phi_e^+)^*(\boldsymbol{t} * \boldsymbol{\eta}) &= \boldsymbol{t}_{\lceil \ell(e^+) \rceil}^{\boldsymbol{a}} u^{\kappa_e} \frac{du}{u} \quad \text{and} \\ (\phi_e^-)^*(\boldsymbol{t} * \boldsymbol{\eta}) &= -\boldsymbol{t}_{\lceil \ell(e^-) \rceil}^{\boldsymbol{a}} \left(v^{-\kappa_e} - r_e(\boldsymbol{t})\right) \frac{dv}{v}. \end{aligned}$$

As these standard coordinates are not unique, we use the prong-matching σ_e to restrict their choice as follows. Given a choice of $\phi_e^\pm$, in these coordinates the prong-matching must be of the form $\sigma_e = \zeta du \otimes dv$, where ζ is some k'th root of unity. We require that our choice of $\phi_e^\pm$ makes this root equal to 1, so that

$$\sigma_e = du \otimes dv. \tag{10.9}$$

In general, the modified differential $\boldsymbol{t} * (\boldsymbol{\eta} + \boldsymbol{\xi})$ does not admit such a simple standard form in a neighborhood of a vertical node. Consider a vertical node with top section $q_e^+ \colon \mathcal{W}_\epsilon \to \mathcal{X}_{\ell(e^+)}$, which is a zero of order $\kappa_e - 1$ of $\boldsymbol{t} * \boldsymbol{\eta}$. Then this zero breaks up into a simple pole and κ_e nearby zeros of the differential $\boldsymbol{t} * (\boldsymbol{\eta} + \boldsymbol{\xi})$. These extraneous nearby zeros should not belong to our plumbed family, so we will construct a family of disks $\mathcal{E}_e^+$ 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* $\Upsilon_e^+ \colon \mathcal{A}_e^+ \to \mathcal{B}_e^+$ putting $\boldsymbol{t} * (\boldsymbol{\eta} + \boldsymbol{\xi})$ into a standard form on a family of annuli over $\mathcal{W}_\epsilon$. These objects are constructed in Theorem 10.4 below. This is the basic analytic ingredient in our plumbing construction. In Section 10.4, we will use these gluing maps to glue in the standard plumbing fixture $\mathbb{V}_e$ defined above.

Adding the modifying differential $\boldsymbol{\xi}$ creates a similar problem at the zero sections $\boldsymbol{z}(h)$ 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 $\mathcal{E}_h$ around $\boldsymbol{z}(h)$, and gluing maps that put $\boldsymbol{t} * (\boldsymbol{\eta} + \boldsymbol{\xi})$ into the standard form on a family of annuli surrounding $\mathcal{E}_h$. In Section 10.4 we will then cut out these disks and glue in a standard family of disks $\mathbb{D}_h$.

Theorem 10.4. *There exists a constant $\delta > 0$ such that for any $R > 0$ there exists a constant $\epsilon > 0$ such that for each vertical edge e and for each half-edge h of Γ there are families of conformal maps of annuli*

$$\begin{aligned} \upsilon_e^+ &\colon \mathcal{W}_\epsilon \times A_{\delta/R,\delta} \to \mathcal{X}_{\ell(e^+)}\,, \\ \upsilon_e^- &\colon \mathcal{W}_\epsilon \times A_{\delta/R,\delta} \to \mathcal{X}_{\ell(e^-)}\,, \quad \text{and} \\ \upsilon_h &\colon \mathcal{W}_\epsilon \times A_{\delta/R,\delta} \to \mathcal{X}_{\ell(h)}\,. \end{aligned}$$

These maps commute with the projections of the source and target to $\mathcal{W}_e$, and have the following properties:

(i) The images of υ_e^+, υ_e^-, υ_h are families of annuli $\mathcal{B}_e^+, \mathcal{B}_e^-$, $\mathcal{B}_h$ that do not contain any zeros or poles of $(\mathcal{X}, \boldsymbol{t} * \boldsymbol{\eta})$. The families of annuli $\mathcal{B}_e^+, \mathcal{B}_e^-$, and $\mathcal{B}_h$ bound families of disks $\mathcal{E}_e^+, \mathcal{E}_e^-$, and $\mathcal{E}_h$, respectively, where

$$Q_e^+ \subset \mathcal{E}_e^+ \subset \mathcal{X}_{\ell(e^+)}\,, \quad Q_e^- \subset \mathcal{E}_e^- \subset \mathcal{X}_{\ell(e^-)}\,, \quad \text{and} \quad \boldsymbol{z}(h) \subset \mathcal{E}_h \subset \mathcal{X}_{\ell(h)}\,.$$

(ii) The pullback of $\boldsymbol{t}*(\boldsymbol{\eta}+\boldsymbol{\xi})$ under each of the maps υ_e^+, υ_e^-, and υ_h, has the standard form on the annulus, that is

$$\begin{aligned} (\upsilon_e^+)^*(\boldsymbol{t} * (\boldsymbol{\eta} + \boldsymbol{\xi})) &= \left(\boldsymbol{t}_{\lceil \ell(e^+)\rceil}^{\boldsymbol{a}} \cdot z^{\kappa_e} - r_e'\right) \frac{dz}{z}\,, \\ (\upsilon_e^-)^*(\boldsymbol{t} * (\boldsymbol{\eta} + \boldsymbol{\xi})) &= \left(-\boldsymbol{t}_{\lceil \ell(e^-)\rceil}^{\boldsymbol{a}} \cdot z^{-\kappa_e} + r_e'\right) \frac{dz}{z}\,, \quad \text{and} \\ \upsilon_h^*(\boldsymbol{t} * (\boldsymbol{\eta} + \boldsymbol{\xi})) &= \boldsymbol{t}_{\lceil \ell(h)\rceil}^{\boldsymbol{a}} \cdot z^{m_h} dz\,. \end{aligned}$$

(iii) The maps υ_e^+, υ_e^-, and υ_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 δ sufficiently small that the maps $\upsilon_e^\pm$ and υ_h are injective and have mutually disjoint images.

We need to allow the constant R to be arbitrarily large to facilitate the proof that $\Omega\,\mathrm{Pl}^v$ is locally injective. See Lemma 10.10, where the choice of R is made.

The location of these annuli is illustrated in the left part of Figure 6. The images of the marked points $p_e^\pm$ and p_h in $\mathcal{X}$ are denoted by $b_e^\pm \in \mathcal{B}^\pm$ and $b_h \in \mathcal{B}_h$, respectively, for each vertical edge or half-edge.

Proof. In the ϕ_e^+-coordinates, the modifying differential $\boldsymbol{t} * \boldsymbol{\xi}$ becomes

$$(\phi_e^+)^*(\boldsymbol{t} * \boldsymbol{\xi}) = \boldsymbol{t}_{\lceil \ell(e^+)\rceil}^{\boldsymbol{a}} \cdot \alpha_e \frac{du}{u}\,,$$

where α_e is a holomorphic function on the product $\mathcal{W}_\epsilon \times \Delta_\epsilon^N \times \Delta_{\delta_1}$ satisfying $\boldsymbol{t}_{\lceil \ell(e^+)\rceil}^{\boldsymbol{a}} \cdot \alpha_e(\boldsymbol{w}, \boldsymbol{t}, 0) = -r_e'(\boldsymbol{w}, \boldsymbol{t})$ and $\alpha_e(\boldsymbol{w}, \boldsymbol{0}, z) \equiv 0$. (Using Corollary 9.4, we see that in fact $\boldsymbol{\xi}$ depends only on the t_i with $i < L$).

Fix a basepoint $b \in A_{\delta/R,\delta}$. By Theorem 4.2, after possibly decreasing ϵ, there is a family of conformal maps

$$\psi_e \colon \mathcal{W}_\epsilon \times A_{\delta/R,\delta} \to \mathcal{W}_\epsilon \times \Delta_{\delta_1} ,$$

which cover the identity on $\mathcal{W}_\epsilon$, fix the section $\mathcal{W}_\epsilon \times \{b\}$, and put $(\phi_e^+)^*(\boldsymbol{t} * (\boldsymbol{\eta} + \boldsymbol{\xi}))$ in the desired standard form as follows:

$$\begin{aligned} (\phi_e^+ \circ \psi_e)^*(\boldsymbol{t} * (\boldsymbol{\eta} + \boldsymbol{\xi})) &= \psi_e^* \left(\boldsymbol{t}^{\boldsymbol{a}}_{\lceil \ell(e^+) \rceil} \cdot z^{\kappa_e} + \boldsymbol{t}^{\boldsymbol{a}}_{\lceil \ell(e^+) \rceil} \cdot \alpha_e(\boldsymbol{w}, \boldsymbol{t}, z) \right) \frac{dz}{z} , \\ &= \left(\boldsymbol{t}^{\boldsymbol{a}}_{\lceil \ell(e^+) \rceil} \cdot z^{\kappa_e} - r'_e \right) \frac{dz}{z} . \end{aligned}$$

Since α_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. Over the locus where $t_{L-1} = \cdots = t_{-N} = 0$, the modifying differential $\boldsymbol{\xi}$ vanishes on level L, which means ψ_e preserves the form $z^{\kappa_e} \frac{dz}{z}$ and fixes the point b, and thus ψ_e is the identity over this locus. We then define $\upsilon_e^+ = \phi_e^+ \circ \psi_e$. The desired family of disks $\mathcal{E}_e^+$ 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 υ_e^- is much simpler at a pole, as then we need only to apply Theorem 4.1 to construct a map υ_e^- putting $\boldsymbol{t} * (\boldsymbol{\eta} + \boldsymbol{\xi})$ 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 υ_h follows from the same technique. In this case, the modifying differential $\boldsymbol{\xi}$ is holomorphic along the zero section $\boldsymbol{z}_h$, so the resulting standard form of $\upsilon_h^*(\boldsymbol{t} * (\boldsymbol{\eta} + \boldsymbol{\xi}))$ has no residue. ◻

10.4. **The vertical plumbing construction.** We now present the basic plumbing construction. The plumbing starts from a family $\mathcal{X} \to \mathcal{W}_\epsilon$ equipped with the family of differentials $\boldsymbol{t} * \boldsymbol{\eta}$ (as defined in Section 10.2) together with a modifying differential $\boldsymbol{\xi}$, and builds a family of meromorphic stable differentials $(\mathcal{Y}^v \to \mathcal{W}_\epsilon, \omega, \boldsymbol{z})$ which vanishes identically on the lower level components over the boundary divisor and is elsewhere holomorphic and nonzero, except for the prescribed zeros and poles $\boldsymbol{z}(h)$.

We define conformal isomorphisms $\Upsilon_e^\pm \colon \mathcal{A}_e^\pm \to \mathcal{B}_e^\pm \subset \mathcal{X}$ by

$$\Upsilon_e^+(\boldsymbol{w}, \boldsymbol{t}, u, v) \;=\; \upsilon_e^+(\boldsymbol{w}, \boldsymbol{t}, u) \quad \text{and} \quad \Upsilon_e^-(\boldsymbol{w}, \boldsymbol{t}, u, v) \;=\; \upsilon_e^-(\boldsymbol{w}, \boldsymbol{t}, v) ,$$

where $\mathcal{B}_e^\pm$ and $\upsilon_e^\pm$ are defined in Theorem 10.4. These maps identify each Ω_e with $\boldsymbol{t} * (\boldsymbol{\eta} + \boldsymbol{\xi})$ 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 10.4 by

$$\Upsilon_h \;=\; \upsilon_h \colon \mathcal{A}_h \to \mathcal{B}_h \subset \mathcal{X} .$$

We let $\mathcal{Y}^v \to \mathcal{W}_\epsilon$ be the family of curves obtained by removing from $\mathcal{X}$ the families of disks $\mathcal{E}_e^\pm$ and $\mathcal{E}_h$ 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 $\mathcal{Y}^v$ inherits a relative one-form ω.

We denote the plumbing annuli as subsurfaces of $\mathcal{Y}^v$ by $\mathcal{C}_e^\pm$ and $\mathcal{C}_h$, denote the middle annuli bounded by the $\mathcal{C}_e^\pm$ as $\mathcal{F}_e$, and denote $c_e^\pm$ the image of the points $p^\pm$ in $\mathcal{C}_e^\pm$.

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 6.

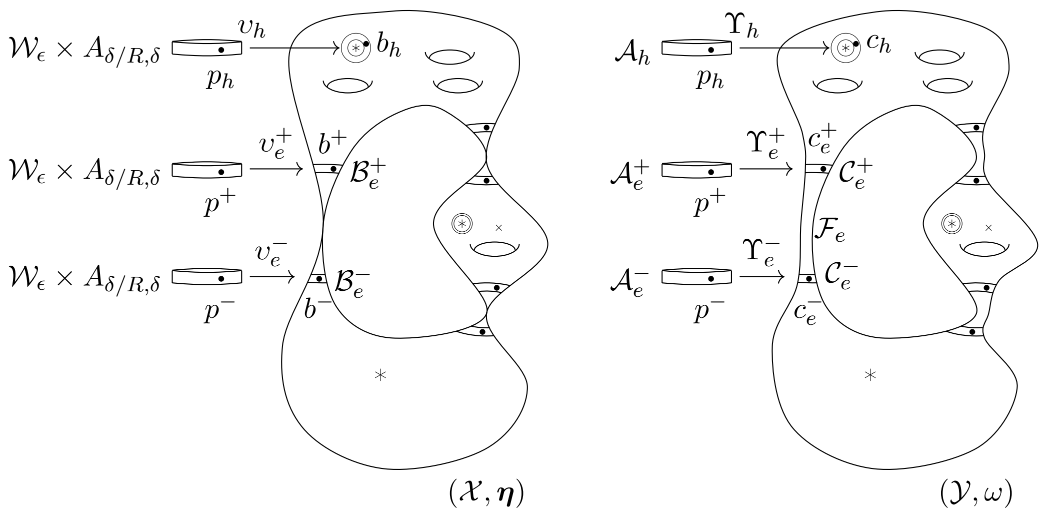

FIGURE 6. The general plumbing construction for our running example of Section 2.6.

We will later use the following consequence of the construction and the fact that the modifying differential $\boldsymbol{\xi}$ on level i depends only on the levels below i and the topological data, see Corollary 9.4.

Proposition 10.5. *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 $i \leq \ell(e^+)$.*

10.5. **The vertical plumbing map.** We now construct the vertical plumbing map $\Omega\,\mathrm{Pl}^v \colon \mathcal{W}_\epsilon \to \Xi\mathcal{D}_\Lambda^{sv}$. Over the generic stratum $\mathcal{W}_\epsilon^\emptyset$, as ω does not vanish identically on any fiber, this map follows immediately from the obvious universal property for the Teichmüller space $\Omega\mathcal{T}_{(\Sigma,\boldsymbol{s})}(\mu)$ of (twisted) differentials.

Now consider a boundary stratum $\mathcal{W}_\epsilon^J$, with $J \subset L(\Lambda)$ a nonempty set of levels, and let $\mathcal{Y}^{v,J} \to \mathcal{W}_\epsilon^J$ be the restriction of $\mathcal{Y}^v$, which is then an equisingular family whose dual graph is the enhanced multicurve $\Lambda_J \rightsquigarrow \Lambda$. To define the plumbing map on this stratum, we restrict appropriate rescalings $\boldsymbol{\omega}^J$ of ω to the irreducible components of $\mathcal{Y}^{v,J}$, construct prong-matchings σ_e^J along the vertical nodes of $\mathcal{Y}^v$, and define a marking of the welding of $\mathcal{Y}^v$ along these prong-matchings.

For each level $i \in L^\bullet(\Lambda)$, define on $\mathcal{X}_{(i)}$ the rescaled form

$$\boldsymbol{t}^{\boldsymbol{a}}_{\lceil i\rceil\setminus J}(\eta_{(i)} + \xi_{(i)}) \;=\; \prod_{k\geq i,\; k\notin J} t_k^{-a_k}\omega_{(i)}, \tag{10.10}$$

where we defined $\boldsymbol{t}_I^{\boldsymbol{a}} = \prod_{i\in I} t_i^{a_i}$ and set $\lceil i\rceil = \{j : j \geq i\}$, extending the definition in (8.3). For each vertical edge e that does not persist in Λ_J, define on $\mathbb{V}_e$ the form

$$\Omega_e^J \;=\; (\boldsymbol{t}^{\boldsymbol{a}}_{\lceil \ell(e^+)\rceil\setminus J}\, u^{\kappa_e} - r'_e)\frac{du}{u} \;=\; -(\boldsymbol{t}^{\boldsymbol{a}}_{\lceil \ell(e^-)\rceil\setminus J}\, v^{-\kappa_e} - r'_e)\frac{dv}{v}\,.$$

The gluing maps for such edges identify these rescaled forms, which defines on $\mathcal{Y}^{v,J}$ a collection of rescaled forms $\boldsymbol{\omega}^J$ that do not vanish identically on any fiber.

For each vertical edge e of Λ_J, we define

$$\text{(10.11)}\qquad \boldsymbol{t}^{\boldsymbol{a}}_{[\ell(e^-),\ell(e^+))\cap J} \;=\; \prod_{\substack{\ell(e^-)\leq k<\ell(e^+)\\ k\notin J}} t_k^{m_{e,k}} \quad\text{and}\quad \sigma_e^J \;=\; \frac{1}{\boldsymbol{t}^{\boldsymbol{a}}_{[\ell(e^-),\ell(e^+))\cap J}}du\otimes dv,$$

which is easily checked to be a prong-matching for $\boldsymbol{\omega}^J$. Note that this agrees with (10.9) when $J = L(\Lambda)$.

We now wish to define markings for the fibers of $\mathcal{Y}^{v,J}$ by transporting the markings of the fibers of $\mathcal{X}$ by almost-diffeomorphisms $f_{\boldsymbol{t}}\colon \overline{Y}_{\boldsymbol{t}} \to \overline{X}_{\boldsymbol{t}}$ between the corresponding welded surfaces. Given a vertical edge e of Λ that does not persist in Λ_J, there is an annulus $E_e \subset \overline{X}_{\boldsymbol{t}}$, obtained by welding the fibers over $\boldsymbol{t}$ in $\mathcal{E}_e^{\pm}$, and a corresponding annulus $F_e \subset \overline{Y}_{\boldsymbol{t}}$, the fiber over $\boldsymbol{t}$ in $\mathcal{F}_e = \mathbb{V}_e \setminus (\mathcal{A}_e^+ \cup \mathcal{A}_e^-)$. We will refer to these F_e as *infinite* if they contain a seam and *finite* otherwise. We denote by $X'_{\boldsymbol{t}} \subset \overline{X}_{\boldsymbol{t}}$ and $Y'_{\boldsymbol{t}} \subset \overline{Y}_{\boldsymbol{t}}$ the complement of these annuli. The plumbing construction defines a canonical conformal map $f_{\boldsymbol{t}}\colon Y'_{\boldsymbol{t}} \to X'_{\boldsymbol{t}}$ which we wish to extend to $\overline{Y}_{\boldsymbol{t}}$.

Let $\gamma_e \subset E_e$ be the continuous curve joining the two boundary components of the annulus E_e obtained by concatenating the two radial curves

$$\gamma_e^+(s) \;=\; \phi_e^+(s) \quad\text{and}\quad \gamma_e^-(s) \;=\; \phi_e^-(sT) \quad\text{for}\quad 0\leq s\leq \delta/R,$$

where $T = \boldsymbol{t}^{\boldsymbol{a}}_{[\ell(e^-),\ell(e^+))\cap J}/|\boldsymbol{t}^{\boldsymbol{a}}_{[\ell(e^-),\ell(e^+))\cap J}|$. Note that γ_e lifts to a continuous curve on $\overline{X}_{\boldsymbol{t}}$ by the choice of prong-matching given by (10.11). Let $\delta_e \subset F_e$ be the curve joining the point $u = \delta/R$ in the upper boundary component to the point $v = T\delta/R$ in the lower boundary component, along which $\arg u$ and $\arg v$ are constant. We extend $f_{\boldsymbol{t}}$ to a homeomorphism $f_{\boldsymbol{t}}\colon \overline{Y}_{\boldsymbol{t}} \to \overline{X}_{\boldsymbol{t}}$ sending each γ_e to δ_e. We then mark $\overline{Y}_{\boldsymbol{t}}$ by composing $f_{\boldsymbol{t}}^{-1}$ with the marking of $\overline{X}_{\boldsymbol{t}}$.

Proposition 10.6. *The almost-diffeomorphisms $f_{\boldsymbol{t}}\colon \overline{Y}_{\boldsymbol{t}} \to \overline{X}_{\boldsymbol{t}}$ have the following properties:*

- *$f_{\boldsymbol{t}}$ identifies $\boldsymbol{\omega}^J$ on $Y'_{\boldsymbol{t}}$ with the rescaled forms* (10.10) *on $X'_{\boldsymbol{t}}$.*
- *For any sequence $\boldsymbol{t}_n \to \boldsymbol{0}$, the maps $f_{\boldsymbol{t}_n}$ are asymptotically turning number preserving.*

Proof. The first statement follows immediately from the definition of the maps and the plumbing construction. For the second statement, it suffices to check that the turning numbers of the γ_e converge to those of the δ_e, and this follows immediately from the fact that the modifying differentials converge to 0 on each component as $\boldsymbol{t} \to \boldsymbol{0}$. ∎

These forms, prong-matchings, and markings give each fiber of $\mathcal{Y}^{v,J}$ the structure of a marked multiscale differential. Applying the standard universal property of the boundary stratum $\Omega\mathcal{B}_{\Lambda_J}$ to the family $(\mathcal{Y}^{v,J} \to \mathcal{W}_\epsilon^J, \boldsymbol{\omega}^J, \boldsymbol{\sigma}^J)$ defines a plumbing map $\Omega\,\mathrm{Pl}^v$ on $\mathcal{W}_\epsilon^J$.

10.6. **Continuity of the plumbing map.** Our next goal is to establish the continuity of the vertical plumbing map. We start with a construction which will be used to extend conformal maps across annuli here and in the discussion of the horizontal plumbing construction.

Consider the round annulus $A = \{r_1 < |z| < r_2\}$ in $\mathbb{C}^*$ (for $0 < r_1 < r_2 < \infty$), which we equip with the flat metric $|du|/u$, together with a C^1 map $\alpha\colon \partial A \to \mathbb{C}^*$. We define the *straight line extension* of α to be the map $F\colon A \to \mathbb{C}^*$ sending the radial geodesic from $r_1e^{i\theta}$ to $r_2e^{i\theta}$ to the geodesic joining $\alpha^-(r_1e^{i\theta})$ to $\alpha^+(r_2e^{i\theta})$, where α^- and α^+ denote the restrictions of α to the respective inner and outer boundary components.

Similarly, given the punctured unit disk Δ^* and a C^1 map $\alpha\colon S^1 \to \mathbb{C}^*$, we define its straight line extension as the map $F\colon \Delta^* \to \mathbb{C}^*$ sending the radial segment from 0 to $e^{i\theta}$ to the radial segment from 0 to $\alpha(e^{i\theta})$.

Lemma 10.7. *Suppose $\alpha_n\colon \partial A \to \mathbb{C}^*$ is a sequence of C^1 maps converging C^1-uniformly to the identity map, and let $F_n\colon A \to \mathbb{C}^*$ be their straight line extensions. Then the F_n also converge C^1-uniformly to the identity and are eventually C^1 diffeomorphisms onto their image.*

The analogous statement holds for the straight line extension of a sequence of maps on S^1.

Proof. We work in logarithmic coordinates z identifying A with the annulus $B = \{0 < \operatorname{Im} z < h\} \subset \mathbb{C}/\mathbb{Z}$. In these coordinates,

$$F_n(x+iy) = \alpha_n^-(x)(1-y/h) + \alpha_n^+(x+ih)y/h,$$

and we have

$$F(x+iy) - (x+iy) = (\alpha_n^-(x) - x)(1-y/h) + (\alpha_n^+(x+ih) - (x+ih))y/h,$$

which easily implies the second claim. It follows that if α_n is sufficiently close to the identity, the Jacobian determinant of F_n is nowhere vanishing, so F_n is a C^1 diffeomorphism onto its image.

The second case is handled similarly. ◻

Proposition 10.8. *The vertical plumbing map $\Omega\,\mathrm{Pl}^v\colon \mathcal{W}_\epsilon \to \Xi\mathcal{D}_\Lambda^{sv}$ constructed above is continuous.*

Proof. We fix a sequence $(X_n, \boldsymbol{\eta}_n, \boldsymbol{t}_n)$ in $\overline{\mathcal{M}D}_\Lambda^s$ that converges to $(X, \boldsymbol{\eta}, \boldsymbol{t})$, and let $(Y_n, \boldsymbol{\omega}_n)$ and $(Y, \boldsymbol{\omega})$ denote the corresponding plumbed surfaces.

Passing to a subsequence, we may assume this sequence belongs to a fixed stratum of $\overline{\mathcal{M}D}_\Lambda^s$. We denote by $g_n\colon \overline{X}_n \to \overline{X}$ the maps exhibiting this convergence as in Proposition 8.3. This means that there are subsurfaces $J_n \subset X_n \setminus \boldsymbol{z}_n$ such that g_n is conformal on J_n, and the images $K_n = g_n(J_n)$ exhaust $X \setminus \boldsymbol{z}$. Moreover, the J_n eventually contain the plumbing annuli $\mathcal{B}_e^\pm$ because these annuli vary continuously in the universal curve, and g_n converges to the identity by Lemma 3.4.

As defined above, we have almost-diffeomorphisms $f_n\colon \overline{Y}_n \to \overline{X}_n$ and $f\colon \overline{Y} \to \overline{X}$ which are conformal on the complements Y_n' and Y' of the plumbing annuli. We define

$h_n = f \circ g_n \circ f_n^{-1}\colon Y_n \to Y$. These maps satisfy all of the requirements to exhibit convergence from Definition 7.4, except that h_n^{-1} is conformal only on Y'. We define $\widetilde{h}_n$ to be the inverse of the straight line extension of the restriction of h_n^{-1} to the boundary of Y_n'. To be precise, $\widetilde{h}_n^{-1}$ may not be continuous at the seams of $\overline{Y}$, but we can modify it on a shrinking neighborhood of each seam, so that it becomes continuous and remains homotopic to h_n^{-1}.

These $\widetilde{\widetilde{h}_n}^{-1}$ are L_n-quasiconformal on an exhaustion of Y with $L_n \to 1$ and satisfy all of the requirements (as modified by Proposition 7.10) to exhibit convergence of $(Y_n, \boldsymbol{\omega}_n)$ to $(Y, \boldsymbol{\omega})$, where on the annuli F_e, we apply Lemma 10.7 to get the necessary weak convergence of forms. ◫

10.7. **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 10.9. *The vertical plumbing map* $\Omega\,\mathrm{Pl}^v : \mathcal{W}_e \to \Xi\mathcal{D}_\Lambda^{sv}$ *is open in a neighborhood of any point in the deepest stratum* $\Omega\mathcal{M}\mathcal{D}_\Lambda^{\Lambda,s}$.

For a surface X that corresponds to a point in the subset $\mathcal{W}_\epsilon$ of the model domain we denote by $X_{(\le i)}^-$ the subsurface consisting of the levels $\le i$, including the plumbing annuli $\mathcal{B}_e^-$ for all e with $\ell(e^-) \le i$ and $\mathcal{B}_h$ for h with $\ell(h) \le i$, but excluding the discs $\mathcal{E}_e^-$ for all e with $\ell(e^+) > i$ and $\mathcal{E}_h$ for h with $\ell(h) = i$. We let $X_{(\le i)}^+$ be the subsurface consisting of the levels $\le i$, including all the plumbing fixtures connecting to higher levels all the way up to the top plumbing annuli $\mathcal{B}_e^+$ for $\ell(e^+) > i$.

Proof. Choose a model differential $(X, \boldsymbol{\eta})$ in the deepest stratum $\Omega\mathcal{M}\mathcal{D}_\Lambda^{\Lambda,s}$ (which we continue to implicitly identify with the deepest stratum of $\Xi\mathcal{D}_\Lambda^{sv}$) and a sequence $(Y_n, \boldsymbol{\omega}_n) \to (X, \boldsymbol{\eta})$ in $\Xi\mathcal{D}_\Lambda^{sv}$. 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.

We can choose representatives of $(X, \boldsymbol{\eta})$ in $\Omega\mathcal{T}_\Lambda^{pm}(\mu)$ and $(Y_n, \boldsymbol{\omega}_n)$ in $\Omega\overline{\mathcal{T}}_{(\Sigma,\boldsymbol{s})}(\mu)$ so that convergence still holds. By the definition of convergence in $\Omega\overline{\mathcal{T}}_{(\Sigma,\boldsymbol{s})}(\mu)$ as given in Section 7, there is a sequence $\boldsymbol{d}_n = \{d_{n,i}\} \in \mathbb{C}^{L^\bullet(\Lambda)}$ and a sequence of almost-diffeomorphisms $g_n\colon Y_n \to \overline{X}$, defined up to isotopy, which are compatible with the markings, asymptotically turning-number-preserving and whose inverses are conformal on an exhaustion $K_{n,(i)}$ of each level $X_{(i)}$. More precisely, this is an exhaustion of the complement of the nodes and marked points z_h in $X_{(i)}$.

We start by choosing the sequence of coordinates $\boldsymbol{t}_n$ defined in terms of these $d_{n,i}$ by

$$t_{n,i} \;=\; \boldsymbol{e}\left(\frac{d_{n,i+1} - d_{n,i}}{a_i}\right). \tag{10.12}$$

Similarly to $\boldsymbol{t}_{\lceil i \rceil}^{\boldsymbol{a}}$ defined in Equation (8.3), we denote $\boldsymbol{t}_{n,\lceil i \rceil}^{\boldsymbol{a}} = \prod_{j \ge i} t_{n,j}^{a_j}$. Since we are not rescaling the top level, $d_{n,0} = 0$ and it follows that $\boldsymbol{t}_{n,\lceil i \rceil}^{\boldsymbol{a}} = \boldsymbol{e}(-d_{n,i})$. By definition of convergence,

$$\boldsymbol{e}(d_{n,i})(g_{n,(i)})_*\omega_n \xrightarrow[n\to\infty]{} \eta_{(i)}, \quad \text{i.e.} \quad \frac{1}{\boldsymbol{t}^{\boldsymbol{a}}_{n,\lceil i\rceil}}(g_{n,(i)})_*\omega_n \xrightarrow[n\to\infty]{} \eta_{(i)}. \tag{10.13}$$

We will now construct inductively the $(X_n, \boldsymbol{\eta}_n)$ such that $\Omega\,\mathrm{Pl}^v(X_n, \boldsymbol{\eta}_n, \boldsymbol{t}_n) = (Y_n, \boldsymbol{\omega}_n)$. Recall that by Proposition 10.5 it makes sense to consider the effect of plumbing only on the bottom part of a surface and to write $\Omega\,\mathrm{Pl}^v(X^{\pm}_{n,(\le i)})$, suppressing the dependence on $(\boldsymbol{\eta}_n, \boldsymbol{t}_n)$ for notational convenience. Note also that the set of connected components of the ϵ_n-thick part of Y_n is eventually the disjoint union of sets of level i components $Y_{n,(i)}$, where $Y_{n,(i)}$ are those components that g_n maps to $X_{n,(i)}$.

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 $\mathcal{W}_\epsilon$ and to construct a conformal map on the bottom level $h_{n,(-N)}\colon \Omega\,\mathrm{Pl}^v(X^-_{n,(\le -N)}) \to Y_{n,(-N)}$ which identifies the two differentials. The second step of the base case is to extend this map by analytic continuation across the plumbing annuli to obtain a conformal map $h^+_{n,(-N)}\colon \Omega\,\mathrm{Pl}^v(X^+_{n,(\le -N)}) \to Y_{n,(\le -N)}$.

The inductive step starts with the map $h^+_{n,(i)}\colon \Omega\,\mathrm{Pl}^v(X^+_{n,(\le i)}) \to Y_{n,(\le i)}$. We choose appropriately $(X_{n,(i+1)}, \eta_{n,(i+1)})$ and construct a conformal map

$$h_{n,(i+1)}\colon \Omega\,\mathrm{Pl}^v\left(X^-_{n,(\le i+1)}, \eta_{n,(\le i+1)}, \boldsymbol{t}_n\right) \to Y_{n,(\le i+1)}$$

which identifies the forms and agrees with $h^+_{n,(i)}$ on its domain. We then analytically continue across the plumbing cylinders and disks to get $h^+_{n,(i+1)}\colon \Omega\,\mathrm{Pl}^v(X^+_{n,(\le i+1)}) \to Y_{n,(\le i+1)}$. This procedure eventually ends at the top level when we have constructed the entire surface X_n together with a conformal isomorphism of $\Omega\,\mathrm{Pl}^v(X_n)$ with Y_n.

We start with the details of the construction at the bottom level. The conformal map $g^{-1}_{n,(-N)}$ is eventually defined on the fixed subsurface K_{-N} containing $X^-_{(-N)}$ but this map only approximately identifies the rescaled differentials, as is indicated in (10.13). We choose a sequence of surfaces $(X_{n,(-N)}, \eta_{n,(-N)})$ of the same topological type as $(X_{(-N)}, \eta_{(-N)})$ such that

$$\mathrm{Per}(X_{n,(-N)}, \eta_{n,(-N)}) \;=\; \frac{1}{\boldsymbol{t}^{\boldsymbol{a}}_{n,(-N)}}\mathrm{Per}_{(-N)}(Y_n, \boldsymbol{\omega}_n) \tag{10.14}$$

and such that $(X_{n,(-N)}, \eta_{n,(-N)})$ converges to $(X_{(-N)}, \eta_{(-N)})$. In this equation $\mathrm{Per}_{(i)} \subset H^1(\Sigma_i \setminus P_i, Z_i; \mathbb{C})$ denotes the relative periods in the level i subsurface. We may choose such a sequence because the period map is open. By convergence in the conformal topology, there exist maps $(g^X_{n,(-N)})\colon X_{n,(-N)} \to X_{(-N)}$ whose inverses are conformal on the same subsurface. We apply Theorem 3.6 to the sequences $(g^X_{n,(-N)})_*(\boldsymbol{t}_n * \boldsymbol{\eta}_n)$ and $(g_{n,(-N)})_*\boldsymbol{\omega}_n$ to produce a conformal map k_n defined on $K_{(-N)}$ which identifies these forms. The composition

$$h_{n,(-N)} \;=\; g^{-1}_{n,(-N)} \circ k_n \circ g^X_{n,(-N)}$$

provides eventually a conformal map

$$h_{n,(-N)}\colon X^-_{n,(-N)} \to Y_n \quad \text{such that} \quad h^*_{n,(-N)}\boldsymbol{\omega}_n \;=\; \boldsymbol{t}^{\boldsymbol{a}}_{n,(-N)}\eta_{n,(-N)}.$$

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 (10.8), 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

$$\nu_{n,e} \;=\; h_n \circ \upsilon^-_{n,e}\colon A^-_{n,e} \to Y_n$$

(where $\upsilon^-_{n,e}$ was defined in Theorem 10.4) identifies the form $\boldsymbol{\omega}_n$ on Y_n with the standard form $\Omega_{n,e}$ on the bottom plumbing annulus. We show in Lemma 10.10 below that for some sufficiently large $R > 0$ the maps $\nu_{n,e}$ can eventually be analytically continued to a conformal map $\nu_{n,e}\colon V'_{n,e} \to Y_n$, where $V'_{n,e} \subset V_{n,e}$ is a round subannulus containing the basepoint $p^+_{n,e} = \delta/R^{1/2}$. *A fortiori* the analytic continuation also identifies $\boldsymbol{\omega}_n$ with $\Omega_{n,e}$. We define marked points $c^+_e = \nu_{n,e}(p^+_e)$ in Y_n.

At this stage, we also analytically continue $h_{n,(-N)}$ across the $\mathcal{E}_h$ at level $-N$.

We now begin the inductive step, assuming that we have constructed conformal maps $h^+_{n,(i)}\colon \Omega\,\mathrm{Pl}^v(X^-_{n,(\le i)}) \to Y_{n,(\le i)}$. 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

$$\mathrm{PPer}(X_{n,(i+1)}, \eta_{n,(i+1)}) \;=\; \frac{1}{\boldsymbol{t}^{\boldsymbol{a}}_{n,(i+1)}}\mathrm{PPer}(Y_{n,(i+1)}, \omega_n)\,. \tag{10.15}$$

The PPer 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 PPer, i.e. we consider the surface cut open at the lower ends and use integration at the "nearby points" c^+_e determined by the induction hypothesis for all cylinders whose top end is on level $(i+1)$. The choice of X_n with the required perturbed periods is eventually possible, since the perturbed period map is open by Proposition 9.7.

To specify $\overline{X}_n$ as a marked surface, we define a marking $\widetilde{f}_n$ of $\overline{X}_n$ by composing the marking of Y_n with the appropriate map $f_n\colon Y_n \to \overline{X}_n$ from Proposition 10.6, which is defined when Y_n is sufficiently close to $\overline{X}$ as an unmarked multi-scale differential.

To complete the proof that $\overline{X}_n \to \overline{X}$ as model differentials, we need to check that the sequence of markings $\widetilde{f}_n$ is asymptotically turning number preserving. Since the sequence of markings of the Y_n is asymptotically turning number preserving by definition, and the sequence f_n is as well by Proposition 10.6, so is their composition $\widetilde{f}_n$. ◻

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.

Given $(Y_n, \omega_n) \to (X, \boldsymbol{\eta})$ as in the above proof, we consider a fixed level i and continue to let $Y_{(i)}$ denote the union of the components at this level of the ε-thick part of Y. For each edge e with $\ell(e^+) = i$, we have a sequence of conformal maps $\upsilon_{n,e}^-\colon A_{n,e}^- \to Y_{n,(<i)}$ of the bottom plumbing annuli of the plumbing fixtures $V(\rho_{n,e})$ such that $(\upsilon_{n,e}^-)^*(\omega/\boldsymbol{t}_{n,\lceil i\rceil}^{\boldsymbol{a}}) = \Omega_e$, where

$$\Omega_e \;=\; (u^\kappa - r_e'(\boldsymbol{t}))\frac{du}{u} \quad \text{and} \quad \rho_{n,e} \;=\; \prod_{j=\ell(e^-)}^{i} t_{n,j}^{m_{e,j}}.$$

Lemma 10.10. *In the above situation, if R is sufficiently large (depending only on the geometry of (X, η)), then the maps $\upsilon_{n,e}^-$ eventually extend to conformal maps $\upsilon_{n,e}$ whose domains contain the annulus $V_{n,e}^\circ = \{\delta/\rho_{n,e} < |u| < \delta/\sqrt{R}\}$, such that $\upsilon_{n,e}^*(\omega/\boldsymbol{t}_{\lceil i\rceil}^{\boldsymbol{a}}) = \Omega_e$, and moreover, the images of $V_{n,e}^\circ$ and $V_{n,e'}^\circ$ are disjoint for any $e \neq e'$.*

Proof of Lemma 10.10. If R is sufficiently large, the subsurface below the outer boundary $\gamma_n = \gamma_{n,e}$ of $A_{n,e}^-$, given by $|v| = \delta/R$, is convex for R sufficiently small, since $r_e'(\boldsymbol{t}_n)/\boldsymbol{t}_{n,\lceil i\rceil}^{\boldsymbol{a}}$ tends to 0. We use orthogonal projection to γ_n to extend $\upsilon_{n,e}^-$. That is, we map the equidistant curve of distance ℓ to γ_n to the equidistant curve of distance ℓ to $\upsilon_{n,e}^-(\gamma_n)$, mapping geodesic rays orthogonal to γ_n to geodesic rays orthogonal to $\upsilon_{n,e}^-(\gamma_n)$.

This gives a well-defined conformal map $\upsilon_{n,e}$ of $V_{n,e}^\circ$ onto its image as long as the Ω_e-distance between the boundary components of $V_{n,e}^\circ$ is smaller than the ω_n-distance from any zero of $Y_{n,(i)}$ to $\upsilon_{n,e}^-(\gamma_n)$. The distance between these boundary components tends to $\delta^\kappa/\kappa R^{\kappa/2}$. Since $(Y_n, \omega_n) \to (X, \boldsymbol{\eta})$, it suffices to take R large enough that $\delta^\kappa/\kappa R^{\kappa/2}$ is smaller than the distance between any zero of $\eta_{(i)}$ and any nodal zero.

Similarly, to ensure the $V_{n,e}^\circ$ are disjoint, it suffices to take R large enough that $\delta^\kappa/\kappa R^{\kappa/2}$ is smaller than half of the distance between any two nodal zeros of $\eta_{(i)}$. ◻

We now deal with the local injectivity of the plumbing map.

Proposition 10.11. *The map $\Omega\,\mathrm{Pl}^v$ is injective in a neighborhood of any point $(X, \boldsymbol{\eta})$ in the deepest stratum $\Omega\mathcal{MD}_\Lambda^{s,\Lambda}$.*

Proof. Consider a sequence Y_n in $\Xi\mathcal{D}_\Lambda^{sv}$ converging to X in the deepest stratum where Y_n is obtained by plumbing two sequences of model differentials X_n^1 and X_n^2, both converging to X. We eventually have almost-diffeomorphisms $f_n^i\colon \overline{Y}_n \to \overline{X}_n^i$, defined in Proposition 10.6, and $h_n = f_n^1 \circ (f_n^2)^{-1}$ is conformal outside the disks $\mathcal{E}_e^\pm$ and $\mathcal{E}_h$. We homotope h_n on these disks to a conformal isomorphism, showing that the X_n^i are in fact the same sequence.

Alternatively, injectivity can be established by analyzing the proof of Proposition 10.9 to show that at every stage of the construction of the sequence X_n and the isomorphism $X_n \to Y_n$, there is eventually a unique choice. This uses in particular that the perturbed period coordinates are injective, established in Proposition 9.7. ◻

Corollary 10.12. *The vertical plumbing map is a local homeomorphism on $\mathcal{W}_\epsilon$.*

Proof. For points P 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. ◻

10.8. **The horizontal plumbing construction.** In this section, we complete the plumbing construction by plumbing the horizontal nodes of the family $(\mathcal{Y}^v, \omega) \to \mathcal{W}_\epsilon$ constructed by the vertical plumbing map $\Omega\,\mathrm{Pl}^v$ to obtain a generically smooth family $\mathcal{Y} \to \mathcal{W}_\epsilon \times \Delta^H$.

We enumerate the horizontal edges as $e_1, \ldots, e_H$ and label the branches through the corresponding nodes q_i by an arbitrary choice of sign. Each q_i is adjacent to two half-infinite cylinders which each contain at least one zero of ω. We denote by z_i^+ and z_i^- an arbitrary choice of a zero in the boundary of each cylinder. We then apply Theorem 4.1 to choose standard coordinates

$$v_i^+ \colon \mathcal{W}_\epsilon \times \Delta_1 \to \mathcal{Y}^v \quad \text{and} \quad v_i^- \colon \mathcal{W}_\epsilon \times \Delta_1 \to \mathcal{Y}^v$$

covering a neighborhood of q_i in the corresponding branch and such that

$$(v_i^+)^*(\omega) \;=\; r_i' \frac{du}{u} \quad \text{and} \quad (v_i^-)^*(\omega) \;=\; -r_i' \frac{dv}{v}.$$

As these standard coordinates are unique up to multiplication by an arbitrary constant, we normalize them by requiring that the zeros z_i^+ and z_i^- correspond to $u = 1$ and $v = 1$ respectively.

We define for each j the (horizontal) plumbing fixture

$$\mathbb{W}_j^\delta \;=\; \left\{(\boldsymbol{w}, \boldsymbol{t}, \boldsymbol{x}, u, v) \in \mathcal{W}_\epsilon \times \Delta_\epsilon^H \times \Delta_\delta^2 : uv = x_j\right\}\,, \tag{10.16}$$

equipped with the relative holomorphic one-form

$$\Omega_j \;=\; -r'_{e_j}(\boldsymbol{t}) \frac{du}{u} \;=\; r'_{e_j}(\boldsymbol{t}) \frac{dv}{v} \tag{10.17}$$

and define two families of annuli $\mathcal{A}_j^\pm \subset \mathbb{W}_j^1$ by removing the upper or lower branches of the singular fibers:

$$\mathcal{A}_j^+ \;=\; \mathbb{W}_j^1 \setminus \{v = 0\} \quad \text{and} \quad \mathcal{A}_j^- \;=\; \mathbb{W}_j^1 \setminus \{u = 0\}\,.$$

We define two families of conformal maps $\Upsilon_j^\pm \colon \mathbb{W}_j^1 \to \mathcal{X}'$ by

$$\Upsilon_j^+(\boldsymbol{w}, \boldsymbol{t}, \boldsymbol{x}, u, v) \;=\; v_j^+(\boldsymbol{w}, \boldsymbol{t}, u) \quad \text{and} \quad \Upsilon_j^-(\boldsymbol{w}, \boldsymbol{t}, \boldsymbol{x}, u, v) \;=\; v_j^-(\boldsymbol{w}, \boldsymbol{t}, v)\,, \tag{10.18}$$

which identify $\mathcal{A}_j^\pm$ with families of annuli $\mathcal{B}_j^\pm \subset \mathcal{Y}^v \times \Delta^H$. Note that, in contrast to the vertical plumbing construction, these families of annuli have moduli tending to ∞ over the singular fibers.

We define the plumbed family $\mathcal{Y} \to \mathcal{W}_\epsilon \times \Delta^H$ by removing from $\mathcal{Y}^v$ the disks bounded by the annuli $\mathcal{B}_j^\pm$ and gluing in the plumbing fixtures $\mathbb{W}_j^1$ by the maps $\Upsilon_j^\pm$. As the gluing maps preserve the one-forms, $\mathcal{Y}$ is equipped with a relative one-form which we continue to denote by ω.

Alternatively, in terms of flat geometry, the family $\mathcal{Y}$ can be obtained by cutting each half-cylinder bounding q_i along a closed geodesic, and gluing the corresponding boundary components by an isometry to obtain a finite annulus C_j. The heights of the core geodesics and the gluing maps are determined by requiring that

$$\int_{\gamma_j} \omega \;=\; r'_j \log x_j,$$

where the integral is along a curve γ_j in C_j from the chosen zero z_j^- to z_j^+.

As for the vertical plumbing construction, standard universal properties applied to the restriction of $\mathcal{Y}$ to the strata $\mathcal{W}_\epsilon^J$ define the plumbing map $\Omega\,\mathrm{Pl}\colon \mathcal{W}_\epsilon \times \Delta^H \to \Xi\mathcal{D}_\Lambda^s$.

Proof of Theorem 10.1. We claim that $\Omega\,\mathrm{Pl}$ is a homeomorphism onto its image, the rest of the statements of the theorem following from the construction and having already been established.

To see this, consider a sequence $(X_n, \boldsymbol{t}_n)$ and point $(X, \boldsymbol{t})$ in $\mathcal{W}_\epsilon \times \Delta^H$ (which we are implicitly identifying with $\Xi\mathcal{D}_\Lambda^{sv} \times \Delta^H$ via the vertical plumbing map), and let Y_n and Y in $\Xi\mathcal{D}_\Lambda^s$ be the corresponding horizontally plumbed surfaces. For each annulus $C_i \subset Y$, fix a subannulus $C'_i \subset C_i$ with geodesic boundary, and let $Y' \subset Y$ be the complement of the C'_i. By the definition of the plumbing construction, Y' is canonically identified with a subsurface $X' \subset X$.

Suppose $(X_n, \boldsymbol{t}_n)$ converges to $(X, \boldsymbol{t})$, which is exhibited by a sequence of maps $g_n\colon K_n \to X_n$ defined on an exhaustion of X. These g_n are eventually defined on X' and define maps $h_n\colon J_n \to Y_n$, defined on an exhaustion J_n of Y'. These h_n satisfy the required properties to exhibit convergence of Y_n to Y, except they need to be extended over the annuli C'_i to be defined on an exhaustion of Y. We do so using the straight line extension of Lemma 10.7.

The pullbacks $h_n^*(du/u)$ then converge to du/u. The extended maps h_n are then quasiconformal on an exhaustion of Y and satisfy the hypotheses of convergence in the quasiconformal topology on forms from Section 3.3. It follows that $Y_n \to Y$ in $\Xi\mathcal{D}_\Lambda^s$, so $\Omega\,\mathrm{Pl}$ is continuous.

If $Y_n \to Y$ in $\Xi\mathcal{D}_\Lambda^s$, convergence of X_n to X is proved similarly by transporting maps exhibiting convergence of Y_n to Y to an exhaustion of X', and then extending across the complementary punctured disks by the straight line extension to obtain quasiconformal maps on an exhaustion of X. Convergence of the $\boldsymbol{t}_n$ is obvious as they are relative periods of a convergent sequence of forms. It follows that $\Omega\,\mathrm{Pl}^{-1}$ is continuous, so $\Omega\,\mathrm{Pl}$ is a local homeomorphism as desired. ◻

10.9. **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 10.1. The desired properties of $\Omega\,\mathrm{Pl}$ were shown in Proposition 10.8 and Corollary 10.12. The equivariance of the map with respect to the group $K_\Lambda = \mathrm{Tw}_\Lambda/\mathrm{Tw}_\Lambda^s$ follows from the construction in Section 10.5, since K_Λ acts on the markings and the simple rescaling parameters only, and they both have been transported from the family over the model domain to the plumbed family. ◻

Proof of Theorem 10.2. We proceed inductively with respect to the partial order induced by undegenerating. The base case of the induction $\Lambda = \emptyset$ is simply $\Omega\mathcal{MD}^s_\emptyset = \Omega\mathcal{B}_\emptyset = \Omega\mathcal{T}^{pm}_\emptyset(\mu) = \Omega\mathcal{T}_{(\Sigma,\boldsymbol{s})}(\mu) = \Xi\mathcal{D}_\emptyset$.

For the induction step we consider 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 $\Xi\mathcal{D}^{\Lambda',s}_{\Lambda,\circ} = \Xi\mathcal{D}^s_{\Lambda'}/(\mathrm{Tw}^s_\Lambda/\mathrm{Tw}^s_{\Lambda'})$ of $\Xi\mathcal{D}^s_\Lambda$ since $\mathrm{Tw}^s_\Lambda/\mathrm{Tw}^s_{\Lambda'}$ acts properly discontinuously. The complex structure on the intersections $\Xi\mathcal{D}^{\Lambda_1,s}_{\Lambda,\circ}$ and $\Xi\mathcal{D}^{\Lambda_2,s}_{\Lambda,\circ}$ agrees, since it stems from the common undegeneration of Λ_1 and Λ_2. So far, we have obtained a complex structure on

$$\bigcup_{\Lambda' \rightsquigarrow \Lambda} \Xi\mathcal{D}^{\Lambda',s}_{\Lambda,\circ} \;=\; \Xi\mathcal{D}^s_\Lambda \setminus \Xi\mathcal{D}^{\Lambda,s}_\Lambda\,.$$

On the other hand, we can cover a neighborhood of the deepest boundary stratum $\Xi\mathcal{D}^{\Lambda,s}_\Lambda$ by open sets of the form $\Omega\,\mathrm{Pl}(U_i)$, where U_i is of the form $\mathcal{W}_\epsilon \times \Delta^H$. The union of all these sets cover all of $\Xi\mathcal{D}^s_\Lambda$. 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 $\Omega\overline{\mathcal{MD}}^s_\Lambda$, 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 $\Omega\mathcal{MD}^s_\Lambda$, its moduli map is holomorphic (as a map to $\Omega\mathcal{B}_\emptyset/\mathrm{Tw}^s_\Lambda$) and this completes the proof.

The complex structure on $\Xi\mathcal{D}_\Lambda$ stems from that on $\Xi\mathcal{D}^s_\Lambda$ and the K_Λ-equivariance of the plumbing map. ◻

10.10. **Horizontal extension of perturbed period coordinates.** We finish this section by extending the perturbed period coordinates discussed in Section 9 to the case with horizontal nodes. These can be regarded as a generalization of the classical period coordinates, giving explicit local coordinates at the boundary of $\Xi\overline{\mathcal{M}}_{g,n}(\mu)$. These coordinates will not be used in the current paper, but the construction is essential for further applications, e.g. in [CMZ20] or [BDG20].

More precisely, 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 (10.4) and (10.18). Let $\mathcal{Y} \to U = \mathcal{W} \times \Delta^H_\epsilon$ be the family that results from plumbing the horizontal nodes of $\mathcal{X} \to \mathcal{W}$ as in Section 10.8. Our goal now is to extend the perturbed period map PPer from Section 9.2 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*

$$\mathrm{ePP} \;=\; \mathrm{PPer} \times \mathrm{Phor}\colon U \to \mathbb{C}^{L^\bullet(\Lambda)} \times \oplus_i \mathcal{R}'_i \times \mathbb{C}^H\,, \tag{10.19}$$

where

$$\text{Phor}\colon \begin{cases} U & \to \quad \mathbb{C}^H, \\ ([X,\boldsymbol{\eta},\boldsymbol{t}],(x_j)) & \mapsto \quad \left(\boldsymbol{e}\left(\frac{\int_{\beta_j}\eta_{(i)}+\xi_{(i)}}{\int_{\alpha_j}\eta_{(i)}+\xi_{(i)}}\right)_{j=1,\dots,H}\right), \end{cases} \tag{10.20}$$

and for integration we use the f-images of the paths in the corresponding fiber of the family $\mathcal{Y}\to U$. 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 $\mathcal{Y}\to U$ 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 10.13. *The perturbed period map* ePP *is a local diffeomorphism in a neighborhood of* $\mathcal{W}\times\mathbf{0}$.

Proof. Using Proposition 9.7 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. ∎

Example 10.14. 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^2_ϵ. 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 $\text{Phor}|_{\mathbf{0}\times\Delta^2_\epsilon}$ is given by

$$(x_1,x_2)\mapsto(k_1x_1,k_2x_2)\,, \tag{10.21}$$

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 (10.16) 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 (10.17) 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 (10.21), where the k_i are the exponentials of $(c_i^1 + c_i^3)/r_i$.

11. Families of multi-scale differentials

In this section we define families of multi-scale differentials, generalizing the definition of a single multi-scale differential in Section 7. Eventually we will give an algebraic description of the stack of multi-scale differential as a blowup. The starting point will be a flat family of pointed stable curves $(\pi\colon \mathcal{X} \to B, \boldsymbol{z})$, over 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 $R_p\colon B_p \to \overline{T}^n_{\Gamma_p}$ to the normalization of the level rotation torus closure (recall the definition and details of this from Section 6.5), 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 11.9 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.

11.1. Germs of families of multi-scale differentials.

We will define all the notions locally first, so until Section 11.2 we will work with a family over the germ B_p of an analytic space B at a point $p \in B$.

Recall e.g. from [ACG11, Proposition X.2.1] that for each node q_e of X_p there is a function $f_e \in \mathcal{O}_{B,p}$, which we call a *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 $\mathcal{O}_{B,p}$. We will write $[f_e] \in \mathcal{O}_{B,p}/\mathcal{O}^*_{B,p}$ for its equivalence class.

Given an enhanced level graph Γ_p, suppose we have a morphism $R\colon B \to \overline{T}^n_{\Gamma_p}$. 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

$$f_e^{\kappa_e} \;=\; s_j \dots s_{i-1}\,. \tag{11.1}$$

Definition 11.1. A *rescaling ensemble* is a morphism $R\colon B \to \overline{T}^n_{\Gamma_p}$ 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 $\pi\colon \mathcal{X} \to B$. △

The s_i will be called the *rescaling parameters* determined by R, in parallel with the notion defined in Section 8.2. The rescaling ensemble R can be thought of as a choice of these parameters that satisfies (11.1) for each edge e of Γ_p, together with the choices of appropriate roots of these which define a lift to $\overline{T}^n_{\Gamma_p}$, see Proposition 6.10 for the precise statement.

Definition 11.2. 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 $X_{p,\le i}$ 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 $r_{e,(i)} \in \mathcal{O}_{B,p}$ to be the period of $\omega_{(i)}$ along the oriented 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 $U_i \cap U_j$ for some $s_k \in \mathcal{O}_{B,p}$ with $s_k(p) = 0$ (where $k = j, \dots, i-1$).

(2) For any edge e joining levels $j < i$ of Γ_p, there are germs of functions u_e, v_e defined on a neighborhood of the corresponding node in $\mathcal{X}$, and a germ of a function f_e on B, such that the family has local normal form $u_e v_e = f_e$, and in these coordinates

$$\omega_{(i)} \;=\; (u_e^{\kappa_e} + f_e^{\kappa_e} r_{e,(i)})\frac{du_e}{u_e} \quad \text{and} \quad \omega_{(j)} \;=\; -(v_e^{-\kappa_e} + r_{e,(j)})\frac{dv_e}{v_e}\,, \tag{11.2}$$

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

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

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 $\boldsymbol{\omega}_p$. △

Some remarks to unravel the meaning of this definition are in order. Condition (2) is often automatic from Theorem 4.3. However, for the case of semipersistent nodes defined below that Theorem does not apply, and the condition needs to be imposed.

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 $j \le k \le i-1$ vanishes in a neighborhood of p on B.

Conditions (4) and (1) together imply the usual global residue condition in each nearby fiber X_q (using the enhanced structure of Γ_{X_q} which we define below by undegenerating from Γ_p). 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 $\ell(e^-) \le j < i$ we have

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

In particular, if $s_j = 0$ for some $\ell(e^-) \le j$, then $r_{e,(i)} = 0$. Consequently, the relations reflect the global residue condition as stated in Proposition 9.2, as we now explore in an Example.

Example 11.3. (*Definition 11.2 extends to fiberwise GRC.*) We consider the level graph given by Figure 7, 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 $B = \operatorname{Spec}\mathbb{C}[t]/(t^3)$, and let $r_{e,(i)}\colon B \to \mathbb{C}$ be as in Definition 11.2. The usual GRC from Section 2.4 states that the residues at the point $t = 0$ satisfy $r_{e_1,(-1)}(0) = r_{e_2,(-1)}(0) = 0$.

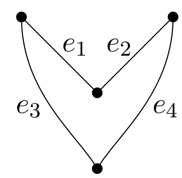

FIGURE 7. An enhanced level graph illustrating versions of the GRC.

Since the vanishing cycles corresponding to the edges e_1 and e_3 are homologous, Definition 11.2 (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 $t = 0$, and imposes a stronger constraint on the higher order terms of the expansion of the residues in t.

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 11.4. Given a germ of a family $\pi\colon \mathcal{X} \to B_p$ 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 $f_e^{\kappa_e} = 0$. △

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 subspace cut out by $u_e = v_e = 0$, so that Q_e can be thought of as the image of a nodal section $B \to \mathcal{X}$. 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 in each branch, 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 (11.2) (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 $\theta_\pm$ 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 subspace $B_e \subset B$ 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 (11.2). Any two possible choices of u_e and v_e are of the form αu_e and $\alpha^{-1} v_e$ for some unit $\alpha \in \mathcal{O}^*_{B,p}$, so the induced prong-matching does not depend on this choice.

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

Definition 11.5. Given a family of pointed stable curves $(\pi\colon \mathcal{X} \to B, \boldsymbol{z})$ and B_p a germ of B at p, the *germ of a family of multi-scale differentials* of type μ over B_p is an equivalence class of the following set of data:

1. the structure of an enhanced level graph on the dual graph Γ_p of the fiber X_p,
2. a rescaling ensemble $R\colon B \to \overline{T}^n_{\Gamma_p}$, compatible with
3. a collection of rescaled differentials $\boldsymbol{\omega} = (\omega_{(i)})_{i \in L^\bullet(\Gamma_p)}$ of type μ, and
4. a collection of prong-matchings $\boldsymbol{\sigma} = (\sigma_e)_{e \in E(\Gamma)^v}$, 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.

The $\mathcal{O}_{B,p}$-valued points of the level rotation torus $T_{\Gamma_p}(\mathcal{O}_B)$ act 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}.

Replacing $T_{\Gamma_p}(\mathcal{O}_{B,p})$ with the extended level rotation torus $T^\bullet_{\Gamma_p}(\mathcal{O}_{B,p})$, the analogous object is called a *germ of family of projectivized multi-scale differentials.* △

Remark 11.6. Note that over a reduced point, this definition of a family of multi-scale differentials agrees with Definition 7.1.

A *morphism between two germs of multi-scale differentials*

$$(\pi'\colon \mathcal{X}' \to B', \boldsymbol{z}', \Gamma_{p'}, \boldsymbol{\omega}', \boldsymbol{\sigma}') \to (\pi\colon \mathcal{X} \to B, \boldsymbol{z}, \Gamma_p, \boldsymbol{\omega}, \boldsymbol{\sigma}) \tag{11.3}$$

is a pair of germs of morphisms $\phi\colon B' \to B$ and $\widetilde{\phi}\colon \mathcal{X}' \to \mathcal{X}$ that jointly define a morphism of families of pointed stable curves (see [ACG11, p. 281]), such that the induced isomorphism of dual graphs $\Gamma_{p'} \to \Gamma_p$ is also an isomorphism of enhanced level graphs and such that $\widetilde{\phi}^*(\boldsymbol{\omega}, \boldsymbol{\sigma})$ is equivalent to $(\boldsymbol{\omega}', \boldsymbol{\sigma}')$.

For later use in the context of marked differential, and as an auxiliary object in the next subsection we define a *simple rescaling ensemble* to be a (germ of a) morphism $R^s\colon B \to \overline{T}^s_{\Gamma_p}$ to the simple level rotation torus closure, such that the composition with $\overline{T}^s_{\Gamma_p} \to \overline{T}^n_{\Gamma_p}$ is a rescaling ensemble in the above sense. Concretely, the map R^s is given by *simple rescaling functions* $t_i\colon B \to \mathbb{C}$ (in terminology consistent with that in

Section 8.2. Given a simple rescaling ensemble R^s, the associated (non-simple) rescaling ensemble is obtained by taking

$$s_i \;=\; t_i^{a_i} \quad \text{and} \quad f_e \;=\; \prod_{i=\ell(e^-)}^{\ell(e^+)-1} t_i^{m_{i,e}}\,, \tag{11.4}$$

similarly to (6.11).

11.2. **Restriction of germs to nearby points.** Now we allow B to be any complex analytic 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 $(\Gamma_p, \boldsymbol{\omega} = (\omega_{(i)}), \boldsymbol{\sigma} = (\sigma_e))$ with $\boldsymbol{\omega}$ compatible with R_p. Let $U \subset B$ be a neighborhood of p over which R_p and all σ_e are 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 $\mathrm{dg}\colon \Gamma_q \rightsquigarrow \Gamma_p$ as in Section 5.1. With notation of that section, this undegeneration is given by $J = \{j_{-1}, \ldots, j_{-M}\}$, where $j_k \in J$ if and only if $s_{j_k}(q) = 0$. Moreover recall that $j_0 = 0$ and $j_{-M-1} = -N-1$ (though they are not part of J). More precisely, the map of dual graphs $\delta\colon \Gamma_p \to \Gamma_q$ is 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 $f_e^{\kappa_e} = s_j \cdots s_{i-1}$, 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 k'th level of Γ_q corresponds to a maximal interval $(j_{k-1}, j_k]$ in $L^\bullet(\Gamma_p)$ such that $s_{j_k}(q) = 0$ for $j_k \in J$ and $s_i(q) \neq 0$ for every smaller i in this interval. The map δ is then compatible with the enhancements of these dual graphs and defines a degeneration dg, as desired.

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

Proposition 11.7. *Given a rescaling ensemble $R_p\colon B \to \overline{T}^n_{\Gamma_p}$, there exists a neighborhood $V \subset B$ of q, a rescaling ensemble $R_q\colon V \to \overline{T}^n_{\Gamma_q}$, and $\tau \in T_{\Gamma_p}$ such that*

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

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

Proof. We take the fiber product of R_p with the finite map $\overline{p}\colon \overline{T}^s_{\Gamma_p} \to \overline{T}^s_{\Gamma_p}/K_{\Gamma_p} = \overline{T}^n_{\Gamma_p}$ to obtain some simple rescaling ensemble $R_p^s\colon B^s \to \overline{T}^s_{\Gamma_p}$, defined on some ramified cover B^s of B. We solve the equation for the corresponding simple objects and then descend. Let $q' \in B^s$ denote some preimage of $q \in B$. Two consecutive 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 $t_i(q') \neq 0$. As in Lemma 6.6 the image of the monomorphism $\overline{\mathrm{dg}}^s_*\colon \overline{T}^s_{\Gamma_{q'}} \to \overline{T}^s_{\Gamma_p}$ of simple level rotation tori is cut out by $t_i = 1$ for all levels i such that the images of level i and $i+1$ are the same in $L^\bullet(\Gamma_q)$. We define $\tau^s \in T^s_{\Gamma_p}$ by

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

This ensures that $\tau^s \cdot R^s_p$ is in the image of dg^s_* and so there exists a simple rescaling ensemble R^s_q defined on a neighborhood of q' in B' such that $\overline{\mathrm{dg}}^s_* \circ R^s_q \;=\; \tau^s \cdot R^s_p$.

The multiplication map τ^s is K_{Γ_p}-equivariant, since the torus $T^s_{\Gamma_p}$ is commutative. Since R^s_p and $\overline{\mathrm{dg}}^s_*$ are K_{Γ_p}-equivariant by construction as fiber product, the map R^s_q descends to the required map R_q, and we let $\tau = \overline{p}(\tau^s)$.

To show the uniqueness of τ up to the action of T_{Γ_q}, observe 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 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}. ◻

Third, we define the collection of rescaled differentials at q. For each $k \in L^\bullet(\Gamma_q)$ let $(j_{k-1}, j_k]$ be its preimage in $L^\bullet(\Gamma_p)$. We act on $(\boldsymbol{\omega}, \boldsymbol{\sigma})$ by the τ from the preceding Proposition. The rescaling ensemble $\overline{\mathrm{dg}}_* \circ R_q$ has $s_i = 1$ for each $i \in (j_{k-1}, j_k)$ and moreover, for any edge e of Γ_p joining two levels in this interval, we have $f_e = 1$. By Condition (1) in Definition 11.2 the restriction of $(\tau \cdot \boldsymbol{\omega})_{(i)}$ to a neighborhood of the fiber over q agree for $i \in (j_{k-1}, j_k]$ on their overlap. So we define $\boldsymbol{\omega}_q$ to be the collection of differentials over q obtained by this gluing for all $k \in L^\bullet(\Gamma_q)$.

The last datum to define is a prong-matching for each edge of Γ_q. An edge e of Γ_p persists in Γ_q exactly when $q \in B_e$, the subscheme defined by $f_e = 0$. The prong-matching σ_e is a section of $\mathcal{P}_e$ and as such restricts to a germ of a section over the neighborhood of q intersected with B_e.

11.3. **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 11.8. Given a family of pointed stable curves $(\pi\colon \mathcal{X} \to B, \boldsymbol{z})$, 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. △

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

Given a family of multi-scale differentials over B and a map $\varphi\colon B' \to B$, 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 $\mathbf{MS}_\mu\colon$ (Analytic spaces) $\to$

(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* $\mathcal{MS}_\mu$ that retains the information of isomorphisms (Section X.12 of [ACG11] provides a textbook introduction, highlighting the difference between $\mathcal{MS}_\mu$ and $\mathbf{MS}_\mu$). 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 11.4.

Proposition 11.9. *If B is a smooth curve, then giving a multi-scale differential of type μ on a family $\mathcal{X} \to B$ 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 11.13 below implies that the family $(\mathcal{X} \to B, \omega)$ is adjustable and hence orderly (see Definitions 11.11 and 11.15 below). The claim then follows from Proposition 11.16. ◻

In contrast to this we observe:

Example 11.10. (*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 $\omega_{(0)}$ may be 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.

11.4. **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 14.1.

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

Definition 11.11. A family of pointed stable differentials $(\pi\colon \mathcal{X} \to B, \omega, \boldsymbol{z})$ 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 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

- there exists a non-zero regular function $h \in \mathcal{O}_{B,p} \setminus \{0\}$ such that $\omega = h\eta$,
- the differentials do not vanish identically on X,
- and $\eta|_X$ 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 $(\mathcal{X}, \omega)$ at the component X, and η is called an *adjusted differential.* △

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 11.16).

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 11.12. (*Adjusting parameters may not exist.*) Recall from [BCGGM18, Example 3.2] that there exist pointed stable differentials whose associated twisted differentials are not unique. Consider such a pointed stable differential (X, ω) and two distinct associated twisted differentials (X, η_1) and (X, η_2). Take two families of generically smooth stable differentials $(\pi_i \colon \mathcal{X}_i \to B_i, \omega_i)$ above smooth curves B_i for $i = 1, 2$, such that the adjusted differentials induce the twisted differentials η_i at the points p_i. Construct a nodal curve B by taking the union of B_1 and B_2 glued at p_1 and p_2. Since the fiber over p_1 of the family of stable differentials $(\mathcal{X}_1, \omega_1)$ coincides with the fiber of $(\mathcal{X}_2, \omega_2)$ over p_2, we can glue $\mathcal{X}_1$ and $\mathcal{X}_2$ to form a family of pointed stable differentials $(\pi \colon \mathcal{X} \to B, \omega)$. This family is not adjustable since the adjusted differentials of the two branches do not coincide over the node of the base curve B.

Next we show that if the base is sufficiently nice, adjusting parameters do exist.

Proposition 11.13. *If the base B is normal, then any family $(\mathcal{X}, \omega)$ satisfying the standing assumption is adjustable. Moreover, any two adjusting parameters for a given point $p \in B$ and a given irreducible component X of X_p differ by multiplication by a unit in $\mathcal{O}_{B,p}$.*

We first recall some terminology. Denote $\mathcal{Z} = \sum_{j=1}^{n} m_j \mathcal{Z}_j$ a divisor on $\mathcal{X}$, where $\mathcal{Z}_j \subset \mathcal{X}$ is the image of the section of the j-th zeros or poles z_j of ω. We call an effective Cartier divisor $V \subset \mathcal{X}$ a *vertical* divisor if the image $\pi(V) \subset B$ is a divisor. Note that any section $B \to \mathcal{X}$ is not vertical because it projects to all of B. 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 $V \subset W$ be the vertical zero divisor of ω in W, i.e. V 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 V contains the generic point of X. Since V is the zero locus of a holomorphic section of a line bundle, it is an effective Cartier divisor (possibly reducible and non-reduced), and we denote by $h \in \mathcal{O}_W \cong \mathcal{O}_{U \times \Delta}$ the local defining equation of V.

We claim that h does not depend on the second factor Δ, i.e. h can be regarded as a function defined on the base U. If this were not the case, then for a generic point $b \in U$ we would be able to solve the equation $h(b, x) = 0$, but then V would map onto U, which contradicts that V is a vertical divisor. We thus conclude that $h \in \mathcal{O}_U$.

Let $W' \subset W$ be the smooth locus of W. Then $(h^{-1}\omega)|_{W'}$ is holomorphic and can have horizontal zeros only, as the vertical zero divisor V is canceled out by h^{-1}.

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's theorem, $(h^{-1}\omega)|_{W'}$ extends to W holomorphically and can still have horizontal zeros only, as the zero locus of $h^{-1}\omega$ must be of codimension one (if not empty). It implies that the zero locus of $h^{-1}\omega$ in W does not contain the generic point of X, and hence $(h^{-1}\omega)|_{X\cap W}$ is holomorphic and not identically zero. Since $X \cap W$ contains the generic point of X, it follows that $h^{-1}\omega$ does not vanish identically on X. Thus h is the desired adjusting parameter for X and $\eta = h^{-1}\omega$ is the corresponding adjusted differential.

Suppose that h_1 is another adjusting parameter for X. Note that $h_1^{-1}\omega = (h/h_1)\eta$. If h/h_1 has a zero or pole at p, then $h_1^{-1}\omega$ would have a zero or pole along the entire 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 $\mathcal{O}_{B,p}$. ∎

In the algebraic situation, we show that adjusting parameters exist étale locally, which will be used in the proof of Theorem 14.8 as a step towards the algebraicity of the moduli space of multi-scale differentials.

Proposition 11.14. *Under the assumptions of Proposition 11.13, if moreover the family $(\mathcal{X} \to B, \omega)$ is algebraic with B irreducible and smooth generic fiber, then there exists an étale base change $B' \to B$ and a preimage p' of p such that the adjusting parameters for the pullback family $(\mathcal{X}' \to B', \omega')$ at p' can be chosen in the algebraic local ring $\mathcal{O}^{\mathrm{alg}}_{B',p'}$.*

Proof. The existence of a function in the local ring after an étale base change of $R = \mathcal{O}^{\mathrm{alg}}_{B,p}$ is by definition equivalent to finding such a function in the (strict) Henselization R^h of R. Let $h_{\mathrm{an}} \in \mathcal{O}^{\mathrm{an}}_{B,p}$ be an analytic adjusting parameter for an irreducible component X of X_p provided by Proposition 11.13. We view h_{an} as an element of the local ring completion $\widehat{R}$. The proof consists of two steps. First we show that there exists an algebraic function $h \in R$ such that h_{an} divides h, as elements of $\widehat{R}$. Secondly, we show that for any factorization $h = h_1 \cdot h_2$ in $\widehat{R}$ there exist $h_1', h_2' \in R^h$ with $h = h_1' \cdot h_2'$ and such that h_i/h_i' is a unit in $\widehat{R}$. Applying this to $h_1 = h_{\mathrm{an}}$ gives the result by taking h_1' as the desired (algebraic) adjusting parameter.

For the first claim, note that the vertical divisor V (as in the proof of Proposition 11.13) is contained in the locus of singular fibers $S \subset \mathcal{X}$, which is the π-preimage of a Cartier divisor $D \subset B$. Let m be the vanishing order of V at a generic point near the component X under consideration. Then h_{an} divides the defining equation h of mS, where h can be regarded as a function in R defining the Cartier divisor mD.

For the second claim, consider the factorization as decomposition of the associated Cartier divisor $D = D_1 \cup D_2$, where $D = V(h)$ and $D_i = V(h_i)$, in $\mathrm{Spec}(\widehat{R})$. Note that a prime ideal of R^h remains prime after lifting in $\widehat{R}$, as a consequence of Artin approximation (see e.g. [Haz00, Section 3A, Proposition 1.9]). It follows that the decomposition of D induces a decomposition $D = D_1' \cup D_2'$ in $\mathrm{Spec}(R^h)$ into (a priori Weil) divisors with $D_i' = V(I_i')$ and $I_i' \otimes_{R^h} \widehat{R} = \langle h_i \rangle$. Since $R^h \to \widehat{R}$ is a faithfully flat morphism, this implies that I_i' is locally free of rank one, too, and their generators h_i' are the elements we wanted to construct. ∎

Suppose for the rest of this section that $(\mathcal{X}, \omega)$ 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 $h_2 \preccurlyeq h_1$ 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 each of the functions f_e and h_v is only defined up to multiplication by some 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 $v \preccurlyeq v'$ if and only if $h_v \preccurlyeq h_{v'}$. In this case, we in fact have

$$h_v/h_{v'} \;=\; f_e^{\kappa_e} \tag{11.5}$$

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 14.12). We will be especially interested in families for which it is a full order.

Definition 11.15. An adjustable family $(\pi\colon \mathcal{X} \to B, \omega)$ is called *orderly* if for every point $p \in B$, the divisibility relation induces a full order on the set of adjusting parameters $(h_v)_{v \in \Gamma_p}$. △

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 11.16. *An orderly family $(\pi\colon \mathcal{X} \to B, \omega, \boldsymbol{z})$ 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 11.5. Namely, such $(\mathcal{X}, \omega, \boldsymbol{z})$ determines a unique family of multi-scale differentials of type μ.*

Proof. The divisibility order of the family $(\mathcal{X}, \omega, \boldsymbol{z})$ 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 $s_0 = h_0 = 1$ and

$$s_i \;:=\; h_i/h_{i+1} \tag{11.6}$$

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

$$\omega_{(i)} := \omega/(s_0 \ldots s_i) = \omega/h_i \tag{11.7}$$

which is generically holomorphic and non-zero on each level i component of X_p, vanishes identically on all lower lever 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 (11.2) 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 11.2.

We will show the existence of a compatible rescaling ensemble $R\colon B \to \overline{T}^n_{\Gamma_p}$ in three steps. First, as a consequence of Theorem 4.3, a map $R'\colon B \to \overline{T}'_{\Gamma_p}$ can always be found by using the tuples s_i and f_e as above such that they satisfy (11.1), where $\overline{T}'_{\Gamma_p}$ denotes the entire torus cut out by Equation (6.10). Next, we want the image of R' to lie in the desired connected component $\overline{T}_{\Gamma_p}$ of $\overline{T}'_{\Gamma_p}$, and this can be done as follows. The torus $\overline{T}'_{\Gamma_p}$ has a map to $(\mathbb{C}^*)^N$ by projection, which is an isogeny since these are two tori of the same dimension. Choose in each connected component of $\overline{T}'_{\Gamma_p}$ an element in the kernel of this 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, and thus by choosing the suitable kernel elements the whole collection can lie in the connected component $\overline{T}_{\Gamma_p}$. Finally, if the base B is normal, the map $R'\colon B \to \overline{T}_{\Gamma_p}$ automatically factors through the normalization of the level rotation torus closure, by the universal property of normalization, and thus gives the rescaling ensemble $R\colon B \to \overline{T}^n_{\Gamma_p}$. ◻

12. Real oriented blowups

The goal of this section is to define a canonical real oriented blowup for a family of multi-scale differentials; see Section 5.3 for a discussion of welding in terms of the real oriented blowup for the case of one Riemann surface. 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 12.2. In fact there are two real oriented blowups, associated with a rescaling ensemble R and with a simple rescaling ensemble R^s.

In Section 12.4 we develop the parallel case of families of marked model differentials.

12.1. **The real blowup construction.**

We start with the local version, which only depends on the rescaling ensemble.

Proposition 12.1. *Let $\pi\colon \mathcal{X} \to B_p$ 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 $\widehat{\pi}\colon \widehat{\mathcal{X}} \to \widehat{B}_p$ of topological spaces with the following properties:*

(i) *There are surjective differentiable maps $\varphi_{\mathcal{X}}\colon \widehat{\mathcal{X}} \to \mathcal{X}$ and $\varphi_B\colon \widehat{B}_p \to B_p$ such that $\pi \circ \varphi_{\mathcal{X}} = \varphi_B \circ \widehat{\pi}$.*
(ii) *All fibers of $\widehat{\pi}$ are vertically welded surfaces (in the sense of Section 5.3).*
(iii) *The fiber of φ_B over the point $p \in B_p$ 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 $B'_p \to B_p$ of the base.

Note that the level-wise real blowup does not modify the neighborhoods of horizontal nodes. Hence in general the fibers of $\widehat{\pi}$ remain nodal.

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 now perform the above construction globally.

Theorem 12.2. *A family of multi-scale differentials $(\boldsymbol{\omega}, \boldsymbol{\sigma})$ on $\pi\colon \mathcal{X} \to B$ 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 $\overline{\pi}_p\colon \overline{\mathcal{X}} \to \overline{B_p}$ the restriction of the local level-wise real blowup to $\overline{B_p}$.*

These germs glue to a global surjective differentiable map $\overline{\pi}\colon \overline{\mathcal{X}} \to \overline{B}$, the (global) level-wise real blowup. *Moreover, the global level-wise real blowup is functorial under pullbacks via maps $B' \to B$ 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 12.1. 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 on B_p. We introduce for each of the variables an S^1-valued partner variable, denoted by the corresponding capital letter. Concretely, we define $\widehat{B}_p \subset B_p \times (S^1)^{E(\Gamma_p)^v} \times (S^1)^{L(\Gamma_p)}$ by the equations

$$F_e|f_e| \;=\; f_e\,, \quad S_i|s_i| \;=\; s_i\,, \quad \text{and} \quad F_e^{\kappa_e} \;=\; S_j \dots S_{i-1}\,, \tag{12.1}$$

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 $S_i \in S^1$ (resp. F_e) is an independent variable, not related to s_i or f_e. The map $\varphi_B : \widehat{B}_p \to B_p$ is then given by the projection onto the first factor.

Next we define the family $\widehat{\pi}\colon \widehat{\mathcal{X}} \to \widehat{B}$ 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

$$U_e^{\pm}|u_e^{\pm}| \;=\; u_e^{\pm} \quad \text{and} \quad U_e^+ U_e^- \;=\; F_e\,. \tag{12.2}$$

The fibers of $\widehat{\pi}$ are not yet smooth in a neighborhood of the preimages of the vertical nodes (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 form the horizontal nodes that can be checked to agree with the one of a welding.

The functoriality of this construction is obvious. ∎

Proof of Theorem 12.2. In view of Remark 11.6, Corollary 6.7 and the definition of $\hat{B}$ given in Equation (12.1) imply 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. ∎

In the special case that all nodes of π are persistent, the base $\overline{B}$ of the real oriented blowup is isomorphic to $B \times (S^1)^{L(\Gamma_p)}$, with parameters $\boldsymbol{S} = (S_i)$. Denoting by $\theta(\boldsymbol{S})$ the argument of the S_i divided by 2π, the fiber of $\overline{\pi}$ over a point $(p, \boldsymbol{S}) \in \overline{B}$ is simply the surface $\mathcal{X}_p$ welded according to the prong-matching $\theta(\boldsymbol{S}) \cdot \boldsymbol{\sigma}$, where this map is defined in (6.2). This also justifies the use of overlines for both constructions.

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

For any fixed subgroup G of the mapping class group $\mathrm{Mod}_{g,n}$ we similarly define the *presheaf of G-markings* $\mathrm{Mark}(\mathcal{Y}/B; G)$ 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 $\mathrm{Mark}(\mathcal{Y}/B; G)$.

We can now define the marked version of families of multi-scale differentials. It starts with germs and glues them by sheafification, as in the unmarked case. Let Λ be an enhanced multicurve and $\Gamma = \Gamma(\Lambda)$ the underlying enhanced level graph.

Definition 12.3. Given a family of pointed stable curves $(\pi\colon \mathcal{X} \to B, \boldsymbol{z})$ and B_p a germ of B at p, the *germ of a Λ-marked family of multi-scale differentials* of type μ over B_p is an equivalence class of the following set of data:

(i) a germ $(\pi\colon \mathcal{X} \to B, \boldsymbol{z}, \Gamma_p, \boldsymbol{\omega}', \boldsymbol{\sigma}')$ of a family of multi-scale differentials, and
(ii) a Tw_Λ-marking f of the level-wise real oriented blowup $\overline{\pi}\colon \overline{\mathcal{X}} \to \overline{B}_p$.

The level rotation torus $T_{\Gamma_p}(\mathcal{O}_B)$ acts on all of the above data (see (6.3) for the action on the marking) and we consider two germs *equivalent* if they differ by the action of an element in $T_{\Gamma_p}(\mathcal{O}_B)$.

A morphism between two germs $(\mathcal{X}', \boldsymbol{\omega}', \boldsymbol{\sigma}', f')$ and $(\mathcal{X}, \boldsymbol{\omega}, \boldsymbol{\sigma}, f)$ is a morphism $(\phi, \widetilde{\phi})$ of the underlying multi-scale differentials such that the induced map $\overline{g}\colon \overline{\mathcal{X}'} \to \overline{\mathcal{X}}$ commutes with the Tw_Λ-marking, up to an isotopy respecting the marked points. △

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

Given a map $\psi\colon B' \to B$, the functoriality of the level-wise real blowup allows to define the pullback of markings along ψ by pulling back the germ of the family as in Section 11 and by restricting the markings along the induced map $\overline{\psi}\colon \overline{B'} \to \overline{B}$. In this way we define a *family of* Λ*-marked multi-scale differentials* by sheafification, just as in Definition (11.8). We have thus defined a moduli functor $\mathbf{MS}_{(\mu,\Lambda)}$ of Λ-marked multi-scale differentials. The notion of a family of Λ-marked multi-scale differentials and this functor has an obvious projectivized version, denoted by $\mathbb{P}\mathbf{MS}_{(\mu,\Lambda)}$.

12.3. **Families of simple marked multi-scale differentials.** There is a similar definition of the notion of a family of simple marked multi-scale differentials of type μ. We aim to require a Tw^s_Λ-marking rather than merely a Tw_Λ-marking but this requires to change the blowup. Suppose $\pi\colon \mathcal{X} \to B$ is a germ of a family of curves with a simple rescaling ensemble R^s. Then a version of Proposition 12.1 holds verbatim, but the torus in its point (iii) is the real torus associated with the simple level rotation torus rather than associated with the level rotation torus. To prove this version of the Proposition we introduce for each of the variables of T_Λ an S^1-valued partner variable, denoted by the corresponding capital letter. Concretely, we define $\widehat{B} \subset B \times (S^1)^{L(\Gamma_p)}$ by the equations

$$T_i \cdot |t_i| \;=\; t_i\,. \tag{12.3}$$

(This torus is a finite cover of the one given in (12.1). The map is given by letting $S_i = T_i^{a_i}$ and $F_e \;=\; T_j^{m_{e,j}} \ldots T_{i-1}^{m_{e,i-1}}$, if e is an edge connecting levels $j < i$ and where the $m_{e,i}$ were defined in (6.7). The remaining construction of the smooth family over this blowup is the same as above.) This procedure should properly be called simple real oriented blowup, but the context (the rescaling ensemble) will always make it clear which version we use.

The zoo of definitions given so far culminates in the following, the moduli functor that will turn out to be indeed represented by a smooth space.

Definition 12.4. Given a family of pointed stable curves $(\pi\colon \mathcal{X} \to B, \boldsymbol{z})$ and B_p a germ of B at p, the *germ of a family of simple* Λ*-marked multi-scale differentials* of type μ over B_p is an equivalence class of the following set of data:

(1) the structure of an enhanced level graph on the dual graph Γ_p of the fiber X_p,
(2) a simple rescaling ensemble $R^s\colon B \to \overline{T}^s_{\Gamma_p}$, compatible with
(3) a collection of rescaled differentials $\boldsymbol{\omega} = (\omega_{(i)})_{i \in L^\bullet(\Gamma_p)}$ of type μ, and
(4) a collection of prong-matchings $\boldsymbol{\sigma} \;=\; (\sigma_e)_{e \in E(\Gamma)^v}$. For the non-semipersistent nodes, these are required to agree with the induced prong-matchings defined before Definition 11.5.
(5) a Tw^s_Λ-marking f of the level-wise simple real oriented blowup $\overline{\pi}\colon \overline{\mathcal{X}} \to \overline{B}$.

The simple level rotation torus $T^s_{\Gamma_p}(\mathcal{O}_B)$ acts on all of the above data and we consider two germs *equivalent* if the differ by the action of an element in $T^s_{\Gamma_p}(\mathcal{O}_B)$. △

With the obvious definition of morphisms, pullbacks and sheafification, this defines the notion of a *family of simple marked multi-scale differentials*. This defines a functor that we denote by $\mathbf{MS}^s_{(\mu,\Lambda)}$ and its projectivized variant by $\mathbb{P}\mathbf{MS}^s_{(\mu,\Lambda)}$.

The group K_Γ acts on germs (and on families where the level graph is an undegeneration of Γ) by post-composing the marking f with the given element in K_Γ. The quotient functor is exactly the functor of (non-simple) marked multi-scale differentials, since in the presence of just the resulting quotient Tw_Λ-marking a simple rescaling ensemble up to the equivalence relation generated by $T^s_{\Gamma_p}(\mathcal{O}_B)$ is the same as a rescaling ensemble up to the equivalence relation generated by $T_{\Gamma_p}(\mathcal{O}_B)$.

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

If the family admits a simple rescaling ensemble R^s*, then there exists a neighborhood* U *of* p *such that* $\pi|_U$ *can be provided with a* Tw^s_Λ*-marking.*

Proof. We need to provide the level-wise real blowup $\overline{\pi}|_U$ with a Tw_Λ-marking f. For this purpose we take U to be simply connected, provide some fiber of $\overline{\pi}$ with a marking and transport the marking along local smooth trivializations of $\overline{\pi}$. 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 12.2 the fibers of $\overline{U} \to U$ 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.

The second statement follows in the same way using the simple version of the real blowup. □

12.4. **Families of marked model differentials.** To highlight similarities and differences, and for further use, we now define (the easier) families of marked and simple marked model differentials. This will be used to verify all the relevant universal properties of the model domain in Section 13.

Recall that 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. While for families of multi-scale differentials we needed to start with a germwise definition to be able to control degeneration, here we can give the global definition right away.

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

Definition 12.6. A *family of twisted differentials* $\boldsymbol{\eta}$ *of type* μ on an equisingular family $\pi\colon \mathcal{X} \to B$ 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 the obvious analogues the conditions in Section 2.4, interpreting the residues as regular functions on the base B. △

Definition 12.7. Let $(\pi\colon \mathcal{X} \to B, \boldsymbol{z})$ be an equisingular family of pointed stable curves. A *family of* Λ*-marked simple model differentials of type* μ *over* B is an equivalence class of the following set of data:

(1) the structure of an enhanced level graph on the dual graph Γ of any fiber of π,
(2) a simple rescaling ensemble $R^s\colon B \to \overline{T}_\Gamma^s$,
(3) a collection $\boldsymbol{\eta} = (\eta_{(i)})_{i\in L^\bullet(\Gamma)}$ of families of twisted differentials of type μ compatible with Γ,
(4) a collection $\boldsymbol{\sigma} = (\sigma_e)_{e\in E(\Gamma)^v}$ of prong-matchings for $\boldsymbol{\eta}$,
(5) a Tw_Λ^s-marking f of the level-wise (simple) real oriented blowup $\overline{\pi}\colon \overline{\mathcal{X}} \to \overline{B}$ defined using (12.3).

The simple level rotation torus $T_\Gamma^s(\mathcal{O}_B)$ acts on the above data, and two elements in the same orbit are defined to be equivalent.

Replacing $T_\Gamma^s(\mathcal{O}_B)$ with the extended level rotation torus $T_\Gamma^{s,\bullet}(\mathcal{O}_B)$, the analogous object is called a *family of marked simple projectivized model differentials of type* μ *over* B. △

The notion of a morphism is derived from morphisms of pointed stable curves as in (11.3). We denote the functor of model differentials by $\mathbf{MD}^s_{(\mu,\Lambda)}$ and its projectivized version by $\mathbb{P}\mathbf{MD}^s_{(\mu,\Lambda)}$.

Remark 12.8. Since a simple rescaling ensemble is simply a collection of functions $\boldsymbol{t} = (t_i)_{i\in L(\Gamma)}$ in $\mathcal{O}_B$ we will denote a family of simple marked model differentials interchangeably by the representatives $(\boldsymbol{\eta}, R^s, \boldsymbol{\sigma}, f)$ or by $(\boldsymbol{\eta}, \boldsymbol{t}, \boldsymbol{\sigma}, f)$ of the $T_\Gamma^s(\mathcal{O}_{B,p})$-orbits.

Similarly we can define (non-simple) marked model differentials by taking the K_Γ-quotient or by using non-simple rescaling examples; we can also define the unmarked versions. Since these other version will not be needed for what follows, we do not give the details.

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.*

To obtain this, we first prove the simple marked version of this 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* $\mathbf{MS}^s_{(\mu,\Lambda)}$ *of simple marked multi-scale differentials.*

Given a family $\pi\colon \mathcal{Y} \to B$ of stable curves with a family of simple marked multi-scale differentials $(\boldsymbol{\omega}, \boldsymbol{\sigma}, f)$, we want to construct functorially a map $m\colon B \to \Xi\mathcal{D}_\Lambda^s$ such that the pullback of the universal family agrees with the given family. Since the complex structure on $\Xi\mathcal{D}^s_\Lambda$ stems from the model domain, we will first establish the universal property of the model domain. The map m will be constructed by using the universal property of the model domain to map there, and then by plumbing using the map $\Omega\,\mathrm{Pl}$ defined in Section 10.

To be able to use the universal property of the model domain, we will need to define an unplumbing construction that takes multi-scale differentials on $\mathcal{Y}$ to model differentials on an equisingular family $\mathcal{X} \to B$. Like the plumbing, the unplumbing construction will depend on several choices, and we will need to carefully arrange the choices consistently on $\mathcal{Y}$ and on the universal family.

13.1. **The universal family over the model domain.** We first exhibit the functor that the space $\Omega\mathcal{T}^{pm}_\Lambda(\mu)$ represents. 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 (compare also to the definitions in Section 12.4).

Definition 13.3. 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 12.6,
(ii) a family of prong-matchings σ, and
(iii) a family of markings $f \in \mathrm{Mark}(\overline{\mathcal{X}}_\sigma/B)$ of the welded family. △

This definition is much simpler than Definition 12.3 or Definition 12.7 and does not require a blowup of the base since there is no equivalence relation by the action of a level-rotation torus, which has a non-trivial fundamental group.

We can now state the universal property of the space $\Omega\mathcal{T}^{pm}_\Lambda(\mu)$. The proof is rather obvious and mainly serves to recall notation.

Proposition 13.4. *Let* $\Lambda \subset \Sigma$ *be a fixed enhanced multicurve. The Teichmüller space of prong-matched twisted differentials* $\Omega\mathcal{T}^{pm}_\Lambda(\mu)$ *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) *over* B.

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 $\mathcal{T}_\Lambda$ of the classical augmented Teichmüller space comes with a universal family $(\pi\colon \mathcal{X} \to \mathcal{T}_\Lambda, \boldsymbol{z}, (f_v)_{v\in V(\Lambda)})$ of pointed stable curves equisingular of type $\Gamma(\Lambda)$, constructed by gluing families of smooth curves $\pi\colon \mathcal{X}_v \to \mathcal{T}_\Lambda$ along the nodes given by the marked sections $q_e^\pm$ corresponding to the edges e of $\Gamma(\Lambda)$. Here $f_v \in \mathrm{Mark}(\mathcal{X}_v/\mathcal{T}_\Lambda)$ 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 $(\mathcal{X}_v, \boldsymbol{z}_v, \boldsymbol{q}_e^\pm, f_v)$.

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

The family of model differentials over the model domain was already constructed in Section 8, and its universal property follows from the construction and from the universal property of the augmented Teichmüller space of flat surfaces.

Proposition 13.5. *The simple model domain $\Omega\overline{\mathcal{MD}}_\Lambda^s$ is the fine moduli space for the functor of simple marked model differentials $\mathbf{MD}^s_{(\mu,\Lambda)}$, and $\overline{\mathcal{MD}}_\Lambda^s$ is the fine moduli space for $\mathbb{P}\mathbf{MD}^s_{(\mu,\Lambda)}$.*

The model domain $\Omega\overline{\mathcal{MD}}_\Lambda$ (considered as quotient stack) is thus isomorphic to the functor of model differentials $\mathbf{MD}_{(\mu,\Lambda)}$.

Proof. Recall that as discussed in Section 8, the family over $\Omega\mathcal{MD}_\Lambda^s$ is simply the Tw_Λ^{sv}-quotient of the family over $\Omega\mathcal{T}_\Lambda^{pm}(\mu)$. We showed in Proposition 13.4 that the latter is the universal family of marked prong-matched twisted differentials. The family over the other strata of $\Omega\overline{\mathcal{MD}}_\Lambda^s$ is constructed by covering the space by charts, and considering the scale comparison. The universal property thus immediately follows from the universal property of $\Omega\mathcal{T}_\Lambda^{pm}(\mu)$. ∎

13.2. **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 then to record the degeneracy of the nodes as the parameters $\boldsymbol{t}$ of model differentials. Technically, we cannot pinch off curves without modifying the differential, due to the presence of non-trivial periods over what we want to be the vanishing (pinching) cycles. This forces us to subtract beforehand some perturbation differentials, whose role is inverse to that of the modifying differentials.

Proposition 13.6. *Given a germ of a family of simple marked multi-scale differentials with all data $(\mathcal{Y} \to B, (\omega_{(i)})_{i\in L^\bullet(\Gamma_p)}, R^s, \boldsymbol{\sigma}, f)$ defined over B, there is an unplumbing construction that produces a family $(\mathcal{X} \to B, (\eta_{(i)})_{i\in L^\bullet(\Gamma_p)}, R^s, \boldsymbol{\sigma}', f)$ of Λ-marked simple model differentials with the following properties:*

(i) *The construction is the identity over the locus B^Λ of all $q \in B$ such that $\Gamma_q = \Gamma(\Lambda)$.*

(ii) *The construction depends only on a finite number of choices of topological data and on a choice of a section near each vertical node.*

(iii) *If B is an open neighborhood of p in the simple Dehn space, then the map $u\colon B \to \Omega\overline{\mathcal{MD}}^s_\Lambda$ induced by the unplumbing of the universal family of model differentials, restricted over B, is a local biholomorphism.*

Proof. The unplumbing construction is level-wise, similarly to how plumbing was defined in Section 10. For simplicity of the exposition, we only treat in detail the case where Γ has two levels, and no horizontal nodes. We may thus write $\omega_{(0)} = s \cdot \omega_{(-1)}$, and $B^\Lambda \subset B$ is then precisely the vanishing locus of s.

For the definition of a perturbation differential we start by choosing some maximal multicurve $\Lambda_{\max} \supseteq \Lambda$. We denote by V the image of Λ in $H_1(\Sigma \setminus P_{\boldsymbol{s}}; \mathbb{Q})$ and, as in Proposition 9.3, we let V' be the subspace generated by curves in $\Lambda_{\max}$ and loops around points in $P_{\boldsymbol{s}}$. Let $\rho\colon B \to \mathrm{Hom}_{\mathbb{Q}}(V, \mathbb{C})$ be the periods of $\omega_{(0)}$ along Λ and let $\widetilde{\rho}$ be the extension of ρ by zero on a subset S of $\Lambda_{\max}$ 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 $\widetilde{\rho}$. A perturbation differential exists and it is uniquely determined by the choice of the topological datum $\Lambda_{\max}$ and the subset S. Since $\widetilde{\rho}$ is divisible by s, the perturbation differential vanishes identically 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. We fix such a section, to be further specified in Section 13.3. 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 disc Δ times the base with a one-form $\eta_{(-1)}$.

Using Theorem 4.1 (more precisely the subsequent remark, and moreover the chosen section, to specify the coordinates uniquely) we 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 $\eta_{(0)}$ does not necessarily have the correct orders of vanishing in the smooth locus. Hence, similarly to what we do in plumbing, 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 finally obtain an equisingular family $\pi\colon \mathcal{X} \to B$ by identifying the points $u = 0$ of $\mathcal{X}_{(0)}$ and $v = 0$ of $\mathcal{X}_{(-1)}$ in each plumbing fixture to form a node.

For an equisingular family, the space of prong-matchings is an unramified cover of the base. Thus, to obtain the prong-matching $\boldsymbol{\sigma}'$ we simply extend the prong-matching $\boldsymbol{\sigma}|_p$ 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 12 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 (12.2).) 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 in both cases they stem from the T^s_Λ-action for the differentials and prong-matchings, and from Tw^s_Λ for the markings.

Finally, to prove (iii) it suffices to show that the tangent map to u is surjective at any point of B^Λ. We argue similarly to the alternative proof of Proposition 10.11, using the fact that perturbed periods give local coordinates on the model domain, as shown in Proposition 9.7. Indeed, by construction, on the restriction of the family $\mathcal{Y}|_{B^\Lambda} \to B^\Lambda$ over B^Λ, the map u is the identity. Working in perturbed periods coordinates on the model domain, it thus suffices to show that the directions corresponding to changing the parameters $\boldsymbol{t}$ of the model differential are in the range of the tangent map to u. This is indeed obvious since those parameters are given by R^s, which is part of the datum of the unplumbed model differential. ∎

13.3. **Consistent unplumbing and the proof of Proposition 13.2.** In order to define the moduli map m we consider the results of applying the unplumbing construction in two situations. First, we apply unplumbing to $\mathcal{Y} \to B$, to obtain a family of model differentials on $\pi\colon \mathcal{X} \to B$. Second, we perform the unplumbing construction for the universal family of multi-scale differentials over $\Xi\mathcal{D}^s_\Lambda$, restricted to a neighborhood $W \subset \Xi\mathcal{D}^s_\Lambda$ of the moduli point of $\mathcal{X}_p$, to obtain a family of model differentials $\pi^{\mathrm{uni}}\colon \mathcal{X}^{\mathrm{uni}} \to W$. We want to perform these two plumbing constructions making all the choices consistently as follows. First, we choose a maximal multicurve $\Lambda_{\max}$ on $\mathcal{Y}_p$ as required for Proposition 13.6, and choose the same maximal multicurve on the universal family over W, which is possible since the surfaces are marked. Second, we choose the normalizing sections for the unplumbing of each node to lie in the neighborhood $V(f_e, \epsilon)$ and in such a way that their relative $\omega_{(\ell(e^-))}$-period to a marked zero on the same level as the lower end of the node is constant in the family. The markings, which are well-defined up to Tw^{sv}_Λ-twists, allow to consistently choose the paths for computing these relative periods.

Let $m'\colon B \to \overline{\mathcal{MD}}^s_\Lambda$ be the moduli map obtained by applying the universal property for the simple model domain to the unplumbed family $\pi\colon \mathcal{X} \to B$. Let u be the moduli map for the universal family as in Proposition 13.6 (iii). We claim that (after possibly shrinking B to fit domains) the composition

$$m \;=\; u^{-1} \circ m'\colon B \to \Xi\mathcal{D}^s_\Lambda$$

will then be the moduli map for B, i.e., that the family of multi-scale differentials $\mathcal{Y} \to B$ is the pullback of the universal family over $\Xi\mathcal{D}^s_\Lambda$ under the map m. By definition there is an isomorphism of families of model differentials $h'\colon (m')^*\mathcal{Y}^{\mathrm{uni}} \to \mathcal{X}$, and we need to exhibit an isomorphism of families of multi-scale differentials $h\colon m^*\mathcal{Y}^{\mathrm{uni}} \to \mathcal{Y}$.

This isomorphism is constructed level by level, and for clarity of exposition we will again only deal with the two-level situation without horizontal nodes, as in the proof of Proposition 13.6. In that setting, 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. This defines the map h on the families of lower subsurfaces. On the upper level subsurfaces a perturbation differential has been added. The consistent choice of $\Lambda_{\max}$ implies that the m'-pullback of the perturbation differential on $\mathcal{Y}_{(0)}^{\mathrm{uni}}$ agrees with that on $\mathcal{Y}_{(0)}$. Similarly, since all choices in both unplumbing constructions are the same, the local modifications near the marked zeros agree under m'-pullback. We thus define the map 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. Lastly, 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 fact that the sections (which determine the normal form uniquely) were chosen using the same relative period of $\omega_{(-1)}$. This completes the proof of the proposition for points p where $\Lambda_p = \Lambda$.

Next we deal with the situation that the pinched multicurve Λ_p is a strict undegeneration of Λ. We use Proposition 12.5 to provide the family with a $\mathrm{Tw}^s_{\Lambda_p}$-marking and the previous argument to obtain a moduli map $B \to \Xi\mathcal{D}^s_{\Lambda_p}$. The composition with the natural maps $\Xi\mathcal{D}^s_{\Lambda_p} \to \Xi\mathcal{D}^{\Lambda_p,s}_{\Lambda} \hookrightarrow \Xi\mathcal{D}^s_{\Lambda}$ then gives the moduli map we want.

Finally, after having constructed $m = m_p$ locally near p, we need to show that the local constructions glue over all of B. The only point that might be not clear is the prong-matching, since $\boldsymbol{\sigma}'$ was constructed in Proposition 13.6 by locally constant extension. However, this might make a difference only if the smoothing parameter f_e is not identically equal to zero, 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 of local neighborhoods. This completes the proof of Proposition 13.2.

13.4. **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 $(\mathcal{Y} \to B, (\omega_{(i)})_{i\in L^\bullet(\Gamma_p)}, R, \boldsymbol{\sigma}, f)$ of marked multi-scale differentials, defined on a neighborhood of p. We proceed similarly to the proof of Proposition 11.7 and let $B^s \to B$ to be the fiber product of $R\colon B \to \overline{T}^n_{\Gamma_p}$ with finite quotient map $\overline{p}\colon \overline{T}^s_{\Gamma_p} \to \overline{T}^s_{\Gamma_p}/K_p = \overline{T}^n_{\Gamma_p}$. The pullback family $\mathcal{Y}^s \to B^s$ comes with a map $R^s\colon B^s \to \overline{T}^s_{\Gamma_p}$ and is thus a family of simple marked multi-scale differentials. By Proposition 13.2 we obtain a moduli map $m^s\colon B^s \to \Xi\mathcal{D}^s_{\Lambda}$. Composing m^s with the quotient map $\Xi\mathcal{D}^s_{\Lambda} \to \Xi\mathcal{D}_{\Lambda}$ we get a map $B^s \to \Xi\mathcal{D}_{\Lambda}$ that is clearly K_{Γ_p}-invariant by construction. It thus descends to the required moduli map $m\colon B \to \Xi\mathcal{D}_{\Lambda}$.

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. Denote by $\mathbb{P}\Omega\overline{\mathcal{M}}^{\mathrm{ninc}}_{g,n}(\mu)$ the normalization of the incidence variety compactification, where the incidence variety is considered as a substack of $\mathbb{P}\Omega\overline{\mathcal{M}}_{g,n}$. In this section we will show that the stack $\mathbb{P}\mathcal{MS}_\mu$ of multi-scale differentials can be obtained from $\mathbb{P}\Omega\overline{\mathcal{M}}^{\mathrm{ninc}}_{g,n}(\mu)$ as the normalization of a certain explicit (complex algebraic, not real oriented) blowup, called the orderly blowup. We will then be able to conclude the proof of Theorem 1.2 and Theorem 1.3, in particular proving algebraicity of $\Xi\overline{\mathcal{M}}_{g,n}(\mu)$. We will conclude this Section with an example that illustrates the orderly blowup and the necessity of the subsequent normalization.

14.1. **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 even $\mathbb{Q}$-factorial, as shown in the following example.

Example 14.1. (*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 the level graph 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 $\mathbb{P}\Omega\overline{\mathcal{M}}^{\mathrm{inc}}_{g,n}(\mu)$ 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 $\mathbb{P}\Omega\overline{\mathcal{M}}^{\mathrm{inc}}_{g,n}(\mu)$. Namely, the map $\mathbb{P}\Xi\overline{\mathcal{M}}_{g,n}(\mu) \to \mathbb{P}\Omega\overline{\mathcal{M}}^{\mathrm{inc}}_{g,n}(\mu)$ 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, and vice versa for $\lambda' = 0$). 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 $\mathbb{P}\Omega\overline{\mathcal{M}}^{\mathrm{inc}}_{g,n}(\mu)$ (as well as its normalization) is not $\mathbb{Q}$-factorial (see e.g. [KM98, Corollary 2.63]).

Given an adjustable but not necessarily orderly family $(\mathcal{X} \to B, \omega)$ (as defined in Section 11.4), 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 $\{h_1, \ldots, h_n\}$ 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 $(h_{i_1}, \ldots, h_{i_k})$, where $\{i_1, \ldots, i_k\}$ 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]). This implies that for any two open subsets U_1 and U_2, we can glue $\widetilde{U_1}$ and $\widetilde{U_2}$ along their common restriction $\widetilde{U_1 \cap U_2}$. 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 $\widetilde{B} \to B$ locally given by $\widetilde{U} \to U$.

Example 14.2. 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 that 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 14.3. *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 $b_1c_1 + \cdots + b_nc_n = a$, which implies that $b_1(c_1/a) + \cdots + b_n(c_n/a) = 1$ 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 $I = (b_i)$. ∎

Lemma 14.4. *Let R be a local ring and let $h_1, \ldots, h_n \in R$ be some elements that are non-zero-divisors. Let $D = \prod(h_{i_1}, \ldots, h_{i_k})$ be the product of ideals where $\{i_1, \ldots, i_k\}$ ranges over all subsets of $\{1, \ldots, n\}$. Then D is a principal ideal (h) with h being a non-zero-divisor if and only if $h_1, \ldots, h_n$ are fully ordered by the divisibility relation.*

Proof. If $h_1, \ldots, h_n$ 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 $(h_{i_1}, \ldots, h_{i_k})$ by Lemma 14.3. Suppose $(h_1, \ldots, h_n) = (b)$ 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 $u_1t_1 + \cdots + u_nt_n = 1$. If all of $t_1, \ldots, t_n$ are not units, then the ideal $(t_1, \ldots, t_n)$ 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. ◻

The orderly blowup construction possesses some functorial property.

Proposition 14.5. *Given an adjustable family of differentials* $(\pi\colon \mathcal{X} \to B, \omega, \boldsymbol{z})$, *the pullback family* $\widetilde{\pi}\colon \widetilde{\mathcal{X}} \to \widetilde{B}$ *over the orderly blowup* $\widetilde{B} \to B$ *is orderly. Moreover, any dominant map* $\pi\colon B' \to B$, *such that the pullback family* $\mathcal{X}' \to B'$ *is orderly, factors through* $\widetilde{B}$.

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 14.4. More precisely, on the orderly blowup, the pullback of $\mathcal{D}_U$ becomes a principal ideal, and hence at every point of $\widetilde{U}$ the pullback family of differentials has adjusting parameters (given by the pullback of the functions h_i) that are fully ordered by divisibility, which implies that the family is orderly over $\widetilde{U}$.

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 $\pi^* h_i$ 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 $\pi^* h_i$ can be used as adjusting parameters for the pullback family over U'. Since the pullback family is orderly, these adjusting parameters $\pi^* h_i$ in U' are fully ordered by divisibility, and consequently the corresponding disorderly ideal $\pi^*\mathcal{D}_U$ in U' is principal (and 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 $\pi\colon U' \to U$ factors through $\widetilde{U}$. ◻

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. This ideal differs from the original definition of the disorderly ideal by a product of principal ideals, and hence the blowup with center D gives the same space as the orderly blowup. 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 warn the reader that the orderly blowup of a normal base may fail to be normal, as illustrated by the following example.

Example 14.6. (*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

$$\left\{(x,y) \times [u,v] \in \mathbb{C}^2 \times \mathbb{P}^1 : x^2 v - y^3 u = 0\right\} . \tag{14.1}$$

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 only isolated singularities.

We are now ready to apply these considerations to the IVC.

Lemma 14.7. *The incidence variety compactification* $\mathbb{P}\Omega\overline{\mathcal{M}}^{\mathrm{inc}}_{g,n}(\mu)$ *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 for the global residue condition (vanishing condition for sums of residues associated with these differentials). ∎

Theorem 14.8. *The moduli stack of multi-scale differentials* $\mathcal{MS}_\mu$ *is equivalent (as an analytic stack) to the normalization* $\widetilde{\Omega\mathcal{M}}^{\mathrm{n}}_{g,n}(\mu)$ *of the orderly blowup of the normalization* $\Omega\overline{\mathcal{M}}^{\mathrm{ninc}}_{g,n}(\mu)$ *of the incidence variety compactification. This orderly blowup* $\widetilde{\Omega\mathcal{M}}_{g,n}(\mu)$*, and thus also* $\mathcal{MS}_\mu$*, is the analytification of an algebraic stack.*

See [ACG11, Chapter XII] for a general introduction to (algebraic) stacks and [Toë99] for analytic stacks and analytification.

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 $\Omega\overline{\mathcal{M}}^{\mathrm{ninc}}_{g,n}(\mu)$ as an analytic stack. Proposition 11.13 ensures that over $\Omega\overline{\mathcal{M}}^{\mathrm{ninc}}_{g,n}(\mu)$ the family of one-forms that are the top level forms of the twisted differential is adjustable. 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 in view of Lemma 14.9 below we can consider the orderly blowup $\widetilde{\Omega\mathcal{M}}^{\mathrm{n}}_{g,n}(\mu)$ as an analytic stack.

Proposition 11.16 then ensures that the resulting orderly family over $\widetilde{\Omega\mathcal{M}}^{\mathrm{n}}_{g,n}(\mu)$ gives a family of multi-scale differentials of type μ. This family induces a map of stacks $\widetilde{\Omega\mathcal{M}}^{\mathrm{n}}_{g,n}(\mu) \to \mathcal{MS}_\mu$.

Conversely, a family in the stack $\mathcal{MS}_\mu$ is orderly by definition, hence by Proposition 14.5 we obtain a map of stacks $\mathcal{MS}_\mu \to \widetilde{\Omega\mathcal{M}}_{g,n}(\mu)$. Since $\mathcal{MS}_\mu$ is normal, this map factors through $\widetilde{\Omega\mathcal{M}}^{\mathrm{n}}_{g,n}(\mu)$, which gives the desired inverse map.

To show that the orderly blowup of $\Omega\overline{\mathcal{M}}^{\mathrm{ninc}}_{g,n}(\mu)$ is an algebraic stack, use that the normalization $\Omega\overline{\mathcal{M}}^{\mathrm{ninc}}_{g,n}(\mu)$ is an algebraic stack and pass from a local quotient groupoid presentation $[U/G]$ to a presentation $[U'/G']$ on a sufficiently large étale cover $U' \to U$ such that adjusting parameters exist (algebraically) on U'. The existence of such a U' is guaranteed by Proposition 11.14. That the blowup is a stack is justified as in the analytic case. Since we could have used $[U'/G']$ also as a presentation in the analytic case, we have justified the last statement of the theorem. ∎

In the preceding proof we used the following general fact.

Lemma 14.9. *Let* X *be an algebraic or analytic variety and let* $\mathcal{I}$ *be a coherent ideal sheaf. Let* $f\colon \widetilde{X} \to X$ *be a finite étale Galois cover with group* G *and let* $\mathcal{J} = f^{-1}\mathcal{I}$

be the pullback (that agrees with the inverse image here). Then there is a natural G-action on $\operatorname{Bl}_{\mathcal{J}} \widetilde{X}$ and an isomorphism $\operatorname{Bl}_{\mathcal{I}} X \cong (\operatorname{Bl}_{\mathcal{J}} \widetilde{X})/G$ (in the algebraic or analytic category).

Proof. The existence of the G-action on $\operatorname{Bl}_{\mathcal{J}} \widetilde{X}$ follows directly from the universal property for $\mathcal{J}$. The universal property of $\mathcal{I}$ (in the version [Har77, Corollary II.7.16]) also gives a map $\operatorname{Bl}_{\mathcal{J}} \widetilde{X} \to \operatorname{Bl}_{\mathcal{I}} X$. This map is equivariant with respect to the G-action on the domain and the trivial action on the target. It thus descends to a map $\operatorname{Bl}_{\mathcal{J}} \widetilde{X}/G \to \operatorname{Bl}_{\mathcal{I}} X$.

Conversely, take the fiber product $F = \widetilde{X} \times_X \operatorname{Bl}_{\mathcal{I}} X$. Since the pullback of $\mathcal{I}$ to F (via $\operatorname{Bl}_{\mathcal{I}} X \to X$) is principal, the pullback of $\mathcal{J}$ to F is principal. The universal property for $\mathcal{J}$ gives a map $F \to \operatorname{Bl}_{\mathcal{J}} \widetilde{X}$, which is obviously G-equivariant. Since $F/G = \operatorname{Bl}_{\mathcal{I}} X$ by definition, as fiber product, the map descends to the desired map $\operatorname{Bl}_{\mathcal{I}} X \to \operatorname{Bl}_{\mathcal{J}} \widetilde{X}/G$ on the G-quotients. ∎

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]). The preceding subsection thus does not answer the question on the projectivity of the coarse space associated with $\mathbb{P}\mathcal{MS}_\mu$.

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 14.10. *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 $\mathbf{MS}_\mu$ of multi-scale differentials of type μ.*

Proof. Given a family of multi-scale differentials $(\pi\colon \mathcal{X} \to B, \boldsymbol{\omega}) \in \mathbf{MS}_\mu(B)$, we want to provide the family locally near any point p with a marking, and then construct the moduli map $m\colon B \to \Xi\overline{\mathcal{M}}_{g,n}(\mu)$ 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 12.5 to provide the family with a marking. The moduli map given by Theorem 13.1 composed with the projection then gives a map $U_p \to \Xi\mathcal{D}_\Lambda \to \Xi\overline{\mathcal{M}}_{g,n}(\mu)$. These maps 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. ∎

Proofs of Theorem 1.2 and Theorem 1.3 completed. For Theorem 1.2, the density (1) and the description of the boundary divisor (2) have been taken care of in Theorem 10.3. Compactness (3) is the content of Theorem 7.12, and the coarse moduli space property (4) has been addressed in Theorem 14.10.

The forgetful map (5) is obvious, e.g. it follows from Theorem 14.8 and its proof.

For Theorem 1.3, the property of being a proper Deligne-Mumford stack carries over from $\overline{\mathcal{M}}_{g,n}$ all the way up through $\mathbb{P}\Omega\overline{\mathcal{M}}_{g,n}$, $\mathbb{P}\Omega\overline{\mathcal{M}}^{\mathrm{ninc}}_{g,n}(\mu)$, and $\widetilde{\mathbb{P}\Omega\mathcal{M}}^{\mathrm{n}}_{g,n}(\mu)$. The isomorphism in the statement of Theorem 1.3 is then 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.

It remains to show the algebraicity and properness of the stack $\mathbb{P}\Xi\overline{\mathcal{M}}_{g,n}(\mu)$. The algebraicity is a consequence of Theorem 14.8 which implies that the coarse moduli space provided by Theorem 14.10 is algebraic. This space is analytically locally covered by quotient stacks $[U_X/G_X]$ near X, where the group G_X is an extension of an automorphisms group of pointed curves by the group K_Γ, by the definition of morphisms of marked simple multi-scale differentials. Here Γ is the level graph compatible with the point X. We need to show that the coarse space is also covered by quotient stacks $[U'_X/G_X]$ étale locally. This is a consequence of Artin's approximation theorem [Art69, Corollary 2.6] that provide an étale neighborhood in the coarse space isomorphic to a neighborhood of the origin in $\mathbb{A}^n/G_X$. Intersecting the preimage in $\mathbb{A}^n$ with its G_X-images to get a G_X-invariant neighborhood gives the U'_X we need. Finally, properness follows from the usual composition rules for proper map, using the fact that the map for a stack to its coarse space is proper. ◻

Remark 14.11. As a consequence of the orderly blowup description, we see that the isotropy group of a point in the stack $\mathcal{MS}_\mu$ has no contribution from the group K_Γ defined in Section 6.4. Consequently, whenever this group is non-trivial, the stacks $\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ and $\mathcal{MS}_\mu$ are not locally isomorphic.

Conversely, whenever K_Γ is trivial, the functors of marked and simple marked multi-scale differentials agree, as we have seen in Section 12.3. By definition $\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ and $\mathcal{MS}_\mu$ are then isomorphic in a neighborhood of such a stratum. This happens of course in the interior $\Omega\mathcal{M}_g$, but also at the generic point of any boundary divisor, by definition of K_Γ.

14.3. **Some moduli spaces in genus zero and cherry divisors.** To illustrate the necessity of both the orderly blowup and the subsequent normalization in the passage from IVC to $\mathcal{MS}_\mu$ 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 14.1.

Example 14.12. (*The cherry requires both the orderly blowup and the normalization*) We consider the incidence variety compactification of $\mathbb{P}\Omega\mathcal{M}_{0,5}(2,1,0,0,-5)$, with two marked regular points on the surface. 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 8 we schematically depict the local structure of $\mathbb{P}\Omega\overline{\mathcal{M}}^{\mathrm{inc}}_{0,5}(2,1,0,0,-5)$ 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.

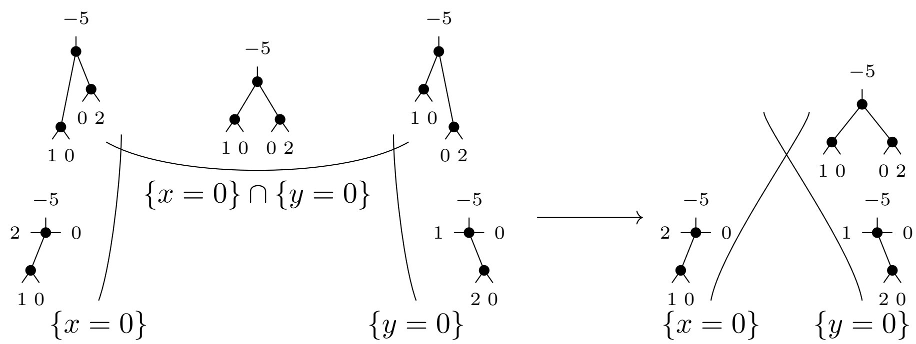

FIGURE 8. The orderly blowup of the incidence variety compactification of $\mathbb{P}\Omega\mathcal{M}_{0,5}(2, 1, 0, 0, -5)$ at a cherry point.

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 14.6 (see (14.1) for the description). We recall that the total space of this blowup 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 8. 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

$$\left\{(x, y) \times [u, v] \in \mathbb{C}^2 \times \mathbb{P}^1 : x^2 v - y^2 u = 0\right\} .$$

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 $\mathrm{GL}_2^+(\mathbb{R})$-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 $\mathrm{GL}_2^+(\mathbb{R})$-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 $\mathrm{GL}_2^+(\mathbb{R})$ 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 contained in the center of of $\mathrm{GL}_2^+(\mathbb{R})$ 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 $\mathrm{GL}_2^+(\mathbb{R})$-action near the boundary. While the constructions have certain similarities, they apparently differ e.g. in the treatment of horizontal nodes.

Theorem 15.1. *The $\mathrm{GL}_2^+(\mathbb{R})$-action on the moduli space $\Omega\mathcal{M}_{g,n}(\mu)$ extends to a continuous $\mathrm{GL}_2^+(\mathbb{R})$-action on the level-wise real blowup $\Xi\widehat{\mathcal{M}}_{g,n}(\mu)$ of the moduli space of multi-scale differentials $\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ along its total boundary divisor.*

In comparison to Section 12, 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 to be by real scaling.

Definition 15.2. A *real multi-scale differential of type* μ on a stable curve X is

(i) a full order $\preccurlyeq$ 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 $\preccurlyeq$, 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 (except the top level) by multiplication by a positive real number. △

To properly work with families of such differentials, we have to leave the category of complex spaces. Recall that *manifolds with corners* are topological spaces locally modeled 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 have to 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 orbifold structure is exclusively due to automorphisms of flat surfaces, as for $\Omega\mathcal{M}_{g,n}$.*

The last statement says that $\Xi\widehat{\mathcal{M}}_{g,n}(\mu)$ resolves the quotient singularites by the groups $K_\Lambda = \mathrm{Tw}_\Lambda/\mathrm{Tw}_\Lambda^s$ of $\Xi\overline{\mathcal{M}}_{g,n}(\mu)$:

Given Theorem 15.3, we define the action of $A \in \mathrm{GL}_2^+(\mathbb{R})$ on $\Xi\widehat{\mathcal{M}}_{g,n}(\mu)$ by

$$A \cdot \big((X_{(i)}), (\omega_{(i)}), \sigma\big) \;=\; \big((A \cdot (X_{(i)}, \omega_{(i)})), A \cdot \sigma\big)\;,$$

where $i \in L^\bullet(\Gamma)$. The first argument is the usual $\mathrm{GL}_2^+(\mathbb{R})$-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 $\boldsymbol{\omega}$) gives a matching of directions of slope $A \cdot \binom{1}{0}$ (for $A \cdot \boldsymbol{\omega}$) 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 11, 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 $\mathrm{GL}_2^+(\mathbb{R})$-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 12.5 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 $\widehat{X} \to \widehat{U}$ constructed in Theorem 12.2 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 $\boldsymbol{\omega} = (\omega_{(i)})_{i \in L^\bullet(\Lambda)}$ on U, and let T_i and F_e be the S^1-valued functions used in the blowup construction (Section 12.1).

Proof of Theorem 15.3. The second statement is an immediate consequence of Theorem 14.10 about the points in $\Xi\overline{\mathcal{M}}_{g,n}(\mu)$ and of Proposition 12.1 (ii).

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), as in the last statement of Theorem 12.2.

The fibers of the level-wise real blowup are a torus consisting of S^1 for each level below zero, see Proposition 12.1. The group $K_\Lambda = \mathrm{Tw}_\Lambda/\mathrm{Tw}_\Lambda^s$ acts freely on this torus. This proves the last statement. ∎

Proof of Theorem 15.1. It remains to justify the continuity of the $\mathrm{GL}_2^+(\mathbb{R})$-action. Consider a sequence $\{\widehat{p}_n\}$ converging to $\widehat{p}$ in $\widehat{U} \subset \widehat{\Xi\mathcal{D}_\Lambda}$. By definition of the topology, this is equivalent (with notations as in Section 12) 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 7.4 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 $\varphi_U(A \cdot \widehat{p}_n) \in U \subset \Xi\overline{\mathcal{M}}_{g,n}(\mu)$. 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 $\mathrm{GL}_2^+(\mathbb{R})$-action on S^1, this will conclude the proof.

For our aim, we use the maps $A \cdot g_n \cdot A^{-1}$, where $A \cdot (\,)$ denotes the induced $\mathrm{GL}_2^+(\mathbb{R})$-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 $\mathrm{GL}_2^+(\mathbb{R})$ on the seam identified with S^1.

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.

$(\Sigma, \boldsymbol{s})$	"Base" compact n-pointed oriented differentiable surface	15
$(X, \boldsymbol{z})$	Pointed stable curve of genus g	13
X_v	Irreducible component of X	13
N_X	Set of nodes of X	13
$X^s = X \setminus N_X$	The smooth part of X	16
N_X^v	Set of vertical nodes of X	13
N_X^h	Set of horizontal nodes of X	13
$f\colon \Sigma \to X$	Marking	15
$P_{\boldsymbol{s}}, Z_{\boldsymbol{s}}$	Subset of $\boldsymbol{s}$ mapped respectively to the poles and zeros of ω	36
$\overline{X}_{\boldsymbol{\sigma}}$	Welded surface associated to the prong-matching $\boldsymbol{\sigma}$	40
$X_{(i)}$	Components of X at level i	12
$\eta_{(i)}$	Restriction of the twisted differential η on $X_{(i)}$	13
$X_{>i}$	Components of X at level $> i$	13
X_ϵ, $(X, \boldsymbol{z})_\epsilon$	ϵ-thick part of X, resp. $X \setminus \boldsymbol{z}$	16

Graphs and Levels.

$\overline{\Gamma} = (\Gamma, \succcurlyeq)$, Γ	Level graph with full order $\succcurlyeq$	12
$V(\Gamma)$	Vertices of Γ	12
$E(\Gamma)$	Edges of Γ	12
$E(\Gamma)^v$, $E(\Gamma)^h$	Set of vertical, resp. horizontal, edges of Γ	12
$\mathrm{val}(v)$	Valence of the vertex v	12
$L^\bullet(\overline{\Gamma})$	Set of levels of the level graph $\overline{\Gamma}$	12
$L(\overline{\Gamma})$	Set of all but the top level of the level graph $\overline{\Gamma}$	12
N	Number of levels strictly below 0	12
$\ell\colon \Gamma \to \underline{N}$	Normalized level function	12
$\overline{\Gamma}_{(i)}$, $\overline{\Gamma}_{>i}$	Subgraph at (resp. above) level i of $\overline{\Gamma}$	12
$\ell(q^\pm)$, $\ell(e^\pm)$	Bottom and top levels of the ends of a node	13
Γ^+, Γ	Enhanced level graph	14

Λ	Multicurve in Σ	35
Λ^+, Λ	Enhanced multicurve	35
$\Gamma^+(\Lambda^+)$, Λ	Enhanced graph associated to the enhanced multicurve Λ	35
$\mathrm{dg}\colon \Lambda_2 \rightsquigarrow \Lambda_1$	Degeneration of the ordered multicurve Λ_2	35
$\delta\colon \underline{N} \twoheadrightarrow \underline{M}$	Map defining a vertical undegeneration	35
$D^h \subseteq \Lambda_1^h$	Subset of horizontal curves inducing a horizontal undegeneration	36
(δ, D^h), δ	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 $\mathbb{P}$.

$\mathcal{T}_{g,n} = \mathcal{T}_{(\Sigma,\boldsymbol{s})}$	Teichmüller space	15
$\mathrm{Mod}_{g,n}$	Classical mapping class group	16
$\overline{\mathcal{T}}_{g,n} = \overline{\mathcal{T}}_{(\Sigma,\boldsymbol{s})}$	Augmented Teichmüller space	16
$\mathcal{D}_\Lambda$	Classical Dehn space	20
$\Omega\mathcal{D}_\Lambda$	Hodge bundle over the Dehn space	21
$\Omega\mathcal{T}_{(\Sigma,\boldsymbol{s})}(\mu)$	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	37
$\Omega\mathcal{T}_\Lambda(\mu)$	Teichmüller space of twisted differentials of type (μ,Λ)	37
$\Omega\mathcal{T}_\Lambda^{pm}(\mu)$	Teichmüller space of prong-matched twisted differentials	41
$\Omega\overline{\mathcal{T}}_{(\Sigma,\boldsymbol{s})}(\mu)$	Augmented Teichmüller space of marked flat surfaces of type μ	53
$\Omega\mathcal{B}_\Lambda$	Λ-boundary stratum	53
$\overline{\mathcal{MD}}_\Lambda^s$, $\overline{\mathcal{MD}}_\Lambda$	(Smooth) model domain	60
$\Xi\mathcal{D}_\Lambda$	Dehn space associated with Λ	70
$\Xi\mathcal{D}_\Lambda^s$	Simple Dehn space associated with Λ	70
$\Xi\mathcal{D}_\Lambda^{sv}$	simple vertical Dehn space	70
$\Xi\widehat{\mathcal{M}}_{g,n}(\mu)$	Level-wise real blowup of $\Xi\overline{\mathcal{M}}_{g,n}(\mu)$	116
$\mathbf{MS}_\mu$	Functor of multi-scale differentials	94
$\mathcal{MS}_\mu$	Grupoid of multi-scale differentials	94
$\mathbf{MS}_{(\mu,\Lambda)}$	Functor of marked multi-scale differentials	101
$\mathbf{MD}_{(\mu,\Lambda)}^s$	Functor of model differentials	103

Families.

$\mathcal{Z}_j$	Image of the section z_j	11
$\mathcal{Z}^0$	Horizontal zero divisor	11
$\mathcal{Z}^\infty$	Horizontal polar divisor	11
$uv = f$	Local equation of a nodal family	29
$u_e^+ u_e^- = f_e$	Local equation of a family near the node q_e	99
$\widehat{\pi}\colon \widehat{\mathcal{X}} \to \widehat{B}$	Level-wise real blowup of a family of multi-scale differentials	98
$\overline{\pi}\colon \overline{\mathcal{X}} \to \overline{B}$	Global level-wise real blowup	99
$\widetilde{\pi}\colon \widetilde{\mathcal{X}} \to \widetilde{B}$	Orderly blowup of an adjustable family of differentials	111

$\boldsymbol{t} = (t_i)_{i \in L}$	Simple rescaling parameters	61
$\boldsymbol{s} = (s_i)_{i \in L}$	Rescaling parameters	61
$R_p \colon B_p \to \overline{T}^n_{\Gamma_p}$	Rescaling ensemble	88
$R^s \colon B \to \overline{T}^s_{\Gamma_p}$	Simple rescaling ensemble	91
x_j	Smoothing parameters for the horizontal nodes	87
f_e	Smoothing parameters	88
h, h_v, h_i	Adjusting parameter	94

Plumbing Constructions.

$\Omega\,\mathrm{Pl}^v$	Vertical plumbing map	77
$\Omega\,\mathrm{Pl}$	Plumbing map	85
$(X_0, \boldsymbol{\eta}_0)$	Base surface in the Λ-boundary stratum	72
$\phi_e^\pm$	Local coordinates at a node e	74
$r_e(\boldsymbol{t})$	Residue of $\boldsymbol{\omega}$ at the node e	74
$\boldsymbol{\xi}$	Family of modifying differentials	63
ρ_i	Period homomorphism	64
A_{δ_1,δ_2}	Annulus of inner radius δ_1 and outer radius δ_2	73
$A^\pm$	Top and bottom plumbing annuli	73
R, δ	Defining constants of the plumbing annuli	73
$V(s)$	Standard plumbing fixture	73
$p^\pm$	Top and bottom marked points	73
$\mathcal{W}_\epsilon$	Base of the plumbing construction	72
Ψ	Local chart center at the base surface	72
$\mathcal{B}^-$, $\mathcal{B}^+$	Bottom and top plumbing annuli in $\mathcal{X}$	76
$\mathcal{C}^-, \mathcal{C}^+$	Bottom and top plumbing annuli in $\mathcal{Y}^v$	76
v_e^+, v_e^-, v_h	Conformal maps on annuli putting $\boldsymbol{t} * (\boldsymbol{\eta} + \boldsymbol{\xi})$ in standard form	75
$b_e^\pm$, b_h	Image of $p^\pm$ in $\mathcal{B}^\pm$	75
$c_e^\pm$	Image of the points $p^\pm$ in $\mathcal{C}_e^\pm$	76

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 $\operatorname{ord}_{q^+} \eta + 1$	14
$\boldsymbol{\sigma}$	Global prong-matching for X	40
$\boldsymbol{d} = (d_i)_{i \in L^\bullet(\Lambda)}$	Tuple in $\mathbb{C}^{L(\Lambda)}$ acting on prong-matched differentials	44
$\cdot$	Action of $\mathbb{C}^{L(\Lambda)}$ on prong-matched twisted differentials	44
$*$	Action of T_Λ on prong-matched twisted differentials	49
$F_{\boldsymbol{d}}$	Fractional Dehn twist	45
a_i, $m_{e,i}$	Defined by $a_i = \operatorname{lcm}_e \kappa_e$ and $m_{e,i} = a_i/\kappa_e$	46
$\boldsymbol{t}^{\boldsymbol{a}}_{\lceil i \rceil}$	Product of the $t_j^{a_j}$ for $j \geq i$	61
P_Γ	Prong rotation group	41
$\mathrm{Tw}_\Lambda^{\mathrm{full}}$	Classical Λ-twist group	20

$\mathbb{Z}^{L^\bullet(\Lambda)}$	Level rotation group	45
$\phi_\Lambda^\bullet$	Map from level rotation group to prong rotation group	45
Tw_Λ^v	Vertical twist group	45
Tw_Λ^h	Horizontal twist group	46
Tw_Λ	Twist group	46
$\mathrm{Tw}_{\Lambda,i}^{sv}$	Simple vertical twist group of level i	46
Tw_Λ^{sv}	Simple vertical twist group	47
K_Λ	Finite group defined by $\mathrm{Tw}_\Lambda^v/\mathrm{Tw}_\Lambda^{sv}$	47
Tw_Λ^s	Simple twist group	47
T_Λ^s, T_Λ	(Simple) level rotation torus	47
H, G	Ramifications groups	49

Other Notations.

$\Delta_r = \{z \in \mathbb{C} : \lvert z\rvert < r\}$	19	$\overline{n} = \{1, \dots, n\}$	11
$\boldsymbol{e}(z) = \exp(2\pi\sqrt{-1}z)$	29	$\underline{N} = \{0, -1, \dots, -N\}$	12

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.

[Art69] M. Artin. *Algebraic approximation of structures over complete local rings*. Inst. Hautes Études Sci. Publ. Math. 36 (1969), pp. 23–58.

[Bar18] I. Barros. *Uniruledness of strata of holomorphic differentials in small genus*. Adv. Math. 333 (2018), pp. 670–693.

[BB02] R. Beattie and H.-P. Butzmann. *Convergence structures and applications to functional analysis*. Kluwer Academic Publishers, Dordrecht, 2002, pp. xiv+264.

[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.

[BDG20] F. Benirschke, B. Dozier, and S. Grushevsky. *Equations of linear subvarieties of strata of differentials*. to appear in Geom. Topol. (2020). arXiv: 2011.11664.

[Ben20a] F. Benirschke. *The boundary of linear subvarietis in strata of differentials*. To appear in J. Eur. Math. Soc. (2020). arXiv: 2007.02502.

[Ben20b] F. Benirschke. *The closure of double ramification loci via strata of exact differentials*. (2020). arXiv: 2012.07703.

[Ber74a] 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.

[Ber74b] L. Bers. "Spaces of degenerating Riemann surfaces". In: *Discontinuous groups and Riemann surfaces (Proc. Conf., Univ. Maryland, College Park, Md., 1973)*. 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.

[CG20] D. Chen and Q. Gendron. *Towards the classification of connected components of the strata of k-differentials*. (2020). arXiv: 2101.01650.

[CMS19] D. Chen, M. Möller, and A. Sauvaget. *Masur-Veech volumes and intersection theory: the principal strata of quadratic differentials, with an appendix by G. Borot, A. Giacchetto, and D. Lewanski*. (2019). arXiv: 1912.02267.

[CMSZ20] D. Chen, M. Möller, A. Sauvaget, and D. Zagier. *Masur-Veech volumes and intersection theory on moduli spaces of Abelian differentials*. Invent. Math. 222.1 (2020), pp. 283–373.

[CMZ19] M. Costantini, M. Möller, and J. Zachhuber. *The area is a good enough metric*. (2019). arXiv: 1910.14151.

[CMZ20] M. Costantini, M. Möller, and J. Zachhuber. *The Euler characteristic of strata*. (2020). arXiv: 2006.12803.

[CW20] D. Chen and A. Wright. *The WYSIWYG compactification*. Journal of the London Mathematical Society 103 (2020), pp. 490–515.

[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.

[Doz20] B. Dozier. *Measure bound for translation surfaces with short saddle connections*. (2020). arXiv: 2002.10026.

[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: *Quasiconformal 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* $\mathrm{SL}(2,\mathbb{R})$ *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* $\mathrm{SL}(2,\mathbb{R})$ *action on moduli space.* Ann. Math. (2) 182.2 (2015), pp. 673–721.

[EMMW20] A. Eskin, C.T. McMullen, R. Mukamel, and A. Wright. *Billiards, quadrilaterals, and moduli spaces.* J. Amer. Math. Soc. 33.4 (2020), pp. 1039–1086.

[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.

[GT21] Q. Gendron and G. Tahar. *Différentielles abéliennes à singularités prescrites.* J. Éc. Polytech., Math. 8 (2021), pp. 1397–1428.

[Har77] R. Hartshorne. *Algebraic geometry.* Graduate Texts in Mathematics, No. 52. New York: Springer-Verlag, 1977, pp. xvi+496.

[Haz00] M. Hazewinkel, ed. *Handbook of algebra. Vol. 2.* Vol. 2. Handbook of Algebra. Elsevier/North-Holland, Amsterdam, 2000, pp. xx+878.

[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.

[HS21] D. Holmes and J. Schmitt. *Infinitesimal structure of the pluricanonical double ramification locus.* Compos. Math. 157.10 (2021), pp. 2280–2337.

[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* **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.

[McM89] C. McMullen. *Amenability, Poincaré series and quasiconformal maps.* Invent. Math. 97.1 (1989), pp. 95–127.

[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.

[Ngu19] D.-M. Nguyen. *Volume form on moduli spaces of d-differentials.* to appear in Geom. Topol. (2019). arXiv: `1902.04830`.

[Pom92] C. Pommerenke. *Boundary behaviour of conformal maps.* Vol. 299. Grundlehren der Mathematischen Wissenschaften. Springer-Verlag, Berlin, 1992, pp. x+300.

[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.

[Sau20] A. Sauvaget. *Volumes of moduli spaces of flat surfaces.* (2020). arXiv: `2004.03198`.

[Sta18] The Stacks project authors. *The Stacks project.* 2018.

[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.

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
Current address: Centro de Investigación en Matemáticas, Guanajuato, Gto., AP 402, CP 36000, México
Email address: quentin.gendron@cimat.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