



A compactification of the moduli space of multiple-spin curves

Emre Can Sertöz¹

Received: 15 October 2022 / Accepted: 21 June 2023 / Published online: 7 July 2023
© The Author(s) 2023

Abstract

We construct a smooth Deligne–Mumford compactification for the moduli space of curves with an m -tuple of spin structures using line bundles on quasi-stable curves as limiting objects, as opposed to line bundles on stacky curves. For all m , we give a combinatorial description of the local structure of the corresponding coarse moduli spaces. We also classify all irreducible and connected components of the resulting moduli spaces of multiple-spin curves.

Keywords Spin curves · Compactification · Moduli · Theta-characteristics · Roots of line bundles

Mathematics Subject Classification (2020) 14H10 · 14D23 · 14B10

1 Introduction

The configuration of the 28 bitangents of a smooth quadric curve in the plane form a beautiful chapter in classical algebraic geometry [11, 34]. A similar structure is observed again with the 120 tritangents of a generic canonical space sextic [4, 7, 12, 30, 31] or, more generally, with the $\binom{2g}{2}$ contact hyperplanes of a generic canonically embedded curve of genus g [9]. The unifying idea is that these hyperplanes correspond to odd spin structures on the curve, i.e., to square roots of the canonical bundle of the curve with a non-zero section.

In the study of algebraic curves, degeneration techniques play an important role. However, there is a technical gap prohibiting the study of “configurations”, or tuples, of spin structures on curves via degeneration: there is no compactification of the relevant moduli spaces that uses curves with line bundles as limiting objects. There are notable exceptions where this gap was partially addressed, or circumvented, with great success. Most prominently, the moduli space of spin curves was compactified by Cornalba [8] which gave rise to a complete Kodaira classification of the resulting spaces [16, 19]. These moduli spaces allow only for the study of a *single* spin structure on each curve. On the other extreme, Caporaso and Sernesi [9] considered the degeneration of *all* odd spin structures to prove that a generic curve is

✉ Emre Can Sertöz
emre@serto.com

¹ Leibniz University Hannover, Welfengarten 1, 30167 Hannover, Germany

determined by its contact hyperplanes. The work [9] could largely avoid the aforementioned technical gap because one can degenerate curves together with all of their odd or even spin structures without constructing a new moduli space. Fan et al. [17] do indeed construct a compactification of the moduli space of curves with more than one spin structure using *stacky* curves as limiting objects.

We find that the quasi-stable curves of Cornalba retain their intuitive appeal in approaching projective geometric problems. In particular, the study of effective limit linear systems keeps their geometric flavor when working with quasi-stable curves. For this reason, we develop here a compactification of the moduli spaces of “multiple spin curves” from the point of view of quasi-stable curves. This construction opens the door to studying multiple and fractional limit linear series on quasi-stable curves. We establish local structural properties of these moduli spaces, such as the smoothness of the moduli stacks and the quotient singularity types of the coarse moduli spaces. We also classify the connected components of these spaces. We describe these results in greater detail in Sect. 1.1.

1.1 Idea of the construction and main results

Let k be an algebraically closed field of characteristic not two. A *spin structure* on a proper smooth curve C/k is a pair $(L, \alpha : \mathcal{L}^{\otimes 2} \xrightarrow{\sim} \omega_{C/k})$ where L is a line bundle on C and $\omega_{C/k}$ is the canonical bundle of C . The tuple (C, L, α) is called a *spin curve*. The moduli space \mathcal{S}_g of spin curves of genus g is quasi-finite over the moduli space \mathcal{M}_g of curves of genus g .

Fix a positive integer m . A sequence $(L_i, \alpha_i)_{i=1}^m$ of spin structures on C will be called a *multiple-spin structure on C* . The tuple $(C, (L_i, \alpha_i)_{i=1}^m)$ will be called a *multiple-spin curve*. For each genus g , the moduli space of multiple-spin curves will be denoted by \mathcal{S}_g^m which is the m -fold fiber product $\mathcal{S}_g \times_{\mathcal{M}_g} \cdots \times_{\mathcal{M}_g} \mathcal{S}_g$. The purpose of this paper is to give a natural compactification of \mathcal{S}_g^m , determine its irreducible components, and describe its basic geometric properties.

There is a compactification $\overline{\mathcal{S}}_g \rightarrow \overline{\mathcal{M}}_g$ of \mathcal{S}_g over the moduli space $\overline{\mathcal{M}}_g$ of stable curves, whose coarse moduli scheme over the complex numbers was originally constructed by Cornalba [8]. Later, the work was completed by Jarvis [26] by the construction of the moduli stack.

Cornalba compactifies \mathcal{S}_g by considering *limit spin curves*, which are tuples of the form $(X, \mathcal{L}, \alpha : \mathcal{L}^{\otimes 2} \rightarrow \omega_{X/k})$ where:

- X is a *quasi-stable* curve of genus g ,
- \mathcal{L} is a line bundle of degree $g - 1$ on X and degree 1 on unstable components,
- α is an isomorphism away from the unstable components and zero on the unstable components.

Note that any unstable component of X is isomorphic to \mathbb{P}^1 and such a component contains exactly two nodes of X . The stabilization of X contracts each unstable component to a point.

Our compactification will be based upon that of Cornalba’s and, in particular, on the product $\overline{\mathcal{S}}_g^{\times m} := \overline{\mathcal{S}}_g \times_{\overline{\mathcal{M}}_g} \cdots \times_{\overline{\mathcal{M}}_g} \overline{\mathcal{S}}_g$. Although the product $\overline{\mathcal{S}}_g^{\times m}$ is compact, the objects it parametrizes are not entirely natural. By definition, an element $(\pi_i : X_i \rightarrow C, \mathcal{L}_i, \alpha_i)_{i=1}^m \in \overline{\mathcal{S}}_g^{\times m}(k)$ is based upon a tuple of curves $(X_i)_{i=1}^m$ whose stabilizations C are all identified but the unstable components of X_i ’s remain distinct. Not only is it unnatural to work with m partially identified curves, but this causes the space $\overline{\mathcal{S}}_g^{\times m}$ to be non-normal. We prove in Proposition 6.10 that the compactification we give normalizes $\overline{\mathcal{S}}_g^{\times m}$.

Roughly, the normalization of $\overline{\mathcal{S}}_g^{\times m}$ could be constructed by adding to each tuple $(\pi_i : X_i \rightarrow C, \mathcal{L}_i, \alpha_i)_{i=1}^m \in \overline{\mathcal{S}}_g^{\times m}$ a master curve X to dominate all X_i (see Figs. 1 and 2). For each $i = 1, \dots, m$ let $\chi_i \subset C$ denote the subset of the nodes over which $\pi_i : X_i \rightarrow C$ is not an isomorphism. We define a “destabilization” $\pi : X \rightarrow C$ of $\chi = \cup_{i=1}^m \chi_i$ by inserting a rational curve to separate the branches at each node in χ (see Definition 2.1). For each $i = 1, \dots, m$ we can factor the map π into $\pi_i \circ \rho_i$ using a partial stabilization map $\rho_i : X \rightarrow X_i$. Denote by $(\mathcal{L}'_i, \alpha'_i)$ the pullback of $(\mathcal{L}_i, \alpha_i)$ via ρ_i . At this point the intermediate objects (π, \mathcal{L}_i) can be forgotten and we may consider the tuple $(X, (\mathcal{L}'_i, \alpha'_i)_{i=1}^m)$ as a natural limit of multiple-spin curves.

We must warn, however, that this construction needs to be refined as there are typically infinitely many choices of maps $(\rho_i)_{i=1}^m$ leading to infinitely many non-equivalent tuples $(\mathcal{L}'_i, \alpha'_i)_{i=1}^m$, even modulo automorphisms of X . In other words, the resulting moduli space of such tuples does not have finite fibers over $\overline{\mathcal{M}}_g$.

We fix the issues of the rough construction by imposing a constraint on the choice of ρ_i . Infinitely many non-equivalent choices of ρ_i 's occur only when there is a pair of indices $i \neq j$ for which X_i and X_j have unstable components over the same node of C , i.e., when $\chi_i \cap \chi_j \neq \emptyset$. In this case, the line bundles \mathcal{L}_i and \mathcal{L}_j both have degree one on the unstable components lying above $\chi_i \cap \chi_j$ and they can be used to stabilize these components.

We will, then, allow only for sequences of maps $(\rho_i)_{i=1}^m$ where for every i, j and every $x \in \chi_i \cap \chi_j$ the line bundles $\rho_i^* \mathcal{L}_i^{\otimes 2}$ and $\rho_j^* \mathcal{L}_j^{\otimes 2}$ are isomorphic in an open neighbourhood of

Fig. 1 The unstable components in the master curve X fully contract to give C . Contracting various subsets of all unstable components of X we can obtain the X_i 's


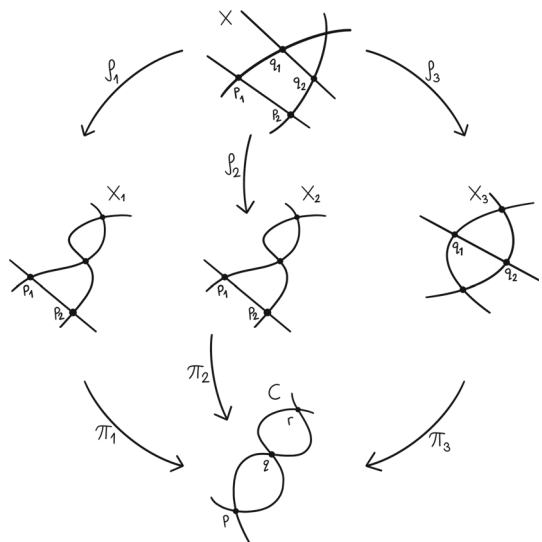


Fig. 2 An example with $m = 3$ illustrating the construction in Fig. 1. The base curve C consists of two smooth components meeting to form three nodes p, q, r . The curves X_1 and X_2 each contract an unstable component \mathbb{P}^1 over p , whereas X_3 contracts a \mathbb{P}^1 over q . The master curve X contracts a \mathbb{P}^1 over p and q each. The morphisms ρ_1 and ρ_2 from X serve to identify the \mathbb{P}^1 's in X_1 and X_2 over p



$\pi^{-1}(x) \subset X$. The resulting tuple $(X, (\mathcal{L}'_i, \alpha'_i)_{i=1}^m)$ will be called a *limit multiple-spin curve*. See Sect. 2.4 for examples and an intrinsic definition which encapsulates *families* of limit multiple-spin curves. We denote the moduli space of limit multiple-spin curves by $\overline{\mathcal{S}}_g^m$.

Theorem 1.1 *The moduli space of limit multiple-spin curves $\overline{\mathcal{S}}_g^m$ is a smooth and proper Deligne–Mumford stack. Moreover, the canonical inclusion $\mathcal{S}_g^m \hookrightarrow \overline{\mathcal{S}}_g^m$ is dense and open. The forgetful map $\overline{\mathcal{S}}_g^m \rightarrow \overline{\mathcal{M}}_g$ is quasi-finite.*

Remark 1.2 This is a special case of our Theorem 6.8. Another special case gives a similar compactification for the moduli space of marked curves with roots of the twisted canonical bundle. Section 6.3 explains these deductions.

For the birational classification of moduli spaces, understanding the nature of the singularities of the coarse moduli spaces associated to the moduli stacks is often a necessary step [6, 18, 24, 32]. The singularities of the coarse moduli space associated to $\overline{\mathcal{S}}_g^m$ are finite quotient singularities. These quotients are described by the action of the automorphism groups of multiple-spin curves on their local deformation functors. In Sect. 5.2 we give a completely combinatorial description of these group actions using a form of dual graph associated to each curve, see Theorem 5.10 and Proposition 5.14.

A complete classification of the connected components of $\overline{\mathcal{S}}_g^m$ is possible. Note that $\overline{\mathcal{S}}_g^m$ is smooth and, therefore, its irreducible and connected components coincide. Moreover, \mathcal{S}_g^m is Zariski dense in $\overline{\mathcal{S}}_g^m$ so the irreducible and connected components of these two spaces coincide.

Given an m -spin structure $(L_i, \alpha_i)_{i=1}^m$ on a genus g curve C we call the associated syzygy relations to be the tuple

$$(h^0(L_i) \pmod 2; h^0(L_j \otimes L_k \otimes L_1^\vee) \pmod 2)_{\substack{1 \leq i \leq m \\ 1 < j < k \leq m}} \in \mathbb{F}_2^{m + \binom{m-1}{2}}. \tag{1.1}$$

Any tuple that can be obtained in this manner will be called an (m, g) -syzygy relation. Let $\text{Gr}(a, b)$ be the Grassmannian of a -planes in \mathbb{F}_2^b . Section 7 contains the proof of the following theorem.

Theorem 1.3 *The irreducible (and connected) components of $\overline{\mathcal{S}}_g^m$ are in natural bijection with tuples (W, \mathfrak{a}) where $W \in \text{Gr}(m - k, m)$, $\mathfrak{a} \in \mathbb{F}_2^{k + \binom{k-1}{2}}$ is a (k, g) -syzygy relation, and $k = 0, \dots, m - 1$. If $g \geq k$ then every $\mathfrak{a} \in \mathbb{F}_2^{k + \binom{k-1}{2}}$ is a (k, g) -syzygy relation and if $g < k$ then Algorithm 7.14 finds all (k, g) -syzygy relations.*

1.2 Outline of the paper

In Sect. 2 we expand the notion of a limit multiple-spin curve defined above and give rigorous definitions. These definitions, employing line bundles on quasi-stable curves, are convenient for solving geometric problems but not for solving the present problem of studying the structure of the relevant moduli spaces. Therefore, in Sect. 3 we give an equivalent definition using torsion-free sheaves on stable curves. In Sect. 4 we prove that the moduli problem is represented by an algebraic stack. In Sect. 5 we undertake a study of the local structure of the moduli stacks. In Sect. 6 we prove basic geometric properties of these moduli stacks. In Sect. 7 we classify the components of \mathcal{S}_g^m . In ‘‘Appendix A’’ we establish some of the technical aspects required for the study of the local deformation functors of limit multiple-spin curves.

2 Families of multiple limit roots

In this section we generalize the construction of limit multiple spin curves given in the introduction to families of curves. As there is nothing special about square roots of the canonical bundle from the point of view of our construction, we will consider the square roots of any line bundle.

2.1 Destabilization of curves

Let k be an algebraically closed field with $\text{char } k \neq 2$. Let C be a connected nodal curve over k . If the relative dualizing sheaf of C is ample then C is said to be *stable*. Let $\chi \in C$ be a subset of the nodes of C and $I \subset \mathcal{O}_C$ the ideal sheaf corresponding to χ .

Definition 2.1 Let $\pi : X = \text{Proj}_C(\text{Sym}^* I) \rightarrow C$. Then the map π , and sometimes X , is called a *destabilization of C at χ* . If C is stable then X is *quasi-stable*.

For each $x \in \chi$, the fiber $\pi^{-1}(x)$ is isomorphic to \mathbb{P}_k^1 . The map π is an isomorphism over $C \setminus \chi$.

Definition 2.2 Let $\pi : X \rightarrow C$ be a destabilization of the nodes $\chi \subset C$. Then for each $x \in \chi$ we will call the fiber $\pi^{-1}(x) \subset X$ an *exceptional component* of X .

2.2 Families of limit roots

We recall the notion of a limit root given in Definition 2.1.1 of [5] but only for *square* roots. All the definitions in this subsection are adapted from *loc. cit.* Fix a line bundle N on the curve C of even degree. Consider a triplet $(\pi : X \rightarrow C, L, \alpha : L^{\otimes 2} \rightarrow \pi^*N)$ where π is a destabilization and L is a line bundle on X of degree $(\deg N)/2$.

Definition 2.3 Suppose α is an isomorphism in the complement of the exceptional components of X . If L has degree 1 on each exceptional component then (π, L, α) is a *limit root of N* . If L has degree 0 or 1 on the exceptional components then (π, L, α) *stabilizes to a limit root of N* .

Let T be a scheme over $\mathbb{Z}[\frac{1}{2}]$, \mathcal{C}/T be a family of stable curves [39, Tag 0E73] and \mathcal{X}/T a family of nodal curves [39, Tag 0C58]. Fix a line bundle \mathcal{N} on \mathcal{C} of relative even degree d .

Definition 2.4 If a morphism $\pi : \mathcal{X} \rightarrow \mathcal{C}$ restricts on each geometric fiber to a destabilization then π is a *destabilization*.

Let $\pi : \mathcal{X} \rightarrow \mathcal{C}$ be a destabilization, \mathcal{L} a line bundle on \mathcal{X} and $\alpha : \mathcal{L}^{\otimes 2} \rightarrow \pi^*\mathcal{N}$ a morphism.

Definition 2.5 If at each geometric fiber of $(\mathcal{X} \xrightarrow{\pi} \mathcal{C}, \mathcal{L}, \alpha)$ is (or stabilizes to) a limit root then $(\mathcal{X} \xrightarrow{\pi} \mathcal{C}, \mathcal{L}, \alpha)$ is (or stabilizes to) a *family of limit roots*.

If $(\pi, \mathcal{L}, \alpha)$ stabilizes to a family of limit roots then there exists a family of limit roots $(\pi' : \mathcal{X}' \rightarrow \mathcal{C}, \mathcal{L}', \alpha')$ and a morphism $\rho : \mathcal{X} \rightarrow \mathcal{X}'$ such that $(\mathcal{L}, \alpha) \simeq \rho^*(\mathcal{L}', \alpha')$. The map ρ is the *partial stabilization map with respect to \mathcal{L}* . A partial stabilization contracts the unstable components of each fiber on which \mathcal{L} has degree 0.

Notation 2.6 If $(\pi, \mathcal{L}, \alpha)$ stabilizes to a family of limit roots then denote by $V(\mathcal{L}) \subset \mathcal{X}$ the largest open set on which the partial stabilization ρ is an isomorphism.

An isomorphism between two families of limit roots $(\mathcal{X}_i \xrightarrow{\pi_i} \mathcal{C}, \mathcal{L}_i, \alpha_i)_{i=1,2}$ is a pair of isomorphisms $(f: \mathcal{X}_1 \rightarrow \mathcal{X}_2, g: \mathcal{L}_1 \rightarrow f^*\mathcal{L}_2)$ such that $\alpha_1 = f^*\alpha_2 \circ g^{\otimes 2}$.

2.3 Families of multiple limit roots

We now want to consider m -tuples of limit roots of \mathcal{N} over \mathcal{C} for a fixed positive integer m .

Definition 2.7 Let $\pi: \mathcal{X} \rightarrow \mathcal{C}$ be a destabilization. Let $\mathfrak{L} := \{\mathcal{L}_i, \alpha_i: \mathcal{L}_i^{\otimes 2} \rightarrow \pi^*\mathcal{N}\}_{i=1}^m$ be such that each $(\mathcal{L}_i, \alpha_i)$ stabilizes to a limit root, but \mathfrak{L} itself is not pulled back from a partial stabilization. Consider a line bundle \mathcal{L} and a sequence of morphisms $\varphi_i: \mathcal{L} \rightarrow \mathcal{L}_i^{\otimes 2}$ satisfying the following:

- $\alpha_i \circ \varphi_i = \alpha_j \circ \varphi_j$ for each i, j .
- Each φ_i restricts to an isomorphism on $V(\mathcal{L}_i)$, see Notation 2.6.

Then, we will call $\mathfrak{F} = (\varphi_i)_{i=1}^2$ a *synchronization data*. The tuple $(\pi, \mathfrak{L}, \mathfrak{F})$ will be called a *multiple limit root*. An isomorphism of multiple limit roots is a sequence of isomorphism of the limit roots commuting with the synchronization data.

Remark 2.8 The definition above agrees with the construction given in the introduction. Indeed, on $V(\mathcal{L}_i) \cap V(\mathcal{L}_j)$ the map $\varphi_j \circ \varphi_i^{-1}$ identifies $\mathcal{L}_i^{\otimes 2}$ and $\mathcal{L}_j^{\otimes 2}$.

We will give an alternative formulation of multiple limit roots in the next section and prove the equivalence of these two notions in Proposition 3.14. To make this equivalence precise, we make the following definition.

Definition 2.9 For $(C \rightarrow T, \mathcal{N})$ and m fixed, let $\mathcal{S}_{\text{lim}}^m(\mathcal{N})$ denote the *category of limit multiple roots of \mathcal{N}* . That is, $\mathcal{S}_{\text{lim}}^m(\mathcal{N}) \rightarrow T$ is fibered in groupoids with fiber over $T' \rightarrow T$ consisting of the multiple limit roots of $(C|_{T'} \rightarrow T', \mathcal{N}|_{T'})$.

2.4 Multiple-spin curves

Take N to be the canonical bundle. In this case, limit roots are called limit spin curves, therefore we will refer to multiple limit roots as a multiple-spin curves. Taking $m = 2$ we will give some examples of multiple-spin curves. For basic results on limit spin curves we refer to [8]. We will say that a spin curve $\xi = (\pi: \tilde{C} \rightarrow C, L, \alpha)$ over C *requires destabilization* if π is not an isomorphism.

The examples above can be readily computed from definitions. Especially the ‘‘Appendix A’’ and Sect. 5 are helpful in constructing more elaborate examples. A form of dual graph is introduced in Sect. 5.2 which gives the general framework to tackle examples like the ones below. For a detailed study of the examples below and additional examples we refer to [37, § II.1.6.2].

Example 2.10 For $i = 1, 2$ let (C_i, p_i) be marked smooth irreducible curves. Consider the nodal curve $C = C_1 \cup_{p_1 \sim p_2} C_2$. Every spin curve over C requires destabilization. Given any pair of spin curves $\xi_i = (\pi_i, L_i, \alpha_i)$, $i = 1, 2$, we can *uniquely* form a limit multiple-spin curve over (ξ_1, ξ_2) upto isomorphisms.

Example 2.11 Let $C = \overline{C}/(p \sim q)$ where $p, q \in \overline{C}$ are distinct points on a smooth irreducible curve \overline{C} . Some spin curves over C requires destabilization and some do not. If we take a pair of spin curves $\xi_i = (\pi_i, L_i, \alpha_i)$, $i = 1, 2$, over C then there is a unique way to form a multiple spin curve, upto isomorphisms, *unless* both spin curves require destabilization. In the latter case, from each pair (ξ_1, ξ_2) we can form two distinct isomorphism classes of multiple-spin curves.

There are three kinds of interactions that can happen over each node between two limit spin structures: both, one or neither may require destabilization. The following demonstrates all three on one curve.

Example 2.12 Consider a smooth irreducible curve \overline{C} with six distinct points $p_{ij} \in \overline{C}$ for $i = 1, 2, 3$ and $j = 1, 2$. Form the curve $C = \overline{C}/(p_{i1} \sim p_{i2})_{i=1}^3$ and let x_i stand for the node $[p_{i1}]$. We will take limit roots ξ_1 and ξ_2 over C such that ξ_1 requires only the node x_1 to be blown-up while ξ_2 requires x_1 and x_2 to be blown-up. We can check that there are two isomorphism classes of limit multiple-spin curves over (ξ_1, ξ_2) .

3 Torsion-free sheaves

The definition given in Sect. 2 is what is intended to be used in geometric applications. However, for the problem of constructing the relevant moduli spaces and studying the local deformation spaces, we gain an advantage by working with stable curves instead of quasi-stable curves. We will now give our “working definition” for limit multiple roots using torsion-free sheaves.

3.1 Torsion-free roots

Let \mathcal{M} be an algebraic stack. The definitions in this subsection are from [26].

Definition 3.1 (*Jarvis*) A torsion-free sheaf on a stable curve $\mathcal{C} \rightarrow \mathcal{M}$ is a coherent $\mathcal{O}_{\mathcal{C}}$ -module \mathcal{E} which is flat and of finite presentation over \mathcal{M} such that over each $s \in \mathcal{M}$ the fiber $\mathcal{E}|_{\mathcal{C}_s}$ has no associated primes of height one.

Note that the smooth locus of the map $\mathcal{C} \rightarrow \mathcal{M}$ is contained in the locus where \mathcal{E} is locally free.

Definition 3.2 (*Deligne, Jarvis*) Let \mathcal{E} be a rank-1 torsion-free sheaf on a curve $\mathcal{C} \rightarrow \mathcal{M}$ and \mathcal{N} a line bundle on \mathcal{C} . Let $\delta: \mathcal{E} \xrightarrow{\sim} \mathcal{N} \otimes \mathcal{E}^\vee$ be an isomorphism. Then the pair (\mathcal{E}, δ) will be called a (square) root of \mathcal{N} .

Definition 3.3 Given a coherent module \mathcal{E} and a line bundle \mathcal{N} on a scheme \mathcal{X} , a homomorphism $b: \mathcal{E}^{\otimes 2} \rightarrow \mathcal{N}$ will be called a bilinear form. A bilinear form induces two maps $b^l, b^r: \mathcal{E} \rightarrow \mathcal{E}^\vee \otimes \mathcal{N}$ where $\mathcal{E}^\vee = \text{hom}(\mathcal{E}, \mathcal{O}_{\mathcal{X}})$, $b^r(e) = b(e, _)$ and $b^l(e) = b(_, e)$. If both b^r and b^l are isomorphisms then b is non-degenerate. If $b^r = b^l$ then b is symmetric, and b factors through the symmetrizing map $\mathcal{E}^{\otimes 2} \rightarrow \text{Sym}^2 \mathcal{E}$.

We will adopt an unusual notational custom and for any A -module E write the d -th symmetric product $\text{Sym}_A^d(E)$ simply as E^d , and given $\mu: E \rightarrow F$ we will denote by μ^d the induced map $E^d \rightarrow F^d$. The same goes for sheaves of modules and morphisms between them. In compensation, we will write out tensor powers and direct sums explicitly as $E^{\otimes d}$ and $E^{\oplus d}$, respectively. We will use the following reformulation of Definition 3.2.

Definition 3.4 Let \mathcal{E} be a rank-1 torsion-free sheaf on a curve $\mathcal{C} \rightarrow \mathcal{M}$ and \mathcal{N} a line bundle on \mathcal{C} . Let $b: \mathcal{E}^2 \rightarrow \mathcal{N}$ be a non-degenerate symmetric form. Then the pair (\mathcal{E}, b) will be called a (*torsion-free*) *root of \mathcal{N} on \mathcal{C}/\mathcal{M}* .

An isomorphism $\mu: (\mathcal{E}, b) \rightarrow (\mathcal{E}', b')$ of roots is defined to be an isomorphism of the underlying sheaf of modules $\mu: \mathcal{E} \xrightarrow{\sim} \mathcal{E}'$ such that $b = b' \circ \mu^2$.

3.2 Relation to limit roots

Suppose (\mathcal{E}, b) is a torsion-free root of \mathcal{N} on \mathcal{C}/\mathcal{M} . Define $\mathbb{P}(\mathcal{E}) := \text{Proj}_{\mathcal{C}}(\text{Sym}^* \mathcal{E})$ with $\pi: \mathbb{P}(\mathcal{E}) \rightarrow \mathcal{C}$ the structure map and $\mathcal{L} = \mathcal{O}_{\mathbb{P}(\mathcal{E})}(1)$ the line bundle corresponding to the Proj construction. Notice that π is a destabilization.

There are natural surjective maps $\pi^* \mathcal{E}^d \rightarrow \mathcal{L}^{\otimes d} = \mathcal{O}_{\mathbb{P}(\mathcal{E})}(d)$ for each $d \geq 0$ and there is the map $\pi^* b: \pi^* \mathcal{E}^2 \rightarrow \pi^* \mathcal{N}$. As is shown in §3.1.3 of [26] there is a natural map α making the following diagram commute:

$$\begin{array}{ccc}
 \pi^* \mathcal{E}^2 & & \\
 \downarrow & \searrow \pi^* b & \\
 \mathcal{L}^2 & \xrightarrow{\alpha} & \pi^* \mathcal{N}
 \end{array}$$

Proposition 3.5 Let (\mathcal{E}, b) be a root of \mathcal{N} . Then $(\mathbb{P}(\mathcal{E}) \xrightarrow{\pi} \mathcal{C}, \mathcal{O}_{\mathbb{P}(\mathcal{E})}(1), \alpha)$, constructed above, is a family of limit roots of \mathcal{N} .

Proof Both Proj and Sym constructions behave well with respect to base change. So we may reduce to $\overline{T} = \text{Spec } k$ where k is an algebraically closed field. Let $L = \mathcal{O}(1)$ and note $\pi_* L \simeq \mathcal{E}$, see Lemma 3.1.4.(2) [26]. To see that L has degree one over any exceptional fiber E over a node x we simply observe $h^0(E, L|_E) = \dim_k \mathcal{E}|_x = 2$. Since $E \simeq \mathbb{P}_k^1$ we are done. The map α is an isomorphism away from the exceptional divisors because b is an isomorphism away from the corresponding nodes. The degrees of L and \mathcal{E} agree because $\pi_* L \simeq \mathcal{E}$. This completes the proof. □

Conversely, given a family of limit roots $(\pi: \mathcal{X} \rightarrow \mathcal{C}, \mathcal{L}, \alpha)$, let $\mathcal{E} := \pi_* \mathcal{L}$. Then, using Lemma 3.1.4.(2) [26] again, we have $\pi_* \mathcal{L}^2 \simeq \mathcal{E}^2$. Using the adjunction map $a: \pi_* \pi^* \mathcal{N} \rightarrow \mathcal{N}$ we may define $b := a \circ \pi_* \alpha: \mathcal{E}^2 \rightarrow \pi_* \pi^* \mathcal{N} \rightarrow \mathcal{N}$.

Proposition 3.6 The tuple (\mathcal{E}, b) obtained in this way is a torsion-free root of \mathcal{N} .

Proof This is similar to the proposition above. The main ingredients are Proposition 3.1.2.(3) and Proposition 3.1.5 of [26] which says that $\pi_* \mathcal{L}$ is torsion-free and b is of the right form respectively. □

As in Definition 2.9 we may define the category of torsion-free roots. The results of this section imply the following.

Corollary 3.7 The category of torsion-free roots of \mathcal{N} is equivalent to the category of limit roots of \mathcal{N} .

3.3 Multiple torsion-free roots

Let $(\mathcal{E}_i, b_i)_{i=1}^m$ be a sequence of torsion-free roots of \mathcal{N} on a family of curves $\mathcal{C} \rightarrow \mathcal{M}$. Now we will phrase the notion of multiple limit roots from Sect. 2.3 in terms of torsion-free roots using the equivalence identified in Sect. 3.2.

For each root we have the destabilization $\mathcal{X}_i := \text{Proj}_{\mathcal{C}}(\text{Sym}^* \mathcal{E}_i) \rightarrow \mathcal{C}$, and the goal is to find a common destabilization $\mathcal{X} \rightarrow \mathcal{C}$. The roots are partially identified with one another after squaring them and \mathcal{X}_i can also be constructed by taking even powers in the symmetric algebra, i.e., $\mathcal{X}_i \simeq \text{Proj}_{\mathcal{C}}(\text{Sym}^{2*} \mathcal{E}_i)$. Our goal is then to isolate the conditions for which the even symmetric algebras are identified wherever possible.

Notation 3.8 Let $V_i \subset \mathcal{C}$ be the open locus (Lemma 4.6) where the rank of \mathcal{E}_i is maximal amongst all \mathcal{E}_j . This is the analogue of Notation 2.6.

Definition 3.9 Let \mathcal{F} be a sheaf of modules on \mathcal{C} and let $(\varphi_i : \mathcal{F} \rightarrow \mathcal{E}_i^2)_{i=1}^m$ be a sequence of maps such that: (1) for all i, j we have $b_i \circ \varphi_i = b_j \circ \varphi_j$, (2) for all i the map $\varphi_i|_{V_i}$ is an isomorphism. Then, we will call $(\varphi_i)_{i=1}^m$ a *pre-sync data* for the sequence of roots $(\mathcal{E}_i, b_i)_{i=1}^m$. The two conditions above are *pre-sync conditions*.

Suppose $(\varphi_i)_{i=1}^m$ is a pre-sync data for $(\mathcal{E}_i, b_i)_{i=1}^m$. On $V_{ij} := V_i \cap V_j$ we can define $\psi_{ij} = \varphi_j|_{V_{ij}} \circ \varphi_i|_{V_{ij}}^{-1} : \mathcal{E}_i^2|_{V_{ij}} \xrightarrow{\sim} \mathcal{E}_j^2|_{V_{ij}}$. Using ψ_{ij} , we get a surjective map

$$\text{Sym}^* \text{Sym}^2 \mathcal{E}_i|_{V_{ij}} \rightarrow \text{Sym}^{2*} \mathcal{E}_j|_{V_{ij}}. \tag{3.1}$$

Lemma 3.10 *The map (3.1) factors through an isomorphism $\text{Sym}^{2*} \mathcal{E}_i|_{V_{ij}} \rightarrow \text{Sym}^{2*} \mathcal{E}_j|_{V_{ij}}$ if and only if $\text{Sym}^2 \text{Sym}^2 \mathcal{E}_i|_{V_{ij}} \rightarrow \text{Sym}^4 \mathcal{E}_j|_{V_{ij}}$ factors through $\text{Sym}^4 \mathcal{E}_i|_{V_{ij}} \rightarrow \text{Sym}^4 \mathcal{E}_j|_{V_{ij}}$.*

Proof The kernel of $\text{Sym}^* \text{Sym}^2 \mathcal{E}_i \rightarrow \text{Sym}^{2*} \mathcal{E}_i$ is generated by the kernel of $\text{Sym}^2 \text{Sym}^2 \mathcal{E}_i \rightarrow \text{Sym}^4 \mathcal{E}_i$. □

Definition 3.11 If the map $\text{Sym}^2 \text{Sym}^2 \mathcal{E}_i|_{V_{ij}} \rightarrow \text{Sym}^4 \mathcal{E}_j|_{V_{ij}}$ factors through $\text{Sym}^4 \mathcal{E}_i|_{V_{ij}} \rightarrow \text{Sym}^4 \mathcal{E}_j|_{V_{ij}}$ then we will call $(\varphi_i)_{i=1}^m$ a *sync data* for the sequence of roots $(\mathcal{E}_i, b_i)_{i=1}^m$. This condition will be called the *sync condition*.

Definition 3.12 The tuple $(\mathcal{E}_i, b_i, \varphi_i)_{i=1}^m$ will be called a *multiple-root of \mathcal{N}* if $(\varphi_i)_{i=1}^m$ satisfy the sync condition.

An isomorphism between a pair of multiple-roots is a sequence of isomorphisms between the underlying roots compatible with the sync data.

Definition 3.13 For $(\mathcal{C} \rightarrow \mathcal{M}, \mathcal{N})$ and m fixed, let $\overline{\mathcal{S}}^m(\mathcal{N})$ denote the *category of multiple-roots of \mathcal{N}* . That is, $\overline{\mathcal{S}}^m(\mathcal{N}) \rightarrow \mathcal{M}$ is fibered in groupoids with fiber over a scheme T/\mathcal{M} consisting of the multiple-roots of $(\mathcal{C}|_T \rightarrow T, \mathcal{N}|_T)$. The sub-category $\mathcal{S}^m(\mathcal{N}) \hookrightarrow \overline{\mathcal{S}}^m(\mathcal{N})$ is defined by taking the objects where all the roots are line bundles.

Proposition 3.14 *The categories $\mathcal{S}_{\text{lim}}^m(\mathcal{N})$ and $\overline{\mathcal{S}}^m(\mathcal{N})$ are equivalent.*

Proof In light of Proposition 3.6 the pushforward of a multiple limit root is clearly a multiple-root with the synchronization data of the former mapping down to the sync data of the latter.

Conversely, let $(\mathcal{E}_i, b_i, \varphi_i)_{i=1}^m$ be a multiple-root of \mathcal{N} on \mathcal{C} . As in Proposition 3.5 let $\mathcal{X}_i \rightarrow \mathcal{C}$ be the destabilization obtained by the symmetric algebra of \mathcal{E}_i and let $(\mathcal{L}_i, \alpha_i)$ be the limit root on \mathcal{X}_i .

Let $D_i = \text{Sym}^{2*} \mathcal{E}_i$ and construct D by gluing D_i on the charts V_i using the sync data $(\varphi_i)_{i=1}^m$, this construction equips D with a map to $\text{Sym}^{2*} \mathcal{N}$. Since D_i are isomorphic to $\text{Sym}^{2*} \mathcal{N}$ on an open neighbourhood containing the complement of V_i , we can construct maps $D \rightarrow D_i$.

Let $\mathcal{X} \rightarrow \mathcal{C}$ be obtained as $\text{Proj}_{\mathcal{C}}(D)$ with $\mathcal{L} := \mathcal{O}_{\mathcal{X}}(1)$. Since \mathcal{X}_i can also be constructed as $\text{Proj}_{\mathcal{C}}(D_i)$, the maps $D \rightarrow D_i$ induce partial contraction maps $\rho_i: \mathcal{X} \rightarrow \mathcal{X}_i$ and maps $\mathcal{L} \rightarrow \rho_i^* \mathcal{L}_i^{\otimes 2}$. It is now clear that the tuple $(\mathcal{X}/\mathcal{C}, \rho_i^* \mathcal{L}_i, \rho_i^* \alpha_i, \mathcal{L} \rightarrow \rho_i^* \mathcal{L}_i^{\otimes 2})$ is a multiple limit root.

These two constructions are inverses to one another and they are functorial with respect to pullback because pushforward and Proj are functorial. □

Remark 3.15 Examples of multiple roots, parallel to the examples given in Sect. 2.4 can be found in [37, § 1.6.1].

3.4 Bounded degree

For technical reasons we need to introduce a boundedness condition on the degree of the line bundles. For the rest of this article we assume our line bundles have absolutely bounded degree in the following sense.

Definition 3.16 If there exists a constant $c \in \mathbb{Z}$ such that on any component Y of any geometric fiber of $\mathcal{C} \rightarrow \mathcal{M}$ we have $\text{deg } \mathcal{N}|_Y \geq c$ then \mathcal{N} will be said to have *absolutely bounded degree*.

This boundedness condition is weak enough that unless \mathcal{M} has infinitely many disconnected components, the condition is automatically satisfied. Even without this condition, the line bundle $\omega_{\mathcal{C}/\mathcal{M}}^{\otimes l}$ for any $l \in \mathbb{Z}$ has absolutely bounded degree (Sublemma 4.1.10 [26] for $l = 1$, the idea readily generalizes to all $l \in \mathbb{Z}$).

4 The moduli space of multiple-roots is algebraic

Let m be a positive integer m . In this section we will show that the moduli space $\overline{\mathcal{S}}^m(\mathcal{N})$ (Definition 3.13) is an algebraic stack and is locally of finite type over the base \mathcal{M} . Note that we assume \mathcal{N} has absolutely bounded degree (Definition 3.16).

4.1 The moduli space of single roots

When $m = 1$ we call $\overline{\mathcal{S}}^m(\mathcal{N})$ the moduli space of single roots and denote it by $\overline{\mathcal{S}}(\mathcal{N})$. The basic properties of $\overline{\mathcal{S}}(\mathcal{N})$ follow directly from the work of Jarvis [26]. We will list these properties, briefly highlighting the differences in proofs.

Theorem 4.1 *The category $\overline{\mathcal{S}}(\mathcal{N})$ is an algebraic stack. Moreover, the morphism $\pi: \mathcal{S}(\mathcal{N}) \rightarrow \mathcal{M}$ is proper, of finite type and quasi-finite. The diagonal of this morphism is finite and unramified. If \mathcal{M} is a Deligne–Mumford stack then so is $\mathcal{S}(\mathcal{N})$.*

Proof When $\mathcal{N} = \omega_{\mathcal{C}/\mathcal{M}}$ these results from those of [26]. However, those proofs apply with little modification to the present case. The only condition required is that of absolutely bounded degree (Definition 3.16) which we assume. The last statement follows from the condition on the diagonal, see also [39, Tag 04YV]. □

4.2 The moduli space of multiple-roots

Lemma 4.2 *For each proper subset $J \subsetneq \{1, \dots, m\}$ there is a natural forgetful functor $\overline{\mathcal{S}}^m(\mathcal{N}) \rightarrow \overline{\mathcal{S}}^{\#J}(\mathcal{N})$ obtained by forgetting all the roots except for those at position $i \in J$ and adjusting the sync data appropriately.*

Proof Let $(C/T, (\mathcal{E}_i, b_i, \varphi_i : \mathcal{F} \rightarrow \mathcal{E}_i^2)_{i=1}^m)$ be a multiple root. After forgetting the roots in the complement of J , use the partial isomorphisms between the remaining roots induced by φ_i 's to glue together a new sheaf \mathcal{F}' and new maps $\varphi'_i : \mathcal{F}' \rightarrow \mathcal{E}_i^2$ for $i \in J$. □

Remark 4.3 It is clear from the proof that there is a natural map $\tau : \mathcal{F} \rightarrow \mathcal{F}'$ commuting with φ_i 's and φ'_i 's.

Theorem 4.4 *The moduli space $\overline{\mathcal{S}}^m(\mathcal{N})$ is an algebraic stack, locally of finite type over \mathcal{M} .*

Proof The proof is by induction on m , with the case $m = 1$ taken care of in Sect. 4.1. Let us write $\mathcal{Y} = \overline{\mathcal{S}}^{m-1}(\mathcal{N}) \times_{\mathcal{M}} \overline{\mathcal{S}}(\mathcal{N})$ which is an algebraic stack, locally of finite type over \mathcal{M} by induction hypothesis. Use the forgetful maps of Lemma 4.2 corresponding to the subsets $J = \{1, \dots, m - 1\}$ and $J = \{m\}$ respectively to obtain a map $\overline{\mathcal{S}}^m(\mathcal{N}) \rightarrow \mathcal{Y}$.

Let $\mathcal{A}' \rightarrow \mathcal{Y}$ be the category consisting of tuples $(C/T, \mathcal{F})$ where \mathcal{F} is a quasi-coherent sheaf on C , finitely presented and T -flat with T -proper support. Using [23] we conclude that \mathcal{A}' is an algebraic stack, locally of finite type over \mathcal{Y} .

A morphism $T \rightarrow \mathcal{A}'$ corresponds to a tuple $(C/T, (\mathcal{E}_i, b_i, \varphi_i : \mathcal{G} \rightarrow \mathcal{E}_i^2)_{i=1}^{m-1}, (\mathcal{E}_m, b_m), \mathcal{F})$. Consider the category \mathcal{A} over \mathcal{A}' of tuples of the form $(T \rightarrow \mathcal{A}', \tau_1, \tau_2)$ where $\tau_1 : \mathcal{F} \rightarrow \mathcal{G}$ and $\tau_2 : \mathcal{F} \rightarrow \mathcal{E}_m^2$ are morphisms such that $b_i \circ \varphi_i \circ \tau_1 = b_m \circ \tau_2$ for each $i = 1, \dots, m - 1$. Since all relevant modules are finitely presented, \mathcal{A} is an algebraic stack. Let us define $\varphi'_i := \varphi_i \circ \tau_1$ for $i = 1, \dots, m - 1$ and $\varphi'_m := \tau_2$. We can define the subcategory \mathcal{A}_1 of \mathcal{A} where $(\varphi'_i)_{i=1}^m$ satisfies the pre-sync condition (Definition 3.9) and the subcategory \mathcal{A}_2 of \mathcal{A}_1 where $(\varphi'_i)_{i=1}^m$ satisfies—in addition—the sync condition (Definition 3.11).

In light of Remark 4.3 the forgetful map $\overline{\mathcal{S}}^m(\mathcal{N}) \rightarrow \mathcal{Y}$ factors through a map $\overline{\mathcal{S}}^m(\mathcal{N}) \rightarrow \mathcal{A}$ and identifies $\overline{\mathcal{S}}^m(\mathcal{N})$ with \mathcal{A}_2 . To conclude the proof, we will show that $\mathcal{A}_1 \hookrightarrow \mathcal{A}$ is an open immersion and $\mathcal{A}_2 \hookrightarrow \mathcal{A}_1$ is a closed immersion. The first of these statements is proven in Proposition 4.9. The second of these statements is a direct application of [22, Corollaire 7.7.8]. □

4.2.1 Pre-sync condition is open

We use the setting in the proof of Theorem 4.4. Let us recall that a morphism $T \rightarrow \mathcal{A}$ corresponds to a tuple $(C/T, (\mathcal{E}_i, b_i, \varphi_i : \mathcal{G} \rightarrow \mathcal{E}_i^2)_{i=1}^{m-1}, (\mathcal{E}_m, b_m), \mathcal{F}, (\tau_1, \tau_2))$. We will prove that the subcategory $\mathcal{A}_1 \subset \mathcal{A}$ where (τ_1, τ_2) induces a pre-sync condition is open in the algebraic stack \mathcal{A} .

Notation 4.5 For any $T \rightarrow \mathcal{A}$ we define the loci $U_i(T), V_i(T) \subset C_T$ for $i = 1, 2$ in the following manner: Let $U_i(T)$ be the locus where τ_i is an isomorphism. Let $V_1(T)$ be the complement of the locus where \mathcal{F}' is free but \mathcal{E}_m is not. Let $V_2(T)$ be the complement of the locus where \mathcal{E}_m is free but \mathcal{F}' is not.

Whenever we are working locally on C , we may assume $m = 2$ since \mathcal{F}' is locally isomorphic to one of \mathcal{E}_i^2 . When $m = 2$ then $\mathcal{F}' = \mathcal{E}_1^2$. With this remark in mind we will assume $m = 2$. In this case, we simply have a pair of maps $(\varphi_i : \mathcal{F} \rightarrow \mathcal{E}_i^2)_{i=1}^2$ satisfying

$b_1 \circ \varphi_1 = b_2 \circ \varphi_2$ and that our goal is to show the second pre-sync condition defines an open locus on the base T . Note that the second pre-sync condition is equivalent to having $U_i(T) = V_i(T)$.

Lemma 4.6 *The loci U_i and V_i are open and respect base change. Precisely, for any $S \rightarrow T \rightarrow \mathcal{A}$ we have $V_i(S) = V_i(T)|_S$ and $U_i(S) = U_i(T)|_S$.*

Proof The complement of V_i is the locus of points for which \mathcal{E}_i is free but \mathcal{E}_j is not ($j \neq i$). This locus is supported on the discriminant locus. We show in Lemma 4.7 that the rank of a root is constant on each connected component of the discriminant locus. Thus V_i^c is a union of components of the discriminant locus, which is closed. Moreover, the condition of being locally free or not behaves well with respect to base change. Therefore it is clear that $V_i(S) = V_i(T)|_S$.

The fact that the U_i 's respect base change is a consequence of the following general fact. Let $\psi : F \rightarrow E$ be a map of finitely presented modules on C_T . Then the set where ψ is an isomorphism is the intersection $\{\ker \psi = 0\} \cap \{\text{coker } \psi = 0\}$. When E is flat over T then for any $T' \rightarrow T$ we have $\{(\ker \psi)|_{T'} = 0\} \cap \{(\text{coker } \psi)|_{T'} = 0\} = \{\ker(\psi|_{T'}) = 0\} \cap \{\text{coker}(\psi|_{T'}) = 0\}$. The zero locus of a finitely generated module is open and respects base change. □

Lemma 4.7 *Let $C \rightarrow T$ be a stable curve and \mathcal{E} a locally self-dual rank-1 torsion-free module on C . Then the rank of \mathcal{E} is constant on each component of the discriminant locus $Z \subset C$.*

Proof Pick any point $\mathfrak{p} \in Z$ and note $\text{rk } \mathcal{E}|_{\mathfrak{p}}$ is either 1 or 2. By semi-continuity, the locus where the rank of \mathcal{E} is 1 is open in C and therefore on Z . It remains to show that the locus in Z where the rank of \mathcal{E} is 2 is also open.

By making an étale base change and passing to an étale neighbourhood U of \mathfrak{p} , we may assume that $T = \text{Spec } R$ and $U = \text{Spec } A$ where $\exists x, y \in A, \pi = xy \in R$ such that $Z|_U$ is defined by the ideal (x, y) . With U chosen appropriately, we may apply Faltings' classification [15] and conclude that $E = \mathcal{E}|_U$ is either free or of the form $E(p, p)$ with $p^2 = \pi$ (see the notation in *loc. cit.* or "Appendix A"). The discriminant locus $Z|_U$ is isomorphic to $\text{Spec } R/(\pi)$ and $E(p, p)$ is free at a point iff p is invertible there. This is impossible on $Z|_U$ since $p^2 = 0 \pmod{\pi}$. □

Lemma 4.8 *If $V_1 \subset U_1$ then $U_2 \subset V_2$. Similarly with the indices swapped.*

Proof Assuming $V_1 \subset U_1$ we have $V_2^c \subset V_1 \subset U_1$ by definitions. Therefore, if $\exists x \in V_2^c \cap U_2$ then $x \in U_1 \cap U_2$. But this is a contradiction, if φ_1 and φ_2 are isomorphisms at x then \mathcal{E}_i 's are isomorphic at x . On the other hand $x \in V_2^c$ implies that the roots have different ranks at x . □

The following proposition proves that $\mathcal{A}_1 \rightarrow \mathcal{A}$ is an open immersion.

Proposition 4.9 *Take a map $\text{Spec } R \rightarrow \mathcal{A}$. Let $\mathfrak{p} \in \text{Spec } R$ be a point such that $U_i(k) = V_i(k)$ for $i = 1, 2$ where k is the residue field of \mathfrak{p} . Then there exists a Zariski open neighbourhood W of \mathfrak{p} such that $U_i(W) = V_i(W)$ for $i = 1, 2$.*

Proof Let $j : C|_{\mathfrak{p}} \hookrightarrow C$ be the inclusion of the fiber over \mathfrak{p} . The fact that $U_i(k) = U_i(R)|_{\text{Spec } k}$ implies that $U_i(R)$ is an open neighbourhood of $j(U_i(k))$, similarly for V_i 's. We know $V_i(k)$'s cover the fiber $C|_{\mathfrak{p}}$ and, by hypothesis, $U_i(k) = V_i(k)$. Therefore the open sets $U_i(R) \cap V_i(R)$ for $i = 1, 2$ cover the fiber $C|_{\mathfrak{p}}$.

Pick a Zariski neighbourhood $W \subset \text{Spec } R$ of \mathfrak{p} such that the preimage of W is covered by $U_i(R) \cap V_i(R)$. Shrink W so that every component of the discriminant locus intersects the fiber over \mathfrak{p} . Let Z be the components of the discriminant locus on which the \mathcal{E}_i 's are both non-free. On $Z|_{\mathfrak{p}}$ the φ_i 's are isomorphisms, hence they will remain an isomorphism in a neighbourhood of $Z \cap C_{\mathfrak{p}} \subset C$. Shrink W one last time so that the φ_i 's are isomorphisms on all of Z .

We claim that $U_i(W) = V_i(W)$. By Lemma 4.8 it will be sufficient to show $V_i(W) \subset U_i(W)$ for $i = 1, 2$. Pick $x \in V_1(W)$ and suppose for a contradiction that $x \notin U_1(W) \cap V_1(W)$. Then x must lie in $U_2(W) \cap V_2(W)$ which implies that either both \mathcal{E}_i 's are free or both \mathcal{E}_i are non-free at x . Furthermore, $x \in U_2(W)$ implies φ_2 is an isomorphism at x . If both the \mathcal{E}_i 's are free then the fact that φ_i 's commute with b_i 's imply that φ_1 is also an isomorphism. Hence $x \in U_1(W)$. If the \mathcal{E}_i 's are both non-free, then $x \in Z$. But, by our construction of W , $x \in Z$ implies that φ_1 is an isomorphism at x . □

5 Local analysis

Fix a family of stable genus g curves $\mathcal{C} \rightarrow \overline{\mathcal{M}}_g$ together with a line bundle \mathcal{N} on \mathcal{C} . The family $\mathcal{C} \rightarrow \overline{\mathcal{M}}_g$ corresponds to a map $\mathcal{M} \rightarrow \overline{\mathcal{M}}_g$. For this section we assume that \mathcal{M} is a Deligne–Mumford stack. We work over an excellent base scheme S defined over $\mathbb{Z}[1/2]$. A natural choice is to take $\mathcal{M} = \overline{\mathcal{M}}_g$ and $S = \text{Spec } \mathbb{Z}[1/2]$. Our main result in this section is the following theorem, whose proof is at the end of Sect. 5.1.

Theorem 5.1 *If the moduli map $\mathcal{M} \rightarrow \overline{\mathcal{M}}_g$ is smooth, then $\overline{\mathcal{S}}^m(\mathcal{N})$ is smooth over the base scheme S .*

In order to facilitate the study of the coarse moduli space associated to $\overline{\mathcal{S}}^m(\mathcal{N})$, we will study the action of the automorphism groups of objects in $\overline{\mathcal{S}}^m(\mathcal{N})$ on their local deformation functors. We end this section with a purely combinatorial description of this action.

5.1 Patching local deformations

For general notions regarding deformation theory using Artin rings we refer to Schlessinger’s original work [35] and Sernesi’s textbook [36]. See “Appendix A” for the deformation theory of nodes and roots near nodes.

Smoothness of $\overline{\mathcal{S}}^m(\mathcal{N}) \rightarrow S$ can be checked around geometric points of S . Fix an algebraically closed field k and a k -valued point ξ of $\overline{\mathcal{S}}^m(\mathcal{N})$. Let $p \in \mathcal{M}(k)$ and $q \in \overline{\mathcal{M}}_g(k)$ be the images of ξ . Then q corresponds to a stable curve X/k and p induces a line bundle \mathcal{N}_X on X . Let $x_1, \dots, x_n \in X$ be the nodes of X and suppose the nodes are indexed so that at least one of the roots of \mathcal{N}_X appearing in ξ is non-free at x_i precisely when $i \leq n'$ for some $n' \leq n$.

Let $s \in S(k)$ be the image of ξ and write Λ for the complete local ring $\hat{\mathcal{O}}_{S,s}$ (Definition A.39). Denote by Art_{Λ} the category of local Artinian Λ -algebras with residue field k . Let us write $D_{\xi} : \text{Art}_{\Lambda} \rightarrow (\text{Sets})$ for the functor of infinitesimal deformations of ξ , this functor is pro-represented by the complete local ring of $\overline{\mathcal{S}}^m(\mathcal{N})$ at ξ . We will compute this pro-representing ring by breaking the deformation functor D_{ξ} into simpler pieces.

Denote ξ by the tuple $(X/k, \mathcal{N}_X, \bar{\mathcal{R}})$ where $\bar{\mathcal{R}} = (\bar{\mathcal{E}}_i, \bar{b}_i, \bar{\varphi}_i)_{i=1}^m$ is a multiple root of \mathcal{N}_X on X . A deformation of ξ over $R \in \text{Art}_{\Lambda}$ begins with a deformation $(\mathcal{X}/R, \mathcal{N}_{\mathcal{X}}, \iota)$ of $(X/k, \mathcal{N}_X)$ where ι is an identification of the central fiber of $\mathcal{X} \rightarrow \text{Spec } R$ with $X \rightarrow \text{Spec } k$.

In addition, we must have a deformation of the root $\bar{\mathcal{R}}$ which will consist of a multiple root $\mathcal{R} := (\mathcal{E}_i, b_i, \varphi_i)_{i=1}^m$ of $\mathcal{N}_{\mathcal{X}}$ and a sequence of isomorphisms $(j_i : \mathcal{E}_i|_X \rightarrow \bar{\mathcal{E}}_i)_{i=1}^m$ compatible with the sync data $(\varphi_i)_{i=1}^m$ and $(\bar{\varphi}_i)_{i=1}^m$.

Let $D_p : \text{Art}_{\Lambda} \rightarrow (\text{Sets})$ be the functor of infinitesimal deformations of p over \mathcal{M} . Forgetting the deformation of the multiple root $\bar{\mathcal{R}}$ we obtain a map $D_{\xi} \rightarrow D_p$ which is the local version of the forgetful map $\bar{\mathcal{S}}^m(\mathcal{N}) \rightarrow \mathcal{M}$.

Let $x_v \in X$ be a node of X and $\hat{x}_v : \text{Spec } \hat{\mathcal{O}}_{X,v} \rightarrow X$ be the complete local neighbourhood of x_v in X . By pulling back the multiple root $\bar{\mathcal{R}}$ via \hat{x}_v we obtain a local multiple root $\hat{\mathcal{R}}_v := \hat{x}_v^* \bar{\mathcal{R}}$ on $\hat{\mathcal{O}}_{X,v}$.

Notation 5.2 Let $G_v : \text{Art}_{\Lambda} \rightarrow (\text{Sets})$ denote the functor of infinitesimal deformations of the node $\hat{\mathcal{O}}_{X,x_v}$ as in Definition A.2. Let $H_v : \text{Art}_{\Lambda} \rightarrow (\text{Sets})$ denote the functor of infinitesimal deformations of the node $\hat{\mathcal{O}}_{X,x_v}$ together with the local multiple root $\hat{\mathcal{R}}_v$ as in Definition A.31.

Remark 5.3 Recall that if a root is free then it deforms trivially so that the forgetful map $H_v \rightarrow G_v$ is an isomorphism if all roots are free at x_v , that is, if $v > n'$. See Lemma A.8.

Given a deformation $(\mathcal{X}/R, \mathcal{N}_{\mathcal{X}}, \iota, (\mathcal{E}_i, b_i, \varphi_i, j_i)_{i=1}^m)$ of ξ , we can first forget the roots and then pass to a local neighbourhood of x_v or we can first pass to a local neighbourhood of x_v in \mathcal{X} and then forget the roots. As a result, we obtain the map

$$D_{\xi} \rightarrow D_p \times_{G_1} H_1 \times_{G_2} H_2 \cdots \times_{G_{n'}} H_{n'}. \tag{5.1}$$

Lemma 5.4 *The natural transformation (5.1) is an isomorphism. In particular, a deformation of a multiple root is completely recovered by its deformations around the nodes.*

Proof Fix an element $(\mathcal{X}/R, \mathcal{N}_{\mathcal{X}}, \iota)$ of $D_p(R)$. For each $1 \leq v \leq n'$ pick a deformation $\mathcal{R}_v := (E_{v,i}, b_{v,i}, \varphi_{v,i}, j_{v,i})_{i=1}^m$ on $\hat{\mathcal{O}}_{\mathcal{X},x_v}$ of the local multiple root $\hat{\mathcal{R}}_v$. We will to show that there exists a unique deformation on \mathcal{X} of the multiple root $\bar{\mathcal{R}}$ which pulls back to \mathcal{R}_v on the complete local neighbourhood of x_v in \mathcal{X} .

Let \mathfrak{m}_R be the maximal ideal of R . Let $R_l := R/\mathfrak{m}_R^l$ and $\mathcal{X}_l := \mathcal{X}|_{R_l}$ for all $l \geq 0$. For each l and $v \leq n'$ we can pullback \mathcal{R}_v to the formal neighbourhood of x_v in \mathcal{X}_l . We will denote this deformation of $\hat{\mathcal{R}}_v$ by $\mathcal{R}_{v,l}$.

Using induction, we fix $N \geq 0$ and suppose that there is a unique deformation of the multiple root $\bar{\mathcal{R}}$ on \mathcal{X}_N such that for all $v \leq n'$ this deformation agrees with $\mathcal{R}_{v,N}$ around the node x_v .

By constructing a lift of this deformation to \mathcal{X}_{n+1} and showing that this lift is unique up to a unique isomorphism will end the proof. We will do this by fpqc-descent on \mathcal{X}_{n+1} . The synchronized roots around the formal neighbourhoods of the nodes are one portion of the descent data. For the rest of the descent data, we will construct the root away from the nodes and then show compatibility.

On the complement W of the nodes $x_1, \dots, x_{n'}$, the roots we have are all free. Use Lemma A.8 and Remark A.9 to conclude that each root deforms uniquely in W . This uniqueness also proves compatibility with the formal neighbourhoods around the nodes. \square

It is well known [10] that $\hat{\mathcal{O}}_{\bar{\mathcal{M}}_g,q} \simeq \Lambda[[t_1, \dots, t_{3g-3}]]$ where the generators t_i to correspond to the deformation of the node x_i for $i = 1, \dots, n$, the labeling of the rest of the generators correspond to deformations of the components of the normalization of X .

With $n' \leq n$ defined as above, let us define a finite extension $\Lambda[[t_1, \dots, t_{3g-3}]] \rightarrow \Lambda[[\tau_1, \dots, \tau_{3g-3}]]$ so that $t_i \mapsto \tau_i^2$ when $i \leq n'$ and $t_i \mapsto \tau_i$ when $i > n'$.

Proposition 5.5 *The complete local ring of $\overline{S}^m(\mathcal{N})$ at ξ is isomorphic over Λ to the tensor product $\hat{\mathcal{O}}_{\mathcal{M},p} \otimes_{\Lambda[[t_1, \dots, t_{3g-3}]]} \Lambda[[\tau_1, \dots, \tau_{3g-3}]]$.*

Proof We computed in ‘‘Appendix A’’ that G_v is pro-represented by $\Lambda[[t_v]]$ and H_v is pro-represented by $\Lambda[[\tau_v]]$ with the map $H_v \rightarrow G_v$ given by $t_v \mapsto \tau_v^2$ if $v \leq n'$ and $t_v \mapsto \tau_v$ if $v > n'$. The isomorphism of the map (5.1) proved in Lemma 5.4 finishes the proof. \square

Now we are ready to prove the main theorem of this section. We will stick to the notation of Proposition 5.5.

Proof of Theorem 5.1 If $\mathcal{M} \rightarrow \overline{\mathcal{M}}_g$ is smooth then $\hat{\mathcal{O}}_{\mathcal{M},p}$ is smooth over $\Lambda[[t_1, \dots, t_{3g-3}]]$. Since smoothness is stable under pullbacks, the tensor product representing the local ring of $\overline{S}^m(\mathcal{N})$ at ξ is smooth over $\Lambda[[\tau_1, \dots, \tau_{3g-3}]]$, which in turn is smooth over Λ . Being smooth is stable under composition and it follows that $\overline{S}^m(\mathcal{N}) \rightarrow S$ is smooth. \square

5.2 Automorphism groups

The complete local ring of $\overline{\mathcal{M}}_g$ at X is $\Lambda[[t_1, \dots, t_{3g-3}]]$. The moduli map $\mathcal{M} \rightarrow \overline{\mathcal{M}}_g$ gives rise to the map $\Lambda[[t_1, \dots, t_{3g-3}]] \rightarrow \hat{\mathcal{O}}_{\mathcal{M},p}$. For each $i = 1, \dots, n$ let us denote the image of t_i in $\hat{\mathcal{O}}_{\mathcal{M},p}$ by \bar{t}_i . As a result, we can reinterpret Proposition 5.5 as follows:

$$\hat{\mathcal{O}}_{\overline{S}^m(\mathcal{N}),\xi} \simeq \hat{\mathcal{O}}_{\mathcal{M},p}[\tau_1, \dots, \tau_{n'}] / (\tau_1^2 - \bar{t}_1, \dots, \tau_{n'}^2 - \bar{t}_{n'}). \tag{5.2}$$

Scaling τ_i by ± 1 gives an automorphisms of the ring $\hat{\mathcal{O}}_{\overline{S}^m(\mathcal{N}),\xi}$ fixing $\hat{\mathcal{O}}_{\mathcal{M},p}$.

There is an action of $\text{Aut}(\xi)$ on $\hat{\mathcal{O}}_{\overline{S}^m(\mathcal{N}),\xi}$ defined as follows. The group $\text{Aut}(\xi)$ acts on the functor of infinitesimal deformations D_ξ of ξ by changing the identification of the central fiber of a deformation through post-composition via an automorphism of the central fiber. Since $\hat{\mathcal{O}}_{\overline{S}^m(\mathcal{N}),\xi}$ pro-represents D_ξ , the group $\text{Aut}(\xi)$ acts on it. Any subgroup of $\text{Aut}(\xi)$ fixing (X, \mathcal{N}_X) over \mathcal{M} will fix $\hat{\mathcal{O}}_{\mathcal{M},p}$.

Definition 5.6 The subgroup of $\text{Aut}(\xi)$ fixing (X, \mathcal{N}_X) in \mathcal{M} will be denoted by $\text{Aut}_0(\xi)$ and will be called the group of *inessential automorphisms*, in accordance with [5, 8]. By the paragraph above, we get a natural morphism

$$\rho: \text{Aut}_0(\xi) \rightarrow \text{Aut}(\hat{\mathcal{O}}_{\overline{S}^m(\mathcal{N}),\xi} / \hat{\mathcal{O}}_{\mathcal{M},p}). \tag{5.3}$$

Let $W = \{x_1, \dots, x_{n'}\} \subset X$ be the set of nodes where at least one root in ξ is non-free. Denote by $X_W \rightarrow X$ the partial normalization of X at W . We will write V for the set of connected components of X_W .

Notation 5.7 Let Γ be the graph with vertex set V and edge set W , the ends of an edge $w \in W$ are the components of the partial resolution $X_W \rightarrow X$ containing a preimage of w .

We will consider the cohomology of the graph Γ with coefficients in the field \mathbb{F}_2 with two elements. Note that we do not need an orientation on the edges of Γ when working with these coefficients. For our applications, it makes sense to identify \mathbb{F}_2 with the set $\{\pm 1\}$ using the bijection $\mathbb{F}_2 \rightarrow \{\pm 1\}; u \mapsto (-1)^u$. Through this bijection we endow the set $\{\pm 1\}$ with a field structure and denote the resulting field by \mathbb{F}_2^\pm .

Let $C^0(\Gamma) = \mathbb{F}_2^\pm(V)$ and $C^1(\Gamma) = \mathbb{F}_2^\pm(W)$ be the 0-chains and 1-chains of Γ respectively. The usual coboundary map $\partial: C^0(\Gamma) \rightarrow C^1(\Gamma)$ gives rise to the cohomology groups $H^0(\Gamma)$

and $H^1(\Gamma)$ which fit into the exact sequence

$$0 \longrightarrow H^0(\Gamma) \longrightarrow C^0(\Gamma) \xrightarrow{\partial} C^1(\Gamma) \longrightarrow H^1(\Gamma) \longrightarrow 0. \tag{5.4}$$

Remark 5.8 For each element $h \in C^1(\Gamma)$ we can define an automorphism of $\hat{\mathcal{O}}_{\bar{g}^m(\mathcal{N}),\xi}$ over $\hat{\mathcal{O}}_{\mathcal{M},p}$ by scaling τ_i by the value of $h(x_i) \in \{\pm 1\}$, this gives an inclusion

$$\sigma : C^1(\Gamma) \hookrightarrow \text{Aut}(\hat{\mathcal{O}}_{\bar{g}^m(\mathcal{N}),\xi} / \hat{\mathcal{O}}_{\mathcal{M},p}). \tag{5.5}$$

Define a map $\psi : \text{Aut}_0(\xi) \rightarrow C^1(\Gamma)$ by the following rule. For each $a \in \text{Aut}_0(\xi)$ we need to define a value $a(x) \in \{\pm 1\}$ for each $x \in W$. Given $x \in W$, there is at least one index $i \leq m$ such that the i -th root $(\tilde{\mathcal{E}}_i, \bar{b}_i)$ is not locally free at x . Pull back $(\tilde{\mathcal{E}}_i, \bar{b}_i)$ to the formal neighbourhood $\hat{\mathcal{O}}_{X,x}$ and apply Lemma A.21 to conclude that the action of the automorphism a on the i -th root around x can be described by a matrix of the form $\begin{pmatrix} \varepsilon_1 & 0 \\ 0 & \varepsilon_2 \end{pmatrix}$ where $\varepsilon_1, \varepsilon_2 \in \{\pm 1\}$. Let us define $\psi(a)(x) = \varepsilon_1 \varepsilon_2$.

Lemma 5.9 *The value of $\psi(a)(x)$ is well defined.*

Proof Although the sign $\varepsilon_1 \varepsilon_2$ obtained from Lemma A.21 is not well defined for an isomorphism between two different roots, it is well defined for an automorphism of a root. As for the choice of i used in defining $\psi(a)(x)$, the sync data forces the sign $\varepsilon_1 \varepsilon_2$ to be independent of this choice, see ‘‘Appendix A.6’’ for more details on local multiple roots. \square

Theorem 5.10 *The map ψ defined above commutes with the natural maps defined in (5.5) and (5.3). That is, the following diagram is commutative:*

$$\begin{array}{ccc} \text{Aut}_0(\xi) & \xrightarrow{\psi} & C^1(\Gamma) \\ & \searrow \rho & \swarrow \sigma \\ & \text{Aut}(\hat{\mathcal{O}}_{\bar{g}^m(\mathcal{N}),\xi} / \hat{\mathcal{O}}_{\mathcal{M},p}) & \end{array} . \tag{5.6}$$

Proof Fix $a \in \text{Aut}_0(\xi)$ and $x_i \in W$. We will denote the parity by $c := \varepsilon_1 \varepsilon_2 = \psi(a)$. Consider the pullback of ξ to $\hat{\mathcal{O}}_{X,x_i}$ and its universal deformation as in ‘‘Appendix A.6’’. We show in Theorem A.38 that this universal deformation has base $\Lambda[[\tau_i]]$. The automorphism a induces an action on this base ring $\Lambda[[\tau_i]]$ defined as in the description leading up to Definition 5.6. It suffices to show that this induced action consists of scaling τ_i by c .

Since $\tau_i \neq 0$ this change in sign is required to accommodate an isomorphism between the two universal families of local multiple roots whose square has parity c , as in Definition A.34. This follows directly from the arguments given in ‘‘Appendix A.6’’. \square

Now we want to express the image of the map $\psi : \text{Aut}_0(\xi) \rightarrow C^1(\Gamma)$ in a more combinatorial fashion. For each $i = 1, \dots, m$ we can forget all but the i -th root in ξ in order to obtain $\xi_i := (X, \mathcal{N}_X, \tilde{\mathcal{E}}_i, \bar{b}_i)$. We will write $W_i \subset W$ for the set of nodes on which $\tilde{\mathcal{E}}_i$ is not free and by V_i the irreducible components of the normalization $X_{W_i} \rightarrow X$ of X at W_i . As in Notation 5.7 we define a graph $\Gamma_i = (V_i, W_i)$. Moreover, the construction of the map ψ applies also when $m = 1$ and in particular to ξ_i giving us maps $\psi_i : \text{Aut}_0(\xi_i) \rightarrow C^1(\Gamma_i)$. We have exact sequences as in (5.4) obtained by replacing Γ with Γ_i .

In $\text{Aut}_0(\xi)$ and $\text{Aut}_0(\xi_i)$ we have a distinguished automorphism which scales all roots by -1 . Let $\{\pm 1\}$ denote the corresponding subgroup generated by this element. It is clear that both ψ and ψ_i contain the subgroup $\{\pm 1\}$ in their kernels, which prompts the following notation.

Notation 5.11 Let $\underline{\text{Aut}}_0(\xi) = \text{Aut}_0(\xi)/\{\pm 1\}$ and $\underline{\text{Aut}}_0(\xi_i) = \text{Aut}_0(\xi_i)/\{\pm 1\}$ for $i = 1, \dots, m$.

Lemma 5.12 [5, 8] For $i = 1, \dots, m$ we have $\text{Aut}_0(\xi_i) \simeq C^0(\Gamma_i)$ and ψ_i establishes an isomorphism $\underline{\text{Aut}}_0(\xi_i) \xrightarrow{\sim} \ker(C^1(\Gamma_i) \rightarrow H^1(\Gamma_i))$.

Proof Let $v_i : X_{W_i} \rightarrow X$ be the partial normalization of $W_i \subset X$. Let $P_i = v_i^*(W_i)$ and $N_i := v_i^*(\mathcal{N}_X)(-P_i)$. Jarvis shows in §4.1.1 [26] that there exists a line bundle L_i on X_{W_i} , and a squaring map $\beta_i : L_i^{\otimes 2} \xrightarrow{\sim} N_i$ such that $\bar{E}_i \simeq v_{i,*}L_i$ and \bar{b}_i is obtained from β_i . Furthermore, he shows that $\text{Aut}(\bar{E}_i, \bar{b}_i) \simeq \text{Aut}(L_i, \beta_i)$.

The automorphism group of (L_i, β_i) is clearly isomorphic to $H^0(X_{W_i}, \mu_2)$ where $\mu_2 \subset \mathcal{O}_{X_{W_i}}^*$ is the kernel of the squaring map. We have an obvious identification $H^0(X_{W_i}, \mu_2) \simeq C^0(\Gamma_i)$ which proves $\text{Aut}_0(\xi_i) \simeq C^0(\Gamma_i)$. Since Γ_i is connected, $H^0(\Gamma_i) \simeq \{\pm 1\}$. Therefore, $\underline{\text{Aut}}_0(\xi_i) = \text{Aut}_0(\xi_i)/\{\pm 1\} \simeq C^0(\Gamma_i)/H^0(\Gamma_i) \simeq \text{im } \partial$. □

Since elements in $\text{Aut}_0(\xi)$ act on individual roots, we have natural maps $\text{Aut}_0(\xi) \rightarrow \text{Aut}_0(\xi_i)$ obtained by restricting the action of an automorphism of ξ to just the i -th root. Any automorphism of multiple roots can be recovered from its action on the individual roots so the combined map $\text{Aut}_0(\xi) \rightarrow \prod_{i=1}^m \text{Aut}_0(\xi_i)$ is an injection.

We can also define maps $C^1(\Gamma) \rightarrow C^1(\Gamma_i)$ obtained by sending all edges in $W \setminus W_i$ to zero. The joint map $C^1(\Gamma) \rightarrow \prod_{i=1}^m C^1(\Gamma_i)$ is an injection since each edge in W appears in at least one W_i .

Lemma 5.13 With the natural maps described above, we obtain the following Cartesian diagram:

$$\begin{array}{ccc} \underline{\text{Aut}}_0(\xi_0) & \hookrightarrow & \prod_{i=1}^m \underline{\text{Aut}}_0(\xi_i) \\ \psi_0 \downarrow & \lrcorner & \downarrow \prod \psi_i \\ C^1(\Gamma) & \hookrightarrow & \prod_{i=1}^m C^1(\Gamma_i) \end{array}$$

Proof The commutativity is immediate since the value of $\psi(a)(x)$ is determined by the action of a on any any root (\bar{E}_i, \bar{b}_i) which is not free on x . To see that the diagram is Cartesian, we observe that a sequence of automorphisms $(a_i)_{i=1}^m$ where a_i acts on ξ_i are compatible with the sync data if the image of $(a_i)_{i=1}^m$ lies in the image of $C^1(\Gamma)$. □

Proposition 5.14 The map ψ identifies $\underline{\text{Aut}}_0(\xi)$ with the kernel of the map $C^1(\Gamma) \rightarrow \prod_{i=1}^m H^1(\Gamma_i)$.

Proof By Lemma 5.13 we conclude that $\underline{\text{Aut}}_0(\xi)$ is the intersection of $C^1(\Gamma)$ with $\prod_{i=1}^m \underline{\text{Aut}}_0(\xi_i)$ in $\prod_{i=1}^m C^1(\Gamma_i)$. Now apply Lemma 5.12. □

Example 5.15 Let us consider Example 2.10. Here $\Gamma = \Gamma_1 = \Gamma_2$ consists of two vertices and one edge between them. Therefore, $C^1(\Gamma) \simeq \mathbb{F}_2$ and $\prod_{i=1}^2 H^1(\Gamma_i) = 0$. We conclude that $\underline{\text{Aut}}_0(\xi) \simeq \mathbb{F}_2$. In (5.2) we have $n' = 1$ so $\hat{\mathcal{O}}_{\mathcal{M}, p}^{\otimes m}[\tau_1]_{\xi} \simeq \hat{\mathcal{O}}_{\mathcal{M}, p}[\tau_1]/(\tau_1^2 - \bar{t}_1)$ and the generator of $\underline{\text{Aut}}_0(\xi)$ acts by $\tau_1 \mapsto -\tau_1$. In particular, the coarse moduli space of double spin curves is not branched around ξ over the moduli space of curves.

Example 5.16 Consider Example 2.11. Here $\Gamma = \Gamma_1 = \Gamma_2$ consists of one vertex and a loop. Therefore, $C^1(\Gamma) \rightarrow \prod_{i=1}^2 H^1(\Gamma_i)$ becomes $\mathbb{F}_2 \rightarrow \mathbb{F}_2^2 : 1 \mapsto (1, 1)$. In particular, $\underline{\text{Aut}}_0(\xi) = 0$. Using (5.2) we see that the locus of such ξ form a divisorial branch locus of the coarse moduli of double spin curves over the moduli space of curves.

6 Fundamental properties of the moduli of multiple-roots

In this section we prove the additional structural results regarding $\overline{\mathcal{S}}^m(\mathcal{N}) \rightarrow \mathcal{M}$ in order to complete the proof of Theorem 1.1. Throughout this section we will assume \mathcal{M} is a Deligne–Mumford stack.

6.1 Proper, Deligne–Mumford compactification

As we proved in Theorem 4.4 that $\overline{\mathcal{S}}^m(\mathcal{N})$ is algebraic, it follows that the relative diagonal $\overline{\mathcal{S}}^m(\mathcal{N}) \rightarrow \mathcal{M}$ is representable by algebraic spaces.

Lemma 6.1 *The diagonal $\Delta: \overline{\mathcal{S}}^m(\mathcal{N}) \rightarrow \overline{\mathcal{S}}^m(\mathcal{N}) \times_{\mathcal{M}} \overline{\mathcal{S}}^m(\mathcal{N})$ is finite and unramified.*

Proof Fix a morphism $B \rightarrow \overline{\mathcal{S}}^m(\mathcal{N}) \times_{\mathcal{M}} \overline{\mathcal{S}}^m(\mathcal{N})$ where B is a scheme. This defines a curve $C \rightarrow B$ and a pair of multiple roots $\mathcal{R} = (\mathcal{E}_i, b_i, \varphi_i)_{i=1}^m$ and $\mathcal{R}' = (\mathcal{E}'_i, b'_i, \varphi'_i)_{i=1}^m$ on the curve C . We claim that the isomorphism functor $\text{Iso}_C(\mathcal{R}, \mathcal{R}')$ is represented by a finite and unramified scheme over B .

An isomorphism of a multiple root is an isomorphism of the underlying sequence of roots compatible with the sync data. Let $\mathcal{R}_i = (\mathcal{E}_i, b_i)$ and $\mathcal{R}'_i = (\mathcal{E}'_i, b'_i)$ be the i -th roots and consider the natural injection $\text{Iso}_C(\mathcal{R}, \mathcal{R}') \rightarrow \prod_{i=1}^m \text{Iso}_C(\mathcal{R}_i, \mathcal{R}'_i)$, where the product of the functors $\text{Iso}_C(\mathcal{R}_i, \mathcal{R}'_i)$ is to be taken over B . This product is finite and unramified since each component $\text{Iso}_C(\mathcal{R}_i, \mathcal{R}'_i)$ is represented by a finite and unramified scheme over B as shown in §4.1.4.3 of [26]. We claim $\text{Iso}_C(\mathcal{R}, \mathcal{R}')$ is a component of the product and we will perform a series of reductions to prove this.

Since the stacks under consideration are locally noetherian and the question is local on the target of Δ we may assume B to be noetherian. Since the diagonal is representable, $\text{Iso}_C(\mathcal{R}, \mathcal{R}')$ is locally noetherian and we may reduce to the case where $B = \text{Spec } R$ where R is a complete discrete valuation ring. Let the η and σ denote the generic and special points of B respectively. To conclude the proof we need only prove that given a sequence of isomorphisms between the roots, the property of being compatible with the sync data both specializes and generalizes. This problem requires studying pairs of roots at a time and therefore we will assume $m = 2$.

Consider a map $h: \text{Spec } R \rightarrow \prod_{i=1}^m \text{Iso}_C(\mathcal{R}_i, \mathcal{R}'_i)$, which gives us a family of curves $\mathcal{X} \rightarrow \text{Spec } R$ and a sequence of isomorphism $(h_i: (\mathcal{E}'_i, b'_i) \rightarrow (\mathcal{E}_i, b_i))_{i=1}^m$ between the roots. We need to check that the sequence $(h_i)_{i=1}^m$ is compatible with the sync data over the generic fiber iff it is compatible with the sync data over the special fiber.

Compatibility is trivially satisfied except around the nodes where both roots may be singular. In this respect, Lemma 5.4 implies that compatibility need only be checked around the formal neighbourhood of the node of the special fiber. In Lemma A.21 we show that in the formal neighbourhood of a node, the isomorphisms h_i between \mathcal{E}_i and \mathcal{E}'_i are of the form $\begin{bmatrix} \varepsilon'_i & 0 \\ 0 & \varepsilon''_i \end{bmatrix}$, with $\varepsilon'_i, \varepsilon''_i \in \{\pm 1\}$. The symmetric square h_i^2 equals $\begin{bmatrix} 1 & 0 & 0 \\ 0 & \varepsilon_i & 0 \\ 0 & 0 & 1 \end{bmatrix}$, where $\varepsilon_i = \varepsilon'_i \varepsilon''_i$.

In this formal neighbourhood of the node, we can replace the sync data with an isomorphism $\psi: \mathcal{E}_1^2 \rightarrow \mathcal{E}_2^2$ on the first pair of roots and ψ' on the second pair of roots (see Remark A.32). From Proposition A.33 we know that $\psi = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \varepsilon & 0 \\ 0 & 0 & 1 \end{bmatrix}$ and $\psi' = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \varepsilon' & 0 \\ 0 & 0 & 1 \end{bmatrix}$ where $\varepsilon, \varepsilon' \in \{\pm 1\}$.

These four isomorphisms give a commuting diagram iff $\varepsilon \varepsilon_1 = \varepsilon' \varepsilon_2$. This equality holds over the special fiber iff it holds over the generic fiber. □

Proposition 6.2 *The moduli space $\overline{\mathcal{S}}^m(\mathcal{N})$ is a Deligne–Mumford stack.*

Proof The proposition above combined with the fact that \mathcal{M} is Deligne–Mumford stack yields this result. □

Lemma 6.3 *The morphism $\overline{\mathcal{S}}^m(\mathcal{N}) \rightarrow \mathcal{M}$ is proper.*

Proof We prove this by the valuative criterion of properness and induction on $m \geq 1$. The result for $m = 1$ is part of Theorem 4.1. Then we need only show that the map $\mathcal{Y} := \overline{\mathcal{S}}^m(\mathcal{N}) \rightarrow \overline{\mathcal{S}}^{m-1}(\mathcal{N}) \times_{\mathcal{M}} \overline{\mathcal{S}}(\mathcal{N})$ constructed in the first paragraph of the proof of Theorem 4.4 is proper. As the diagonal is locally noetherian we restrict to checking the valuative criterion using complete DVRs.

Let R be a complete DVR, with residue field K . Consider a 2-commutative diagram:

$$\begin{array}{ccc}
 \text{Spec } K & \longrightarrow & \overline{\mathcal{S}}^m(\mathcal{N}) \\
 \downarrow & & \downarrow \\
 \text{Spec } R & \longrightarrow & \mathcal{Y}
 \end{array}$$

This means that we have a synchronized $(m - 1)$ -tuple of roots and an m -th root over the curve $C_R \rightarrow \text{Spec } R$. Furthermore, these roots are all synchronized over the general fiber. However, as we demonstrated in the proof of Proposition 6.1 a synchronization on the generic fiber over a complete DVR extends to the entire family uniquely. □

We now prove that $\overline{\mathcal{S}}^m(\mathcal{N})$ is indeed a ‘closure’ of $\mathcal{S}^m(\mathcal{N})$ in the sense we would expect.

Lemma 6.4 *If $\mathcal{C} \rightarrow \mathcal{M}$ is generically smooth then $\mathcal{S}^m(\mathcal{N}) \rightarrow \overline{\mathcal{S}}^m(\mathcal{N})$ is a dense open immersion.*

Proof Since all roots on a smooth curve are locally free, it suffices to show that any root can be deformed onto a smooth curve. Provided that any singular curve X can be deformed to a smooth curve over \mathcal{M} , it follows immediately from the local deformation functors discussed in Sect. 5.1 that any tuple of roots on X can also be deformed onto a smooth curve over \mathcal{M} . In particular, this is an immediate consequence of Lemma 5.4. □

Remark 6.5 If $\mathcal{C} \rightarrow \mathcal{M}$ is not assumed to be generically smooth the result will certainly not hold, even when $m = 1$. For example one could take $\mathcal{M} = \text{Spec } k$. In this case, the isomorphism classes of roots of a fixed line bundle form a discrete set. If \mathcal{C} is singular, then some of these roots will not be free.

6.2 The coarse moduli space

Lemma 6.6 *The map $\overline{\mathcal{S}}^m(\mathcal{N}) \rightarrow \mathcal{M}$ is quasi-finite.*

Proof We will build on the fact that $\overline{\mathcal{S}}(\mathcal{N}) \rightarrow \mathcal{M}$ is quasi-finite which is a part of Theorem 4.1. Fix a geometric point of the m -fold product $\overline{\mathcal{S}}^m(\mathcal{N}) \times_{\mathcal{M}} \cdots \times_{\mathcal{M}} \overline{\mathcal{S}}(\mathcal{N})$. Our goal is to show that there are finitely many synchronizations on the corresponding sequence of roots. Using Lemma 5.4 we see that we need to synchronize the roots only around the nodes. Now apply Remark A.32 to express synchronizations as a sequence of isomorphisms of the local multiple roots. Now apply Lemma A.21 to see that there are only finitely many isomorphisms. □

The proof below is adapted from Proposition 3.1.1 [25].

Proposition 6.7 *If the coarse moduli space of \mathcal{M} is projective over S then the coarse moduli space of $\overline{\mathcal{S}}^m(\mathcal{N})$ is projective over S .*

Proof It is well known that separated Deligne–Mumford stacks are coarsely represented by algebraic spaces (e.g. Corollary 1.3.1 [28]). So we let $X = \text{coarse}(\overline{\mathcal{S}}^m(\mathcal{N}))$ and $Y = \text{coarse}(\mathcal{M})$ be these coarse moduli spaces with $f : X \rightarrow Y$ the natural map between them.

This map f is proper because the corresponding map between the stacks is proper. Also f is quasi-finite by Lemma 6.6. Therefore f is finite and hence projective. When $Y \rightarrow S$ is projective then so is $X \xrightarrow{f} Y \rightarrow S$. □

6.3 Proof of Theorem 1.1

We will first prove a general version of Theorem 1.1 below. To that end, let us recall all assumptions. Take $S \rightarrow \text{Spec } \mathbb{Z}[1/2]$ to be an excellent scheme. Suppose that \mathcal{M} is a proper Deligne–Mumford stack over S and $\mathcal{C} \rightarrow \mathcal{M}$ is a stable curve of genus $g \geq 2$. Further assume that the generic fiber of $\mathcal{C} \rightarrow \mathcal{M}$ over each irreducible component of \mathcal{M} is smooth. Take a line bundle \mathcal{N} on \mathcal{C} of absolutely bounded degree (Definition 3.16).

Theorem 6.8 *With the conditions above the moduli space $\overline{\mathcal{S}}^m(\mathcal{N})$ is a smooth and proper Deligne–Mumford stack over the base scheme S . Furthermore, the inclusion $\mathcal{S}^m(\mathcal{N}) \hookrightarrow \overline{\mathcal{S}}^m(\mathcal{N})$ is dense and open while the forgetful map $\overline{\mathcal{S}}^m(\mathcal{N}) \rightarrow \mathcal{M}$ is quasi-finite.*

Proof The moduli space $\overline{\mathcal{S}}^m(\mathcal{N})$ is smooth by Theorem 5.1 and Deligne–Mumford by Proposition 6.2. The inclusion $\mathcal{S}^m(\mathcal{N}) \hookrightarrow \overline{\mathcal{S}}^m(\mathcal{N})$ is a dense open immersion by Lemma 6.4. The map $\overline{\mathcal{S}}^m(\mathcal{N}) \rightarrow \mathcal{M}$ is proper by Lemma 6.3 and therefore $\overline{\mathcal{S}}^m(\mathcal{N}) \rightarrow S$ is proper since $\mathcal{M} \rightarrow S$ is assumed to be proper. Lemma 6.6 establishes the claim on quasi-finiteness. □

Take $S = \text{Spec } \mathbb{Z}[1/2]$ or one of its geometric points. Let $\mathcal{M} = \overline{\mathcal{M}}_{g,n}$ be the moduli space of marked stable curves and take the universal curve $\overline{\mathcal{C}}_{g,n} \rightarrow \overline{\mathcal{M}}_{g,n}$ over it with sections $\sigma_1, \dots, \sigma_n : \overline{\mathcal{M}}_{g,n} \rightarrow \overline{\mathcal{C}}_{g,n}$. Fix a tuple $(a_1, \dots, a_n) \in \mathbb{Z}^n$ and define the line bundle \mathcal{N} to be the twisted relative dualizing sheaf $\omega_{\overline{\mathcal{C}}_g/\overline{\mathcal{M}}_g}(\sum_{i=1}^n a_i \sigma_i)$. With these conditions the moduli space $\overline{\mathcal{S}}^m(\mathcal{N})$ satisfies the hypotheses of Theorem 6.8.

Definition 6.9 The moduli space of multiple-spin curves $\overline{\mathcal{S}}_g^m$ is the space $\overline{\mathcal{S}}^m(\mathcal{N})$ where we are using the setup above with $n = 0$ and $\mathcal{N} = \omega_{\mathcal{C}/\mathcal{M}_g}$.

Theorem 1.1 is a special case of Theorem 6.8.

6.4 Normalization of the product

Take $\mathcal{M} = \overline{\mathcal{M}}_g$ with the universal curve $\overline{\mathcal{C}}_g \rightarrow \overline{\mathcal{M}}_g$ and take as the line bundle \mathcal{N} the relative dualizing sheaf $\omega_{\overline{\mathcal{C}}_g/\overline{\mathcal{M}}_g}$. The resulting moduli space $\overline{\mathcal{S}}^m(\mathcal{N})$ comes with a map $\overline{\mathcal{S}}^m(\mathcal{N}) \rightarrow \overline{\mathcal{S}}_g^{\times m}$ which forgets the synchronizations.

Proposition 6.10 *The moduli space $\overline{\mathcal{S}}_g^{\times m}$ is not normal for $m \geq 2$ and the forgetful map $\overline{\mathcal{S}}^m(\mathcal{N}) \rightarrow \overline{\mathcal{S}}_g^{\times m}$ is the normalization.*

Proof Cornalba [8] proves that there is a unique divisorial component of the ramification locus of the map $\overline{\mathcal{S}}_g \rightarrow \overline{\mathcal{M}}_g$ on each of the two components of $\overline{\mathcal{S}}_g$. The two components are defined by the parity of the spin structure which plays no role for the nature of our problem.

The general member of the ramification locus is a limit spin structure on an irreducible curve with a single node, with the spin structure acquiring a singularity at the node. Therefore, $\overline{\mathcal{S}}_g^{\times m}$ has divisorial singularities precisely along the closure of the locus of curves whose generic member is an irreducible curve with a single node and at least two of the m roots are singular at the node. For the rest of the proof we may assume $m = 2$, the argument being similar for $m \geq 2$.

Using Example 5.4.(b) of [8] we can describe the map $\overline{\mathcal{S}}_g^{\times 2} \rightarrow \overline{\mathcal{M}}_g$ locally around the non-normal locus, which in local coordinates is given by the map

$$k[[t_1, \dots, t_{3g-3}]] \rightarrow k[[\tau_{11}, \tau_{12}, t_2, \dots, t_{3g-3}]]/(\tau_{11}^2 - t_1, \tau_{12}^2 - t_1). \tag{6.1}$$

The normalization of the ring on the right is given by

$$k[[\tau_{11}, \tau_{12}, t_2, \dots, t_{3g-3}]]/(\tau_{11}^2 - t_1, \tau_{12}^2 - t_1) \rightarrow (k[[\tau, t_2, \dots, t_{3g-3}]]/(\tau^2 - t_1))^{\times 2}$$

$$\begin{bmatrix} \tau_{11} \\ \tau_{12} \end{bmatrix} \mapsto \begin{bmatrix} (\tau, \tau) \\ (\tau, -\tau) \end{bmatrix}. \tag{6.2}$$

The two components of the normalization correspond to the choice of parity (Definition A.34) of the local sync data around the node. The corresponding formal neighbourhoods in $\overline{\mathcal{S}}^2(\mathcal{N})$ are given exactly as in the normalized components on the left. This follows from the description of the local deformation functors (Proposition 5.5) and the fact that a generic irreducible nodal curve has no automorphisms and, in this case, there are no inessential automorphisms of the roots that act non-trivially on the base of the local deformation functors (Proposition 5.14). When there are $m \geq 2$ roots and $n \leq m$ of them are singular around the node, we would get 2^{n-1} components from the normalization which agrees with the number of choices of parity. □

7 Components of the moduli space of multiple-spin curves

Let \mathcal{S}_g/k be the moduli space of spin curves, as described in the introduction, over an algebraically closed field k of characteristic not 2. In this section, we will classify the irreducible components of $\mathcal{S}_g^m = \mathcal{S}_g \times_{\mathcal{M}_g} \dots \times_{\mathcal{M}_g} \mathcal{S}_g$ for any $m \geq 1$. In doing so, we will also classify the irreducible components of its compactification $\overline{\mathcal{S}}_g^m$ (Definition 6.9).

Lemma 7.1 *The connected components of \mathcal{S}_g^m are irreducible. The number of (connected and irreducible) components of \mathcal{S}_g^m does not depend on the field k . In fact, for any k these components are in natural bijection with the components of $\mathcal{S}_g^m/\mathbb{C}$. These statements are also true for $\overline{\mathcal{S}}_g^m$ in place of \mathcal{S}_g^m .*

Proof The inclusion $\mathcal{S}_g^m \hookrightarrow \overline{\mathcal{S}}_g^m$ is dense and open by Theorem 1.1 so that the irreducible components of \mathcal{S}_g^m and $\overline{\mathcal{S}}_g^m$ coincide.

Theorem 6.8 implies that we can define $\overline{\mathcal{S}}_g^m$ over $\mathbb{Z}[1/2]$, where it is still proper and smooth. Apply Theorem 4.17.(iii) of [10] to conclude that for any algebraically closed field k of char(k) $\neq 2$, the number of connected components of $\overline{\mathcal{S}}_g^m|_k$ is the same as the number of connected components of $\overline{\mathcal{S}}_g^m|_{\mathbb{C}}$. The bijection of the components are attained by taking

the closure of the components over $\overline{\mathcal{S}_g^m}|_{\mathbb{C}}$. Alternatively, the classification of the components given in Sect. 7.1 is characteristic independent and makes the bijection more explicit.

Finally, since the spaces in question are smooth, their connected components and irreducible components coincide. \square

We have thus reduced the problem to classifying the components over $k = \mathbb{C}$. We will take our base field to be \mathbb{C} for the rest of this section. The components of \mathcal{S}_g^m over \mathbb{C} will be classified via the monodromy action on the fibers of $\mathcal{S}_g^m \rightarrow \mathcal{M}_g$.

7.1 Deformation invariants of spin structures

In this subsection we will describe a natural set of deformation invariants of an m -tuple of spin structures on a curve. The main theorem of this section states that these deformation invariants fully classify the components of $\mathcal{S}_g^m = \mathcal{S}_g \times_{\mathcal{M}_g} \cdots \times_{\mathcal{M}_g} \mathcal{S}_g$. The proofs given in this subsection are outlines with proofs completed in the rest of this section.

Let X be smooth proper genus g curve over \mathbb{C} and let S_X denote the fiber of $\mathcal{S}_g \rightarrow \mathcal{M}_g$ over X , which we will view as the isomorphism classes of spin structures $(\eta, \alpha : \eta^{\otimes 2} \xrightarrow{\sim} \omega_{X/\mathbb{C}})$ on X . Up to isomorphisms, we can drop the map α . In other words, the set S_X is in bijection with the isomorphism classes of line bundles η such that $\eta^{\otimes 2} \xrightarrow{\sim} \omega_X$. These isomorphism classes (and, by abuse of notation, the line bundles themselves) are called *theta characteristics*.

The only deformation invariant of a spin curve (X, η) is the *parity* of the theta characteristic η [33], which is defined to be the parity of the dimension of the global sections of η . We will denote the parity by

$$P(\eta) := h^0(\eta) \pmod{2} \in \mathbb{F}_2. \tag{7.1}$$

In other words, \mathcal{S}_g consists of two components and the component to which (X, η) belongs is determined by the parity of η .

Given three theta characteristics $\eta_1, \eta_2, \eta_3 \in S_X$, we can build a fourth one $\eta_1 \otimes \eta_2 \otimes \eta_3 \otimes \omega_X^\vee = \eta_1 \otimes \eta_2 \otimes \eta_3^\vee$. In general, the parity of η_i 's alone do not determine the parity of $\eta_1 \otimes \eta_2 \otimes \eta_3 \otimes \omega_X^\vee$ which, therefore, introduces another deformation invariant.

Definition 7.2 For an m -tuple $\eta_\bullet = (\eta_1, \dots, \eta_m) \in S_X^m$ we define the *syzygy relations* $\text{syzy}(\eta_\bullet)$ to be the parities of all η_i and all $\eta_j \otimes \eta_k \otimes \eta_l^\vee$:

$$\text{syzy}(\eta_\bullet) := \left(P(\eta_i); P(\eta_j \otimes \eta_k \otimes \eta_l^\vee) \right)_{\substack{1 \leq i \leq m \\ 1 < j < k \leq m}} \in \mathbb{F}_2^{m + \binom{m-1}{2}}. \tag{7.2}$$

Once we account for the relations between the theta characteristics, the syzygy relations will give a complete set of deformation invariants. Moreover, we will see that the syzygy relations will be a minimal set of deformation invariants when $g \geq m$ and when there are no relations between the theta characteristics.

Definition 7.3 A vector $c_\bullet \in \mathbb{F}_2^m$ with $\sum c_i = 0$ is a *relation* for a tuple $\eta_\bullet \in S_X^m$ if

$$\eta_1^{c_1} \otimes \cdots \otimes \eta_m^{c_m} \simeq \omega_X^k \quad \text{for some } k \in \mathbb{Z}. \tag{7.3}$$

The set of relations $R(\eta_\bullet) \subset \mathbb{F}_2^m$ of η_\bullet is a vector space. If $R(\eta_\bullet) = 0$ then η_\bullet (as well as (X, η_\bullet)) will be called *non-degenerate*.

Remark 7.4 Note that the condition $\sum c_i = 0$ is redundant if $g \neq 1$ by degree considerations. Also the lift of $c_i \in \mathbb{F}_2$ to \mathbb{Z} in taking a power of η_i influences only the value of k , but is otherwise immaterial.

Remark 7.5 If c_\bullet is a relation for η_\bullet and $c_i \neq 0$ then

$$\eta_i \simeq \omega_X^{k'} \bigotimes_{j \neq i} \eta_j^{c_j} \tag{7.4}$$

for some k' . In other words, if $R(\eta_\bullet) \neq 0$ then η_\bullet can be reconstructed from a non-degenerate sub-tuple of η_\bullet .

A portion of the main result of this section can be stated as follows.

Theorem 7.6 *The pair $(R(\eta_\bullet), \text{syzy}(\eta_\bullet))$ is a complete deformation invariant of (X, η_\bullet) . That is, two elements $(X, \eta_\bullet), (X', \eta'_\bullet) \in \mathcal{S}_g^m$ lie in the same component of \mathcal{S}_g^m if and only if $R(\eta_\bullet) = R(\eta'_\bullet)$ and $\text{syzy}(\eta_\bullet) = \text{syzy}(\eta'_\bullet)$.*

Proof The deformation invariance of (R, syzy) is Proposition 7.19. In particular, we can deform (X', η'_\bullet) to some (X, η''_\bullet) while keeping (R, syzy) constant. Now we use the monodromy action of $\mathcal{S}_g^m \rightarrow \mathcal{M}_g$ over X to argue that (X, η_\bullet) and (X, η''_\bullet) lie on the same component if and only if their space of relations and their syzygies coincide. The last step combines Sect. 7.3 (Corollary 7.21) and Sect. 7.4 (Corollary 7.27). \square

The theorem above can be strengthened by describing exactly which syzygy relations can occur. We begin by shortening the syzygy relations by using a non-degenerate sub-tuple.

Definition 7.7 Given $\eta_\bullet \in \mathcal{S}_X^m$, there is a unique non-degenerate sub-tuple $\tilde{\eta}_\bullet = (\eta_{i_1}, \dots, \eta_{i_k})$ where $k = m - \dim R(\eta_\bullet)$ and $i_1 < \dots < i_k$ is the lexicographically smallest index set. Define the *reduced syzygy relations* $\widetilde{\text{syzy}}(\eta_\bullet)$ to be the syzygy relations $\text{syzy}(\tilde{\eta}_\bullet)$ of this sub-tuple.

Let $\text{Gr}(a, m)$ be the Grassmannian of a -dimensional subspaces in \mathbb{F}_2^m .

Definition 7.8 Given $R \in \text{Gr}(m - k, m)$ and $\mathfrak{a} \in \mathbb{F}_2^{k + \binom{k-1}{2}}$ define

$$\mathcal{S}_g^{(R, \mathfrak{a})} := \{(X, \eta_\bullet) \in \mathcal{S}_g^m \mid R(\eta_\bullet) = R, \widetilde{\text{syzy}}(\eta_\bullet) = \mathfrak{a}\}. \tag{7.5}$$

If $R = 0$ then we denote $\mathcal{S}_g^{(0, \mathfrak{a})}$ by $\mathcal{S}_g^\mathfrak{a}$ and call it a *non-degenerate component* of \mathcal{S}_g^m . When $R \neq 0$, the component $\mathcal{S}_g^{(R, \mathfrak{a})}$ is *degenerate*.

For any $R \in \text{Gr}(m - k, m)$ and $\mathfrak{a} \in \mathbb{F}_2^{k + \binom{k-1}{2}}$ we can consider two spaces $\mathcal{S}_g^\mathfrak{a} \subset \mathcal{S}_g^k$ and $\mathcal{S}_g^{(R, \mathfrak{a})} \subset \mathcal{S}_g^m$. It is clear that there is an isomorphism $\mathcal{S}_g^\mathfrak{a} \xrightarrow{\sim} \mathcal{S}_g^{(R, \mathfrak{a})}$ which maps (X, η_\bullet) to $(X, \omega_X^{k'} \bigotimes_j \eta_j^{c_{i,j}})$ where $k_i \in \mathbb{Z}$ and $c_{i,j} \in \mathbb{F}_2$ depend only on R . In particular, $\mathcal{S}_g^{(R, \mathfrak{a})}$ is non-empty iff $\mathcal{S}_g^\mathfrak{a}$ is non-empty, regardless of R .

As a consequence, $\widetilde{\text{syzy}}(\eta_\bullet)$ determines $\text{syzy}(\eta_\bullet)$ and Theorem 7.6 holds with the syzygy relations syzy replaced by reduced syzygy relations $\widetilde{\text{syzy}}$. Hence, the components of \mathcal{S}_g^m are indexed by a subset of

$$\bigcup_{k=0}^{m-1} \text{Gr}(m - k, m) \times \mathbb{F}_2^{k + \binom{k-1}{2}}. \tag{7.6}$$

Definition 7.9 Let $\text{Syzy}_g^m \subset \mathbb{F}_2^{m + \binom{m-1}{2}}$ be the set of all *admissible syzygy relations*,

$$\text{Syzy}_g^m = \{\mathfrak{a} \in \mathbb{F}_2^{m + \binom{m-1}{2}} \mid \mathcal{S}_g^\mathfrak{a} \neq \emptyset\}. \tag{7.7}$$

Since $\mathcal{S}_g^{(R, \alpha)} \simeq \mathcal{S}_g^\alpha$, the decomposition of \mathcal{S}_g^m into its (non-empty) irreducible components is given by

$$\mathcal{S}_g^m = \bigcup_{k=0}^{m-1} \bigcup_{R \in \text{Gr}(m-k, m)} \bigcup_{\alpha \in \text{Syzy}_g^k} \mathcal{S}_g^{(R, \alpha)}. \tag{7.8}$$

To complete the classification of components of \mathcal{S}_g^m , it remains to describe the set Syzy_g^m for any m .

Theorem 7.10 *If $g \geq m$ then $\text{Syzy}_g^m = \mathbb{F}_2^{m+\binom{m-1}{2}}$, i.e., $\forall \alpha \in \mathbb{F}_2^{m+\binom{m-1}{2}}$ the space \mathcal{S}_g^α is non-empty. If $g < \frac{m-1}{2}$ then $\text{Syzy}_g^m = \emptyset$. In general, for given (α, g) , Algorithm 7.14 decides if \mathcal{S}_g^α is empty or not.*

Proof We use the equivalence of theta characteristics and quadratic forms (Sect. 7.2) so that the problem is about the existence of a non-degenerate m -tuple of quadratic forms with syzygy relations α on a $2g$ -dimensional non-singular symplectic \mathbb{F}_2 -vector space. This problem is addressed in Sect. 7.5, in particular by Lemma 7.29 and Proposition 7.32. See the last paragraph of Sect. 7.5 to see Algorithm 7.14 rephrased in that context. \square

Remark 7.11 The bounds in the theorem are optimal. If $g < m$, some \mathcal{S}_g^α will be empty. In fact, the sequence α consisting solely of 1s will give an empty component whenever $g < m$. If $g \geq \frac{m-1}{2}$ then some \mathcal{S}_g^α will be non-empty.

Example 7.12 For $g \geq 2$, \mathcal{S}_g^2 has 4 irreducible components: $\mathcal{S}_g^{(0,0)}$, $\mathcal{S}_g^{(1,0)}$, $\mathcal{S}_g^{(0,1)}$, $\mathcal{S}_g^{(1,1)}$. If $g = 1$ then $\mathcal{S}_g^{(1,1)}$ is empty. Indeed, an elliptic curve has only one odd ($P = 1$) theta characteristic.

Example 7.13 For $g \geq 3$, \mathcal{S}_g^3 has 16 components of the form $\mathcal{S}_g^{(a_1, a_2, a_3; a_{23})}$. If $g = 2$ then $\mathcal{S}_g^{(1,1,1;1)}$ is empty. To see this directly, take a genus 2 hyperelliptic curve C with Weierstrass points w_1, \dots, w_6 . Up to relabeling, any three distinct odd theta characteristics can be written as $\eta_i \equiv w_i$ for $i = 1, 2, 3$. But then $\eta_1 \otimes \eta_2 \otimes \eta_3^{-1} \equiv w_1 + w_2 + w_3 - g_2^1$ is an even ($P = 0$) theta characteristic.

Algorithm 7.14 Given $\alpha = (a_i; a_{ij}) \in \mathbb{F}_2^{m+\binom{m-1}{2}}$ and $g \in \mathbb{N}$ this algorithm decides if \mathcal{S}_g^α is empty or not. For $i, j \in \{2, \dots, m\}$ let $e_{ii} = 0$ and $e_{ij} = a_{ij} + a_i + a_j + a_1$ for $i \neq j$. Let $E = (e_{ij})_{2 \leq i, j \leq m}$ and $K = \ker E \subset \mathbb{F}_2^{m-1}$. Consider the quadratic form $q = \sum_{2 \leq i \leq m} (a_i + a_1)x_i^2 + \sum_{2 \leq i < j \leq m} e_{ij}x_i x_j$ as a function on \mathbb{F}_2^{m-1} with coordinates x_2, \dots, x_m . Let $U \subset \mathbb{F}_2^{m-1}$ be a space complementary to K , i.e., $U \oplus K = \mathbb{F}_2^{m-1}$. If $q|_K \not\equiv 0$ or if the Arf invariant of $q|_U$ is a_1 then let $c = 0$, otherwise let $c = 1$. The equality $2g \geq (m-1) + \dim K + 2c$ holds iff \mathcal{S}_g^α is non-empty.

For small (m, g) we can count the number of non-degenerate components $\mathcal{S}_g^\alpha \subset \mathcal{S}_g^m$ using this algorithm. Let $N(m, g)$ be the cardinality of Syzy_g^m . Theorem 7.10 implies that $N(m, \geq m) = 2^{m+\binom{m-1}{2}}$ and $N(2k+1, < k) = 0$. The Table 1 lists the values of $N(m, g)$ for $m, g \leq 7$ obtained by applying Algorithm 7.14 to each possible α .

Remark 7.15 It is easy to prove using Sect. 7.5 that $N(m, m-1) = N(m, m) - 1$. The other extreme is also easy to describe: $N(2k+1, k)$ counts the number of non-singular quadratic

Table 1 The number $N(m, g)$ of non-degenerate components of S_g^m

$m \setminus g$	0	1	2	3	4	5	6	7
1	1	2	2	2	2	2	2	2
2	0	3	4	4	4	4	4	4
3	0	4	15	16	16	16	16	16
4	0	0	84	127	128	128	128	128
5	0	0	448	1876	2047	2048	2048	2048
6	0	0	0	41664	64852	65535	65536	65536
7	0	0	0	888832	3819200	4191572	4194303	4194304

If $g \geq m$ then $N(m, g) = 2^{m+\binom{m-1}{2}}$. If $2g < m - 1$ then $N(m, g) = 0$

forms on \mathbb{F}_2^{2k} . We can evaluate this number by summing the orbit sizes of an even and odd quadratic form:

$$N(2k + 1, k) = \frac{\# \text{GL}(\mathbb{F}_2^{2k})}{\# O^+(2k)} + \frac{\# \text{GL}(\mathbb{F}_2^{2k})}{\# O^-(2k)}, \tag{7.9}$$

where $O^\pm(2k) \subset \text{GL}(\mathbb{F}_2^{2k})$ is the orthogonal group preserving a non-singular even (+) or odd (-) quadratic form. The cardinality of these groups are well known [27]:

$$\# \text{GL}(\mathbb{F}_2^{2k}) = \prod_{j=0}^{2k-1} (2^{2k} - 2^j), \quad \# O^\pm(2k) = 2^{k^2-k+1} \cdot (2^k \mp 1) \prod_{j=1}^{k-1} (2^{2j} - 1).$$

7.2 Theta characteristics and quadratic forms

This section translates the set-up from the previous section on theta characteristics into the language of quadratic forms, which will be more convenient for the proofs. See [20] for our point of view and a survey of the connection between quadratic forms, theta characteristics and Arf invariants.

We recall our notation from the previous section: X/\mathbb{C} is a smooth proper curve of genus $g \geq 0$ and S_X is the set of theta characteristics on X , i.e., isomorphism classes of line bundles η on such that $\eta^{\otimes 2} \simeq \omega_X$.

Let V_X be the set of 2-torsion line bundles of X up to isomorphism, i.e., for $\varepsilon \in V_X$ we have $\varepsilon^{\otimes 2} \simeq \mathcal{O}_X$. We will make the identification $V_X = H^1(X, \mathbb{F}_2)$ where the right hand side uses Betti cohomology. This identification comes from viewing Pic_X^0 as the quotient $H^1(X, \mathcal{O}_X)/H^1(X, \mathbb{Z})$ from which it follows that the 2-torsion elements are $H^1(X, \mathbb{F}_2) \simeq H^1(X, \frac{1}{2}\mathbb{Z})/H^1(X, \mathbb{Z})$. Note that $V_X \simeq \mathbb{F}_2^{2g}$ and S_X is a V_X torsor.

Using the intersection product on Betti cohomology (or equivalently the Weil pairing on 2-torsion line bundles [1, p. 282]) we see that V_X is a non-singular symplectic vector space: $\forall v \in V_X, v \cdot v = 0$ and $\forall v \in V_X, \exists u \in V_X$ such that $u \cdot v \neq 0$. Let $\text{Sp}(V_X)$ be the symplectic linear group acting on V_X .

Denote by $Q_X \subset \text{Sym}^2 V_X^\vee$ the set of quadratic forms $V_X \rightarrow \mathbb{F}_2$ associated to the inner product on V_X . In characteristic 2 this means that for any $q \in Q_X$ we have

$$\forall u, v \in V_X, \quad q(u + v) + q(u) + q(v) = u \cdot v. \tag{7.10}$$

The set Q_X is an affine space (i.e., torsor) over V_X . Given $v \in V_X$ and $q \in Q_X$ we write $q + v$ for the quadratic form defined by $(q + v)(u) = q(u) + v \cdot u$. Conversely, given two

quadratic forms $q, q' \in Q_X$ we denote by $q - q'$ the element of V_X translating q' to q . The only $\text{Sp}(V_X)$ -invariant of a quadratic form $q \in Q_X$ is its Arf invariant $\text{Arf}(q) \in \mathbb{F}_2$, see [2]. The value $\text{Arf}(q)$ coincides with the value which q attains most frequently on V_X .

The spaces Q_X and S_X are canonically identified, see [1, p. 290] and [33]. We will denote this identification by $S_X \xrightarrow{\sim} V_X: \eta \mapsto q_\eta$ where q_η is defined by

$$\forall v \in V_X, \quad q_\eta(v) = h^0(\eta \otimes v) - h^0(\eta) \pmod{2}. \tag{7.11}$$

The Arf invariant of the quadratic form q_η coincides with the parity of η , [33],

$$\text{Arf}(q_\eta) = h^0(\eta) \pmod{2} = P(\eta). \tag{7.12}$$

As we are in characteristic 2, we have $q_i - q_j = q_j - q_i$ so we define $q_i + q_j := q_i - q_j = q_j - q_i$. It is standard that with this operation $V_X \oplus Q_X$ becomes a vector space with \mathbb{F}_2 -grading.

Definition 7.16 A relation of q_\bullet is an element $c_\bullet \in \mathbb{F}_2^m$ satisfying $\sum c_j q_j = 0$. The \mathbb{F}_2 -grading implies that $\sum c_j = 0$. We denote the subspace of relations of q_\bullet by $R(q_\bullet) \subset \mathbb{F}_2^m$. When $R(q_\bullet) = 0$ we say q_\bullet is non-degenerate.

Let $\eta_\bullet \in S_X^m$ be an m -tuple of theta characteristics and $q_\bullet \in Q_X^m$ be the corresponding sequence of quadratic forms ($q_i := q_{\eta_i}$).

Lemma 7.17 *The space of relations for η_\bullet and q_\bullet coincide, i.e., $R(q_\bullet) = R(\eta_\bullet) \subset \mathbb{F}_2^m$.*

Proof Using (7.11) and (7.10) it can be checked that $q_i + q_j = q_{\eta_i \otimes \eta_j \otimes \omega_X^\vee}$. It is now clear that the relations of η_\bullet as defined in Definition 7.3 coincides with the relations of q_\bullet . \square

Remark 7.18 Using $q_i + q_j = q_{\eta_i \otimes \eta_j \otimes \omega_X^\vee}$ we find $q_1 + q_2 + q_3 = q_{\eta_1 \otimes \eta_2 \otimes \eta_3 \otimes \omega_X^\vee}$. Therefore, the syzygy relations of a sequence of theta characteristics on X (Definition 7.2) and the syzygy relations of the corresponding quadratic forms (Definition 7.23) coincide.

Proposition 7.19 *For any deformation of (X, η_\bullet) the space of relations $R(\eta_\bullet) \subset \mathbb{F}_2^m$ and the syzygy relations $\text{syzy}(\eta_\bullet)$ remains unchanged.*

Proof As (X, η_\bullet) varies in a smooth family, the space V_X remains analytic locally constant (Ehresmann’s fibration theorem) and, therefore, so does Q_X . As the parity of a theta characteristic is a deformation invariant [33] the identification of η_\bullet with quadratic forms q_\bullet can be done in families using (7.11). In particular, the syzygy relations remain constant during deformations.

As V_X and Q_X are discrete, the sections corresponding to η_\bullet and q_\bullet are locally constant. Then the space of relations $R(q_\bullet)$ is locally constant. But the ambient space \mathbb{F}_2^m of $R(q_\bullet)$ is globally constant and then so is $R(q_\bullet)$. Now use Lemma 7.17. \square

7.3 Monodromy action on theta characteristics

We can deduce the connected components of S_g^m by studying the monodromy action of the finite cover $\pi : S_g^m \rightarrow \mathcal{M}_g$. The references for the standard results in this section are [13, 21] and [1, Appendix B].

Let X be a smooth proper complex curve of genus $g \geq 1$, S_X the fiber of π over X , $V_X = H^1(X, \mathbb{F}_2) \simeq \mathbb{F}_2^{2g}$, and Q_X is the affine space of quadratic forms on V_X compatible with the symplectic intersection product.

We recall that the orbifold fundamental group of \mathcal{M}_g is isomorphic to the mapping class group Γ_g . The monodromy action of the mapping class group Γ_g on $V_X = H^1(X, \mathbb{F}_2)$ preserves the intersection product and induces a *surjective* map $s: \Gamma_g \rightarrow \text{Sp}(V_X)$. Let $\text{Sp}(V_X)$ act on Q_X from the left by inverse precomposition, that is $\text{Sp}(V_X) \times Q_X \rightarrow Q_X: (\gamma, q) \mapsto q \circ \gamma^{-1}$. The following statement appears to be well known [38] and, in any case, is not hard to prove (see [1, p. 294]).

Proposition 7.20 *The monodromy action of Γ_g on S_X factors through the natural surjective map $s: \Gamma_g \rightarrow \text{Sp}(V_X)$. With respect to the identification $S_X \mapsto Q_X: \eta \mapsto q_\eta$, the monodromy action on S_X coincides with the precomposition action of $\text{Sp}(V_X)$ on Q_X . More precisely, $\forall u \in \Gamma_g, \forall \eta \in S_X$ we have $q_{u \cdot \eta} = q_\eta \circ s(u)^{-1}$.*

Combining Proposition 7.20 with the irreducibility of \mathcal{M}_g we can determine the irreducible components of the m -fold product $\mathbb{S}_g^m = \mathbb{S}_g \times_{\mathcal{M}_g} \cdots \times_{\mathcal{M}_g} \mathbb{S}_g$.

Corollary 7.21 *The irreducible (and connected) components of \mathbb{S}_g^m are in bijection with the $\text{Sp}(V_X)$ -orbits of $Q_X^m = Q_X \times \cdots \times Q_X$ acting by precomposition, where $\gamma \in \text{Sp}(V_X)$ maps (q_1, \dots, q_m) to $(q_1 \circ \gamma^{-1}, \dots, q_m \circ \gamma^{-1})$.*

Proof The morphism $\pi: \mathbb{S}_g^m \rightarrow \mathcal{M}_g$ is an orbifold finite cover. Therefore, the components of \mathbb{S}_g^m are in bijection with the orbits of any fiber of π under the monodromy action. Proposition 7.20 states that the monodromy action of $\pi_1(\mathcal{M}_g, X) \simeq \Gamma_g$ on $S_X \simeq Q_X$ factors through a surjective map $\Gamma_g \rightarrow \text{Sp}(V_X)$ followed by the precomposition action of $\text{Sp}(V_X)$ on Q_X . \square

7.4 Affine geometry of quadratic forms

The goal of this subsection is to classify the orbits of m -tuples of quadratic forms on a symplectic \mathbb{F}_2 -vector space under the action of the symplectic group. Via Corollary 7.21 this will allow us to classify the components of \mathbb{S}_g^m . The standard text book [3] offers a comprehensive treatment of quadratic spaces in characteristic not 2, the basic treatment in characteristic 2 is similar [29].

Let V be a $2g$ -dimensional vector space over \mathbb{F}_2 with a non-singular symplectic inner product $f: V \times V \rightarrow \mathbb{F}_2$. Let Q be the set of quadratic forms on V associated to the inner product f . The set Q is an affine space with space of translations V . The translation action is given by $V \times Q \rightarrow V: (v, q) \mapsto q + f(v, \cdot)$. We will simply write the latter quadratic form as $q + v$. Moreover, if $q, q' \in Q$ are such that $q' = q + v$ then we will express v as $q' - q$, or as $q' + q$ since the characteristic is 2.

Let $\text{Sp}(V) := \text{Sp}(V, f)$ be the symplectic group of V preserving f . There is a natural action of $\text{Sp}(V)$ on Q via precomposition. In other words, there is a pairing $\text{Sp}(V) \times Q \rightarrow Q: (\gamma, q) \mapsto \gamma_* q = q \circ \gamma^{-1}$. Notice that $\gamma_*(q + v) = \gamma_* q + \gamma(v)$ since f is invariant under γ . The Arf invariant $\text{Arf}: Q \rightarrow \mathbb{F}_2$ is the only invariant under the action of $\text{Sp}(V)$ [2].

To classify the $\text{Sp}(V)$ -orbits of $Q^m = Q \times \cdots \times Q$ under precomposition we could remove the diagonals from Q^m and use induction, since the orbits contained in a diagonal correspond to orbits in Q^{m-1} . However, we can do better: the affine linear structure on Q allows us to reduce the number of elements in a tuple of quadratic forms to an affine linearly independent subset. As a consequence, we need only consider affine linearly independent tuples from Q^m .

Definition 7.22 A sequence $q_\bullet = (q_1, \dots, q_m)$ of quadratic forms on V is called *non-degenerate* if the affine span generated by q_\bullet in Q is of dimension $m - 1$. (This definition agrees with the one in Definition 7.16.)

For non-degenerate tuples, we will see that the “syzygy relations” as defined below classify the $\text{Sp}(V)$ -orbits.

Definition 7.23 Given a sequence of quadratic forms $q_\bullet = (q_1, \dots, q_m)$, let the *syzygy relation of q_\bullet* to be the tuple

$$\text{syzy}(q_\bullet) := (\text{Arf}(q_i); \text{Arf}(q_i + q_j + q_1))_{\substack{1 \leq i \leq m \\ 1 < j < k \leq m}} \in \mathbb{F}_2^{m + \binom{m-1}{2}}. \tag{7.13}$$

Notation 7.24 Given $q_\bullet \in Q^m$ define $W(q_\bullet) = (W, q)$ where W is obtained by translating the affine span of q_\bullet by q_1 and the quadratic form q is the restriction of q_1 to $W \subset V$. The obvious generating vectors for W will be denoted by $v_i = q_i + q_1$ for $i = 2, \dots, m$. Let $e_i = q(v_i)$, $e_{ij} = f(v_i, v_j)$.

Lemma 7.25 *If $\text{syzy}(q_\bullet) = (a_i; a_{ij})$ then $q(v_i) = a_i + a_1$ and $f(v_i, v_j) = a_{ij} + a_i + a_j + a_1$.*

Proof There is a relation between the Arf invariant of a quadratic form and the Arf invariant of its translate, which give

$$\text{Arf}(q_i) = \text{Arf}(q_1 + v_i) = \text{Arf}(q_1) + q_1(v_i). \tag{7.14}$$

Hence we deduce $q(v_i) = a_1 + a_i$. To evaluate $f(v_i, v_j)$ we combine the relation above and the bilinearity relation as in (7.10). □

Witt’s Lemma (Theorem 3.3 and Exercise 3.31 in [40]) Let (V, q) and (V', q') be isometric non-singular quadratic spaces. Let $W \subset V$ and $W' \subset V'$ be subspaces and $\mu : (W, q|_W) \xrightarrow{\sim} (W', q'|_{W'})$ be an isometry. Then, there is an isometry $\gamma : (V, q) \rightarrow (V', q')$ such that $\gamma|_W = \mu$.

Proposition 7.26 *Two sequences of non-degenerate quadratic forms on V are in the same $\text{Sp}(V)$ -orbit if and only if their syzygy relations are equal.*

Proof Let $q_\bullet, q'_\bullet \in Q^m$ be non-degenerate. Since $\text{Sp}(V)$ preserves the Arf invariant, if there is a $\gamma \in \text{Sp}(V)$ such that $\gamma_*q_\bullet = q'_\bullet$ then the associated syzygy relations are equal, that is, $\text{syzy}(q_\bullet) = \text{syzy}(q'_\bullet)$.

Conversely, suppose that $\text{syzy}(q_\bullet) = \text{syzy}(q'_\bullet)$. Since, $\text{Arf}(q_1) = \text{Arf}(q'_1)$ the non-singular quadratic spaces (V, q_1) and (V, q'_1) are isometric. Let $W(q_\bullet) = (W, q)$ and $W(q'_\bullet) = (W', q')$ be given as in Notation 7.24. The obvious generating elements form a basis since q_\bullet and q'_\bullet are non-degenerate. We denote these bases by $W = \langle v_2, \dots, v_m \rangle$ and $W' = \langle v'_2, \dots, v'_m \rangle$.

Define a linear map $\mu : W \rightarrow W' : v_i \mapsto v'_i$. Lemma 7.25 and the equality of the syzygy relations implies that $q(v_i) = q'(v'_i)$ and $f(v_i, v_j) = f(v'_i, v'_j)$. Therefore, $\mu : (W, q|_W) \rightarrow (W', q'|_{W'})$ is an isometry.

Applying Witt’s Lemma above we conclude that there is an isometry $\gamma : (V, q_1) \rightarrow (V, q'_1)$ extending μ . Necessarily $\gamma \in \text{Sp}(V)$ and $\gamma_*q_1 = q'_1$. Furthermore, $\gamma_*(q_i) = \gamma_*(q_1 + v_i) = \gamma_*(q_1) + \gamma(v_i) = q'_i$. Thus $\gamma_*q_\bullet = q'_\bullet$. □

Let $q_\bullet \in Q^m$ be a possibly degenerate m -tuple. We will use the space of relations $R(q_\bullet) \subset \mathbb{F}_2^m$ from Definition 7.16 and the obvious adaptation of the reduced syzygy relations $\widetilde{\text{syzy}}(q_\bullet)$ from Definition 7.7. The following follows immediately from Proposition 7.26.

Corollary 7.27 *Two tuples $q_\bullet, q'_\bullet \in Q^m$ are in the same $\text{Sp}(V)$ -orbit if and only if $R(q_\bullet) = R(q'_\bullet)$ and $\widetilde{\text{syzy}}(q_\bullet) = \widetilde{\text{syzy}}(q'_\bullet)$.*

7.5 Prescribing syzygy relations

In this section we will give an algorithm to decide if given $\mathbf{a} = (a_i; a_{ij}) \in \mathbb{F}_2^{m+\binom{m-1}{2}}$ there is a non-degenerate sequence $q_\bullet \in \mathcal{Q}^m$ with $\text{syzy}(q_\bullet) = \mathbf{a}$.

With Notation 7.24 and Lemma 7.25 in mind we define what would be an abstract copy of $W(q_\bullet)$ if non-degenerate q_\bullet corresponding to \mathbf{a} exists.

Notation 7.28 Let $(W_\mathbf{a}, q_\mathbf{a})$ be such that $W_\mathbf{a} = \mathbb{F}_2^{m-1}$ with basis v_2, \dots, v_m and dual basis x_2, \dots, x_m . Let

$$q_\mathbf{a} = \sum_{i=2}^m e_i x_i^2 + \sum_{2 \leq i < j \leq m} e_{ij} x_i x_j, \tag{7.15}$$

where $e_i = a_i + a_1$ and $e_{ij} = a_{ij} + a_i + a_j + a_1$.

Lemma 7.29 *There exists a non-degenerate sequence q_\bullet of quadratic forms on V with syzygy relations \mathbf{a} if and only there is a quadratic form q_V on V with $\text{Arf}(q_V) = a_1$ and an isometric immersion $\iota: (W_\mathbf{a}, q_\mathbf{a}) \hookrightarrow (V, q_V)$.*

Proof If there is an immersion ι , let $q_1 = q_V$ and $q_i = q_1 + \iota(v_i)$ for $i = 2, \dots, m$. Then $q_\bullet = (q_1, \dots, q_m)$ is non-degenerate with syzygy relations \mathbf{a} . Conversely, given $q_\bullet = (q_1, \dots, q_m)$ let $q_V = q_1$ and consider the subspace $W(q_\bullet)$ of V as defined in Notation 7.24. By design $(W_\mathbf{a}, q_\mathbf{a})$ is isomorphic to $W(q_\bullet)$, giving the map ι . □

We now give an easy criterion to check for the existence of immersions of the kind required by Lemma 7.29. Let (V, q_V) and (W, q_W) be quadratic spaces with V non-singular symplectic and W possibly singular symplectic. The intersection pairing on W defines a linear map $W \rightarrow W^\vee$ with kernel K . Choose a subspace $U \subset W$ complementary to K , i.e., $W = K \oplus U$. Note that the intersection pairing on U is non-degenerate and the value of $\text{Arf}(q_W|_U)$ does not depend on our choice of U .

Notation 7.30 For $a \in \mathbb{F}_2$ let $c(q_W, a) \in \{0, 1\}$ be the *correction term* defined as follows:

$$c(q_W, a) = \begin{cases} 0 : & q_W|_K \not\equiv 0 \\ 0 : & q_W|_K \equiv 0, \text{Arf}(q_W|_U) = a \\ 1 : & q_W|_K \equiv 0, \text{Arf}(q_W|_U) \neq a \end{cases}$$

Remark 7.31 In evaluating the correction term for $K = 0$ or $U = 0$, use the convention that the quadratic form on the zero space is zero and has Arf invariant zero.

Proposition 7.32 *There is an isometric immersion $(W, q_W) \hookrightarrow (V, q_V)$ iff the following inequality is satisfied:*

$$\dim V \geq \dim W + \dim K + 2c(q_W, \text{Arf}(q_V)). \tag{7.16}$$

Proof Recall our decomposition $W = K \oplus U$. Let K' be an isomorphic copy of K and define $W' = W \oplus K'$. Extend the symplectic pairing on W to W' so that K' is orthogonal to U and is dual to K . Naturally, W' is a non-singular symplectic space and any embedding $W \hookrightarrow V$ will extend to an embedding $W' \hookrightarrow V$. This forces the inequality $\dim V \geq \dim W + \dim K$.

Pick a basis $\gamma_1, \dots, \gamma_s$ of K and its dual basis $\gamma'_1, \dots, \gamma'_s \in K'$. Any extension of the quadratic form q_W to a quadratic form $q_{W'}$ on W' requires only the values $q_{W'}(\gamma'_i)$ for

$i = 1, \dots, s$. Using the obvious hyperbolic decomposition of $K \oplus K'$ defined by our choice of bases, we note that:

$$\text{Arf}(q_{W'}) = \text{Arf}(q_W|_U) + \sum_{i=1}^s q_W(\gamma_i)q_{W'}(\gamma'_i).$$

If $q_W|_K \neq 0$ then we can choose an extension $q_{W'}$ of either Arf invariant. However, if $q_W|_K \equiv 0$ then any extension $q_{W'}$ will necessarily have $\text{Arf}(q_{W'}) = \text{Arf}(q_W|_U)$. To change the parity, we would have to join an odd plane to W' , increasing the dimension by 2 and obtaining, say $(W'', q_{W''})$. This forces the refined inequality $\dim V \geq \dim W + \dim K + 2c(q_W, \text{Arf}(q_V))$.

However, this inequality is also sufficient. Construct $(W', q_{W'})$ (or the larger $(W'', q_{W''})$ if necessary) such that $\text{Arf}(q_{W'}) = \text{Arf}(q_V)$ (or $\text{Arf}(q_{W''}) = \text{Arf}(q_V)$). Using the hyperbolic decompositions of V and W' (or W'') it is clear that we can find an embedding $W' \hookrightarrow V$ (or $W'' \hookrightarrow V$). □

Corollary 7.33 For any $\mathfrak{a} \in \mathbb{F}_2^{m+\binom{m-1}{2}}$ there exists a non-degenerate sequence $q_\bullet = (q_1, \dots, q_m)$ of quadratic forms on V with syzygy relations \mathfrak{a} if

$$g \geq m, \quad \text{where } g = \frac{\dim V}{2}. \tag{7.17}$$

Proof Observe that $m - 1 = \dim W_\mathfrak{a} \geq \dim K$ and $1 \geq c(q_\mathfrak{a}, \bullet)$ so that the right hand side of (7.16) is at most $2m$. Since $\dim V \geq 2m$, Proposition 7.32 implies the existence of an isometric immersion $(W_\mathfrak{a}, q_\mathfrak{a}) \hookrightarrow (V, q_V)$. We conclude by Lemma 7.29. □

Remark 7.34 Note that the inequality (7.17) is optimal. Indeed, if $g < m$ then the tuple \mathfrak{a} consisting entirely of 1s does not constitute a syzygy relation. In this case, the intersection product on $W_\mathfrak{a}$ is identically zero. Therefore, $K = W, U = 0$, and $\text{Arf}(q_\mathfrak{a}|_U) = 0 \neq 1 = \text{Arf}(q_V)$. The inequality (7.16) is then violated.

In general, we can decide if \mathfrak{a} is the syzygy relation of a non-degenerate m -tuple of quadratic forms on a $2g$ -dimensional non-singular symplectic space as follows. The tuple \mathfrak{a} allows us to construct $q_\mathfrak{a}$ as in (7.15). The bilinear pairing on $W_\mathfrak{a} = \mathbb{F}_2^{m-1}$ induced by $q_\mathfrak{a}$ is given by the coefficients e_{ij} of $x_i x_j$ ($i \neq j$). In the context of Lemma 7.29 we can thus evaluate the correction term $c(q_\mathfrak{a}, a_1)$ to check if the inequality in Proposition 7.32 holds. In this setting, this inequality specializes to

$$2g \geq (m - 1) + \dim K + c(q_\mathfrak{a}, a_1). \tag{7.18}$$

The inequality holds if and only if the syzygy relations are realized. This is Algorithm 7.14.

Acknowledgements It is my pleasure to thank my adviser Gavril Farkas for generously sharing his insight into research as well as giving me financial and academic support during the course of my PhD. I would like to thank my co-adviser Gerard van der Geer for numerous discussions during my stay in Amsterdam. In addition, thanks to Lenny Taelman and David Holmes for providing helpful suggestions at key moments. Special thanks go to Fabio Tonini for helping me with stacks and to Klaus Altmann for helping me extend my scholarship. Finally, I thank Özde Bayer Sertöz for the help with the picture. This research constitutes a chapter in my PhD thesis. My PhD was funded by the Berlin Mathematical School and Graduiertenkolleg 1800 of the Deutsche Forschungsgemeinschaft. I am grateful to the referee for their careful reading and insightful comments.

Funding Open Access funding enabled and organized by Projekt DEAL.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Appendix A. The universal deformation of a rooted node

In this section we describe the universal deformation of a node together with a root of a line bundle, that is, of a *rooted node*. This amounts to bringing together the results available in literature, specifically [15, 26]. Faltings [15] studies, essentially, square roots of vector bundles whereas Jarvis [26] studies r -th roots of line bundles. We are interested in the intersection of the two, the square roots of line bundles. Because of its simplicity, a treatment of this special case is quite revealing.

A.1. Conventions

In this section we are solely concerned with local, or infinitesimal, deformation functors. Therefore, at times, we will simply refer to these as deformation functors. Infinitesimal deformations of an affine scheme are always affine [36] and therefore we will work in the dual category of algebras instead of schemes.

For the rest of this section k is an algebraically closed field of characteristic $\neq 2$ and Λ is a complete noetherian local ring with residue field k . The category of Artinian local Λ -algebras with residue field k is denoted by Art_Λ . We denote by $\hat{\text{Art}}_\Lambda$ the category of complete noetherian local Λ -algebras (R, \mathfrak{m}_R) such that for each $n \geq 1$ the quotient R/\mathfrak{m}_R^n belongs to Art_Λ . Every $R \in \hat{\text{Art}}_\Lambda$ comes equipped with a natural map $R \rightarrow k = R/\mathfrak{m}_R$.

A functor $F: \text{Art}_\Lambda \rightarrow (\text{Sets})$ is said to be *pro-represented* by $R \in \hat{\text{Art}}_\Lambda$ if F is isomorphic to the restriction of the functor $\text{hom}_{\hat{\text{Art}}_\Lambda}(R, \cdot)$ to Art_Λ , see [35] for more details.

We will use equality between two objects $X = Y$ to mean that there exists a unique isomorphism between X and Y .

A.2. Deformations of a node

Definition A.1 The k -algebra $\bar{A} := k[[x, y]]/(xy)$ is called *the standard node*. A tuple $(A/R, \iota: A \rightarrow \bar{A})$ where A is a complete local flat R -algebra and ι factors through an isomorphism $A \otimes_R k \rightarrow \bar{A}$ is called a *deformation of the node over R* . Two deformations $(A/R, \iota)$ and $(A'/R, \iota')$ are isomorphic if there is an R -isomorphism of A and A' commuting with the maps ι and ι' .

Definition A.2 The *functor of infinitesimal deformations of the node* is the functor $G: \text{Art}_\Lambda \rightarrow (\text{Sets})$ which maps R to the set of isomorphism classes of deformations of the node over R .

The following theorem is folklore. See [39, Tag 0CBX] for the idea of proof.

Theorem A.3 *The deformation $(\Lambda[[x, y, t]]/(xy - t)/\Lambda[[t]], j: t \mapsto 0)$ is universal. That is, $\Lambda[[\tau]]$ pro-represents G , which for any deformation $(A/R, \iota) \in G(R)$ defines a unique map $\Lambda[[t]] \rightarrow R$ giving rise to an isomorphism $A \simeq R[[x, y]]/(xy - \pi)$.*

A.3. Deformations of a root

Definition A.4 Let $(A/R, \iota)$ be a deformation of the node and let E be an R -flat and R -relatively torsion-free rank-1 A -module. Suppose $b: E^{\otimes 2} \rightarrow A$ is a non-degenerate bilinear form (Definition 3.3). Then the pair (E, b) will be called a *root*. If E is a free A -module, and hence $E \simeq A$, then (E, b) is called a *free root*. An isomorphism between two roots (E, b) and (E', b') is an isomorphism $\mu: E \rightarrow E'$ such that $b' \circ \mu^2 = b$.

Notation A.5 For a root (E, b) on $(A/R, \iota)$ we will write ι_*E for $E \otimes_A \bar{A}$ and ι_*b for the map $\iota_*E^2 \rightarrow \bar{A}$ induced from b .

Definition A.6 Let (\bar{E}, \bar{b}) be a root on the standard node and let (E, b) be a root on a deformation of the node $(A/R, \iota)$. An isomorphism $j: \iota_*E \xrightarrow{\sim} \bar{E}$ such that $\bar{b} \circ j^2 = \iota_*b$ is called a *restriction map*. The tuple (E, b, j) is a *deformation of the root* (\bar{E}, \bar{b}) . An isomorphism of deformations is an isomorphism of roots commuting with the restriction maps.

Definition A.7 The standard node together with a root $(\bar{A}/k, \bar{E}, \bar{s})$ will be called a *rooted node*. A deformation of the standard node together with a deformation of the root (\bar{E}, \bar{s}) is a tuple $(A/R, \iota, (E, b, j))$, which will be called a *deformation of the rooted node over R* . Isomorphisms are defined in the obvious way.

Lemma A.8 *If (\bar{E}, \bar{s}) is a free root on \bar{A}/k , then for any $R \in \hat{\text{Art}}_\Lambda$ there exists a unique deformation of $(\bar{A}/k, \bar{E}, \bar{s})$ over R , up to unique isomorphisms.*

Proof Since A is complete with respect to $\mathfrak{m}_R \cdot A$ we just have to show that there exists a unique lift of a free root from R/\mathfrak{m}_R^n to R/\mathfrak{m}_R^{n+1} . Existence of the lift is clear. In order to conclude that there exists a *unique* isomorphisms between any two lifts, we observe that any two lifts of square roots of an invertible element in R/\mathfrak{m}_R^n are equal. \square

Remark A.9 The proof does not require the presence of a node and the argument works just as well around a smooth point of a curve.

Definition A.10 Given a root (\bar{E}, \bar{s}) on \bar{A}/k let $F: \text{Art}_\Lambda \rightarrow (\text{Sets})$ be the functor of isomorphism classes of deformations of the rooted node $(\bar{A}/k, \bar{E}, \bar{s})$.

Recall that $G: \text{Art}_\Lambda \rightarrow (\text{Sets})$ is the functor of infinitesimal deformations of the node \bar{A}/k . There is a natural transformation $F \rightarrow G$ obtained by forgetting the root. Lemma A.8 implies that if (\bar{E}, \bar{s}) is a free root then the forgetful functor $F \rightarrow G$ is an isomorphism. Complementing Theorem A.3 we have the following result.

Theorem A.11 *If (\bar{E}, \bar{s}) is not free then F is pro-represented by $\Lambda[[\tau]]$. The natural transformation $F \rightarrow G$ obtained by forgetting the root corresponds to the map $\Lambda[[t]] \rightarrow \Lambda[[\tau]]: t \mapsto \tau^2$.*

We will prove this theorem by constructing a universal family over $\Lambda[[\tau]]$, see Theorem A.26. To do this we need Faltings’ classification of torsion-free modules.

A.4. Classification of roots

Fix $R \in \hat{\text{Art}}_\Lambda$ and $A = R[[x, y]]/(xy - \pi)$ for some $\pi \in \mathfrak{m}_R$. Define $\iota: A \rightarrow \bar{A} = k[[x, y]]/(xy)$ using $R \rightarrow R/\mathfrak{m}_R = k$. For the rest of this section, we will be interested in non-free roots.

A.4.1. Faltings' construction

Let $p, q \in R$ be such that $pq = \pi$. Define 2×2 matrices with entries in A :

$$\alpha = \begin{pmatrix} x & p \\ q & y \end{pmatrix}, \quad \beta = \begin{pmatrix} y & -p \\ -q & x \end{pmatrix}$$

To avoid confusion, we may write $\alpha(p, q)$ and $\beta(p, q)$ instead. Clearly $\alpha\beta = \beta\alpha = 0$ but, moreover, we get an exact infinite periodic complex (see [15]):

$$\dots \rightarrow A^{\oplus 2} \xrightarrow{\alpha} A^{\oplus 2} \xrightarrow{\beta} A^{\oplus 2} \xrightarrow{\alpha} A^{\oplus 2} \xrightarrow{\beta} A^{\oplus 2} \rightarrow \dots$$

Definition A.12 Define $E(p, q) \subset A^{\oplus 2}$ to be the image of α or, equivalently, the kernel of β . Truncating the complex above we get a free resolution of $E(p, q)$. Whenever we refer to the standard resolution of $E(p, q)$ this is the one we mean.

Remark A.13 The module $E(p, q)$ is relatively torsion-free and R -flat as shown in Construction 3.2 of [15].

Remark A.14 If either p or q is invertible, then $E(p, q)$ is free. As we are interested in non-free roots, from now on we assume $p, q \in \mathfrak{m}_R$. This forces π to be in \mathfrak{m}_R^2 .

There is a natural isomorphism between $E(p, q)^\vee = \text{hom}(E(p, q), A)$ and $E(q, p)$. In particular, when $p = q$ the module $E(p, p)$ is self-dual and we get a natural pairing $s_p: E(p, p)^2 \rightarrow A$.

Definition A.15 The natural pairing $s_p: E(p, p)^2 \rightarrow A$ will be called the *standard map*.

Definition A.16 Let us refer to $(E(p, p), s_p)$ as a *standard root on $(A/R, \iota)$* .

A.4.2. Properties of standard roots

Given any bilinear map b on $E(p, q)$ we can lift it to $\text{Sym}^2 A^{\oplus 2} \rightarrow \text{Sym}^2 E(p, q)$ to get a morphism $\tilde{b}: \text{Sym}^2 A^{\oplus 2} \rightarrow A$. Letting e_1, e_2 be the standard generators of $A^{\oplus 2}$ and e_1^2, e_1e_2, e_2^2 the corresponding generators of $\text{Sym}^2 A^{\oplus 2}$ we may uniquely identify b with the values $b_0 := \tilde{b}(e_1^2), b_1 := \tilde{b}(e_1e_2), b_2 := \tilde{b}(e_2^2)$. By abuse of notation we will write $b = (b_0, b_1, b_2)$.

Lemma A.17 For a standard root $(E(p, p), s_p)$ we have $s_p = (x, p, y)$.

Proof Let $\langle \cdot, \cdot \rangle$ denote the natural pairing $A^{\oplus 2} \times A^{\oplus 2} \rightarrow A$. The identification of $E(p, p)^\vee$ with $E(p, p)$ makes it clear that if $e, f \in E(p, p)$ and $u, v \in A^{\oplus 2}$ are such that $e = \alpha(u)$ and $f = \alpha(v)$ then we have $s_p(e, f) = \langle u, \alpha(v) \rangle = \langle \alpha(u), v \rangle$. Now, direct computation yields the result. □

Lemma A.18 Any root $(E(p, p), b)$ on A/R is isomorphic to $(E(p, p), s_p)$.

Proof Lemma 5.4.10 [26] states that $b = (ax, b_1, awy)$ where $a \in A^*$ and $w \in R^*$ such that $wp = p$. Note here that as we are working with square roots of line bundles, the hypothesis of the cited lemma is satisfied (as stated in Corollary 5.4.9 *loc.cit.*).

Let v be a square root of w and consider the isomorphism $\mu : E \rightarrow E$ which descends from multiplication by $\begin{pmatrix} v & 0 \\ 0 & v^{-1} \end{pmatrix}$ on $A^{\oplus 2}$. Clearly $\mu^*b = a(vb_0, b_1, v^{-1}b_2) = va(x, v^{-1}b_1, y)$. By scaling E we may now assume $va = 1$ and $b = (x, b_1, y)$ where we changed b_1 .

Since $\alpha(y, 0) = \alpha(0, p)$ and $\alpha(0, x) = \alpha(p, 0)$ we see that $pb_2 = yb_1$ and $pb_0 = xb_1$. Which means $y(b_1 - p) = x(b_1 - p) = 0$ (we used $wp = p$). But $\text{Ann}_A(x, y) = 0$ hence $b_1 = p$. □

Theorem A.19 (Faltings) For any root (E, b) on A , $\exists p \in \mathfrak{m}_R$ such that $(E, b) \xrightarrow{\sim} (E(p, p), s_p)$.

Proof In [15, Theorem 3.7] Faltings classifies non-degenerate quadratic forms on E , or equivalently bilinear forms $b: E^2 \rightarrow A$. The conclusion is that $(E, b) \simeq (E(p, p), b')$ for some $p \in \mathfrak{m}_R$ and b' . Apply Lemma A.18 to change b' to the standard map s_p . □

Notation A.20 If $\mu \in \text{hom}(E(p, p), E(q, q))$ lifts to a map $\begin{pmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{pmatrix} : A^{\oplus 2} \rightarrow A^{\oplus 2}$ then we will write $\mu = \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix}$.

Lemma A.21 $\text{Iso}((E(p, p), s_p), (E(q, q), s_q)) = \left\{ \begin{bmatrix} \varepsilon_1 & 0 \\ 0 & \varepsilon_2 \end{bmatrix} \mid \varepsilon_1, \varepsilon_2 \in \{\pm 1\}, q = \varepsilon_1\varepsilon_2p \right\}$

Proof We adapt the proof of Proposition 4.1.12 of [26]. An easy observation is that we can choose a lift of μ of the form

$$\begin{pmatrix} u_+(x) & v_+(x) \\ v_-(y) & u_-(y) \end{pmatrix}$$

where $u_+, v_+ \in R[[x]] \subset A$ and $u_-, v_- \in R[[y]] \subset A$. Now we simply have to calculate what it means to have $\mu^*s_p = s_q$ in terms of u_{\pm}, v_{\pm} . Using that x (resp. y) does not annihilate $R[[x]]$ (resp. $R[[y]]$) we see immediately that $v_{\pm} = 0, u_{\pm} \in \{\pm 1\}$ is forced. Then $q = u_+u_-p$. □

Notation A.22 On the standard node \bar{A}/k we will define $\bar{E} := E(0, 0)$ and $\bar{s} := s_0$.

Definition A.23 On $(A/R, \iota)$ there is a natural restriction map from $E(p, q)$ to $\bar{E} = E(0, 0)$ which is the map r completing the diagram below:

$$\begin{array}{ccc} A^{\oplus 2} & \xrightarrow{\iota} & \bar{A}^{\oplus 2} \\ \downarrow \alpha(p,q) & & \downarrow \alpha(0,0) \\ E(p, q) & \xrightarrow{r} & \bar{E} \end{array}$$

Remark A.24 Lemma A.21 implies that choosing a restriction map r rigidifies the root. That is, $\text{Aut}(E(p, p), s_p, r) = 1$. The following result takes this observation one step further.

Proposition A.25 Suppose that (E, b, j) is a deformation of (\bar{E}, \bar{s}) . Then there exists precisely one $p \in R$ such that $(E, b, j) \simeq (E(p, p), s_p, r)$. Moreover, this isomorphism is unique.

Proof Uniqueness of the isomorphism follows from Remark A.24. By Theorem A.19 we know that $(E, b) \simeq (E(p, p), s_p)$ for some $p \in R$. Picking one such isomorphism we may assume $(E, b, j) = (E(p, p), s_p, j)$ for some j . However, with our choice of identification, j is not necessarily equal to the natural restriction r .

Let $\gamma = j \circ r^{-1} : (\bar{E}, \bar{s}) \rightarrow (\bar{E}, \bar{s})$. Then γ is uniquely defined by $\gamma(0) \in \{\pm \text{id}, \pm \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}\}$. An isomorphism $\mu : (E(p, p), s_p) \xrightarrow{\sim} (E(q, q), s_q)$ commutes with j and r iff $\iota_*\mu : (\bar{E}, \bar{s}) \xrightarrow{\sim} (\bar{E}, \bar{s})$ is the inverse of γ . Having classified such μ in Lemma A.21 we know that there exists precisely one q and one μ which will restrict to γ^{-1} . □

A.5. Universal deformation of rooted node

Take a *non-free* root (\bar{E}, \bar{s}) on \bar{A}/k and let $F : \text{Art}_\Lambda \rightarrow (\text{Sets})$ be the deformation functor of $(\bar{A}/k, \bar{E}, \bar{s})$ (Definition A.10).

Theorem A.26 *The ring $\Lambda[[\tau]]$ pro-represents F via the universal family*

$$(\Lambda[[\tau]] \rightarrow \Lambda[[x, y, \tau]]/(xy - \tau^2), \tau \mapsto 0, E(\tau, \tau), s_\tau, r).$$

Proof Given any deformation of the rooted node $(A/R, \iota, E, b, j)$ we wish to show that there exists a unique map $\varphi : \Lambda[[\tau]] \rightarrow R$ such that A is canonically isomorphic to $\Lambda[[x, y, \tau]]/(xy - \tau^2) \otimes_{\Lambda[[\tau]]} R$ and $\varphi_*(E(\tau, \tau), s_\tau, r) \simeq (E, b, j)$. Furthermore, that this isomorphism is unique.

Proposition A.25 shows that there exists a *unique* $p \in R$ such that (E, b, j) is (uniquely) isomorphic to $(E(p, p), s_p, r)$, moreover this implies $A = R[[x, y]]/(xy - \pi)$ with $\pi = p^2$. Define φ by $\tau \mapsto p$. Since the maps s_p and r are natural, the pushforward of $(E(\tau, \tau), s_\tau, r)$ is (uniquely) isomorphic to $(E(p, p), s_p, r)$.

Choosing any other map $\tau \mapsto q$ would give a root that is not isomorphic to (E, b, j) . Thus we have proven the existence and uniqueness of the map φ of the desired form. □

A.6. Deformations of nodes with multiple roots

Taking a multiple root in the sense of Definition 3.12 and restricting it to the formal neighbourhood of a node in a family of curves warrants the study of the objects defined here. Fix a positive integer m , let $(A/R, \iota)$ be a deformation of the node as in Definition A.1 and let $(E_i, b_i)_{i=1}^m$ be a sequence of roots on A/R .

Notation A.27 We will reorder the roots so that we may assume there is an $m' \in \{0, \dots, m\}$ such that a root E_i is free if and only if $i > m'$.

Definition A.28 A sequence of maps $(\varphi_i : F \rightarrow E_i^2)_{i=1}^m$ which satisfy the following conditions are called a *local pre-sync data*: (1) if E_i is not free then φ_i is an isomorphism, (2) if E_i are free then all φ_i are isomorphisms, (3) for all i, j we have $b_i \circ \varphi_i = b_j \circ \varphi_j$.

Remark A.29 It will be often more convenient to identify the sequence $(\varphi_i)_{i=1}^m$ with a sequence of isomorphisms $(E_1^2 \rightarrow E_i^2)_{i=1}^{m'}$. The free roots do not need additional synchronization as the maps b_i already identify their squares with the target.

Definition A.30 Let $(\varphi_i)_{i=1}^{m'}$ be a local pre-sync data. If for every $i, j \leq m'$ the natural maps $\text{Sym}^2 \text{Sym}^2 E_i \rightarrow \text{Sym}^4 E_j$ factor through an isomorphism $\text{Sym}^4 E_i \rightarrow \text{Sym}^4 E_j$ then $(\varphi_i)_{i=1}^{m'}$ is called a *local sync data*.

Believing this can be no source of confusion, we will abuse notation and call a tuple $(E_i, b_i, \varphi_i)_{i=1}^m$ a *multiple-root on A*, where $(\varphi_i)_{i=1}^m$ is a local sync data for the roots $(E_i, b_i)_{i=1}^m$. An *isomorphism between two multiple-roots* $(E_i, b_i, \varphi_i : F \rightarrow E_i^2)_{i=1}^m$ and $(E'_i, b'_i, \varphi'_i : F' \rightarrow (E'_i)^2)_{i=1}^m$ is a sequence of isomorphisms $(h_i : E_i \rightarrow E'_i)_{i=1}^m$ for which there exists an isomorphism $f : F \rightarrow F'$ satisfying $h_i^2 \circ \varphi_i = \varphi'_i \circ f$.

Using the map $\iota : A \rightarrow \bar{A}$ we can define the push-forward $\iota_*(E_i, b_i, \varphi_i)_{i=1}^m$ to be the tuple $(\iota_*E_i, \iota_*b_i, \iota_*\varphi_i)_{i=1}^m$ with $\iota_*E_i = E_i \otimes_A \bar{A}$ and ι_*b_i as in Definition A.5 and $\iota_*\varphi_i : F \otimes_A \bar{A} \rightarrow \iota_*E_i$ the natural restriction of φ_i .

Fix a multiple-root $\bar{\xi} := (\bar{E}_i, \bar{b}_i, \bar{\varphi}_i)_{i=1}^m$ on the standard node \bar{A} . Suppose $\xi := (E_i, b_i, \varphi_i)_{i=1}^m$ is a multiple-root on A such that there exists a sequence of isomorphism $h_i : \iota_*E_i \rightarrow \bar{E}_i$ giving rise to an isomorphism between the multiple-roots $\iota_*\xi$ and $\bar{\xi}$. Then, the sequence of maps $j_i : E_i \rightarrow \iota_*E_i \xrightarrow{h_i} \bar{E}_i$ will be called a *restriction map* of ξ to $\bar{\xi}$. The tuple $(A/R, \iota, (E_i, b_i, \varphi_i, j_i)_{i=1}^m)$ will be called a *deformation of* $(\bar{A}/k, \bar{\xi})$. An isomorphism between two such deformations must commute with the restriction maps.

Definition A.31 Let $H : \text{Art}_\Lambda \rightarrow (\text{Sets})$ be the functor assigning to each R the set of isomorphism classes of deformations of $(\bar{A}/k, \bar{\xi})$.

Remark A.32 Given a sync data $(\varphi_i)_{i=1}^m$, we can replace it with a sequence of isomorphisms $(E_1^2 \rightarrow E_i^2)_{i \leq m'}$.

In order to construct a universal family for the functor H , we will reformulate the definition of local sync data by analyzing isomorphisms between squares of roots. Let $(E(p, p), s_p)$ and $(E(q, q), s_q)$ be standard roots on A/R as in Definition A.16. Suppose $\varphi : E(p, p)^2 \xrightarrow{\sim} E(q, q)^2$ is an isomorphism satisfying $s_p = s_q \circ \varphi$.

Proposition A.33 *The map $\tilde{\varphi} : \text{Sym}^2 \text{Sym}^2 E(p, p) \rightarrow \text{Sym}^4 E(q, q)$ induced by φ factors through an isomorphism $\text{Sym}^4 E(p, p) \rightarrow \text{Sym}^4 E(q, q)$ if and only if there is an isomorphism $\mu : E(p, p) \rightarrow E(q, q)$ such that $\varphi = \mu^2$.*

Proof Let $A^{\oplus 2} \rightarrow E(p, p)$ be the first step in the standard resolution (Definition A.12). Denote the images of $(1, 0)$ and $(0, 1)$ by $\xi_1, \xi_2 \in E(p, p)$ respectively. Similarly define $\zeta_1, \zeta_2 \in E(q, q)$. We need to characterize the condition that

$$\tilde{\varphi}(\xi_1^2 \xi_2^2 - (\xi_1 \xi_2)^2) = 0. \tag{A.1}$$

Identify $\text{Sym}^2 A^{\oplus 2}$ with $A^{\oplus 3}$. Using standard arguments (§A2.3 [14]) we find presentations of $E(p, p)^2$ and $E(q, q)^2$, for instance:

$$E(p, p)^2 = \text{coker} (L(p, p) : A^{\oplus 4} \rightarrow A^{\oplus 3}) \tag{A.2}$$

where

$$L(p, p) = \begin{pmatrix} y & -p & 0 & 0 \\ -p & x & y & -p \\ 0 & 0 & -p & x \end{pmatrix}. \tag{A.3}$$

We want to choose a simple lift of φ to a map $A^{\oplus 3} \rightarrow A^{\oplus 3}$. Using the relations provided by $L(p, p)$ and $L(q, q)$ we may construct a lift $A^{\oplus 3} \rightarrow A^{\oplus 3}$ such that the corresponding 3×3 matrix contains no terms involving y in the first row, x or y in the second row, x in the third

row. In fact, such a lift is unique and we will denote it by $\hat{\varphi}$. It is easily seen that φ commutes with s_p and s_q iff the lift can be expressed as

$$\hat{\varphi} = \begin{pmatrix} 1 & 0 & 0 \\ a_1 & u & a_2 \\ 0 & 0 & 1 \end{pmatrix}, \tag{A.4}$$

where $u \in R^\times$ is such that $p = uq$ and $a_1, a_2 \in \text{Ann}(q) = \text{Ann}(p)$.

In terms of the entries of this matrix we may now write

$$\tilde{\varphi}(\xi_1^2 \xi_2^2 - (\xi_1 \xi_2)^2) = \zeta_1^2 \zeta_2^2 + a_1 \zeta_2^2 (\zeta_1 \zeta_2) + a_2 \zeta_1^2 (\zeta_1 \zeta_2) + (a_1 a_2 - u^2) (\zeta_1 \zeta_2)^2.$$

As before, we can calculate a presentation of $E(q, q)^4$. This presentation looks similar to Eq. A.3 but with an extra block. It follows that the equality (A.1) is satisfied iff $a_1 = a_2 = 0$ and $u^2 = 1$. Since R is a complete local ring of characteristic not two, we have $u^2 = 1$ iff $u = \pm 1$.

Supposing (A.1) holds, we conclude $p = uq = \pm q$. By Lemma A.21 we can find isomorphisms μ between the two roots. Moreover, μ admits a representation $\begin{bmatrix} \varepsilon_1 & 0 \\ 0 & \varepsilon_2 \end{bmatrix}$ where $\varepsilon_1, \varepsilon_2 \in \{-1, 1\}$ and $\varepsilon_1 \varepsilon_2 = u$. It is now clear that $\varphi = \mu^2$. Conversely, if $\varphi = \mu^2$ then the representation $\hat{\varphi}$ is of the required form and (A.1) holds. \square

Definition A.34 Let φ be the square of an isomorphism $\mu: (E(p, p), s_p) \rightarrow (E(q, q), s_q)$. Then the entry $u \in \{\pm 1\}$ appearing in (A.4) will be called the *parity* of φ . If $p = q$ and $\mu = \pm \text{id}$ then $u = 1$ otherwise $u = -1$.

Remark A.35 Observe that the parity of φ can be determined by restricting φ to the central fiber \bar{A}/k . In particular, the root with multiple roots $(\bar{A}/k, \bar{\xi})$ completely determines the local sync data for any of its deformations.

Notation A.36 For our fixed root $\bar{\xi}$ on \bar{A}/k let us identify the non-free roots with the standard root $(\bar{E}, \bar{s}) = (E(0, 0), s_0)$. We can and will choose this identification so that the parity of each of the isomorphisms $(\bar{E}_1^2 \rightarrow \bar{E}_i^2)_{i=1}^{m'}$ induced by the local sync data $(\bar{\varphi}_i)_{i=1}^{m'}$ is 1.

The following lemma shows how we can put a deformations of $\bar{\xi}$ into standard form. In light of the Remark A.35 above, we will omit the local sync data of these deformations as they are already determined by the local sync data of $\bar{\xi}$.

Lemma A.37 Let $(A/R, \iota, (E_i, b_i, \varphi_i, j_i)_{i=1}^{m'})$ be a deformation of $(\bar{A}/k, \bar{\xi})$. Then $\exists! p \in R$ such that for all $i \leq m'$ we have $(E_i, b_i, j_i) = (E(p, p), s_p, r)$ with r the natural restriction map from Definition A.23.

Proof By Proposition A.25 we know that $\forall i \leq m' \exists! p_i \in R$ such that $(E_i, b_i, j_i) = (E(p_i, p_i), s, r)$. Let $p = p_1$. From Proposition A.33 we observe that all roots (E_i, b_i) must be isomorphic and by Lemma A.21 we conclude $p_i \in \{\pm p\}$.

If $p = 0$ then there is nothing more to prove so assume $p \neq 0$. Then the sign of p_i is completely determined by the local sync data and in particular by the sign of u in (A.4). We choose the identification of the roots on the central fiber so that this sign is always 1. \square

Since free roots have no non-trivial deformations and they face no obstructions in deforming, if all m roots in $\bar{\xi}$ are free ($m' = 0$) then the deformation functor H of $(\bar{A}/k, \bar{\xi})$ agrees with the deformation functor of \bar{A}/k alone (Lemma A.8). Therefore, we will now assume $m' > 0$.

Theorem A.38 *The functor of infinitesimal deformations H of $(\bar{A}/k, \bar{\xi})$ is pro-represented by $\Lambda[[\tau]]$ with the universal deformation given by $(\Lambda[[\tau]] \rightarrow \Lambda[[\tau, x, y]]/(xy - \tau^2), \tau \mapsto 0)$ with the first m' roots equal to $(E(\tau, \tau), s_\tau, r)$ and the remaining roots free. The local sync data are trivially obtained from those of $\bar{\xi}$, which by our Notation A.36 correspond to the squares of the identity maps.*

Proof Let $(A/R, \iota)$ be a deformation of the node and let $(\mathcal{R}, j) = (E_i, b_i, j_i)_{i=1}^{m_1}$ together with the local sync data $(\varphi_i)_{i=1}^{m_1}$ be a deformation of $(\bar{A}/k, \bar{\xi})$. Any map $\nu: \Lambda[[\tau]] \rightarrow R$ is uniquely defined by the choice of $p \in R$ for which $\tau \mapsto p$. Lemma A.37 tells us that there is a unique $p \in R$ for which $\nu_*(E(\tau, \tau), s_\tau, r) = (E_i, b_i, j_i)$. This proves the existence and uniqueness of ν provided we show that the synchronizations agree. This is done by reducing to the central fiber, where the compatibility of the synchronizations is immediate. \square

A.7. Formal neighbourhoods in general

Let k be any field. Then by Cohen structure theorem there exists a *universal coefficient ring*, which we will denote by \mathfrak{o}_k , so that any complete local ring with residue field k contains a copy of \mathfrak{o}_k . If $\text{char } k = 0$ then $\mathfrak{o}_k = k$.

Let S be scheme over k and let $s: \text{Spec } k \rightarrow S$ be k -valued point. The usual complete local ring $\hat{\mathcal{O}}_{S,s}$ pro-represents the functor $\text{Art}_{\mathfrak{o}_k} \rightarrow (\text{Sets})$ defined by $A \mapsto \text{hom}_s(A, S)$, where the subscript s indicates that the morphisms must restrict to s on the residue field. This allows one to generalize the definition of a complete local ring to more general situations, for instance to geometric points of Deligne–Mumford stacks.

If k' is any field and $s': \text{Spec } k' \rightarrow S$ is a k' -valued point of S then we can still define the functor $\mathcal{Q}_{s'}: \text{Art}_{\mathfrak{o}_{k'}} \rightarrow (\text{Sets})$ via the rule $A \mapsto \text{hom}_{s'}(A, S)$. If s' factors through a k -valued point $s: \text{Spec } k \rightarrow S$ then $\mathcal{Q}_{s'}$ is pro-represented by the complete local ring $\hat{\mathcal{O}}_{S,s} \otimes_{\mathfrak{o}_k} \mathfrak{o}_{k'}$.

For a Deligne–Mumford stack \mathcal{M} and a point $p: \text{Spec } k \rightarrow \mathcal{M}$ the functor $\mathcal{Q}_p: \text{Art}_{\mathfrak{o}_k} \rightarrow (\text{Sets})$ can be defined as above. This functor is seen to be pro-representable by using any étale chart.

Definition A.39 Let \mathcal{M} be a Deligne–Mumford stack and let $p: \text{Spec } k \rightarrow \mathcal{M}$ be any k -valued point. The complete local ring pro-representing \mathcal{Q}_p will be denoted by $\hat{\mathcal{O}}_{\mathcal{M},p}$. The functor \mathcal{Q}_p is called the *local deformation functor* of p and $\text{Spec } \mathcal{O}_{\mathcal{M},p}$ is called the *formal neighbourhood* of p .

References

1. Arbarello, E., Cornalba, M., Griffiths, P.A., Harris, J.: *Geometry of Algebraic Curves. Volume I* (Grundlehren der Mathematischen Wissenschaften), vol. 267, p. xvi+386. Springer, New York (1985)
2. Arf, C.: Untersuchungen über quadratische Formen in Körpern der Charakteristik 2. I. J. Reine Angew. Math. **183**, 148–167 (1941)
3. Artin, E.: *Geometric Algebra* (Wiley Classics Library), p. x+214. Wiley, New York (1988). Reprint of the 1957 original, A Wiley-Interscience Publication
4. Bruin, N., Sertöz, E.C.: Prym varieties of genus four curves. Trans. Am. Math. Soc. **373**(1), 149–183 (2020)
5. Caporaso, L., Casagrande, C., Cornalba, M.: Moduli of roots of line bundles on curves. Trans. Am. Math. Soc. **359**(8), 3733–3768 (2007)
6. Chiodo, A., Farkas, G.: Singularities of the moduli space of level curves. J. Eur. Math. Soc. (JEMS) **19**(3), 603–658 (2017)

7. Coble, A.B.: Algebraic Geometry and Theta Functions (American Mathematical Society Colloquium Publications), vol. 10, p. vii+282. American Mathematical Society, Providence (1982). Reprint of the 1929 edition
8. Cornalba, M.: Moduli of curves and theta-characteristics. In: Lectures on Riemann Surfaces (Trieste, 1987), pp. 560–589. World Scientific Publishing, Teaneck (1989)
9. Caporaso, L., Sernesi, E.: Characterizing curves by their odd theta-characteristics. *J. Reine Angew. Math.* **562**, 101–135 (2003)
10. Deligne, P., Mumford, D.: The irreducibility of the space of curves of given genus. *Inst. Hautes Études Sci. Publ. Math.* **36**, 75–109 (1969)
11. Dolgachev, I.V.: Classical Algebraic Geometry, A Modern View, p. xii+639. Cambridge University Press, Cambridge (2012)
12. Del Centina, A., Recillas, S.: Some projective geometry associated with unramified double covers of curves of genus 4. *Ann. Mat. Pura Appl.* **4**(133), 125–140 (1983)
13. Earle, C.J., Fowler, R.S.: Holomorphic families of open Riemann surfaces. *Math. Ann.* **270**(2), 249–273 (1985)
14. Eisenbud, D.: Commutative Algebra, With a View Toward Algebraic Geometry (Graduate Texts in Mathematics), vol. 150, p. xvi+785. Springer, New York (1995)
15. Faltings, G.: Moduli-stacks for bundles on semistable curves. *Math. Ann.* **304**(3), 489–515 (1996)
16. Farkas, G.: The birational type of the moduli space of even spin curves. *Adv. Math.* **223**(2), 433–443 (2010)
17. Fan, H., Jarvis, T., Ruan, Y.: The Witten equation, mirror symmetry, and quantum singularity theory. *Ann. Math.* **2**(1), 1–106 (2013). (178)
18. Farkas, G., Ludwig, K.: The Kodaira dimension of the moduli space of Prym varieties. *J. Eur. Math. Soc. (JEMS)* **12**(3), 755–795 (2010)
19. Farkas, G., Verra, A.: Moduli of theta-characteristics via Nikulin surfaces. *Math. Ann.* **354**(2), 465–496 (2012)
20. Gross, B.H., Harris, J.: On some geometric constructions related to theta characteristics. In: Contributions to Automorphic Forms, Geometry, and Number Theory, pp. 279–311. Johns Hopkins University Press, Baltimore (2004)
21. Grothendieck, A.: Techniques de construction en géométrie analytique. X. Construction de l'espace de Teichmüller. *Séminaire Henri Cartan* **13**, 2 (1960)
22. Grothendieck, A.: Éléments de géométrie algébrique?: III. étude cohomologique des faisceaux cohérents, seconde partie. *Publications Mathématiques de l'IHÉS* **17**, 5–91 (1963)
23. Hall, J.: Openness of versality via coherent functors. *J. Reine Angew. Math.* **722**, 137–182 (2017)
24. Harris, J., Mumford, D.: On the Kodaira dimension of the moduli space of curves. *Invent. Math.* **67**(1), 23–88 (1982)
25. Jarvis, T.J.: Geometry of the moduli of higher spin curves. *Int. J. Math.* **11**(5), 637–663 (2000)
26. Jarvis, T.J.: Torsion-free sheaves and moduli of generalized spin curves. *Compos. Math.* **110**(3), 291–333 (1998)
27. Kleidman, P., Liebeck, M.: The Subgroup Structure of the Finite Classical Groups (London Mathematical Society Lecture Note Series), vol. 129, p. x+303. Cambridge University Press, Cambridge (1990)
28. Keel, S., Mori, S.: Quotients by groupoids. *Ann. Math. (2)* **145**(1), 193–213 (1997)
29. Knebusch, M.: Specialization of Quadratic and Symmetric Bilinear Forms (Algebra and Applications), vol. 11, p. xiv+192. Springer, London (2010). Translated from the German by Thomas Unger
30. Lehavi, D.: Effective reconstruction of generic genus 4 curves from their theta hyperplanes. *Int. Math. Res. Not. IMRN* **19**, 9472–9485 (2015)
31. Lehavi, D.: Effective reconstruction of generic genus 5 curves from their theta hyperplanes. *Exp. Math.* **25**, 1–15 (2022)
32. Ludwig, K.: On the geometry of the moduli space of spin curves. *J. Algebraic Geom.* **19**(1), 133–171 (2010)
33. Mumford, D.: Theta characteristics of an algebraic curve. *Ann. Sci. École Norm. Sup.* **4**(4), 181–192 (1971)
34. Riemann, B.: Zur theorie der abelschen funktionen für den fall $p = 3$, in *Gesammelte Mathematische Werke und wissenschaftlicher Nachlass*. Teubner V, Leipzig (1876)
35. Schlessinger, M.: Functors of Artin rings. *Trans. Am. Math. Soc.* **130**, 208–222 (1968)
36. Sernesi, E.: Deformations of Algebraic Schemes (Grundlehren der Mathematischen Wissenschaften), vol. 334, p. xii+339. Springer, Berlin (2006)
37. Sertöz, E.C.: Enumerative geometry of double spin curves. Ph.D. dissertation, Humboldt-Universität zu Berlin, Mathematisch-Naturwissenschaftliche Fakultät (2017). <https://doi.org/10.18452/18455>

38. Sipe, P.L.: Roots of the canonical bundle of the universal Teichmüller curve and certain subgroups of the mapping class group. *Math. Ann.* **260**(1), 67–92 (1982)
39. Stacks, T.: Project Authors, Stacks project. <http://stacks.math.columbia.edu> (2020)
40. Wilson, R.A.: *The Finite Simple Groups* (Graduate Texts in Mathematics), vol. 251, p. xvi+298. Springer, London (2009)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.