

ON THE SEVERI PROBLEM IN ARBITRARY CHARACTERISTIC[★]

by KARL CHRIST, XIANG HE, and ILYA TYOMKIN 

ABSTRACT

In this paper, we show that Severi varieties parameterizing irreducible reduced planar curves of a given degree and geometric genus are either empty or irreducible in any characteristic. Following Severi's original idea, this gives a new proof of the irreducibility of the moduli space of smooth projective curves of a given genus in positive characteristic. It is the first proof that involves no reduction to the characteristic zero case. As a further consequence, we generalize Zariski's theorem to positive characteristic and show that a general reduced planar curve of a given geometric genus is nodal.

CONTENTS

1. Introduction	1
2. Severi varieties	6
3. Tropical curves	9
4. Tropicalization	15
5. Degeneration via point constraints	27
6. The local geometry of Severi varieties	35
7. The proof of the Main Theorem	40
Acknowledgements	43
References	43

1. Introduction

In the current paper, we study the geometry of Severi varieties over an algebraically closed field \mathbf{K} of arbitrary characteristic. Recall that the *Severi variety* $V_{g,d}$ is defined to be the locus of degree d reduced plane curves of geometric genus g in the complete linear system $|\mathcal{O}_{\mathbf{P}^2}(d)|$. We denote by $V_{g,d}^{\text{irr}}$ the union of the irreducible components parameterizing irreducible curves. The main goal of this paper is to prove that the Severi varieties $V_{g,d}^{\text{irr}}$ are either empty or irreducible.

These varieties were introduced by Severi in 1921 in order to prove the irreducibility of the moduli space M_g of genus g compact Riemann surfaces [Sev21]. He noticed that for d large enough, there exists a natural surjective map from $V_{g,d}^{\text{irr}}$ to M_g since any Riemann surface of genus g admits an immersion as a plane curve of degree d . Therefore, the irreducibility of M_g follows once one proves that $V_{g,d}^{\text{irr}}$ itself is irreducible.

The irreducibility (or rather connectedness) of M_g was first asserted by Klein in 1882, who deduced it from the connectedness of certain spaces parameterizing branched coverings of the Riemann sphere, nowadays called *Hurwitz schemes*, see [Kle82, Hur91, EC85, Ful69]. Severi's goal was to provide an algebraic treatment of Klein's assertion.

[★] IT is partially supported by the Israel Science Foundation (grant No. 821/16). KC is supported by the Israel Science Foundation (grant No. 821/16) and by the Center for Advanced Studies at BGU. XH was supported by the ERC Consolidator Grant 770922 – BirNonArchGeom.



However, his proof of irreducibility of $V_{g,d}^{\text{irr}}$ contained a gap, and the question, of whether Severi varieties are irreducible, remained open for more than 60 years.

In 1969, Deligne and Mumford proved the irreducibility of the moduli space \mathcal{M}_g of smooth projective genus g curves in arbitrary characteristic [DM69]; see also [Ful69]. The proof uses reduction to characteristic zero and Teichmüller theory over the field of complex numbers. As a result, it is based on transcendental methods. Only in 1982 Fulton gave a purely algebraic proof of the irreducibility of \mathcal{M}_g in an appendix to the paper [HM82] of Harris and Mumford, who constructed a good compactification of Hurwitz schemes. However, Fulton's proof still proceeds by a reduction to the characteristic zero case.

Although the question of the irreducibility of \mathcal{M}_g was settled, the Severi problem remained open and attracted quite a lot of attention. In 1982, Zariski gave a dimension-theoretic characterization of Severi varieties in characteristic zero [Zar82]. Namely, he proved that if $V \subseteq |\mathcal{O}_{\mathbf{P}^2}(d)|$ parameterizes reduced curves of geometric genus g , then the dimension of V is bounded by $3d + g - 1$. Moreover, if equality holds, then a general curve parameterized by V is necessarily nodal. An alternative proof of this result can be given using the techniques developed by Arbarello and Cornalba; [AC81a, AC81b]. Both approaches, however, fail in positive characteristic.

Soon after, Harris settled the Severi problem in characteristic zero [Har86]. Harris' proof of irreducibility follows the general line of Severi's approach. He first proves that any component of $V_{g,d}^{\text{irr}}$ contains a rational nodal curve and then shows that such a component is unique. The second step is relatively simple and uses a monodromy argument. However, the first step is rather involved: it uses Zariski's theorem, a careful analysis of degenerations of curves, and a remarkable trick based on studying the deformation spaces of tacnodes.

In positive characteristic, much less was known about the Severi problem and Zariski's theorem. In 2013, the third author proved that the bound in Zariski's theorem holds true in arbitrary characteristic. However, the assertion about the nodality of a general curve of genus g was shown to be wrong, at least on certain weighted projective planes; see [Tyo13]. In the planar case, however, it remained an intriguing open question, which we also settle in the current paper. Notice that in characteristic zero, both parts of Zariski's theorem generalize naturally to many rational surfaces, including weighted projective planes, see, e.g., [KS13].

1.1. The main result. — Let $d \geq 1$ and $1 - d \leq g \leq \binom{d-1}{2}$ be integers. We first prove that $V_{g,d}$ is locally closed (Lemma 2.6) of pure dimension $3d + g - 1$ (Proposition 2.7). The main result is then the following theorem, which settles the Severi problem in arbitrary characteristic.

Main Theorem. — Let $d \geq 1$ and $0 \leq g \leq \binom{d-1}{2}$ be integers. Then the Severi variety $V_{g,d}^{\text{irr}}$ is irreducible of dimension $3d + g - 1$ over any algebraically closed ground field.

As an intermediate step, we establish Zariski's theorem in arbitrary characteristic (Corollary 5.3):

Theorem (Zariski's Theorem). — *Let $V \subset V_{g,d}$ be an irreducible subvariety. Then,*

- (1) $\dim(V) \leq 3d + g - 1$, and
- (2) *if $\dim(V) = 3d + g - 1$, then for a general $[C] \in V$, the curve C is nodal.*

Note that unlike in the proof of Harris in characteristic zero, we do not use Zariski's theorem to establish the degeneration part of the argument but rather obtain Zariski's theorem as a consequence of our Degeneration Theorem (Theorem 5.1).

As an immediate corollary of the main result, we obtain the first proof of the irreducibility of the moduli spaces \mathcal{M}_g in arbitrary characteristic that involves *no* reduction to characteristic zero. Indeed, since there is a dense open subset U of $V_{g,d}^{\text{irr}}$ that parameterizes nodal curves, the universal family of curves over U is equinormalizable and hence induces a map $U \rightarrow \mathcal{M}_g$. However, since for any d large enough, any smooth genus g curve admits a birational immersion as a planar degree d curve, it follows that the map $U \rightarrow \mathcal{M}_g$ is dominant. Thus, the irreducibility of \mathcal{M}_g follows from the Main Theorem in arbitrary characteristic.

1.2. The techniques and the idea of the proof. — Our proof of the Main Theorem also follows the general strategy of Severi's original approach. However, the tools we use are completely different and combine classical algebraic geometry and tropical geometry.

As a first step, we prove the Degeneration Theorem asserting that the closure $\bar{V} \subseteq |\mathcal{O}_{\mathbf{P}^2}(d)|$ of any irreducible component $V \subseteq V_{g,d}$ contains $V_{1-d,d}$. We proceed by induction, i.e., we prove that \bar{V} contains an irreducible component of $V_{g-1,d}$, and therefore a component of $V_{1-d,d}$. But the variety $V_{1-d,d}$ is irreducible since it parameterizes unions of d distinct lines, which allows us to conclude that \bar{V} contains the whole of $V_{1-d,d}$. This is the step where the new tools from tropical geometry are involved. To do tropical geometry, we assume that the ground field is the algebraic closure of a complete discretely valued field, which is harmless because the irreducibility property is stable under field extensions, and any field is a subfield of the algebraic closure of the field of Laurent power series with coefficients in the given field.

To prove that \bar{V} contains an irreducible component of $V_{g-1,d}$, we consider the intersection Z of V with the space of curves passing through $3d + g - 2$ points in general position, chosen such that their tropicalizations are *vertically stretched*. By carefully analyzing the tropicalization of the corresponding one-parameter family of curves of genus g , we show that there exists $[C] \in \bar{Z}$ such that C is a reduced curve of geometric genus $g - 1$. Since such a C passes through $3d + g - 2$ points in general position, it follows that the intersection $\bar{V} \cap V_{g-1,d}$ has the same dimension as $V_{g-1,d}$ itself and therefore contains one of its irreducible components.

Our degeneration argument is based on the study of the tropicalization of general one-parameter families of algebraic curves. For such a family, we investigate the

induced map $\alpha: \Lambda \rightarrow M_{g,n,\nabla}^{\text{trop}}$ from the tropicalization of the base to the moduli space of parameterized tropical curves. We prove that α is a piecewise integral affine map and in good cases, it satisfies the *harmonicity* and *local combinatorial surjectivity* properties, see Definition 3.4 and Theorem 4.6. These properties allow us to describe the intersection of the image of α with two types of strata of $M_{g,n,\nabla}^{\text{trop}}$: the *nice strata* – parameterizing weightless tropical curves whose underlying graph is 3-valent, and the *simple walls* – parameterizing such curves but with a unique 4-valent vertex. We show that for a nice stratum, the image of α is either disjoint from it or intersects it along a straight interval whose boundary does not belong to the stratum. For a simple wall, we show that the image of α is either disjoint from it or intersects *all* nice strata in the star of the wall.

Next, we use these general properties to reduce the assertion of Theorem 5.1 to a combinatorial game with floor decomposed tropical curves, the latter being a very convenient tool in tropical geometry introduced by Brugallé and Mikhalkin [BM09]. The goal of the game is to prove that the image of α necessarily intersects a nice stratum parameterizing tropical curves having a contracted edge of varying length. Indeed, if this is the case, then Λ contains a leg parameterizing tropical curves for which the length of the contracted edge is growing to infinity. Since the legs of Λ correspond to marked points of the base, it follows that there is a closed point $[C] \in \bar{Z}$, such that the generalized tropicalization of the normalization of C has an edge of infinite length, i.e., C has geometric genus less than g . With a minor extra effort, we show that C is necessarily reduced, see Step 3 in the proof of Theorem 5.1.

On the tropical side of the game, we show that for a point $q \in \Lambda$ corresponding to a tropical curve passing through $3d + g - 1$ vertically stretched points, the stratum M_Θ of $M_{g,n,\nabla}^{\text{trop}}$ containing $\alpha(q)$ is nice, and the boundary of the interval $\text{Im}(\alpha) \cap M_\Theta$ belongs to simple walls adjacent to M_Θ . We then use the *floor-elevator* structure of floor decomposed curves to describe the nice strata adjacent to these simple walls. We show that by traveling along the nice strata and by crossing the corresponding simple walls, we can find a leg of Λ such that α maps it to a nice stratum parameterizing tropical curves having a contracted edge of varying length, see Lemma 5.5, and the figures within its proof.

To complete the proof of the Main Theorem we follow the ideas of [Ty07], and consider the so-called *decorated Severi varieties* $U_{d,\delta}$ that parameterize reduced curves with $\delta := \binom{d-1}{2} - g$ marked nodes. We prove that these varieties are smooth and equidimensional and that the natural map $\bar{U}_{d,\delta} \rightarrow |\mathcal{O}_{\mathbf{P}^2}(d)|$ is generically finite and has $\bar{V}_{g,d}$ as its image. Finally, we show that the variety $U_{d,\delta}^{\text{irr}}$ parameterizing irreducible reduced curves of degree d with δ marked nodes is irreducible (Theorem 7.1), which implies the Main Theorem.

The analysis of the decorated Severi varieties allows us to generalize the well-known description of the local geometry of the closure of the Severi variety $\bar{V}_{g,d}$ along $V_{1-d,d}$ to the case of arbitrary characteristic (Theorem 6.1). We show that as in the case of characteristic zero, at a general $[C_0] \in V_{1-d,d}$, the variety $\bar{V}_{g,d}$ consists of smooth

branches parameterized by the subsets μ of δ nodes of C_0 . Furthermore, the intersections of branches are reduced, have expected dimension, and the branch corresponding to μ belongs to $\overline{V}_{g,d}^{\text{irr}}$ if and only if $C_0 \setminus \mu$ is connected.

1.3. Generalizations and applications. — The ideas and the techniques introduced in this paper admit a variety of generalizations and applications, some of which we investigate in our follow-up papers. In particular, our Key Lemma (Lemma 5.5) is a new tool providing good control over the skeletons of curves and the way they vary as curves move in tropically general one-parameter families. Therefore, we expect that further applications of our tropical machinery to studying families of curves will be found soon. Let us describe two generalizations we already explored.

A natural generalization considered in [CHT22a] is the case of polarized toric surfaces corresponding to *h-transverse polygons*, e.g., Hirzebruch surfaces. These are the types of polygons for which the technique of floor-decomposed curves applies. It turns out that if a polygon Δ is *h-transverse*, then any component of the Severi variety $V_{g,\Delta}^{\text{irr}}$ contains $V_{0,\Delta}^{\text{irr}}$ in its closure. However, unlike in the planar case, to prove this in full generality, one must deal with walls corresponding to elliptic components. This brings the geometry of the moduli spaces of elliptic curves with a level structure into play and restricts the validity of the results for the general *h-transverse polygons* to characteristic zero (or large enough characteristic). In some cases, such as Hirzebruch surfaces, one can avoid elliptic walls and prove the degeneration result in arbitrary characteristic. The degeneration result combined with the recent monodromy result of Lang [Lan20] gives rise to the irreducibility of the Severi varieties on most such surfaces. Notice, however, that there are examples of reducible Severi varieties even on polarized toric surfaces corresponding to certain *h-transverse polygons*; see [Tyo14] for a simple example.

In [CHT22b], we develop further the tropical tools introduced in the current paper, and deduce the irreducibility of Hurwitz schemes in small positive characteristic. Remarkably, our tools allow to replace the monodromy argument with a purely tropical one. To the best of our knowledge, this problem has been open since Hurwitz schemes were introduced by Fulton in 1969 [Ful69].

1.4. The structure of the paper. — The first two sections are preliminaries from algebraic and tropical geometry. In Section 2, we define the Severi varieties and prove that they are equidimensional of the expected dimension. In Section 3, we recall the notion of parameterized tropical curves, their moduli spaces, and floor decomposed curves.

In Section 4, we study the tropicalization of a general one-parameter family of smooth curves with a map to the projective plane. The results of this section are the main technical tool in the proof of the Degeneration Theorem, which is proved in Section 5 together with the generalization of Zariski's theorem to the case of positive characteristic (Corollary 5.3).

The main result of Section 6 is Theorem 6.1 describing the local geometry of $\overline{V}_{g,d}$ along $V_{1-d,d}$. We also define and study the decorated Severi varieties in this section. Finally, in Section 7, we prove the irreducibility of decorated Severi varieties and deduce the Main Theorem.

2. Severi varieties

In this section, we recall the notion of Severi varieties on algebraic surfaces and prove their basic properties. We show that in the case of toric surfaces, Severi varieties are locally closed, and their closures are either empty or equidimensional of the expected dimension. Let S be a projective surface over an algebraically closed ground field \mathbf{K} , and \mathcal{L} a line bundle on S . We write $[C] \in |\mathcal{L}|$ for the \mathbf{K} -point in the complete linear system $|\mathcal{L}|$ parameterizing a curve C .

Definition 2.1. — *The Severi varieties $V_{g,\mathcal{L}}$ and $V_{g,\mathcal{L}}^{\text{irr}}$ are defined to be the following loci in $|\mathcal{L}|$:*

$$V_{g,\mathcal{L}} := \{[C] \in |\mathcal{L}| \mid C \text{ is reduced, } C \cap S^{\text{sing}} = \emptyset, p_g(C) = g\},$$

where $p_g(C)$ denotes the geometric genus of C , and

$$V_{g,\mathcal{L}}^{\text{irr}} := \{[C] \in V_{g,\mathcal{L}} \mid C \text{ is irreducible}\}.$$

We denote by $\overline{V}_{g,\mathcal{L}}$ and $\overline{V}_{g,\mathcal{L}}^{\text{irr}}$ the closures in $|\mathcal{L}|$ of $V_{g,\mathcal{L}}$ and $V_{g,\mathcal{L}}^{\text{irr}}$, respectively. Notice that since the locus of integral curves is open in the linear system, the variety $\overline{V}_{g,\mathcal{L}}^{\text{irr}}$ is a union of irreducible components of $\overline{V}_{g,\mathcal{L}}$. If $(S, \mathcal{L}) = (\mathbf{P}^2, \mathcal{O}_{\mathbf{P}^2}(d))$, we will often use the classical notation $V_{g,d} := V_{g,\mathcal{L}}$ and $V_{g,d}^{\text{irr}} := V_{g,\mathcal{L}}^{\text{irr}}$.

Definition 2.2. — *A family of curves $\mathcal{X} \rightarrow \mathbf{B}$ is called generically equinormalizable if the normalization of the total space, $\mathcal{X}^{\nu} \rightarrow \mathbf{B}$, is a family of curves with smooth generic fiber.*

Lemma 2.3. — *[dJ96, §5.10] Let $\mathcal{X} \rightarrow \mathbf{B}$ be a projective morphism of integral schemes with reduced curves as fibers. Then there exist maps*

$$(2.1) \quad \begin{array}{ccc} \mathcal{X}' & \longrightarrow & \mathcal{X} \\ \downarrow & & \downarrow \\ \mathbf{B}' & \longrightarrow & \mathbf{B} \end{array}$$

such that $\mathbf{B}' \rightarrow \mathbf{B}$ is finite and surjective, \mathcal{X}' is an irreducible component of $(\mathcal{X} \times_{\mathbf{B}} \mathbf{B}')_{\text{red}}$ dominating \mathcal{X} and $\mathcal{X}' \rightarrow \mathbf{B}'$ is a generically equinormalizable family of reduced projective curves.

Remark 2.4. — In characteristic zero, any family $\mathcal{X} \rightarrow \mathbf{B}$ like in Lemma 2.3 is generically equinormalizable by [DPT80, Page 80] and [CHL06]. This fails in characteristic p , as the classic example below shows. Furthermore, the map $\mathbf{B}' \rightarrow \mathbf{B}$ of Lemma 2.3 may fail to be generically étale, which is one of the pitfalls in a naïve attempt to generalize the known characteristic-zero approaches to Zariski's theorem and Harris' theorem to the case of positive characteristic.

Example 2.5. — In characteristic $p > 2$, consider the family over $\mathbf{A}_{\mathbf{K}}^1 = \text{Spec}(\mathbf{K}[t])$, given in an affine chart $\text{Spec}(\mathbf{K}[x, y, t])$ by $x^p + y^2 + t = 0$. It has normal total space isomorphic to $\mathbf{A}_{\mathbf{K}}^2$. However, none of its fibers is smooth. In this case, the necessary base change in Lemma 2.3 is given by adjoining $\sqrt[p]{t}$. In particular, the base change is nowhere étale.

Lemma 2.6. — The Severi varieties $V_{g, \mathcal{L}}^{\text{irr}}$ and $V_{g, \mathcal{L}}$ are locally closed subsets of $|\mathcal{L}|$.

Proof. — Since $V_{g, \mathcal{L}}^{\text{irr}}$ is the intersection of $V_{g, \mathcal{L}}$ with the open locus of irreducible curves in $|\mathcal{L}|$, it is sufficient to prove the assertion for $V_{g, \mathcal{L}}$. Since the universal curve $\mathcal{X}_{|\mathcal{L}|} \rightarrow |\mathcal{L}|$ is flat and proper, the set $U_1 := \{[C] \in |\mathcal{L}| \mid C \text{ is reduced}\}$ is open in $|\mathcal{L}|$; see [Sta20, Tag 0C0E].

Let us first show that $V_{g, \mathcal{L}}$ is constructible. We can inductively define a finite stratification of U_1 by geometric genera as follows: Consider the restriction $\mathcal{X}_{U_1} \rightarrow U_1$ of the universal curve to U_1 . Assume first that the total space \mathcal{X}_{U_1} is irreducible. Let $\mathcal{X}'_{U_1} \rightarrow U'_1$ be a family associated to $\mathcal{X}_{U_1} \rightarrow U_1$ as in Lemma 2.3. Then there is a dense open subset V'_1 of U'_1 over which $\mathcal{X}'_{U_1} \rightarrow U'_1$ has constant geometric genus. Since the map $U'_1 \rightarrow U_1$ is finite, the image Z of the complement of V'_1 is closed and nowhere dense in U_1 . Set $V_1 := U_1 \setminus Z$. Then V_1 is open and dense in U_1 . Furthermore, $\mathcal{X}_{U_1} \rightarrow U_1$ has constant geometric genus over V_1 since so does $\mathcal{X}'_{U_1} \rightarrow U'_1$ over V'_1 , and the pullback of V_1 is a subset of V'_1 . If the total space \mathcal{X}_{U_1} is reducible, we apply the above reasoning to each irreducible component of \mathcal{X}_{U_1} separately and define V_1 to be the intersection of the dense open sets obtained in this way.

For an irreducible component U_2 of the complement $U_1 \setminus V_1$, the same procedure gives rise to a dense open subset $V_2 \subset U_2$ such that the restriction of $\mathcal{X}_{U_2} \rightarrow U_2$ to V_2 has fibers of constant geometric genus. Repeating this process for all irreducible components of the $U_i \setminus V_j$, gives a finite stratification of U_1 by the locally closed V_j 's. By construction, each $V_{g, \mathcal{L}}$ is a union of V_j 's, hence a constructible subset of $|\mathcal{L}|$ as asserted.

It remains to prove that the constructible subsets $\bigcup_{g' \leq g} V_{g', \mathcal{L}}$ and $\bigcup_{g' < g} V_{g', \mathcal{L}}$ are closed in U_1 , since then $V_{g, \mathcal{L}}$ is locally closed. Let \mathbf{R} be a discrete valuation ring, and $\phi: \text{Spec}(\mathbf{R}) \rightarrow U_1$ a map taking the generic point $\eta \in \text{Spec}(\mathbf{R})$ to $[C_\eta] \in V_{g', \mathcal{L}}$ for some $g' \leq g$. After a quasi-finite base change, we may assume that the normalization of C_η is smooth over η and admits a stable model over $\text{Spec}(\mathbf{R})$. Let $[C] \in U_1$ be the image of the closed point of $\text{Spec}(\mathbf{R})$. Then C is a reduced curve dominated by a semi-stable reduction of the normalization C_η^ν . Thus, $p_g(C) \leq p_g(C_\eta) = g'$. Therefore $\bigcup_{g' \leq g} V_{g', \mathcal{L}}$ and

$\bigcup_{g' < g} V_{g', \mathcal{L}}$ are closed in U_1 by the valuative criterion [Sta20, Tag 0ARK], and we are done. \square

In general, Severi varieties may have components of different dimensions; see [CC99]. However, in the case of toric surfaces, they are equidimensional, as the following proposition shows:

Proposition 2.7. — *If S is a toric surface, then the Severi varieties $V_{g, \mathcal{L}}^{\text{irr}}$ and $V_{g, \mathcal{L}}$ are either empty or equidimensional of dimension $-\mathcal{L} \cdot \mathbf{K}_S + g - 1$. Furthermore, a general $[C] \in V_{g, \mathcal{L}}$ corresponds to a curve C intersecting the boundary divisor transversally.*

Proof. — Since $V_{g, \mathcal{L}}^{\text{irr}}$ is a union of irreducible components of $V_{g, \mathcal{L}}$, it is sufficient to prove the assertion for $V_{g, \mathcal{L}}$. Let $V \subseteq V_{g, \mathcal{L}}$ be an irreducible component, and $[C] \in V$ a general point. Suppose that the curve C has m irreducible components C_1, \dots, C_m . Set $g_i := p_g(C_i)$ and $\mathcal{L}_i := \mathcal{O}_S(C_i)$, and consider the map $\prod_{i=1}^m V_{g_i, \mathcal{L}_i} \rightarrow \overline{V}_{g, \mathcal{L}}$. Its fiber over $[C]$ is finite, and the map is locally surjective. Thus, $\dim_{[C]}(V) = \sum_{i=1}^m \dim_{[C_i]}(V_{g_i, \mathcal{L}_i})$. However, $\dim_{[C_i]}(V_{g_i, \mathcal{L}_i}) \leq -\mathcal{L}_i \cdot \mathbf{K}_S + g_i - 1$ by [Tyol3, Theorem 1.2], and hence

$$\dim_{[C]}(V) \leq \sum_{i=1}^m (-\mathcal{L}_i \cdot \mathbf{K}_S + g_i - 1) = -\mathcal{L} \cdot \mathbf{K}_S + g - 1.$$

Furthermore, if C intersects the boundary divisor non-transversally, then the inequality is strict. Next, let us show the opposite inequality, implying both assertions of the proposition.

Let X be the normalization of C , and $f: X \rightarrow S$ the corresponding map. It defines a point $[X, f]$ in the Kontsevich moduli space $\mathcal{M}_{g,n}(S, |\mathcal{L}|)$ of stable maps to S with image in $|\mathcal{L}|$; see [Kon95] and [dJHS11, Theorem 5.7]. Since f is birational onto its image, it is automorphism-free, and hence $\mathcal{M}_{g,n}(S, |\mathcal{L}|)$ is a quasi-projective scheme in a neighborhood of $[X, f]$. Let (\mathbf{R}, \mathbf{m}) be the formal completion of the algebra of functions on $\mathcal{M}_{g,n}(S, |\mathcal{L}|)$ at $[X, f]$. By [III71, Théorème III.2.1.7], the tangent space to $\mathcal{M}_{g,n}(S, |\mathcal{L}|)$ at $[X, f]$ is given by $(\mathbf{m}/\mathbf{m}^2)^* = H^0(X, \mathcal{N}_f)$, where $\mathcal{N}_f = \text{Coker}(T_X \rightarrow f^*T_S)$ is the normal sheaf to f , and the obstruction space is a subspace in $H^1(X, \mathcal{N}_f)$.

Let us express the algebra (\mathbf{R}, \mathbf{m}) as a quotient of a ring of formal power series (A, \mathbf{n}) by an ideal $I \subset \mathbf{n}$ such that $\mathbf{n}/\mathbf{n}^2 \simeq \mathbf{m}/\mathbf{m}^2$. Then, by a standard argument in deformation theory, the ideal I is generated by at most $h^1(X, \mathcal{N}_f)$ elements, and hence the dimension of any local germ of $\mathcal{M}_{g,n}(S, |\mathcal{L}|)$ at $[X, f]$ is bounded from below by $\chi(X, \mathcal{N}_f)$; cf. [Mor79, Proposition 3]. Since $c_1(\mathcal{N}_f) = -\mathcal{L} \cdot \mathbf{K}_S + 2g - 2$, it follows from the Riemann-Roch theorem that

$$\dim_{[X, f]} \mathcal{M}_{g,n}(S, |\mathcal{L}|) \geq \chi(X, \mathcal{N}_f) = c_1(\mathcal{N}_f) + 1 - g = -\mathcal{L} \cdot \mathbf{K}_S + g - 1.$$

Consider the projection from a small neighborhood of $[X, f]$ to $|\mathcal{L}|$ given by sending a point $[X', f']$ to $[f'(X')]$. Since the fiber over $[C]$ is finite, it follows that

$$\dim_{[C]}(\mathcal{V}) \geq \dim_{[X, f]} \mathcal{M}_{g,n}(S, |\mathcal{L}|) \geq -\mathcal{L} \cdot \mathbf{K}_S + g - 1,$$

and we are done. \square

Remark 2.8. — If \mathbf{M} is a lattice, and (S, \mathcal{L}) is the polarized toric surface corresponding to a convex lattice polygon $\Delta \subset \mathbf{M}_{\mathbf{R}}$, then $-\mathcal{L} \cdot \mathbf{K}_S = |\partial \Delta \cap \mathbf{M}|$. Hence the dimension of the Severi varieties can be expressed combinatorially in terms of Δ and g . In particular, $\dim(\mathcal{V}_{g,d}) = 3d + g - 1$.

3. Tropical curves

In this section, we review the theory of (parameterized) tropical curves and recall the notions of floor decomposed curves and moduli spaces of parameterized tropical curves. We mainly follow [Mik05, GM07, BM09, Tyo12, ACP15], to which we refer for further details. We also discuss families of parameterized tropical curves, cf. [CCUW20, Ran19], and introduce the notions of harmonicity and local combinatorial surjectivity of the induced map to the moduli space. Throughout the section, \mathbf{M} and \mathbf{N} denote a pair of dual lattices.

3.1. Families and parameter spaces for tropical curves.

3.1.1. Abstract tropical curves. — The graphs we consider have half-edges, called *legs*. A *tropical curve* is a weighted metric graph $\Gamma = (\mathbf{G}, \ell)$ with ordered legs, i.e., \mathbf{G} is a (connected) graph with ordered half-edges and equipped with a *weight (or genus) function* $g: \mathbf{V}(\mathbf{G}) \rightarrow \mathbf{Z}_{\geq 0}$, and a *length function* $\ell: \mathbf{E}(\mathbf{G}) \rightarrow \mathbf{R}_{>0}$. Here $\mathbf{V}(\mathbf{G})$ and $\mathbf{E}(\mathbf{G})$ denote the sets of vertices and edges of \mathbf{G} , respectively. We denote by $\mathbf{L}(\mathbf{G})$ the set of legs of \mathbf{G} and extend ℓ to it by declaring the lengths of the legs to be infinite. Set $\bar{\mathbf{E}}(\mathbf{G}) := \mathbf{E}(\mathbf{G}) \cup \mathbf{L}(\mathbf{G})$. We will view tropical curves as polyhedral complexes by identifying the edges of \mathbf{G} with bounded closed intervals of corresponding lengths in \mathbf{R} and identifying the legs of \mathbf{G} with semi-bounded closed intervals in \mathbf{R} .

For $e \in \bar{\mathbf{E}}(\mathbf{G})$ we denote by e° the interior of e , and use \vec{e} to indicate a choice of orientation on e . If $e \in \mathbf{L}(\mathbf{G})$ is a leg, then it will *always* be oriented away from the vertex. Bounded edges will be considered with both possible orientations. For $v \in \mathbf{V}(\mathbf{G})$, we denote by $\text{Star}(v)$ the *star* of v , i.e., the collection of oriented edges and legs having v as their tail. In particular, $\text{Star}(v)$ contains two oriented edges for every loop based at v . The number of edges and legs in $\text{Star}(v)$ is called the *valency* of v and is denoted by $\text{val}(v)$. By abuse of notation, we will often identify $\text{Star}(v)$ with its geometric realization — the vertex v with germs of the oriented edges $\vec{e} \in \text{Star}(v)$ attached to it.

The *genus* of Γ is defined to be $g(\Gamma) = g(\mathbf{G}) := 1 - \chi(\mathbf{G}) + \sum_{v \in V(\mathbf{G})} g(v)$, where $\chi(\mathbf{G}) := b_0(\mathbf{G}) - b_1(\mathbf{G})$ denotes the Euler characteristic of \mathbf{G} . Finally, a tropical curve $\Gamma = (\mathbf{G}, \ell)$ is called *stable* if so is \mathbf{G} , that is, if every vertex of weight zero has valency at least three, and every vertex of weight one has valency at least one.

3.1.2. Parameterized tropical curves. — A *parameterized tropical curve* is a pair (Γ, h) , where $\Gamma = (\mathbf{G}, \ell)$ is a tropical curve, and $h: \Gamma \rightarrow \mathbf{N}_{\mathbf{R}}$ is a map such that:

- (a) for any $e \in \bar{E}(\mathbf{G})$, the restriction $h|_e$ is an integral affine function; and
- (b) (Balancing condition) for any vertex $v \in V(\mathbf{G})$, we have $\sum_{\vec{e} \in \text{Star}(v)} \frac{\partial h}{\partial \vec{e}} = 0$.

Note that the slope $\frac{\partial h}{\partial \vec{e}} \in \mathbf{N}$ is not necessarily primitive, and its integral length is the stretching factor of the affine map $h|_e$. We call it the *multiplicity* of h along \vec{e} . In particular, $\frac{\partial h}{\partial \vec{e}} = 0$ if and only if h contracts e . A parameterized tropical curve (Γ, h) is called *stable* if so is Γ . Its *combinatorial type* Θ is defined to be the weighted underlying graph \mathbf{G} with ordered legs equipped with the collection of slopes $\frac{\partial h}{\partial \vec{e}}$ for $e \in \bar{E}(\mathbf{G})$. We define the *extended degree* $\bar{\nabla}$ to be the sequence of slopes $\left(\frac{\partial h}{\partial \vec{l}} \right)_{\vec{l} \in L(\mathbf{G})}$, and the *degree* ∇ to be the subsequence $\left(\frac{\partial h}{\partial \vec{l}_i} \right)$, where \vec{l}_i 's are the legs not contracted by h . The (extended) degree clearly depends only on the combinatorial type Θ . An *isomorphism* of parameterized tropical curves $h: \Gamma \rightarrow \mathbf{N}_{\mathbf{R}}$ and $h': \Gamma' \rightarrow \mathbf{N}_{\mathbf{R}}$ is an isomorphism of metric graphs $\varphi: \Gamma \rightarrow \Gamma'$ such that $h = h' \circ \varphi$. Similarly, an *isomorphism* of combinatorial types Θ and Θ' is an isomorphism of the underlying graphs respecting the order of the legs, the weight function, and the slopes of the edges and legs. We denote the group of automorphisms of a combinatorial type Θ by $\text{Aut}(\Theta)$, and the isomorphism class of Θ by $[\Theta]$.

3.1.3. Families of tropical curves. — Next we define families of (parameterized) tropical curves; cf. [CCUW20, §3.1] and [Ran19, §2.2]. We will adapt these more general definitions to the special case we need – families of (parameterized) tropical curves over a tropical curve Λ without stacky structure. Since tropical curves are nothing but parameterized tropical curves with $\mathbf{N} = 0$, we discuss only families of parameterized tropical curves.

Let $\Lambda = (\mathbf{G}_{\Lambda}, \ell_{\Lambda})$ be a tropical curve *without loops*. Loosely speaking, a family of parameterized tropical curves over Λ is a continuous family such that for any $e \in \bar{E}(\mathbf{G}_{\Lambda})$, the combinatorial type of the fiber is constant along e° and the lengths of the edges of the fibers, as well as the images of the vertices of the fibers in $\mathbf{N}_{\mathbf{R}}$, form integral affine functions on e . To define this notion formally, consider a datum (\dagger) consisting of the following:

- an extended degree $\bar{\nabla}$;
- a combinatorial type $\Theta_e = \left(\mathbf{G}_e, \left(\frac{\partial h}{\partial \vec{\gamma}} \right) \right)$ of extended degree $\bar{\nabla}$ for each $e \in \bar{E}(\mathbf{G}_{\Lambda})$;
- an integral affine function $\ell(\gamma, \cdot): e \rightarrow \mathbf{R}_{\geq 0}$ for each $e \in \bar{E}(\mathbf{G}_{\Lambda})$ and $\gamma \in E(\mathbf{G}_e)$;
- an integral affine function $h(u, \cdot): e \rightarrow \mathbf{N}_{\mathbf{R}}$ for each $e \in \bar{E}(\mathbf{G}_{\Lambda})$ and $u \in V(\mathbf{G}_e)$;
- a parameterized tropical curve $h_w: \Gamma_w \rightarrow \mathbf{N}_{\mathbf{R}}$ of extended degree $\bar{\nabla}$ for each $w \in V(\mathbf{G}_{\Lambda})$;

- a weighted contraction $\varphi_{\vec{e}}: \mathbf{G}_e \rightarrow \mathbf{G}_w$ preserving the order of the legs for each $w \in V(\mathbf{G}_\Lambda)$ and $\vec{e} \in \text{Star}(w)$.

For any $e \in \bar{E}(\mathbf{G}_\Lambda)$ and $q \in e^\circ$, set $\Gamma_q := (\mathbf{G}_e, \ell_q)$, where $\ell_q: E(\mathbf{G}_e) \rightarrow \mathbf{R}_{\geq 0}$ is the function defined by $\ell_q(\gamma) := \ell(\gamma, q)$. Denote by $h_q: \Gamma_q \rightarrow \mathbf{N}_{\mathbf{R}}$ the unique piecewise affine map for which $h_q(u) = h(u, q)$ for all $u \in V(\mathbf{G}_e)$ and $\left(\frac{\partial h_q}{\partial t}\right) = \bar{\nabla}$.

Definition 3.1. — *Let Λ be a tropical curve without loops. We say that a datum (\dagger) is a family of parameterized tropical curves over Λ if the following compatibilities hold for all $w \in V(\mathbf{G}_\Lambda)$, $\vec{e} \in \text{Star}(w)$, and $q \in e^\circ$:*

- (1) (Γ_q, h_q) is a parameterized tropical curve of combinatorial type Θ_e ;
- (2) the length of $\varphi_{\vec{e}}(\gamma)$ in Γ_w is $\ell(\gamma, w)$ for all $\gamma \in E(\mathbf{G}_e)$;
- (3) $h_w(\varphi_{\vec{e}}(u)) = h(u, w)$ for all $u \in V(\mathbf{G}_e)$.

A family of parameterized tropical curves over Λ will be denoted by $h: \Gamma_\Lambda \rightarrow \mathbf{N}_{\mathbf{R}}$.

3.1.4. The parameter space. — We denote by $M_{g,n,\nabla}^{\text{trop}}$ the parameter space of isomorphism classes of genus g stable parameterized tropical curves of degree ∇ with exactly n contracted legs. We assume that the contracted legs are l_1, \dots, l_n . Similarly to the case of the parameter space of abstract tropical curves studied in [ACP15, Section 2], we consider $M_{g,n,\nabla}^{\text{trop}}$ as a *generalized polyhedral complex*, i.e. $M_{g,n,\nabla}^{\text{trop}}$ is glued from “orbifold” quotients of polyhedra \bar{M}_Θ factored by finite groups of automorphisms $\text{Aut}(\Theta)$. In particular, the integral affine structure on the quotient is determined by that of M_Θ , and the “new” faces of the geometric quotient $\bar{M}_\Theta/\text{Aut}(\Theta)$ are not considered to be faces of the orbifold quotient; see *loc. cit.* for more details. More explicitly, $M_{g,n,\nabla}^{\text{trop}}$ is constructed as follows.

A parameterized tropical curve of type Θ defines a point in $\mathbf{N}_{\mathbf{R}}^{|\mathbf{V}(\mathbf{G})|} \times \mathbf{R}_{>0}^{|\mathbf{E}(\mathbf{G})|}$, and the set of parameterized tropical curves of type Θ gets identified with the interior M_Θ of a convex polyhedron $\bar{M}_\Theta \subseteq \mathbf{N}_{\mathbf{R}}^{|\mathbf{V}(\mathbf{G})|} \times \mathbf{R}_{\geq 0}^{|\mathbf{E}(\mathbf{G})|}$; see, e.g., [Mik05, Proposition 2.13] or [GM07, § 3] for details. The lattice $\mathbf{N}^{|\mathbf{V}(\mathbf{G})|} \times \mathbf{Z}^{|\mathbf{E}(\mathbf{G})|} \subset \mathbf{N}_{\mathbf{R}}^{|\mathbf{V}(\mathbf{G})|} \times \mathbf{R}^{|\mathbf{E}(\mathbf{G})|}$ defines an integral affine structure on \bar{M}_Θ .

For each subset $E \subseteq E(\mathbf{G})$, the intersection of \bar{M}_Θ with the locus

$$\{(n_u, x_e) \in \mathbf{N}_{\mathbf{R}}^{|\mathbf{V}(\mathbf{G})|} \times \mathbf{R}_{\geq 0}^{|\mathbf{E}(\mathbf{G})|} \mid x_e = 0 \text{ if and only if } e \in E\}$$

is either empty or can be identified naturally with M_{Θ_E} , where Θ_E is the type of degree ∇ and genus g , obtained from Θ by the weighted edge contraction of $E \subset E(\mathbf{G})$. Plainly, the corresponding inclusion $\iota_{\Theta,E}: \bar{M}_{\Theta_E} \hookrightarrow \bar{M}_\Theta$ respects the integral affine structures.

Next, notice that an isomorphism $\alpha: \Theta_1 \rightarrow \Theta_2$ of combinatorial types induces an isomorphism $\mathbf{N}_{\mathbf{R}}^{|\mathbf{V}(\mathbf{G}_1)|} \times \mathbf{R}^{|\mathbf{E}(\mathbf{G}_1)|} \rightarrow \mathbf{N}_{\mathbf{R}}^{|\mathbf{V}(\mathbf{G}_2)|} \times \mathbf{R}^{|\mathbf{E}(\mathbf{G}_2)|}$ that takes M_{Θ_1} to M_{Θ_2} and also respects the integral affine structures. We thus obtain isomorphisms $\iota_\alpha: \bar{M}_{\Theta_1} \rightarrow \bar{M}_{\Theta_2}$. In particular, the group $\text{Aut}(\Theta)$ acts naturally on \bar{M}_Θ .

The space $\mathbf{M}_{g,n,\nabla}^{\text{trop}}$ is defined to be the colimit of the diagram, whose entries are $\overline{\mathbf{M}}_{\Theta}$'s for all combinatorial types Θ of genus g , degree ∇ curves having exactly n contracted legs l_1, \dots, l_n ; and arrows are the inclusions $\iota_{\Theta,E}$'s and the isomorphisms ι_{α} 's described above. By the construction, for each type Θ of degree ∇ and genus g with exactly n contracted legs l_1, \dots, l_n , we have a finite map $\mathbf{M}_{\Theta} \rightarrow \mathbf{M}_{g,n,\nabla}^{\text{trop}}$, whose image $\mathbf{M}_{\Theta}/\text{Aut}(\Theta)$ is denoted by $\mathbf{M}_{[\Theta]}$.

Definition 3.2. — *Let Λ be a tropical curve, and $\alpha: \Lambda \rightarrow \mathbf{M}_{g,n,\nabla}^{\text{trop}}$ a continuous map. We say that α is piecewise integral affine if for any $e \in \overline{\mathbf{E}}(\Lambda)$ the restriction $\alpha|_e$ lifts to an integral affine map $e \rightarrow \overline{\mathbf{M}}_{\Theta}$ for some combinatorial type Θ .*

Proposition 3.3. — *Let $h: \Gamma_{\Lambda} \rightarrow \mathbf{N}_{\mathbf{R}}$ be a family of parameterized tropical curves of degree ∇ . Then the induced map $\alpha: \Lambda \rightarrow \mathbf{M}_{g,n,\nabla}^{\text{trop}}$ is piecewise integral affine.*

Proof. — Follows immediately from the definitions. □

Definition 3.4. — *Let Λ be a tropical curve, $\alpha: \Lambda \rightarrow \mathbf{M}_{g,n,\nabla}^{\text{trop}}$ a piecewise integral affine map, w a vertex of Λ , and Θ a type such that $\alpha(w) \in \mathbf{M}_{[\Theta]}$.*

- (1) *Suppose $\alpha(\text{Star}(w)) \subset \mathbf{M}_{[\Theta]}$. We say that α is harmonic at w if $\alpha|_{\text{Star}(w)}$ lifts to a harmonic map $\alpha_w: \text{Star}(w) \rightarrow \mathbf{M}_{\Theta}$, i.e., α_w is piecewise integral affine, and $\sum_{\tilde{e} \in \text{Star}(w)} \frac{\partial \alpha_w}{\partial \tilde{e}} = 0$.*
- (2) *Suppose $\alpha(\text{Star}(w)) \not\subset \mathbf{M}_{[\Theta]}$. We say that α is locally combinatorially surjective at w if for any combinatorial type Θ' with an inclusion $\iota_{\Theta',E}: \overline{\mathbf{M}}_{\Theta} \hookrightarrow \overline{\mathbf{M}}_{\Theta'}$, there exists a germ \tilde{e} of $\text{Star}(w)$ such that $\alpha(\tilde{e}) \cap \mathbf{M}_{[\Theta']} \neq \emptyset$.*

3.1.5. Regularity. — A parameterized tropical curve (Γ, h) , its type Θ , and the corresponding stratum $\mathbf{M}_{[\Theta]}$ are called *regular* if \mathbf{M}_{Θ} has the expected dimension,

$$\text{expdim}(\mathbf{M}_{\Theta}) := |\nabla| + n + (\text{rank}(\mathbf{N}) - 3)\chi(\mathbf{G}) - \text{ov}(\mathbf{G}),$$

where \mathbf{G} , as usual, denotes the underlying graph, $\chi(\mathbf{G})$ its Euler characteristic, and $\text{ov}(\mathbf{G})$ the overvalency of \mathbf{G} , i.e., $\text{ov}(\mathbf{G}) := \sum_{v \in \mathbf{V}(\Gamma)} \max\{0, \text{val}(v) - 3\}$. By [Mik05, Proposition 2.20], the expected dimension provides a lower bound on the dimension of a stratum \mathbf{M}_{Θ} . Note that in *loc. cit.*, the expected dimension is stated to be $\text{expdim}(\mathbf{M}_{\Theta}) - c$, where c denotes the number of edges contracted by h . This discrepancy in the formulae for the expected dimension appears because in *loc. cit.* only deformations of the image $h(\Gamma)$ are considered. The following is Mikhalkin's [Mik05, Proposition 2.23] stated in terms of the current paper:

Proposition 3.5. — *Assume that $\text{rank}(\mathbf{N}) = 2$ and h is an immersion away from the contracted legs. If \mathbf{G} is weightless and 3-valent, then (Γ, h) is regular, and $\dim(\mathbf{M}_{\Theta}) = |\nabla| + n + g - 1$. On the other hand, if $\text{ov}(\mathbf{G}) > 0$, then $\dim(\mathbf{M}_{\Theta}) \leq |\nabla| + n - \chi(\mathbf{G}) - 1 \leq |\nabla| + n + g - 2$.*

3.1.6. *The evaluation map.* — Consider the natural evaluation map

$$\mathrm{ev}: \mathbf{M}_{g,n,\nabla}^{\mathrm{trop}} \rightarrow \mathbf{N}_{\mathbf{R}}^n,$$

defined by $\mathrm{ev}(\Gamma, h) := (h(l_1), \dots, h(l_n))$. Let Θ be a combinatorial type, $(q_1, \dots, q_n) \in \mathbf{N}_{\mathbf{R}}^n$ any tuple, and ev_{Θ} the composition $\overline{\mathbf{M}}_{\Theta} \rightarrow \mathbf{M}_{g,n,\nabla}^{\mathrm{trop}} \rightarrow \mathbf{N}_{\mathbf{R}}^n$. Notice that ev_{Θ} is the restriction of a linear projection $\mathbf{N}_{\mathbf{R}}^{|\mathbf{V}(\mathbf{G})|} \times \mathbf{R}^{|\mathbf{E}(\mathbf{G})|} \rightarrow \mathbf{N}_{\mathbf{R}}^n$. Thus, $\mathrm{ev}_{\Theta}^{-1}(q_1, \dots, q_n)$ is a polyhedron cut out in $\overline{\mathbf{M}}_{\Theta}$ by an affine subspace, and hence its boundary is disjoint from \mathbf{M}_{Θ} unless $\mathrm{ev}_{\Theta}^{-1}(q_1, \dots, q_n)$ is a point.

Definition 3.6. — We say that a tuple $(q_1, \dots, q_n) \in \mathbf{N}_{\mathbf{R}}^n$ is general with respect to ∇ and g if for any combinatorial type Θ of degree ∇ and genus g with exactly n contracted legs, the intersection $\mathrm{ev}_{\Theta}^{-1}(q_1, \dots, q_n) \cap \mathbf{M}_{\Theta}$ is either empty, or has codimension $n \cdot \mathrm{rank}(\mathbf{N})$ in \mathbf{M}_{Θ} .

In what follows, tuples of points will be called general, without mentioning ∇ and g , which we tacitly assume to be fixed. Since there are only finitely many isomorphism classes of types of fixed degree and genus, the set of non-general tuples is a nowhere dense closed subset in $\mathbf{N}_{\mathbf{R}}^n$.

Remark 3.7. — Other notions of general position can be found in the literature. For instance, Mikhalkin uses a different definition in case $\mathrm{rank}(\mathbf{N}) = 2$, [Mik05, Definitions 4.2 and 4.7]. Let $n = |\nabla| + g - 1$ and suppose that $(q_1, \dots, q_n) \in \mathbf{N}_{\mathbf{R}}^n$ are in general position in the sense of Mikhalkin. Then, by Mikhalkin's definition, any curve (Γ, h) of genus g and degree ∇ passing through the q_i 's is necessarily 3-valent, weightless, and all slopes of h are non-zero (except for the slopes of the contracted legs). In our setting, similar properties hold true as the following proposition shows.

Proposition 3.8. — Assume that $\mathrm{rank}(\mathbf{N}) = 2$, $n = |\nabla| + g - 1$, and $(q_1, \dots, q_n) \in \mathbf{N}_{\mathbf{R}}^n$ is a general tuple. Let Θ be a combinatorial type of degree ∇ and genus g with n contracted legs and let \mathbf{G} be its underlying weighted graph. If $\mathrm{ev}_{\Theta}^{-1}(q_1, \dots, q_n) \neq \emptyset$, then \mathbf{G} is 3-valent and weightless, and all slopes are non-zero, except for the slopes of the contracted legs.

Proof. — Let $\xi \in \mathrm{ev}_{\Theta}^{-1}(q_1, \dots, q_n)$ be any point. Since weightless 3-valent graphs are not contractions of other stable weighted graphs, we may assume that $\xi \in \mathbf{M}_{\Theta}$. The map ev_{Θ} is the restriction of a linear projection $\mathbf{N}_{\mathbf{R}}^{|\mathbf{V}(\mathbf{G})|} \times \mathbf{R}^{|\mathbf{E}(\mathbf{G})|} \rightarrow \mathbf{N}_{\mathbf{R}}^n$ to the polyhedral cone $\overline{\mathbf{M}}_{\Theta} \subset \mathbf{N}_{\mathbf{R}}^{|\mathbf{V}(\mathbf{G})|} \times \mathbf{R}^{|\mathbf{E}(\mathbf{G})|}$. By definition of points in general position, $\mathrm{ev}_{\Theta}^{-1}(q_1, \dots, q_n)$ has codimension $2n$ in \mathbf{M}_{Θ} . Since $\mathbf{N}_{\mathbf{R}}^n$ has dimension $2n$, it follows that ev_{Θ} is a submersion at any point of $\mathrm{ev}_{\Theta}^{-1}(q_1, \dots, q_n) \cap \mathbf{M}_{\Theta}$, hence also in a small open neighborhood of ξ . By [Mik05, Corollary 4.12], the tuples of points in tropically general position in the sense of Mikhalkin are dense, and hence there exists a small perturbation $\xi' \in \mathrm{ev}_{\Theta}^{-1}(q'_1, \dots, q'_n)$ of ξ , where (q'_1, \dots, q'_n) is a tuple in tropically general position in the sense of Mikhalkin. Thus, the assertion holds by definition (see Remark 3.7). \square

Definition 3.9. — A stratum $M_{[\Theta]} \subset M_{g,n,\nabla}^{\text{trop}}$ is called nice if it is regular, and the underlying graph \mathbf{G} is weightless and 3-valent. A stratum $M_{[\Theta]}$ is called a simple wall if it is regular, and \mathbf{G} is weightless and 3-valent except for a unique 4-valent vertex.

Lemma 3.10. — Assume that $(q_1, \dots, q_n) \in \mathbf{N}_{\mathbf{R}}^n$ is general, $\text{ev}_{\Theta}^{-1}(q_1, \dots, q_n) \neq \emptyset$, and $\text{rank}(\mathbf{N}) = 2$. Then,

- (1) If $n = |\nabla| + g - 1$, then $M_{[\Theta]}$ is nice;
- (2) If $n = |\nabla| + g - 2$ and $M_{[\Theta]}$ is a simple wall, then $\text{ev}_{\Theta}^{-1}(q_1, \dots, q_n)$ is a point in M_{Θ} ;
- (3) If $n = |\nabla| + g - 2$ and $M_{[\Theta]}$ is nice, then $\text{ev}_{\Theta}^{-1}(q_1, \dots, q_n) \subset \overline{M}_{\Theta}$ is an interval, whose boundary is disjoint from M_{Θ} .

Proof. — Assertion (1) follows from Proposition 3.8. Thus, we may assume that $n = |\nabla| + g - 2$. If $M_{[\Theta]}$ is a simple wall, then it is regular, and the graph \mathbf{G} is weightless and has overvalency one. Therefore,

$$\dim(\text{ev}_{\Theta}^{-1}(q_1, \dots, q_n)) = |\nabla| + n + g - 2 - 2n = 0,$$

and hence $\text{ev}_{\Theta}^{-1}(q_1, \dots, q_n)$ is a point in M_{Θ} , as asserted in (2). Similarly, if $M_{[\Theta]}$ is nice, then the dimension count shows that $\text{ev}_{\Theta}^{-1}(q_1, \dots, q_n)$ is an interval, which implies assertion (3). \square

3.2. Floor decomposed curves. — Next, we recall the notion of a floor decomposed curve as introduced by Brugallé and Mikhalkin [BM09]. We will follow the presentation in [BIMS15, §4.4], see also [Rau17, §2.5]. For the rest of this section, we will assume $\mathbf{M} = \mathbf{N} = \mathbf{Z}^2$ and $\Delta \subset \mathbf{M}_{\mathbf{R}} = \mathbf{R}^2$ is an h -transverse polygon (cf. [BM09, §2]), e.g., the triangle Δ_d with vertices $(0, 0)$, $(0, d)$ and $(d, 0)$. We associate to Δ the reduced degree ∇ of tropical curves dual to Δ , i.e., ∇ consists of primitive outer normals to the sides of Δ , and the number of slopes outer normal to a given side is equal to its integral length.

Definition 3.11. — Let $\lambda \in \mathbf{R}$ be a positive real number. A point configuration q_1, \dots, q_n in \mathbf{R}^2 is called vertically λ -stretched if for any pair of distinct points $q_i = (x, y)$ and $q_j = (x', y')$ in the configuration the following holds: $|y - y'| > \lambda \cdot |x - x'|$.

Definition 3.12. — A parameterized tropical curve (Γ, h) is called floor decomposed if for any edge e of Γ , the slope $\frac{\partial h}{\partial e}$ is either $(\pm 1, *)$ or $(0, *)$.

Let (Γ, h) be a floor decomposed parameterized tropical curve. If the image of a non-contracted edge (resp. leg) is vertical, then the edge (resp. leg) is called an *elevator*. After removing the interiors of all elevators in Γ , we are left with a disconnected graph. The non-contracted connected components of this graph are called the *floors* of Γ . In Figure 1, the elevators are the red edges, and the blue components are the floors.




FIG. 1. — A collection of vertically stretched points, and a floor decomposed cubic tropical curve passing through them and having the reduced tropical degree associated to Δ_3

Proposition 3.13. — [BM09, §5.1] *Let g be an integer, and set $n := |\nabla| + g - 1$. Then there exists λ such that for any generic configuration $q_1, \dots, q_n \subset \mathbf{R}^2$ of vertically λ -stretched points, every parameterized tropical curve (Γ, h) in $\text{ev}^{-1}(q_1, \dots, q_n) \subset \mathbf{M}_{g,n,\nabla}^{\text{trop}}$ is floor decomposed. Furthermore, each floor and each elevator of Γ contains exactly one of the q_i 's in its image.*

By abuse of language, we call a collection of points as in Proposition 3.13 *vertically stretched*.

Remark 3.14. — If $\Delta = \Delta_d$, then there is a unique elevator adjacent to the top floor, and, by the balancing condition, for every floor, there is a downward elevator adjacent to it. The elevator adjacent to the top floor and the downward elevators adjacent to the bottom floor – all have multiplicity one. If, in addition, the points q_1, \dots, q_n belong to the line given by $y = -\mu x$ for $\mu \gg 1$, then the x -coordinate of the marked point belonging to the image of the top floor is smaller than the x -coordinate of any other q_i .

Remark 3.15. — In the original definition of [BM09], an elevator either connects two floors or is a leg. In our convention, an elevator of [BM09] may split into a couple of elevators separated by a vertex mapped to one of the q_i 's.

4. Tropicalization

In this section, we construct the tropicalization of one-parameter families of algebraic curves with a map to a toric variety (Theorem 4.6). As a result, we obtain a tropical tool (Corollary 4.9) for studying algebraic degenerations that plays a central role in the proof of the Main Theorem.

4.1. Notation and terminology. — Throughout this section, we assume that \mathbf{K} is the algebraic closure of a complete discretely valued field. We denote the valuation by $v: \mathbf{K} \rightarrow \mathbf{R} \cup \{\infty\}$, the ring of integers by \mathbf{K}^0 , and its maximal ideal by \mathbf{K}^{00} . We fix a pair of dual lattices \mathbf{M} and \mathbf{N} and a toric variety \mathbf{S} with lattice of monomials \mathbf{M} . In this paper,

we will mostly be interested in the case $S = \mathbf{P}^2$, but the tropicalization we are going to describe works in general.

By a *family of curves*, we mean a flat, projective morphism of finite presentation and relative dimension one. By a collection of *marked points* on a family of curves, we mean a collection of disjoint sections contained in the smooth locus of the family. A family of curves with marked points is *prestable* if its fibers have at-worst-nodal singularities; cf. [Sta20, Tag 0E6T]. It is called (*semi-*)*stable* if so are its geometric fibers. A prestable curve with marked points over a field is called *split* if the irreducible components of its normalization are geometrically irreducible and smooth, and the preimages of the nodes in the normalization are defined over the ground field. A family of prestable curves with marked points is called *split* if all of its fibers are so; cf. [dJ96, § 2.22]. If $U \subset Z$ is open, and (Y, σ_\bullet) is a family of curves with marked points over U , then a *model* of (Y, σ_\bullet) is a family of curves with marked points over Z , whose restriction to U is (Y, σ_\bullet) .

4.2. Tropicalization of curves. — We first recall the canonical tropicalization construction for a fixed (parameterized) curve over \mathbf{K} . This is well established, and we refer to [BPR13, Tyo12] for details. The normalization of signs is chosen such that the algebraic definition is compatible with the standard tropical pictures. A *parameterized curve* in S is a smooth projective curve with marked points (X, σ_\bullet) and a map $f: X \rightarrow S$ such that $f(X)$ does not intersect orbits of codimension greater than one, and the image of $X \setminus (\bigcup_i \sigma_i)$ under f is contained in the dense torus $T \subset S$.

Let $f: X \rightarrow S$ be a parameterized curve, and $X^0 \rightarrow \text{Spec}(\mathbf{K}^0)$ a prestable model. Denote by \tilde{X} the fiber of X^0 over the closed point of $\text{Spec}(\mathbf{K}^0)$. As usual, a point $s \in \tilde{X}$ is called *special* if it is either a node or a marked point of \tilde{X} . Let $D = \sum_j D_j$ be the boundary divisor of S . Set $f^*(D_j) := \sum_r d_{j,r} p_{j,r}$ with $d_{j,r} \in \mathbf{N}$ and $p_{j,r} \in X$. By definition, $\{p_{j,r}\} \subseteq \{\sigma_i\}$. The collection of multisets $\{d_{j,r}\}_r$ for each j is called the *tangency profile* of $f: X \rightarrow S$. We say that the tangency profile is trivial if $d_{j,r} = 1$ for all j and r .

The *tropicalization* $\text{trop}(X)$ of X with respect to the model X^0 is the tropical curve $\Gamma = (\mathbf{G}, \ell)$ defined as follows: The underlying graph \mathbf{G} is the dual graph of the central fiber \tilde{X} , i.e., the vertices of \mathbf{G} correspond to irreducible components of \tilde{X} , the edges – to nodes, the legs – to marked points, and the natural incidence relation holds. For a vertex v of \mathbf{G} , its weight is defined to be the geometric genus of the corresponding component of the reduction \tilde{X}_v . As for the length function, if $e \in E(\mathbf{G})$ is the edge corresponding to a node $z \in \tilde{X}$, then $\ell(e)$ is defined to be the valuation of λ , where $\lambda \in \mathbf{K}^{00}$ is such that étale locally at z , the total space of X^0 is given by $xy = \lambda$. Although λ depends on the étale neighborhood, its valuation does not, and hence the length function is well-defined. Finally, notice that the order on the set of marked points induces an order on the set of legs of Γ . By abuse of notation, we will not distinguish between the tropical curve $\text{trop}(X)$ and its geometric realization.

Next, we explain how to construct the parameterization $h: \text{trop}(X) \rightarrow \mathbf{N}_{\mathbf{R}}$. Let \tilde{X}_v be an irreducible component of \tilde{X} . Then, for any $m \in M$, the pullback $f^*(x^m)$ of

the monomial x^m is a non-zero rational function on X^0 , since the preimage of the big orbit is dense in X . Thus, there is $\lambda_m \in \mathbf{K}^\times$, unique up to an element invertible in \mathbf{K}^0 , such that $\lambda_m f^*(x^m)$ is an invertible function at the generic point of \tilde{X}_v . The function $h(v)$, associating to $m \in M$ the valuation $v(\lambda_m)$, is clearly linear, and hence $h(v) \in \mathbf{N}_{\mathbf{R}}$. The parameterization $h: \text{trop}(X) \rightarrow \mathbf{N}_{\mathbf{R}}$ is defined to be the unique piecewise integral affine function with values $h(v)$ at the vertices of $\text{trop}(X)$, whose slopes along the legs satisfy the following: for any leg l_i and $m \in M$ we have $\frac{\partial h}{\partial l_i}(m) = -\text{ord}_{\sigma_i} f^*(x^m)$, where σ_i is the marked point corresponding to l_i . Then $h: \text{trop}(X) \rightarrow \mathbf{N}_{\mathbf{R}}$ is a parameterized tropical curve, by [Tyo12, Lemma 2.23]. The curve $\text{trop}(X)$ (resp. $h: \text{trop}(X) \rightarrow \mathbf{N}_{\mathbf{R}}$) is called the *tropicalization* of X (resp. $f: X \rightarrow S$) with respect to the model X^0 . Plainly, the tropical curve $\text{trop}(X)$ is independent of the parameterization and depends only on X^0 . If the family $X \rightarrow S$ is stable and X^0 is the stable model, then the corresponding tropicalization is called simply *the tropicalization* of X (resp. $f: X \rightarrow S$). Plainly, $\text{trop}(X)$ (resp. $h: \text{trop}(X) \rightarrow \mathbf{N}_{\mathbf{R}}$) is stable in this case.

Remark 4.1. — Let l_i be the leg corresponding to a marked point σ_i . If σ_i is mapped to the boundary divisor D , then, by definition, $f(\sigma_i)$ belongs to a unique component of D , and in particular, to the smooth locus of D . Thus, the slope $\frac{\partial h}{\partial l_i}$ is completely determined by the irreducible component of D containing the image of σ_i . Furthermore, the multiplicity of $\frac{\partial h}{\partial l_i}$ is the multiplicity of σ_i in f^*D . Note also that $\frac{\partial h}{\partial l_i} = 0$ if and only if $f(\sigma_i)$ is contained in the dense orbit $T \subset S$.

Next, let us recall the *tropicalization map* $\text{trop}: X(\mathbf{K}) \setminus \bigcup_i \sigma_i \rightarrow \text{trop}(X)$ from the set of \mathbf{K} -points of X to its tropicalization. The image $\text{trop}(\eta)$ of a \mathbf{K} -point $\eta \in X(\mathbf{K}) \setminus \bigcup_i \sigma_i$ is defined as follows. We temporarily add η to the collection of marked points of X and consider the minimal modification $(X^0)' \rightarrow X^0$ such that $(X^0)'$ is a prestable model of $(X, \sigma_\bullet, \eta)$. Then $\text{trop}(X)'$ is obtained from $\text{trop}(X)$ by attaching a leg either to an existing vertex or to a new two-valent vertex splitting an edge or a leg of $\text{trop}(X)$. Either way, the geometric realization of $\text{trop}(X)'$ is obtained from $\text{trop}(X)$ by attaching a leg to some point $q \in \text{trop}(X)$. The tropicalization map $\text{trop}: X(\mathbf{K}) \setminus \bigcup_i \sigma_i \rightarrow \text{trop}(X)$ sends η to the point q .

Remark 4.2. — The tropicalization defined above is compatible with the natural tropicalization of the torus in the sense that the following diagram is commutative:

$$\begin{array}{ccc} X \setminus f^{-1}(D) = f^{-1}(T) & \xrightarrow{\text{trop}} & \text{trop}(X) \\ \downarrow f & & \downarrow h \\ T & \xrightarrow{\text{trop}} & \mathbf{N}_{\mathbf{R}} \end{array}$$

where the map $\text{trop}: T \rightarrow \mathbf{N}_{\mathbf{R}}$ is defined by $p \mapsto (m \mapsto -v(x^m(p)))$. In particular, in the notation of Remark 4.1, if h contracts l_i , then $h(l_i) = \text{trop}(f(\sigma_i))$.

Below we will need a more explicit description of $\text{trop}(\eta) \in \text{trop}(\mathbf{X})$. Let $\psi: \text{Spec}(\mathbf{K}) \rightarrow \mathbf{X}$ be the immersion of the point η . Since $\mathbf{X}^0 \rightarrow \text{Spec}(\mathbf{K}^0)$ is proper, ψ admits a unique extension $\psi^0: \text{Spec}(\mathbf{K}^0) \rightarrow \mathbf{X}^0$ with image $\bar{\eta}$. Let $s \in \text{Spec}(\mathbf{K}^0)$ be the closed point. Its image in $\tilde{\mathbf{X}}$ under the map ψ^0 is called the *reduction* of η , and is denoted by $\text{red}(\eta)$. If the reduction of η is a non-special point of a component $\tilde{\mathbf{X}}_v$, then $(\mathbf{X}^0)' = \mathbf{X}^0$ and $\text{trop}(\eta) = v$. If the reduction of η is a node $z \in \tilde{\mathbf{X}}$, then $\text{trop}(\eta)$ belongs to the edge e corresponding to z , and we shall specify the distance from $\text{trop}(\eta)$ to the two vertices v and w adjacent to e . Let $\lambda \in \mathbf{K}^{00}$ be such that \mathbf{X}^0 is given by $ab = \lambda$ étale locally at z , and let $\tilde{\mathbf{X}}_v, \tilde{\mathbf{X}}_w$ be the branches of $\tilde{\mathbf{X}}$ associated to the vertices v and w , respectively. Without loss of generality, $\tilde{\mathbf{X}}_v$ is given locally by $a = 0$ and $\tilde{\mathbf{X}}_w$ by $b = 0$.

Lemma 4.3. — *The distance from $\text{trop}(\eta)$ to w in e is given by $v(\psi^*(a))$.*

Proof. — The model $(\mathbf{X}^0)'$ is a blow up of \mathbf{X}^0 at the point z , and its local charts are given by $at = \mu$ and $bt^{-1} = \lambda\mu^{-1}$. In the first chart, the exceptional divisor is given by $a = 0$ and $\tilde{\mathbf{X}}_w$ by $t = 0$. Furthermore, $\psi^*(t) \in \mathbf{K}^0$ is invertible since η specializes to a non-special point of the exceptional divisor. Thus, $v(\psi^*(a)) = v(\mu)$ is the distance from $\text{trop}(\eta)$ to w in e . \square

Finally, if the reduction of η is the specialization of a marked point σ_i , and l_i is the leg corresponding to σ_i , then $\text{trop}(\eta)$ belongs to the leg l_i . To specify the distance from $\text{trop}(\eta)$ to the vertex w adjacent to l_i , let $a = 0$ be an étale local equation of σ_i^0 in \mathbf{X}^0 around $\psi^0(s)$. Then,

Lemma 4.4. — *The distance from $\text{trop}(\eta)$ to w in l_i is given by $v(\psi^*(a))$.*

Proof. — The proof is completely analogous to that of Lemma 4.3. \square

4.3. Tropicalization of one-parameter families of curves. — In this section, we describe the tropicalization procedure for one-dimensional families of stable curves with a map to a toric variety. We expect that the main existence statement – the first part of Theorem 4.6 – is known to experts; see, in particular, [AGGS20, CCUW20, Ran19] for related constructions. We use the local (toroidal) structure of the moduli space of pointed curves and the universal curve over it; cf. [DM69, Theorem 5.2] and [Knu83, Theorem 2.7].

Definition 4.5. — *A family of parameterized curves in \mathbf{S} consists of the following data:*

- (1) *a smooth, projective base curve with marked points $(\mathbf{B}, \tau_\bullet)$,*
- (2) *a family of stable marked curves $(\mathcal{X} \rightarrow \mathbf{B}, \sigma_\bullet)$, smooth over $\mathbf{B}' := \mathbf{B} \setminus (\bigcup_i \tau_i)$, and*
- (3) *a rational map $f: \mathcal{X} \dashrightarrow \mathbf{S}$, defined over \mathbf{B}' , such that for any closed point $b \in \mathbf{B}'$, the restriction $f_b: \mathcal{X}_b \rightarrow \mathbf{S}$ is a parameterized curve.*

Theorem 4.6. — *Let $f: \mathcal{X} \dashrightarrow S$ be a family of parameterized curves, B^0 a prestable model of the base curve (B, τ_\bullet) , and assume that the family $(\mathcal{X} \rightarrow B, \sigma_\bullet)$ admits a split stable model $(\mathcal{X}^0 \rightarrow B^0, \sigma_\bullet^0)$. Let $\Lambda := \text{trop}(B)$ be the tropicalization of (B, τ_\bullet) with respect to B^0 and assume that Λ has no loops. Then there exists a family of parameterized tropical curves $h: \Gamma_\Lambda \rightarrow \mathbb{N}_{\mathbf{R}}$ such that for any \mathbf{K} -point $\eta \in B' = B \setminus (\bigcup_i \tau_i)$, the fiber of (Γ_Λ, h) over $\text{trop}(\eta)$ is the tropicalization of $f: \mathcal{X}_\eta \rightarrow S$ with respect to the model $\mathcal{X}^0|_{\bar{\eta}}$. Furthermore, the induced map $\alpha: \Lambda \rightarrow M_{g,n,\nabla}^{\text{trop}}$ is either harmonic or locally combinatorially surjective at any vertex $w \in V(\Lambda)$ for which the curve Γ_w is weightless and 3-valent except for at most one 4-valent vertex.*

Definition 4.7. — *The family $h: \Gamma_\Lambda \rightarrow \mathbb{N}_{\mathbf{R}}$ constructed in Theorem 4.6 is called the tropicalization of $f: \mathcal{X} \dashrightarrow S$ with respect to $\mathcal{X}^0 \rightarrow B^0$.*

Remark 4.8. — Without the splitness assumption, the dual graphs \mathbf{G}_w can be defined only up to an automorphism. In general, one needs to consider families $\Gamma_\Lambda \rightarrow \Lambda$ with a *stacky structure*, following [CCUW20].

Corollary 4.9. — *Let $f: \mathcal{X} \dashrightarrow S$ be a family of parameterized curves, and $h: \Gamma_\Lambda \rightarrow \mathbb{N}_{\mathbf{R}}$ its tropicalization with respect to $\mathcal{X}^0 \rightarrow B^0$. Let $\tau \in B(\mathbf{K})$ be a marked point, $l \in L(\Lambda)$ the associated leg, and \mathbf{G}_l the underlying graph of the tropical curves parameterized by l . Assume that the lengths of all but one edge of \mathbf{G}_l are constant in the family Γ_Λ and the map h is constant on all vertices of \mathbf{G}_l . Then,*

- (1) \mathcal{X}_τ has exactly one node. In particular, the geometric genus of \mathcal{X}_τ is one less than the geometric genus of a general fiber of $\mathcal{X} \rightarrow B$;
- (2) The rational map $f: \mathcal{X} \dashrightarrow S$ is defined on the fiber \mathcal{X}_τ . Furthermore, f maps the generic point(s) of \mathcal{X}_τ to the dense orbit $T \subset S$, and $f: \mathcal{X}_\tau \rightarrow S$ has the same tangency profile as the general fiber $f: \mathcal{X}_b \rightarrow S$.

Remark 4.10. — Under the assumptions of the corollary, the edge of varying length gets necessarily contracted by the map h .

The rest of the section is devoted to the proof of Theorem 4.6 and Corollary 4.9. Recall that \mathbf{K} is the algebraic closure of a complete discretely valued field, which we denote by F . Any \mathbf{K} -scheme of finite type is defined over a finite extension F' of F , which is also a complete discretely valued field since so is F . Thus, we may view a \mathbf{K} -scheme of finite type as the base change of a scheme over F' . In the proofs below, we will work over F' to have a well-behaved total space of the families we consider. To ease the notation, we will assume that all models and points we are interested in are defined already over F^0 but will not assume that the valuation of the uniformizer π is one.

In particular, we may assume that $B^0, \mathcal{X}^0, \mathcal{X}^0 \rightarrow B^0$, and $\bar{\eta}$ are all defined over F^0 . In the proof, we will use the following notation: $\psi: \text{Spec}(F) \rightarrow B$ will denote the immersion of the point η , $s := \text{red}(\eta) \in \tilde{B}$ its reduction, and $\tilde{\mathcal{X}}_s$ the corresponding fiber. We set $D_{B^0} := \tilde{B} \cup (\bigcup_i \tau_i)$ and $D_{\mathcal{X}^0} := \tilde{\mathcal{X}} \cup (\bigcup_i \mathcal{X}_{\tau_i}) \cup (\bigcup_j \sigma_j)$, where \mathcal{X}_{τ_i} are the fibers

over the marked points $\tau_i \in \mathbf{B}(\mathbf{F})$. Then the pullback of any monomial function $f^*(x^m)$ is regular and invertible on $\mathcal{X}^0 \setminus \mathbf{D}_{\mathcal{X}^0}$ by the assumptions of the theorem.

4.4. Proof of Theorem 4.6.

Step 1: the tropicalization $h_\eta: \text{trop}(\mathcal{X}_\eta) \rightarrow \mathbf{N}_{\mathbf{R}}$ depends only on $\text{trop}(\eta) \in \Lambda$. — We set $q := \text{trop}(\eta)$. After modifying \mathbf{B}^0 , we may assume that $s \in \tilde{\mathbf{B}}$ is non-special and, by definition, this modification depends only on $\text{trop}(\eta)$, cf. Section 4.2. Let $\tilde{\mathbf{B}}_w$ be the component containing s , and $w = q$ the corresponding vertex of Λ . By definition, the underlying graph of $\text{trop}(\mathcal{X}_\eta)$ is the dual graph \mathbf{G}_s of $\tilde{\mathcal{X}}_s$. Let z be a node of $\tilde{\mathcal{X}}_s$. Then there exist an étale neighborhood \mathbf{U} of s in \mathbf{B}^0 and a function $g_z \in \mathcal{O}_{\mathbf{U}}(\mathbf{U})$ vanishing at s such that the family $\mathcal{X}^0 \times_{\mathbf{B}^0} \mathbf{U}$ is given by $xy = g_z$ étale locally near z . After shrinking \mathbf{U} , we may assume that the latter is true in a neighborhood of any node of $\tilde{\mathcal{X}}_s$ over \mathbf{U} . Furthermore, since $s \in \tilde{\mathbf{B}}$ is non-special, we may choose \mathbf{U} such that the pullback of $\mathbf{D}_{\mathbf{B}^0}$ to \mathbf{U} is $\tilde{\mathbf{B}}_w$.

By assumption, \mathcal{X} is smooth over \mathbf{B}' , and hence each g_z is invertible on the complement of $\mathbf{D}_{\mathbf{B}^0}$, i.e., away from $\tilde{\mathbf{B}}_w$. However, \mathbf{U} is normal, and $g_z(s) = 0$ for all nodes z . Thus, all the g_z 's vanish identically along $\tilde{\mathbf{B}}_w$, which implies that the dual graph \mathbf{G}_s of $\tilde{\mathcal{X}}_s$ is étale locally constant over $\tilde{\mathbf{B}}_w$. We claim that the lengths of the edges of \mathbf{G}_s are also étale locally constant. Indeed, pick a node $z \in \tilde{\mathcal{X}}_s$. Since g_z is invertible away from $\tilde{\mathbf{B}}_w$, and π vanishes to order one along $\tilde{\mathbf{B}}_w$, there exists $k_w \in \mathbf{N}$ such that $\pi^{-k_w} g_z$ is regular and invertible in codimension one, and hence, by normality, it is regular and invertible on \mathbf{U} . Thus, the length of the edge of \mathbf{G}_s corresponding to z is given by $\nu(\psi^* g_z) = k_w \nu(\pi)$, cf. Section 4.2, which is étale locally constant around s . Finally, since $\mathcal{X}|_{\tilde{\mathbf{B}}_w} \rightarrow \tilde{\mathbf{B}}_w$ is split, the identifications of graphs \mathbf{G}_s with the dual graph \mathbf{G}_w of the generic fiber is canonical. Hence the tropicalization $\text{trop}(\mathcal{X}_\eta)$ depends only on q , and we set $\Gamma_q := \text{trop}(\mathcal{X}_\eta)$.

It remains to check that the parameterization h_η also depends only on q . Since the tangency profile of $\mathcal{X}_\eta \rightarrow \mathbf{S}$ is independent of η , so are the slopes of the legs of $\text{trop}(\mathcal{X}_\eta)$, cf. Remark 4.1. Let u be a vertex of \mathbf{G}_w , and $\tilde{\mathcal{X}}_u \subset \tilde{\mathcal{X}}$ the component corresponding to u that dominates $\tilde{\mathbf{B}}_w$. Pick any $m \in \mathbf{M}$. Since $f^*(x^m)$ is regular and invertible away from $\mathbf{D}_{\mathcal{X}^0}$, \mathcal{X}^0 is normal, and π vanishes to order one along $\tilde{\mathcal{X}}_u$, it follows that there is an integer $k_u \in \mathbf{Z}$ such that $\pi^{k_u} f^*(x^m)$ is regular and invertible at the generic point of $\tilde{\mathcal{X}}_u$. Thus, $h_\eta(u)(m) = k_u \nu(\pi)$, and hence h_η depends only on q . We set $h_q := h_\eta$, and obtain a parameterized tropical curve $h_q: \Gamma_q \rightarrow \mathbf{N}_{\mathbf{R}}$ that is canonically isomorphic to the tropicalization of $f: \mathcal{X}_\eta \rightarrow \mathbf{S}$ for any \mathbf{K} -point $\eta \in \mathbf{B}'$ satisfying $\text{trop}(\eta) = q$.

Step 2: $h: \Gamma_\Lambda \rightarrow \mathbf{N}_{\mathbf{R}}$ is a family of parameterized tropical curves. — We need to show that the fiberwise tropicalizations constructed in Step 1 form a family of parameterized tropical curves; that is, we need to specify a datum (\dagger) as in § 3.1.3 that satisfies the conditions of Definition 3.1. Since the tropicalizations of \mathbf{K} -points give rise only to rational points in the tropical curve, we will work with rational points $\Lambda_{\mathbf{Q}} \subset \Lambda$, and in the very end, extend the family by linearity to the non-rational points of Λ .

By Step 1, the extended degree of $h_q: \Gamma_q \rightarrow \mathbf{N}_{\mathbf{R}}$ is independent of $q \in \Lambda$, and will be denoted by $\overline{\mathbf{V}}$. Furthermore, if q is an inner point of some $e \in \overline{\mathbf{E}}(\mathbf{G}_{\Lambda})$, and $\text{trop}(\eta) = q$, then s is the special point of $\widetilde{\mathbf{B}}$ corresponding to e , and the underlying graph of Γ_q is the dual graph of the reduction $\widetilde{\mathcal{X}}_s$. Therefore, the underlying graph of Γ_q depends only on e and will be denoted by \mathbf{G}_e . We will see below that in this case, the slopes of the bounded edges of Γ_q , also depend only on e . Hence so does the combinatorial type of (Γ_q, h_q) .

Next, we specify the contraction maps. Let $s' \in \widetilde{\mathbf{B}}_w$ be a node of $\widetilde{\mathbf{B}}$ corresponding to an edge e of Λ , and $\mathbf{G}_e := \mathbf{G}_{s'}$ the corresponding weighted graph. Any degeneration of stable curves corresponds to a weighted edge contraction on the level of dual graphs. Thus, for the étale local branch of $\widetilde{\mathbf{B}}_w$ at s' corresponding to $\vec{e} \in \text{Star}(w)$, we get the contraction $\varphi_{\vec{e}}: \mathbf{G}_e \rightarrow \mathbf{G}_w$ between the associated dual graphs. Analogously, for the reduction $s' \in \widetilde{\mathbf{B}}_w$ of a marked point with corresponding leg $l \in \mathbf{L}(\Lambda)$, we obtain the desired contraction $\varphi_l: \mathbf{G}_l \rightarrow \mathbf{G}_w$.

Finally, let us define ℓ and h . Let $w \in \mathbf{V}(\Lambda)$ be a vertex, and $\vec{e} \in \text{Star}(w)$ an edge or a leg. We define the functions $\ell(\gamma, \cdot): e \cap \mathbf{Q} \rightarrow \mathbf{R}_{\geq 0}$ and $h(u, \cdot): e \cap \mathbf{Q} \rightarrow \mathbf{N}_{\mathbf{R}}$ for each $\gamma \in \mathbf{E}(\mathbf{G}_e)$ and $u \in \mathbf{V}(\mathbf{G}_e)$ as follows: if $q \in e^\circ \cap \mathbf{Q}$, then we set $\ell(\gamma, q) := \ell_q(\gamma)$ and $h(u, q) := h_q(u)$, and if q is the tail of \vec{e} , then $\ell(\gamma, q) := \ell_q(\varphi_{\vec{e}}(\gamma))$ and $h(u, q) := h_q(\varphi_{\vec{e}}(u))$. To finish the proof of Step 2, it remains to show that the functions $\ell(\gamma, \cdot)$ and $h(u, \cdot)$ are restrictions of integral affine functions on e . Indeed, if this is the case, then the family extends to the irrational points of e . Furthermore, it follows that the slope of h_q along any $\gamma \in \overline{\mathbf{E}}(\mathbf{G}_e)$ is continuous on e° , and obtains values in \mathbf{N} at the rational points $q \in e^\circ \cap \mathbf{Q}$. Therefore, it is necessarily constant on e° , and hence so is the combinatorial type.

Notice that for the function h , the assertion can be verified separately for the evaluations $h(u, \cdot)(m)$ of $h(u, \cdot)$ for each $m \in \mathbf{M}$. For the rest of the proof, we fix an $m \in \mathbf{M}$, an edge $\gamma \in \mathbf{E}(\mathbf{G}_q)$, and a vertex $u \in \mathbf{V}(\mathbf{G}_q)$. Let $z \in \widetilde{\mathcal{X}}_s$ denote the node corresponding to γ and $\widetilde{\mathcal{X}}_{s,u}$ the irreducible component of $\widetilde{\mathcal{X}}_s$ corresponding to u ; see Figure 2. There are two cases to consider.

Case 1: $s \in \widetilde{\mathbf{B}}$ is a node. Let e be the edge of Λ corresponding to s . Then $q \in e^\circ$ and $\mathbf{G}_q = \mathbf{G}_e$. After shrinking \mathbf{U} , we may assume that it is given by $ab = \pi^{k_s}$ for some positive integer k_s , and $\widetilde{\mathbf{B}}$ has two components $\widetilde{\mathbf{B}}_v$ and $\widetilde{\mathbf{B}}_w$ in \mathbf{U} given by $a = 0$ and $b = 0$, respectively. Thus, the length $\ell_{\Lambda}(e)$ is given by $k_s \nu(\pi)$.

We start with the function ℓ . Since g_z is regular and vanishes only along $\widetilde{\mathbf{B}}_v \cup \widetilde{\mathbf{B}}_w$ in \mathbf{U} , there exist $k_a, k_b, k_{\pi} \in \mathbf{N}$ such that $a^{-k_a} b^{-k_b} \pi^{-k_{\pi}} g_z$ is regular and invertible on \mathbf{U} . Therefore,

$$\begin{aligned} (4.1) \quad \ell(\gamma, q) &= \nu(\psi^* g_z) = k_a \nu(\psi^* a) + k_b \nu(\psi^* b) + k_{\pi} \nu(\pi) \\ &= (k_a - k_b) \nu(\psi^* a) + k_b \ell_{\Lambda}(e) + k_{\pi} \nu(\pi). \end{aligned}$$

By Lemma 4.3, $\nu(\psi^*(a))$ is the distance of q from w in e . Thus, (4.1) defines an integral affine function of q on e with slope $k_a - k_b$; see Figure 3 for an illustration.




FIG. 2. — The left picture shows the components of $D_{\mathcal{X}^0}$. On the right, we indicate the irreducible components of $\tilde{\mathcal{X}}$

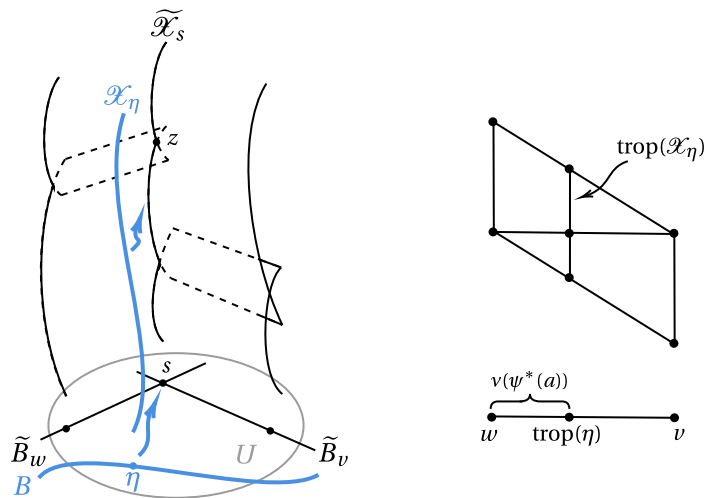


FIG. 3. — The tropicalization of $\mathcal{X} \rightarrow B$ near a node of \tilde{B}

It remains to show that the value of (4.1) for $\nu(\psi^*(a)) = 0$ is $\ell_w(\varphi_{\tilde{e}}(\gamma))$, where the orientation on e is such that w is its tail. By Step 1, $\ell_w(\varphi_{\tilde{e}}(\gamma)) = k_w \nu(\pi)$. Since a is invertible at the generic point of \tilde{B}_w , the order of vanishing k_w of g_z at the generic point of \tilde{B}_w is equal to the order of vanishing of $b^{k_b} \pi^{k_\pi}$, which in turn is the order of vanishing of $\pi^{k_b k_s + k_\pi}$. Thus,

$$(4.2) \quad \ell_w(\varphi_{\tilde{e}}(\gamma)) = k_w \nu(\pi) = (k_b k_s + k_\pi) \nu(\pi) = k_b \ell_\Lambda(e) + k_\pi \nu(\pi),$$

as needed.

Next we consider the function h . For the two orientations \vec{e} and \bar{e} on e , let $\varphi_{\vec{e}}: \mathbf{G}_e \rightarrow \mathbf{G}_w$ and $\varphi_{\bar{e}}: \mathbf{G}_e \rightarrow \mathbf{G}_v$ be the weighted edge contractions defined above. As in Step 1, we have irreducible components $\widetilde{\mathcal{X}}_{\varphi_{\vec{e}}(u)}$ and $\widetilde{\mathcal{X}}_{\varphi_{\bar{e}}(u)}$ of $\widetilde{\mathcal{X}}$; they are the components of $\widetilde{\mathcal{X}}$ containing $\widetilde{\mathcal{X}}_{s,u}$ and supported over $\widetilde{\mathbf{B}}_v$ and $\widetilde{\mathbf{B}}_w$, respectively; see Figure 2. Since $f^*(x^m)$ is regular and invertible outside $\mathbf{D}_{\mathcal{X}^0}$, there exist $r_a, r_b, r_\pi \in \mathbf{Z}$ such that $a^{r_a} b^{r_b} \pi^{r_\pi} f^*(x^m)$ is regular and invertible at the generic points of $\widetilde{\mathcal{X}}_{\varphi_{\vec{e}}(u)}$ and $\widetilde{\mathcal{X}}_{\varphi_{\bar{e}}(u)}$, and hence, by normality of \mathcal{X} , it is regular and invertible in a neighborhood of the generic point of $\widetilde{\mathcal{X}}_{s,u}$. Therefore,

$$(4.3) \quad \begin{aligned} h(u, q)(m) &= r_a v(\psi^*(a)) + r_b v(\psi^*(b)) + r_\pi v(\pi) \\ &= (r_a - r_b) v(\psi^*(a)) + r_b \ell_\Lambda(e) + r_\pi v(\pi). \end{aligned}$$

By Lemma 4.3, $v(\psi^*(a))$ is the distance of q from w in e . Thus, (4.3) defines an integral affine function on e with slope $r_a - r_b$. Furthermore, the value of this function for $v(\psi^*(a)) = 0$ is $h_w(\varphi_{\vec{e}}(u))$. Indeed, since a is regular and invertible at the generic point of $\widetilde{\mathbf{B}}_w$, it is regular and invertible at the generic point of $\widetilde{\mathcal{X}}_{\varphi_{\vec{e}}(u)}$. Thus, by the definition of h_w ,

$$h_w(\varphi_{\vec{e}}(u)) = v(b^{r_b} \pi^{r_\pi}) = v(\pi^{k_s r_b + r_\pi}) = (k_s r_b + r_\pi) v(\pi) = r_b \ell_\Lambda(e) + r_\pi v(\pi),$$

as needed. Similarly, the value of (4.3) for $v(\psi^*(a)) = \ell_\Lambda(e)$ is $h_v(\varphi_{\bar{e}}(u))$.

Case 2: $s \in \widetilde{\mathbf{B}}$ is the reduction of a marked point. Let $\tau \in \mathbf{B}$ be the marked point with reduction s , l the associated leg of Λ , and $\widetilde{\mathbf{B}}_w$ the component of $\widetilde{\mathbf{B}}$ containing s . Then $q \in l^\circ$ and $\mathbf{G}_q = \mathbf{G}_l$. After shrinking \mathbf{U} , we may assume that $\widetilde{\mathbf{B}} \cup \tau$ in \mathbf{U} is given by $\pi a = 0$.

Again, we start with the function ℓ . Since g_z is regular in \mathbf{U} and vanishes only along $\widetilde{\mathbf{B}} \cup \tau$, there exist $k_w, k_\tau \in \mathbf{N}$ such that $\pi^{-k_w} a^{-k_\tau} g_z$ is regular and invertible on \mathbf{U} . The length $\ell(\gamma, q) = v(\psi^* g_z)$ is thus given by


$$(4.4) \quad \ell(\gamma, q) = k_w v(\pi) + k_\tau v(\psi^*(a)).$$

By Lemma 4.4, the valuation $v(\psi^*(a))$ is the distance of q from w in l . Thus, (4.4) defines an integral affine function on l with slope k_τ and, by Step 1, its value at $v(\psi^*(a)) = 0$ is $k_w v(\pi) = \ell_w(\varphi_l(\gamma))$, as required; see Figure 4 for an illustration.

We proceed with the function h . Let $\varphi_l: \mathbf{G}_l \rightarrow \mathbf{G}_w$ be the edge contraction defined above. As before, we have a unique component $\widetilde{\mathcal{X}}_{\varphi_l(u)}$ of $\widetilde{\mathcal{X}}$ containing $\widetilde{\mathcal{X}}_{s,u}$ and supported over $\widetilde{\mathbf{B}}_w$. Similarly, there is a unique component $\mathcal{X}_{\tau,u}$ of \mathcal{X}_τ containing $\widetilde{\mathcal{X}}_{s,u}$ in its closure. Since $f^*(x^m)$ is regular and invertible away from $\mathbf{D}_{\mathcal{X}^0}$, there exist $r_w, r_\tau \in \mathbf{Z}$ such that $\pi^{r_w} a^{r_\tau} f^*(x^m)$ is regular and invertible on \mathbf{U} . Therefore,

$$(4.5) \quad h(u, q)(m) = r_\tau v(\psi^*(a)) + r_w v(\pi).$$

By Lemma 4.4, the valuation $v(\psi^*(a))$ is the distance of q from w in l . Thus, (4.5) defines an integral affine function on l with slope r_τ , and the value $h_w(\varphi_l(u))(m)$ at w , since w is given by $v(\psi^*(a)) = 0$, cf. Step 1.

FIG. 4. — The tropicalization of $\mathcal{X} \rightarrow \mathbb{B}$ near a marked point of $\tilde{\mathbb{B}}$

Step 3: the harmonicity and the local combinatorial surjectivity of the map α . — Let $w \in V(\Lambda)$ be a vertex, Θ the combinatorial type of $\alpha(w)$, and $\mathcal{C} \rightarrow \overline{\mathcal{M}}_{g,n+|\nabla|}$ the universal curve. Consider the natural map to the coarse moduli space $\chi : \mathbb{B}^0 \rightarrow \overline{\mathbb{M}}_{g,n+|\nabla|}$. There are two cases to consider:

Case 1: χ contracts $\tilde{\mathbb{B}}_w$. Set $p := \chi(\tilde{\mathbb{B}}_w)$. Since $\tilde{\mathcal{X}} \rightarrow \tilde{\mathbb{B}}$ is split, it follows that the restriction of $\tilde{\mathcal{X}}$ to $\tilde{\mathbb{B}}_w$ is the product $\tilde{\mathbb{B}}_w \times \mathcal{C}_p$. Furthermore, \mathbf{G}_w is the dual graph of \mathcal{C}_p , and α maps $\text{Star}(w)$ to $\mathbb{M}_{[\Theta]}$. Since α is a piecewise integral affine map, it lifts to a map to \mathbb{M}_Θ , which we denote by α_w . Let us show that α is harmonic at w , i.e.,

$$\sum_{\tilde{z} \in \text{Star}(w)} \frac{\partial \alpha_w}{\partial \tilde{z}} = 0.$$

The latter equality can be verified coordinatewise. Recall that the integral affine structure on \mathbb{M}_Θ is induced from $\mathbf{N}_{\mathbf{R}}^{|\mathbf{V}(\mathbf{G}_w)|} \times \mathbf{R}^{|\mathbf{E}(\mathbf{G}_w)|}$. Let $\gamma \in \mathbf{E}(\mathbf{G}_w)$ be an edge corresponding to a node $z \in \mathcal{C}_p$, and assume for simplicity that γ is not a loop. The case of a loop can be treated similarly, and we leave it to the reader. Let $u, u' \in \mathbf{V}(\mathbf{G}_w)$ be the vertices adjacent to γ , \mathcal{C}_u and $\mathcal{C}_{u'}$ the corresponding components of \mathcal{C}_p , and $\tilde{\mathcal{X}}_u, \tilde{\mathcal{X}}_{u'}, \mathbf{Z}$ the pullbacks of $\mathcal{C}_u, \mathcal{C}_{u'}, z$ to $\tilde{\mathcal{X}}$, respectively. The universal curve \mathcal{C} is given étale locally at z by $xy = m_z$, where m_z is defined on an étale neighborhood of p and vanishes at p . Thus, \mathcal{X}^0 is given by $xy = g_z := \chi^* m_z$ in an étale neighborhood of \mathbf{Z} . Notice that we constructed a function defined in a neighborhood of the whole family of nodes \mathbf{Z} , which globalizes the local construction of Step 1 in the case we consider here.

To prove harmonicity with respect to the coordinate x_γ corresponding to γ , we now consider $(\pi^{-k_w} g_z)|_{\tilde{\mathbb{B}}_w}$, where, as in Step 1, k_w denotes the order of vanishing of g_z at the generic point of $\tilde{\mathbb{B}}_w$. Thus, $(\pi^{-k_w} g_z)|_{\tilde{\mathbb{B}}_w}$ is a non-zero rational function on $\tilde{\mathbb{B}}_w$, and

hence the sum of orders of zeroes and poles of this function is zero. We claim that this is precisely the harmonicity condition we are looking for. Indeed, since \mathcal{X}^0 is smooth over B' , the function g_Z is regular and invertible away from D_{B^0} , and hence, as usual, $(\pi^{-k_w} g_Z)|_{\tilde{B}_w}$ has zeroes and poles only at the special points of \tilde{B}_w . For $\vec{e} \in \text{Star}(w)$, let s be the corresponding special point. If s is a node of \tilde{B} , and \tilde{B}_v is the second irreducible component containing s , then pick $k_s \in \mathbf{N}$ as in Case 1 of Step 2, i.e., such that B^0 is given étale locally at s by $ab = \pi^{k_s}$. Then, using Case 1 of Step 2,

$$\frac{\partial x_\gamma}{\partial \vec{e}} = \frac{\ell_v(\gamma) - \ell_w(\gamma)}{\ell_\Lambda(e)} = \frac{k_v - k_w}{k_s},$$

which is the order of vanishing of $(\pi^{-k_w} g_Z)|_{\tilde{B}_w}$ at s . Similarly, if s is the specialization of a marked point, $\frac{\partial x_\gamma}{\partial \vec{e}}$ is again the order of vanishing of $(\pi^{-k_w} g_Z)|_{\tilde{B}_w}$ at s .

Next, let us show harmonicity with respect to a coordinate n_u for $u \in V(\mathbf{G}_w)$, that is, we need to show that

$$\sum_{\vec{e} \in \text{Star}(w)} \frac{\partial h(u, \cdot)(m)}{\partial \vec{e}} = 0$$

for any $m \in M$. Let $k_u \in \mathbf{Z}$ be such that $\pi^{k_u} f^*(x^m)$ is regular and invertible at the generic point of $\tilde{\mathcal{X}}_u$. Then the divisor $D_{u,m}$ of $(\pi^{k_u} f^*(x^m))|_{\tilde{\mathcal{X}}_u}$ has horizontal components supported on the special points of the fibers of $\tilde{\mathcal{X}} \rightarrow \tilde{B}$, and vertical components supported on the preimages of the special points of \tilde{B}_w . Pick a general point $c \in \mathcal{C}_u$, and consider the horizontal curve $\tilde{B}_{w,c} := \tilde{B}_w \times \{c\} \subset \tilde{\mathcal{X}}_u$. It intersects no horizontal components of $D_{u,m}$, and intersects its vertical components transversally.

Let $s \in \tilde{B}_w$ be a special point, and $\tilde{\mathcal{X}}_{s,u}$ the corresponding vertical component of $D_{u,m}$. Its multiplicity in $D_{u,m}$ is equal to the multiplicity of the point (s, c) in the divisor of $(\pi^{k_u} f^*(x^m))|_{\tilde{B}_{w,c}}$, and hence the sum over all special points of \tilde{B}_w of these multiplicities vanishes. On the other hand, we claim that the multiplicity of $\tilde{\mathcal{X}}_{s,u}$ in $D_{u,m}$ is nothing but $\frac{\partial h(u, \cdot)(m)}{\partial \vec{e}}$, where $\vec{e} \in \text{Star}(w)$ corresponds to the special point s . Indeed, if s is the specialization of a marked point τ , then $\frac{\partial h(u, \cdot)(m)}{\partial \vec{e}} = r_\tau$, where r_τ is the order of pole of $f^*(x^m)$ at τ as in Case 2 of Step 2. Since π does not vanish at τ , we conclude that $\frac{\partial h(u, \cdot)(m)}{\partial \vec{e}}$ is the multiplicity of the reduction $\tilde{\mathcal{X}}_{s,u}$ of \mathcal{X}_τ in $D_{u,m}$, as asserted. A similar computation shows that if s is a node of \tilde{B} , then $\frac{\partial h(u, \cdot)(m)}{\partial \vec{e}}$ is again the multiplicity of $\tilde{\mathcal{X}}_{s,u}$ in $D_{u,m}$. We leave the details to the reader.

Case 2: χ does not contract \tilde{B}_w . By the assumptions of the theorem, the graph \mathbf{G}_w is weightless and 3-valent except for at most one 4-valent vertex. Since rational curves with three special points have no moduli, it follows that \mathbf{G}_w has a 4-valent vertex, which we denote by $u \in \mathbf{G}_w$. In this case we will show that the map α is locally combinatorially surjective at w ; see Figure 5 for an illustration.

Consider $\tilde{\mathcal{X}}_u$ as above. By construction, $\tilde{\mathcal{X}}_u \rightarrow \tilde{B}_w$ is a family of rational curves and since $\tilde{\mathcal{X}} \rightarrow \tilde{B}$ is split, the four marked points in each fiber define sections of the




Fig. 5. — An illustration for Case 2 of Step 3 in the proof of Theorem 4.6. Since Γ_Λ is not embeddable in \mathbf{R}^3 , the picture is a “cartoon”. The valency of w may be greater than three, but for each resolution of the 4-valent vertex, there is at least one germ in $\text{Star}(w)$ as in the picture

family. Let $\xi: \tilde{\mathbf{B}}_w \rightarrow \overline{\mathbf{M}}_{0,4} \simeq \mathbf{P}^1$ be the induced map. Since χ does not contract $\tilde{\mathbf{B}}_w$, it follows that ξ is not constant, and hence surjective. We conclude that $\tilde{\mathbf{B}}_w$ contains points, the fibers over which have dual graphs corresponding to the three possible splittings of the 4-valent vertex u into a pair of 3-valent vertices joined by an edge. In particular, α does not map $\text{Star}(w)$ into $\mathbf{M}_{[\Theta]}$, and we need to show that for each Θ' with an inclusion $\mathbf{M}_\Theta \hookrightarrow \overline{\mathbf{M}}_{\Theta'}$ there is $\vec{e} \in \text{Star}(w)$ such that $\alpha(\vec{e}) \cap \mathbf{M}_{[\Theta']} \neq \emptyset$. To see this, notice that any such polyhedron $\mathbf{M}_{\Theta'}$ corresponds to one of the three possible splittings of the 4-valent vertex u since, by balancing, the slope of the new edge is uniquely determined. \square

4.5. Proof of Corollary 4.9. — Set $s' := \text{red}(\tau) \in \tilde{\mathbf{B}}$, and let z be the node of the fiber $\tilde{\mathcal{X}}_{s'}$ corresponding to the edge of varying length $\gamma \in \mathbf{E}(\mathbf{G}_l)$. Let $\tilde{\mathbf{B}}_w$ be the component of $\tilde{\mathbf{B}}$ containing s' . Recall that we set $\mathbf{D}_{\mathcal{X}^0} := \tilde{\mathcal{X}} \cup \left(\bigcup_i \mathcal{X}_{\tau_i} \right) \cup \left(\bigcup_j \sigma_j \right)$, where σ_j are the marked points of $\mathcal{X} \rightarrow \mathbf{B}$.

(1) As in Case 2 of Step 2 in the proof of Theorem 4.6, in an étale neighborhood \mathbf{U} of s' , $\tilde{\mathbf{B}} \cup \tau$ is given by $\pi a = 0$, and the family $\mathcal{X}^0 \rightarrow \mathbf{B}^0$ over \mathbf{U} is given étale locally near z by $xy = g_z$ for some $g_z \in \mathcal{O}_{\mathbf{U}}(\mathbf{U})$. The length of γ is an integral affine function on l , and by (4.4), its slope is given by the order of vanishing k_τ of g_z at τ . Since the slope of $\ell(\gamma, \cdot)$ along l is not constant, $k_\tau > 0$, i.e., g_z vanishes at τ . Thus, \mathcal{X}_τ has a node as asserted. Vice versa, any node of \mathcal{X}_τ specializes to some node z' of $\mathcal{X}_{s'}$ corresponding to an edge $\gamma' \in \mathbf{E}(\mathbf{G}_l)$. Then étale locally at z' , the family $\mathcal{X}^0 \rightarrow \mathbf{B}^0$ is given by $xy = g_{z'}$, and $g_{z'}$ vanishes at τ . By (4.4) we conclude that the slope of $\ell(\gamma', \cdot)$ is not constant along l . Thus, $\gamma = \gamma'$ by the assumption of the corollary, and hence $z' = z$.

(2) Let $\mathcal{X}'_\tau \subseteq \mathcal{X}_\tau$ be an irreducible component. First, we show that the pullback $f^*(x^m)$ is regular and invertible at the generic point of \mathcal{X}'_τ for any $m \in \mathbf{M}$. Hence the rational map f is defined at the generic point of \mathcal{X}'_τ and maps it to the dense orbit \mathbf{T} . Pick a vertex u of \mathbf{G}_l such that $\mathcal{X}_{s,u}$ belongs to the closure of \mathcal{X}'_τ . By (4.5), the slope of $h(u, \cdot)(m)$ along l is given by the order of pole of $f^*(x^m)$ along \mathcal{X}'_τ . However, $h(u, \cdot)(m)$ is constant along l by the assumptions of the corollary, and hence $f^*(x^m)$ has neither zero

nor pole at the generic point of \mathcal{X}'_τ . Second, notice that since \mathcal{X}^0 is normal, $f^*(x^m)$ is regular and invertible away from $\mathbf{D}_{\mathcal{X}^0}$, and it is regular and invertible in codimension one on \mathcal{X}_τ , it follows that f is defined on $\mathcal{X}_\tau \setminus \left(\bigcup_j \sigma_j\right)$. Pick any σ_j , and let us show that f is defined at $\sigma_j(\tau) \in \mathcal{X}_\tau$, too. Let $S_j \subseteq S$ be the affine toric variety consisting of the dense torus orbit and the orbit of codimension at most one containing the image of $\sigma_j(b)$ for a general $b \in B$. Then the pullback of any regular monomial function $x^m \in \mathcal{O}_{S_j}(S_j)$ is regular in codimension one in a neighborhood of $\sigma_j(\tau)$, and hence regular at $\sigma_j(\tau)$ by normality of \mathcal{X}^0 . Thus, f is defined at $\sigma_j(\tau)$, and maps it to S_j . The last assertion is clear. \square

5. Degeneration via point constraints

In this section, $(S, \mathcal{L}) = (\mathbf{P}^2, \mathcal{O}_{\mathbf{P}^2}(d))$, $M = N = \mathbf{Z}^2$, and $\nabla = \nabla_d$ is the reduced degree of tropical curves associated to the triangle Δ_d , i.e., ∇ consists of $3d$ vectors: $(1, 1)$, $(-1, 0)$, $(0, -1)$, each appearing d times. Recall that for an integer $1 - d \leq g \leq \binom{d-1}{2}$, the Severi variety parameterizing curves of degree d and geometric genus g is a locally closed subset $V_{g,d} \subseteq |\mathcal{O}_{\mathbf{P}^2}(d)|$ of pure dimension $3d + g - 1$ by Lemma 2.6 and Proposition 2.7. In this section, we prove the Degeneration Theorem and its corollary, which generalizes Zariski's Theorem to arbitrary characteristic.

Theorem 5.1 (Degeneration Theorem). — *Let $d \in \mathbf{N}$ and $1 - d \leq g \leq \binom{d-1}{2}$ be integers, and $V \subseteq V_{g,d}$ an irreducible component. Then \bar{V} contains $V_{1-d,d}$.*

Remark 5.2. — If the geometric genus of a degree d curve C is $1 - d$, then C is necessarily a union of d lines. Therefore, $V_{1-d,d}$ is dominated by $(\mathbf{P}^2)^*{}^d$. In particular, $V_{1-d,d}$ is irreducible, and its general element corresponds to a nodal curve.

Corollary 5.3 (Zariski's Theorem). — *Let $V \subset V_{g,d}$ be an irreducible subvariety. Then,*

- (1) $\dim(V) \leq 3d + g - 1$, and
- (2) if $\dim(V) = 3d + g - 1$, then for a general $[C] \in V$, the curve C is nodal.

Proof. — Assertion (1) follows from Proposition 2.7. If $\dim(V) = 3d + g - 1$, then V is an irreducible component of $V_{g,d}$ by Proposition 2.7, and hence $V_{1-d,d} \subseteq \bar{V}$ by Theorem 5.1. Thus, there exists $[C] \in \bar{V}$ such that C is nodal. However, the locus of nodal curves is open in \bar{V} by [Sta20, Tag 0DSC], and V is irreducible. Hence for a general $[C] \in V$, the curve C is nodal. \square

Remark 5.4. — Assertion (1) in arbitrary characteristic and for any toric surface was already proved in [Tyol3, Theorem 1.2] by the third author. Examples of toric surfaces for which assertion (2) fails can also be found *loc. cit.*

The rest of the section is devoted to the proof of Theorem 5.1, which proceeds by induction on (d, g) with the lexicographic order. The base of induction, $(d, g) = (1, 0)$, is clear. To prove the induction step, let $(d, g) > (1, 0)$ be a pair of integers such that $1 - d \leq g \leq \binom{d-1}{2}$, and assume that for all $(d', g') < (d, g)$ the assertion is true. Let us prove that it holds true also for the pair (d, g) . If $g = 1 - d$, then there is nothing to prove since the variety $V_{1-d,d}$ is irreducible by Remark 5.2. Thus, we may assume that $g > 1 - d$.

Step 1: the reduction to the case of irreducible curves. — We claim that it is enough to prove the assertion for $V \subseteq V_{g,d}^{\text{irr}}$. Indeed, let $[C] \in V$ be a general point and assume that C is reducible. Denote the degrees of the components by d_i , and their geometric genera by g_i . Then $d_i < d$ for all i , $d = \sum d_i$, and $g - 1 = \sum (g_i - 1)$. Consider the natural finite map $\prod \bar{V}_{g_i, d_i}^{\text{irr}} \rightarrow \bar{V}_{g,d}$, whose image contains $[C]$. By Proposition 2.7, the dimension of each irreducible component of $\bar{V}_{g_i, d_i}^{\text{irr}}$ is equal to $3d_i + g_i - 1$, and since

$$\sum_{i=1}^m (3d_i + g_i - 1) = 3d + g - 1 = \dim(V),$$

it follows that \bar{V} is dominated by a product of irreducible components $\bar{V}_i \subseteq \bar{V}_{g_i, d_i}^{\text{irr}}$. By the induction assumption, $V_{1-d_i, d_i} \subseteq \bar{V}_i$ for all i . But a general point of $\prod \bar{V}_{1-d_i, d_i}$ corresponds to a union of $d = \sum d_i$ general lines, and hence the map $\prod \bar{V}_{g_i, d_i}^{\text{irr}} \rightarrow \bar{V}_{g,d}$ takes it to a general point of $\bar{V}_{1-d,d}$. It follows now that $V_{1-d,d} \subseteq \bar{V}$. From now on, we assume that $V \subseteq V_{g,d}^{\text{irr}}$.

Our goal is to prove that the locus of curves of geometric genus $g - 1$ in \bar{V} has dimension $3d + g - 2$, since then \bar{V} necessarily contains a component of $V_{g-1,d}$ by Proposition 2.7, and hence also $V_{1-d,d}$ by the induction assumption. We proceed as follows: set $n := 3d + g - 1$, and pick $n - 1$ points $\{p_i\}_{i=1}^{n-1} \subset \mathbf{P}^2$ in general position. For each i , let $H_i \subset |\mathcal{O}_{\mathbf{P}^2}(d)|$ be the hyperplane parameterizing curves passing through the point p_i . Since $\dim(V) = n$ and the points $\{p_i\}_{i=1}^{n-1} \subset \mathbf{P}^2$ are in general position, the intersection $Z := V \cap \left(\bigcap_{i=1}^{n-1} H_i \right)$ has pure dimension one and a general $[C] \in Z$ corresponds to an integral curve C of geometric genus g that intersects the boundary divisor transversally by Proposition 2.7. It is sufficient to show that there exists $[C'] \in \bar{Z}$ such that C' is reduced and has geometric genus $g - 1$. Indeed, such a C' passes through a general collection of $n - 1$ points $\{p_i\}_{i=1}^{n-1} \subset \mathbf{P}^2$, and therefore the locus of curves of geometric genus $g - 1$ in \bar{V} has dimension $n - 1 = 3d + g - 2$.

To apply the results of the previous section, we assume that the field \mathbf{K} is the algebraic closure of a complete discretely valued field, which we may do by the Lefschetz principle. Furthermore, we may assume that the points $\{p_i\}_{i=1}^{n-1} \subset \mathbf{P}^2$ tropicalize to distinct vertically stretched points $\{q_i\}_{i=1}^{n-1}$ in \mathbf{R}^2 . This assumption allows us to work on the

tropical side with floor decomposed curves, which are very convenient for controlling the degenerations of curves parameterized by Z . The rest of the proof proceeds as follows. In Step 2, we modify the base curve Z and the family of curves over it so that both Z and the general fiber of the family become smooth. Hence, the tropicalization results of Section 4.3 apply. And in Step 3, we investigate the tropicalization of the modified family and prove that it necessarily contains a tropical curve that corresponds to an algebraic fiber of genus $g - 1$.

Step 2: constructing a family of parameterized curves. — This step follows the ideas and the techniques of de Jong [dJ96] based on the results of Deligne [Del85]. The goal is to construct a family $f: \mathcal{X} \dashrightarrow \mathbf{P}^2$ of parameterized curves over a smooth base curve (B, τ_\bullet) , a prestable model B^0 of (B, τ_\bullet) whose reduction \tilde{B} has smooth irreducible components, and a finite morphism $B \rightarrow \bar{Z}$ that satisfy the following properties: (i) $(\mathcal{X} \rightarrow B, \sigma_\bullet)$ extends to a split family of stable marked curves over B^0 , and (ii) for a general $[C] \in Z$, the fibers of $\mathcal{X} \rightarrow B$ over the preimages of $[C]$ are the normalization X of C equipped with the natural map to \mathbf{P}^2 and with $3d + n - 1$ marked points such that the first $n - 1$ of them are mapped to p_1, \dots, p_{n-1} , and the rest – to the boundary divisor; cf. Definition 4.5.

Let us start with the normalization $B \rightarrow \bar{Z}$ and with the pullback $\mathcal{X} \rightarrow B$ of the tautological family to B equipped with the natural map $f: \mathcal{X} \rightarrow \mathbf{P}^2$. We will replace B with finite coverings, dense open subsets, and compactifications several times, but to simplify the presentation, we will use the same notation B , \mathcal{X} , and f . First, we apply Lemma 2.3. After replacing B with B' and \mathcal{X} with \mathcal{X}' as in the lemma, we may assume that the family $\mathcal{X} \rightarrow B$ is a generically equinormalizable family of projective curves. After shrinking B , we may further assume that $\mathcal{X} \rightarrow B$ is equinormalizable, the pullback of the boundary divisor of \mathbf{P}^2 on each fiber is reduced, and the preimages of p_i 's are smooth points of the fibers.

Second, we label the points of \mathcal{X} that are mapped to $\{p_i\}_{i=1}^{n-1}$ and to the boundary divisor, which results in a finite covering of the base curve B . After replacing B with this covering, and \mathcal{X} with the normalization of the pullback, we equip the family $\mathcal{X} \rightarrow B$ with marked points σ_\bullet such that the first $n - 1$ of them are mapped to p_1, \dots, p_{n-1} , and the rest – to the boundary divisor. In particular, we obtain a 1-morphism $B \rightarrow \mathcal{M}_{g, 3d+n-1}$ that induces the family $(\mathcal{X} \rightarrow B, \sigma_\bullet)$. Notice that the natural morphism $f: \mathcal{X} \rightarrow \mathbf{P}^2$ satisfies Property (ii).

By [dJ96, §2.24], the compactification $\overline{\mathcal{M}}_{g, 3d+n-1}$ admits a finite surjective morphism from a projective scheme \overline{M} . Thus, after replacing B with an irreducible component of $B \times_{\overline{\mathcal{M}}_{g, 3d+n-1}} \overline{M}$, we may extend the family $(\mathcal{X} \rightarrow B, \sigma_\bullet)$ to a family of stable curves with marked points over a smooth projective base, i.e., we may assume that B is projective. Furthermore, the family over B is the pullback of the universal family over \overline{M} along $B \rightarrow \overline{M}$. Plainly, the induced rational map $f: \mathcal{X} \dashrightarrow \mathbf{P}^2$ still satisfies Property (ii).

Let B_{st} be the stable model of B , and B^0 the closure of the graph $B \rightarrow \overline{M}$ in the product $B_{\text{st}} \times \overline{M}$. It is a projective integral model of B over which the family $(\mathcal{X} \rightarrow B, \sigma_\bullet)$

naturally extends. Furthermore, the extension is obtained by pulling back the universal family from $\overline{\mathbf{M}}$. After replacing \mathbf{B}^0 with a prestable model dominating \mathbf{B}^0 , whose reduction has smooth irreducible components, we obtain a model of the base over which the family extends to a family of stable marked curves satisfying Property (ii). It remains to achieve splitness. To do so, we proceed as in [dj96, § 5.17].

We begin by adding sections such that the nodes of the geometric fibers are contained in the sections. Such sections exist by [dj96, Lemma 5.3]. We temporarily add these sections as marked points, replace \mathbf{B} with an appropriate finite covering, and construct a prestable model of the base over which the family of curves with the extended collection of marked points admits a stable model. By construction, the irreducible components of the geometric fibers of the new family are smooth. By applying [dj96, Lemmata 5.2 and 5.3] once again, we may assume that the new family admits a tuple of sections such that for any geometric fiber, any component contains a section, and any node is contained in a section. This implies that the new family is split; cf. [dj96, § 5.17]. Finally, we remove the temporarily marked points σ_i 's and stabilize. The obtained family is still split.

To summarize, we constructed a projective curve with marked points $(\mathbf{B}, \tau_\bullet)$, its integral model \mathbf{B}^0 , a family of marked curves $(\mathcal{X} \rightarrow \mathbf{B}, \sigma_\bullet)$, and a rational map $f: \mathcal{X} \dashrightarrow \mathbf{P}^2$ that satisfy Properties (i) and (ii).

Step 3: the conclusion.

Lemma 5.5. — *Let $h: \Gamma_\Lambda \rightarrow \mathbf{R}^2$ be the tropicalization of $f: \mathcal{X} \dashrightarrow \mathbf{P}^2$ with respect to $\mathcal{X}^0 \rightarrow \mathbf{B}^0$. Then there exists a leg l of Λ such that the lengths of all but one edge γ_l of \mathbf{G}_l are constant in the family Γ_Λ and the map h is constant on all vertices of \mathbf{G}_l , where \mathbf{G}_l is the underlying graph of the tropical curves parameterized by l . Furthermore, each connected component of $\mathbf{G}_l \setminus \{\gamma_l\}$ contains a non-contracted leg.*

We postpone the proof of the lemma and first deduce the theorem. Let $\tau \in \mathbf{B}(\mathbf{K})$ be the marked point corresponding to the leg l as in Lemma 5.5. By Corollary 4.9, the geometric genus of \mathcal{X}_τ is $g - 1$. Furthermore, the map f is defined on \mathcal{X}_τ and maps its generic points to the dense orbit. We claim that $f|_{\mathcal{X}_\tau}$ is birational onto its image, and hence $[f(\mathcal{X}_\tau)] \in \mathbf{Z}$ is a reduced curve of genus $g - 1$, which completes the proof of Theorem 5.1.

To show the birationality of $f|_{\mathcal{X}_\tau}$, first note that f does not contract irreducible components of \mathcal{X}_τ . Indeed, the tropicalization of a component must contain a non-contracted leg by Lemma 5.5 and therefore, its image intersects the boundary divisor. However, its generic point is mapped to the dense orbit, which implies that the component is not contracted by f . If $f|_{\mathcal{X}_\tau}$ is not birational onto its image, then $f(\mathcal{X}_\tau)$ (with the reduced structure) is a curve of geometric genus at most $g - 1$ and degree at most $d - 1$. However, $f(\mathcal{X}_\tau)$ passes through the $n - 1 = 3d + g - 2$ points in general position p_1, \dots, p_{n-1} , which is a contradiction by Proposition 2.7. \square

Proof of Lemma 5.5. — To prove the lemma, we shall first analyze the image of the induced map $\alpha : \Lambda \rightarrow \mathbf{M}_{g,n-1,\nabla}^{\text{trop}}$. The strata of the moduli space we will deal with admit no automorphisms. Thus, throughout the proof, one can think about $\mathbf{M}_{g,n-1,\nabla}^{\text{trop}}$ as a usual polyhedral complex rather than a generalized one. In particular, we will use the standard notion of *star* of a given stratum in a polyhedral complex.

Notice that α is not constant, since by the construction of \mathbf{B} , for any $p \in \mathbf{P}^2$, there exists $b \in \mathbf{B}$ such that the curve $f(\mathcal{X}_b)$ passes through p , and hence $h(\Gamma_{\text{trop}(b)})$ contains $\text{trop}(p)$, which can be any point in $\mathbf{Q}^2 \subset \mathbf{R}^2$, cf. Remark 4.2. We will also need the following key properties of α , which we prove next: (a) if α maps a vertex $v \in V(\Lambda)$ to a simple wall \mathbf{M}_Θ , see Definition 3.9, then there exists a vertex $w \in V(\Lambda)$ such that $\alpha(w) = \alpha(v)$ and the map α is locally combinatorially surjective at w ; and (b) if $\mathbf{M}_{\Theta'}$ is nice and $\alpha(\Lambda) \cap \mathbf{M}_{\Theta'} \neq \emptyset$, then

$$\alpha(\Lambda) \cap \mathbf{M}_{\Theta'} = \mathbf{M}_{\Theta'} \cap \text{ev}_{\Theta'}^{-1}(q_1, \dots, q_{n-1}),$$

where $\text{ev}_{\Theta'}$ denotes the evaluation map defined in Section 3.1.6. In particular, $\alpha(\Lambda) \cap \mathbf{M}_{\Theta'}$ is an interval, whose boundary is disjoint from $\mathbf{M}_{\Theta'}$ by Lemma 3.10 (3).

We start with property (a). Notice that $\mathbf{M}_\Theta \cap \text{ev}_\Theta^{-1}(q_1, \dots, q_{n-1})$ is a point by Lemma 3.10 (2). However, α is not constant, and hence there exists a vertex w such that $\alpha(\text{Star}(w)) \not\subseteq \mathbf{M}_\Theta$ and $\alpha(w) = \alpha(v)$. Thus, α is not harmonic at w , and hence it is locally combinatorially surjective at w by Theorem 4.6.

Property (b) follows from harmonicity. Indeed, since $\mathbf{M}_{\Theta'} \cap \text{ev}_{\Theta'}^{-1}(q_1, \dots, q_{n-1})$ is an interval, $\emptyset \neq \alpha(\Lambda) \cap \mathbf{M}_{\Theta'} \subseteq \overline{\mathbf{M}_{\Theta'} \cap \text{ev}_{\Theta'}^{-1}(q_1, \dots, q_{n-1})}$, and α is affine on the edges and legs, it follows that $\alpha(\Lambda) \cap \overline{\mathbf{M}_{\Theta'}}$ is a finite union of closed intervals. However, by Theorem 4.6, α is harmonic at the vertices $w \in V(\Lambda)$ that are mapped to $\mathbf{M}_{\Theta'}$. Thus, $\alpha(\Lambda) \cap \overline{\mathbf{M}_{\Theta'}}$ is a single closed interval, whose boundary is disjoint from $\mathbf{M}_{\Theta'}$, which implies (b).

Pick a point $p_n \in \mathbf{P}^2$ with tropicalization $q_n \in \mathbf{R}^2$ such that the collection $\{p_i\}_{i=1}^n$ is in general position, and the configuration $\{q_i\}_{i=1}^n$ is vertically stretched, see Section 3.2. Assume further that the points $\{q_i\}_{i=1}^n$ belong to a line defined by $y = -\mu x$ for some $\mu \gg 1$. Let $b \in \mathbf{B}(\mathbf{K})$ be a point such that $f(\mathcal{X}_b)$ passes through p_n . Then the tropicalization (Γ, h) of (\mathcal{X}_b, f) is a floor decomposed curve by Proposition 3.13. By Proposition 3.8, if we marked the point p_n on \mathcal{X}_b , the tropicalization $\text{trop}(\mathcal{X}_b; p_1, \dots, p_n)$ would be weightless and 3-valent, and the map to \mathbf{R}^2 would be an immersion away from the contracted legs. The tropicalization $\text{trop}(\mathcal{X}_b; p_1, \dots, p_{n-1})$ is obtained from $\text{trop}(\mathcal{X}_b; p_1, \dots, p_n)$ by removing the leg contracted to q_n and stabilizing. Thus, $\alpha(\text{trop}(b))$ belongs to a nice stratum $\mathbf{M}_{\Theta'}$.

Denote the floors from the bottom to the top by F_1, \dots, F_d and the elevator adjacent to the top floor by E . Recall that E has multiplicity one, cf. Remark 3.14. Let F_k , $k < d$, be the non-top floor adjacent to E . We may assume that q_n belongs to the image of the elevator E , and q_i to the image of the floor F_i for all $1 \leq i \leq d$. Indeed, pick a permutation $\sigma : \{1, \dots, n\} \rightarrow \{1, \dots, n\}$ such that $q_{\sigma(n)} \in h(E)$ and $q_{\sigma(i)} \in h(F_i)$ for all $1 \leq i \leq d$.

Set $p'_i := p_{\sigma(i)}$, and consider the curve B' associated to $\{p'_i\}_{i=1}^{n-1}$ and the corresponding family of parameterized curves $f': \mathcal{X}' \dashrightarrow \mathbf{P}^2$. By construction, there exists $b' \in B'(\mathbf{K})$ such that $(\mathcal{X}'_{b'}, f'|_{\mathcal{X}'_{b'}}) = (\mathcal{X}_b, f|_{\mathcal{X}_b})$. It remains to replace B with B' , (\mathcal{X}, f) with (\mathcal{X}', f') , and p_i 's with p'_i 's.

Denote the x -coordinate of $E \cap F_k$ by x and let E' be the downward elevator adjacent to F_k whose x -coordinate x' is the closest to x . Notice that by Remark 3.14, the x -coordinate of the marked point q_d does not belong to the interval joining x and x' . Without loss of generality, we may assume that $x' > x$. We will call a point of F_k *special* if it is either a marked point or a vertex. Denote the x -coordinates of the special points of F_k that belong to the interval $[x, x']$ by $x = x_0 < x_1 < \dots < x_r = x'$. If x_i belongs to an elevator, then the elevator will be denoted by E_i . Plainly, $E = E_0$ and $E_r = E'$.

Set $q_n(t) := q_n + t(x_1 - x_0, 0)$, $0 \leq t \leq 1$, and consider the continuous family of tropical curves $(\Gamma_t, h_t) \in \overline{M}_{\Theta'} \cap \text{ev}_{\Theta'}^{-1}(q_1, \dots, q_{n-1})$, where $(\Gamma_0, h_0) = (\Gamma, h)$, and such that $q_n(t) \in h_t(E)$ for all t . Since $q_1, \dots, q_{n-1}, q_n(t)$ are vertically λ -stretched for a very large value of λ , the floors of the curve remain disjoint in the deformation (Γ_t, h_t) . Notice also that since the points q_1, \dots, q_{n-1} are fixed, the x -coordinate of each elevator except E remains the same in the deformation. Furthermore, the x -coordinates of all elevators adjacent to F_i are fixed in the deformation for any $i \neq k, d$, as well as the position of the marked point q_i on F_i . Thus, the restriction of the parameterization h to F_i is also fixed. Finally, since the x -coordinate of q_d does not belong to $[x, x']$, it follows that (Γ_t, h_t) belongs to the intersection $\overline{M}_{\Theta'} \cap \text{ev}_{\Theta'}^{-1}(q_1, \dots, q_{n-1})$ for any $0 \leq t \leq 1$. Furthermore, the curve (Γ_1, h_1) belongs to a simple wall M_{Θ} , in which the elevator $E = E_0$ gets adjacent to a 4-valent vertex on the floor F_k together with either the elevator E_1 or the leg contracted to the marked point q_k .




FIG. 6. — On the left, the elevator E reaches a special point. On the right – the local pictures of typical elements in the three nice cones in $\text{Star}(M_{\Theta})$. The top row illustrates the case when the special point is a vertex, and the bottom – when it is a marked point




FIG. 7. — The case in which the 4-valent vertex is adjacent to two elevators of multiplicity one. The right picture illustrates the fact that in this case, the curves parameterized by $M_{\Theta'''} \cap \text{ev}_{\Theta'''}^{-1}(q_1, \dots, q_{n-1})$ have constant image in \mathbf{R}^2 and contain an edge of varying length contracted to u

Let us describe $\text{Star}(M_{\Theta})$ explicitly. It consists of three nice strata $M_{\Theta'}$, $M_{\Theta''}$, $M_{\Theta'''}$, cf. Case 2 of Step 3 in the proof of Theorem 4.6. The stratum $M_{\Theta''}$ parameterizes curves in which the elevator $E = E_0$ has x -coordinate larger than that of E_1 (resp. q_k), and $M_{\Theta'''}$ parameterizes curves in which E_0 and E_1 (resp. q_k) get adjacent to a common vertex u in the perturbation of the 4-valent vertex of (Γ_1, h_1) ; see Figure 6. By Property (b), there exists a vertex v of Λ , such that $\alpha(v) \in M_{g,n-1,\nabla}^{\text{trop}}$ is the isomorphism class of (Γ_1, h_1) , and hence by Property (a), there exists a vertex w such that $\alpha(w) = \alpha(v)$ and α is locally combinatorially surjective at w .

We proceed by induction on (k, r) with the lexicographic order and start with the extended base of induction: $r = 1$ and the multiplicity of E' is one. Since ∇ is reduced, this is the case in the actual base of induction $(k, r) = (1, 1)$. Let $e \in \text{Star}(w)$ be an edge such that $\alpha(e) \subset M_{\Theta'''}$. Since the elevators E and E' both have multiplicity one, it follows from the balancing condition that the third edge adjacent to u gets contracted by the parameterization map h ; see Figure 7. Thus, the locus $M_{\Theta'''} \cap \text{ev}_{\Theta'''}^{-1}(q_1, \dots, q_{n-1})$ is an unbounded interval. Furthermore, all curves in this locus factor through (Γ_1, h_1) . By Property (b), $M_{\Theta'''} \cap \text{ev}_{\Theta'''}^{-1}(q_1, \dots, q_{n-1}) = \alpha(\Lambda) \cap M_{\Theta'''}$. Hence there exists a leg $l \in L(\Lambda)$ such that $\alpha(l) \subset M_{\Theta'''} \cap \text{ev}_{\Theta'''}^{-1}(q_1, \dots, q_{n-1})$ is not bounded. The leg l satisfies the first assertion of the lemma. Note that here the edge γ_l is just the third edge adjacent to u (considered as a vertex of \mathbf{G}_l) other than E and E' , mentioned above. As γ_l connects F_k and u , and u is connected to the top floor F_d by E , each component of $\mathbf{G}_l \setminus \{\gamma_l\}$ contains either the non-contracted legs in F_k or the non-contracted legs in F_d . This fulfils the second assertion. Next, let us prove the induction step. We distinguish between two cases.

Case 1: $r > 1$. Let $e \in \text{Star}(w)$ be an edge such that $\alpha(e) \subset M_{\Theta''}$, and $(\Gamma_{1+\epsilon}, h_{1+\epsilon})$ the tropical curve parameterized by any point in e° . Then the corresponding invariant is $(k, r - 1)$, and hence, by the induction assumption, there exists a leg $l \in L(\Lambda)$ satisfying the assertion of the lemma. See Figure 8 for an illustration.

Case 2: $r = 1$ and $k > 1$. We already checked the case when the multiplicity of E' is one in the extended induction base. Thus, we may assume that the multiplicity $\mathfrak{w}(E')$ of E' is greater than one. Let q_j be the marked point contained in $h_1(E')$, and E'' the second elevator for which $q_j \in h_1(E'')$. Let $F_{k'}$ be the floor adjacent to E'' . Then $k' < k$. Let $e \in \text{Star}(w)$ be an edge such that $\alpha(e) \subset M_{\Theta'''}$, where $M_{\Theta'''}$ denotes the nice cone in the star of M_{Θ} in which the elevators E and E' get attached to a 3-valent vertex u


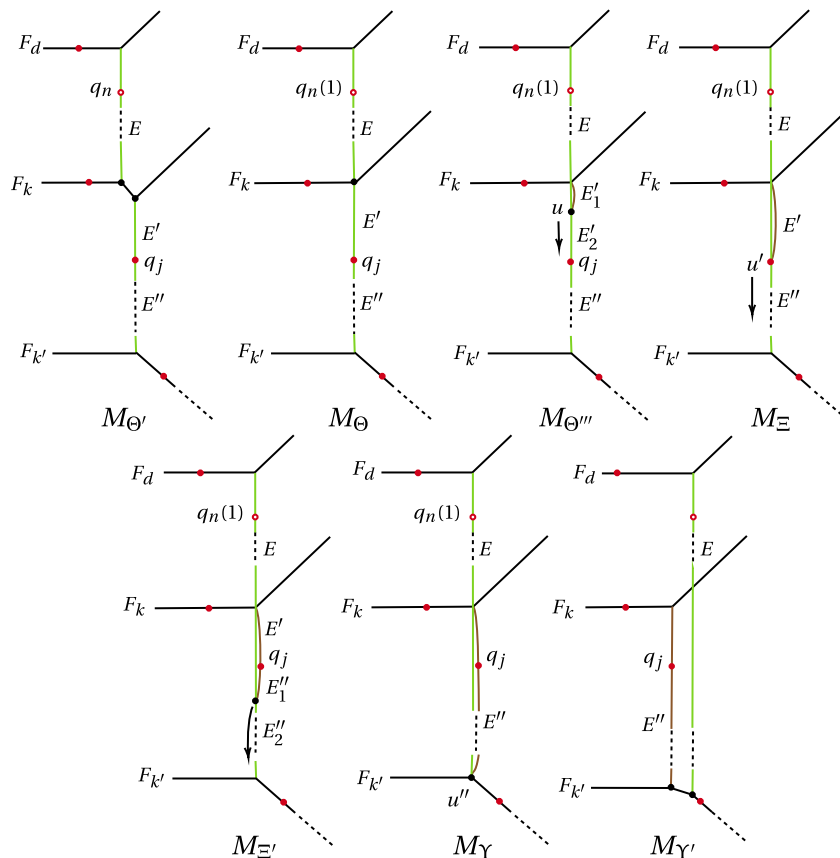


FIG. 8. — Moving the elevator E along the floor F_k . The top row illustrates the wall-crossing corresponding to a marked point, the bottom – to a downward elevator

obtained from the perturbation of the 4-valent vertex of (Γ_1, h_1) . We denote the two elevators corresponding to E' in this cone by E'_1 and E'_2 , see Figure 9 for an illustration.

It follows from the balancing condition that the slope of the elevator E'_1 in $\text{Star}(u)$ is $(0, \mathfrak{w}(E') - 1)$. Thus, the locus $M_{\Theta''} \cap \text{ev}_{\Theta''}^{-1}(q_1, \dots, q_{n-1})$ is a bounded interval, whose second boundary point belongs to a simple wall M_{Ξ} , which parameterizes curves with a 4-valent vertex u' adjacent to the elevators E , E' , E'' and to the leg l_j contracted to q_j .

As before, there exists a vertex w' of Λ such that $\alpha(w') \in M_{\Xi}$ and α is locally combinatorially surjective at w' by Property (b). Let $M_{\Xi'} \in \text{Star}(M_{\Xi})$ be the nice stratum in which the 4-valent vertex u' is perturbed to a 3-valent vertex adjacent to E and the downward elevator E''_2 , and a 3-valent vertex adjacent to E' and l_j . By Property (a), the image of α intersects $M_{\Xi'}$ non-trivially. Furthermore, $M_{\Xi'} \cap \text{ev}_{\Xi'}^{-1}(q_1, \dots, q_{n-1})$ is a bounded interval, whose second boundary point belongs to the simple wall M_{Υ} parameterizing curves with a 4-valent vertex u'' adjacent to the floor F_k and to the elevators E and E'' . As usual, $M_{\Xi'} \cap \text{ev}_{\Xi'}^{-1}(q_1, \dots, q_{n-1})$ belongs to the image of α , and hence all three nice strata of $\text{Star}(M_{\Upsilon})$ intersect the image of α non-trivially. In particular, there exists an edge e' of Λ parameterizing curves with invariant $(k', *) < (k, 1)$, and hence, by the induction assumption, there exists a leg $l \in L(\Lambda)$ satisfying the assertion of the lemma, which completes the proof. \square


 FIG. 9. — Going down with an elevator E' of multiplicity greater than one, and the corresponding wall-crossings

6. The local geometry of Severi varieties

In this section, we describe the local geometry of $\overline{V}_{g,d}$ and $\overline{V}_{g,d}^{\text{irr}}$ along $V_{1-d,d}$.

Theorem 6.1. — *The germ of $\overline{V}_{g,d}$ at a general $[C_0] \in V_{1-d,d}$ is a union of smooth branches indexed by subsets $\mu \subseteq C_0^{\text{sing}}$ of cardinality $\delta := \binom{d-1}{2} - g$. Moreover,*

- (1) *The scheme-theoretic intersection $\text{Br}(\mu) \cap \text{Br}(\mu')$ is smooth of codimension $|\mu' \setminus \mu|$ in $\text{Br}(\mu)$;*
- (2) *The branch $\text{Br}(\mu)$ belongs to $\overline{V}_{g,d}^{\text{irr}}$ if and only if $C_0 \setminus \mu$ is connected.*

Throughout this section, we fix the degree $d \geq 1$, the genus $1 - d \leq g \leq \binom{d-1}{2}$, and a general union of lines $[C_0] \in V_{1-d,d}$. Recall that for a projective algebraic curve C and a point $p \in C$, the δ -invariant of C at p is defined to be $\delta(C; p) := \dim((\mathcal{O}_{C^v} / \mathcal{O}_C) \otimes \mathcal{O}_{C,p})$, where C^v denotes the normalization of C . Plainly, $\delta(C; p) \neq 0$ if and only if p is singular.

The total δ -invariant of \mathbf{C} is then defined to be $\delta(\mathbf{C}) := \dim(\mathcal{O}_{\mathbf{C}^v}/\mathcal{O}_{\mathbf{C}}) = \sum_{p \in \mathbf{C}} \delta(\mathbf{C}; p)$. Recall also, that $\delta(\mathbf{C}) = p_a(\mathbf{C}) - p_g(\mathbf{C})$ is the difference between the arithmetic and the geometric genera; see, e.g., [Ros52, Theorem 8]. We set $\delta := \binom{d-1}{2} - g$, which is the total δ -invariant of an irreducible curve of degree d and geometric genus g in \mathbf{P}^2 . In order to analyze the local geometry of $\overline{\mathbf{V}}_{g,d}$ along $\mathbf{V}_{1-d,d}$, it is convenient to consider the decorated Severi varieties as introduced in [Tyo07].

Definition 6.2. — *The decorated Severi variety is the incidence locus*

$$\mathbf{U}_{d,\delta} \subset |\mathcal{O}_{\mathbf{P}^2}(d)| \times (\mathbf{P}^2)^\delta$$

consisting of the tuples $[\mathbf{C}; p_1, \dots, p_\delta]$, in which \mathbf{C} is a reduced curve of degree d , and p_1, \dots, p_δ are distinct nodes of \mathbf{C} . The union of the irreducible components $\mathbf{U} \subseteq \mathbf{U}_{d,\delta}$, for which \mathbf{C} is irreducible, is also called a decorated Severi variety, and is denoted by $\mathbf{U}_{d,\delta}^{\text{irr}}$.

The following is a version of [Tyo07, Proposition 2.11].

Proposition 6.3. — *Let $\phi: |\mathcal{O}_{\mathbf{P}^2}(d)| \times (\mathbf{P}^2)^\delta \rightarrow |\mathcal{O}_{\mathbf{P}^2}(d)|$ be the natural projection, and $\mathbf{C} := [\mathbf{C}; p_1, \dots, p_\delta] \in \mathbf{U}_{d,\delta}$ any point. Then,*

- (1) *The restriction of $d\phi$ to the tangent space $\mathbf{T}_{\mathbf{C}}(\mathbf{U}_{d,\delta})$ is injective;*
- (2) *The variety $\mathbf{U}_{d,\delta}$ is smooth of pure dimension $\dim(\mathbf{U}_{d,\delta}) = \binom{d+2}{2} - 1 - \delta = 3d + g - 1$;*
- (3) *$\phi(\overline{\mathbf{U}}_{d,\delta}) = \overline{\mathbf{V}}_{g,d}$, and $(\phi|_{\overline{\mathbf{U}}_{d,\delta}})^{-1}(\overline{\mathbf{V}}_{g,d}^{\text{irr}}) = \overline{\mathbf{U}}_{d,\delta}^{\text{irr}}$.*

Proof. — The proof is a rather straight-forward computation. After removing a general line from \mathbf{P}^2 , we get an open affine plane $\mathbf{A}^2 = \text{SpecK}[x, y] \subset \mathbf{P}^2$ containing all the p_l 's. We denote the coordinates x, y on the l -th copy of \mathbf{P}^2 in $(\mathbf{P}^2)^\delta$ by x_l, y_l , and identify $\mathbf{H}^0(\mathbf{P}^2, \mathcal{O}_{\mathbf{P}^2}(d))$ with the space of polynomials of degree at most d in the variables x and y . Then the decorated Severi variety $\mathbf{U}_{d,\delta}$ is given locally at \mathbf{C} by the system of 3δ (homogeneous in the coefficients of \mathbf{F}) equations:

$$\mathbf{F}(x_l, y_l) = \mathbf{F}_x(x_l, y_l) = \mathbf{F}_y(x_l, y_l) = 0,$$

where $\mathbf{F} = \sum_{i,j} a_{ij} x^i y^j \in \mathbf{H}^0(\mathbf{P}^2, \mathcal{O}_{\mathbf{P}^2}(d))$, and $\mathbf{F}_x, \mathbf{F}_y$ are its partial derivatives.

Let $\mathbf{G} \in \mathbf{H}^0(\mathbf{P}^2, \mathcal{O}_{\mathbf{P}^2}(d))$ be a polynomial defining \mathbf{C} , and for each l , let (λ_l, μ_l) be the coordinates of the point p_l . Then the tangent space $\mathbf{T}_{\mathbf{C}}(\mathbf{U}_{d,\delta})$ is given by the system of 3δ equations

$$(6.1) \quad \forall l \quad \begin{cases} \sum_{i,j} \lambda_l^i \mu_l^j da_{ij} + \mathbf{G}_x(p_l) dx_l + \mathbf{G}_y(p_l) dy_l = 0 \\ \sum_{i,j} i \lambda_l^{i-1} \mu_l^j da_{ij} + \mathbf{G}_{xx}(p_l) dx_l + \mathbf{G}_{xy}(p_l) dy_l = 0 \\ \sum_{i,j} j \lambda_l^i \mu_l^{j-1} da_{ij} + \mathbf{G}_{xy}(p_l) dx_l + \mathbf{G}_{yy}(p_l) dy_l = 0 \end{cases}$$

Notice that the kernel $\text{Ker}(d\phi)$ is given by $da_{ij} = 0$ for all i and j . However, since $p_l \in \mathbb{C}$ is a node for each l , the matrices

$$\begin{bmatrix} G_{xx}(p_l) & G_{xy}(p_l) \\ G_{xy}(p_l) & G_{yy}(p_l) \end{bmatrix}$$

are invertible, and $G_x(p_l) = G_y(p_l) = 0$. Therefore, the intersection $T_{\mathbf{C}}(\mathbb{U}_{d,\delta}) \cap \text{Ker}(d\phi)$ is zero, and assertion (1) follows. Furthermore, the dimension of the space of solutions of (6.1) is equal to the dimension of the space of solutions of the following system of δ equations:

$$(6.2) \quad \forall l \quad \sum_{i,j} \lambda_l^i \mu_l^j da_{ij} = 0.$$

Using the canonical identification $T_{[\mathbb{C}]|\mathcal{O}_{\mathbf{P}^2}(d)} \cong H^0(\mathbb{C}, \mathcal{O}_{\mathbb{C}}(d))$, we conclude that the differential $d\phi$ induces an isomorphism $T_{\mathbf{C}}(\mathbb{U}_{d,\delta}) \rightarrow H^0(\mathbb{C}, \mathcal{I}(d))$, where \mathcal{I} is the ideal sheaf of the scheme of nodes $Z := \bigcup_{l=1}^{\delta} p_l$. To proceed, we need the following lemma, whose proof we postpone.

Lemma 6.4. — $h^1(\mathbb{C}, \mathcal{I}(d)) = 0$.

It follows from the lemma that the sequence

$$0 \rightarrow H^0(\mathbb{C}, \mathcal{I}(d)) \rightarrow H^0(\mathbb{C}, \mathcal{O}_{\mathbb{C}}(d)) \rightarrow H^0(Z, \mathcal{O}_Z) \rightarrow 0$$

is exact. Thus, $h^0(\mathbb{C}, \mathcal{O}_{\mathbb{C}}(d)) - h^0(\mathbb{C}, \mathcal{I}(d)) = h^0(Z, \mathcal{O}_Z) = \delta$, and hence the system (6.2) has full rank. Therefore, $\mathbb{U}_{d,\delta}$ is smooth of pure dimension $\binom{d+2}{2} - 1 - \delta$ by the Jacobian criterion, and assertion (2) follows.

To prove (3), notice that the fibers of $\phi|_{\mathbb{U}_{d,\delta}}$ are finite. Thus, $\phi(\overline{\mathbb{U}}_{d,\delta})$ is pure dimensional of dimension $\binom{d+2}{2} - 1 - \delta = 3d + g - 1$. However, by the very definition, $\phi(\mathbb{U}_{d,\delta}) \subseteq \bigcup_{g' \leq g} \overline{V}_{g',d}$, and $\dim(V_{g',d}) = 3d + g' - 1 < 3d + g - 1$ for any $g' < g$ by Proposition 2.7 and Remark 2.8. Thus, $\phi(\overline{\mathbb{U}}_{d,\delta})$ is a union of irreducible components of $\overline{V}_{g,d}$. On the other hand, by Corollary 5.3, any irreducible component $V \subseteq \overline{V}_{g,d}$ admits a dense open subset that parameterizes nodal curves, and hence belongs to $\phi(\mathbb{U}_{d,\delta})$. Therefore, $\overline{V}_{g,d} \subseteq \phi(\overline{\mathbb{U}}_{d,\delta})$, and hence $\phi(\overline{\mathbb{U}}_{d,\delta}) = \overline{V}_{g,d}$. The last assertion of (3) now follows from the definitions of $\overline{V}_{g,d}^{\text{irr}}$ and $\overline{\mathbb{U}}_{d,\delta}^{\text{irr}}$. \square

Proof of Lemma 6.4. — The lemma is identical to [Tyo07, Claim 2.12], and the proof given in *loc. cit.* works in arbitrary characteristic. For the completeness of the presentation, we include a variation of this proof that proceeds by induction on the number of irreducible components of \mathbb{C} .

If \mathbb{C} is irreducible, consider the conductor ideal $\mathcal{I}^{\text{cond}} \subseteq \mathcal{O}_{\mathbb{C}}$, i.e., the annihilator of $\mathcal{O}_{\mathbb{C}^\nu}/\mathcal{O}_{\mathbb{C}}$, where $\mathbb{C}^\nu \rightarrow \mathbb{C}$ denotes the normalization of \mathbb{C} . Notice that $\mathcal{I}^{\text{cond}}$ is an

ideal also in \mathcal{O}_{C^\vee} under the natural embedding $\mathcal{O}_C \subseteq \mathcal{O}_{C^\vee}$. Since the fibers of $C^\vee \rightarrow C$ are zero-dimensional, $H^i(C, \mathcal{I}^{\text{cond}}(d)) = H^i(C^\vee, \mathcal{I}^{\text{cond}}(d))$ for all i . It follows from the definition that the vanishing locus of $\mathcal{I}^{\text{cond}}$ is C^{sing} , and hence $\mathcal{I}^{\text{cond}} \subseteq \mathcal{I}$. Consider the exact sequence of cohomology

$$H^1(C, \mathcal{I}^{\text{cond}}(d)) \rightarrow H^1(C, \mathcal{I}(d)) \rightarrow H^1(C, \mathcal{I}/\mathcal{I}^{\text{cond}} \otimes \mathcal{O}_C(d)).$$

Since the quotient $\mathcal{I}/\mathcal{I}^{\text{cond}}$ is a torsion sheaf, the group $H^1(C, \mathcal{I}/\mathcal{I}^{\text{cond}} \otimes \mathcal{O}_C(d))$ vanishes, and hence the map $H^1(C, \mathcal{I}^{\text{cond}}(d)) \rightarrow H^1(C, \mathcal{I}(d))$ is surjective. However, by [Ros52, Theorem 14], the degree of the invertible sheaf $\mathcal{I}^{\text{cond}}(d)$ on C^\vee is given by

$$\begin{aligned} c_1(\mathcal{I}^{\text{cond}}(d)) &= d^2 - 2\delta(C) = d^2 + 2p_g(C^\vee) - (d-2)(d-1) \\ &= 3d + 2g - 2 > 2g - 2, \end{aligned}$$

and hence $h^1(C, \mathcal{I}^{\text{cond}}(d)) = h^1(C^\vee, \mathcal{I}^{\text{cond}}(d)) = 0$, which implies $h^1(C, \mathcal{I}(d)) = 0$.

To prove the induction step, let $C_1 \subset C$ be an irreducible component of degree $d_1 < d$, and C_2 be the union of the other irreducible components. Set $d_2 := d - d_1$, and $Z_i := (Z \cap C_i) \setminus C_{3-i}$ for $i = 1, 2$. Let $\mathcal{I}_i \subset \mathcal{O}_{C_i}$ be the ideal sheaf of Z_i and $F_i \in H^0(C, \mathcal{O}_C(d_i))$ a polynomial defining C_i . Finally, denote by ι the closed immersion $C_1 \hookrightarrow C$. Consider the exact sequence

$$(6.3) \quad 0 \rightarrow \iota_* \mathcal{I}_1(d_1) \rightarrow \mathcal{I}(d) \rightarrow \mathcal{F}(d) \rightarrow 0,$$

where the first map is multiplication by F_2 . Then \mathcal{F} is supported on C_2 . Furthermore, it is the ideal sheaf of Z_2 union with the zero-dimensional scheme defined by the ideal $\langle F_1, F_2 \rangle$. In particular, there is a natural embedding $\mathcal{I}_2(d_2) \hookrightarrow \mathcal{F}(d)$ given by multiplication by F_1 , whose cokernel \mathcal{G} is a torsion sheaf. We conclude that there is an exact sequence of cohomology

$$H^1(C_2, \mathcal{I}_2(d_2)) \rightarrow H^1(C_2, \mathcal{F}(d)) \rightarrow H^1(C_2, \mathcal{G}),$$

in which the first group vanishes by the induction assumption and the last group vanishes since \mathcal{G} is a torsion sheaf. Thus, $H^1(C, \mathcal{F}(d)) = H^1(C_2, \mathcal{F}(d)) = 0$. Similarly, we get an exact sequence of cohomology

$$H^1(C, \iota_* \mathcal{I}_1(d_1)) \rightarrow H^1(C, \mathcal{I}(d)) \rightarrow H^1(C, \mathcal{F}(d))$$

associated to (6.3). The last group vanishes, and the first group vanishes too by the induction assumption, since it is isomorphic to $H^1(C_1, \mathcal{I}_1(d_1))$. Therefore $h^1(C, \mathcal{I}(d)) = 0$, which completes the proof. \square

Definition 6.5. — A δ -marking on C_0 is an ordered subset $\vec{\mu} \subseteq C_0^{\text{sing}}$ of cardinality δ . A δ -marking $\vec{\mu}$ is called irreducible if $C_0 \setminus \vec{\mu}$ is connected.

Remark 6.6. — Since δ and C_0 are fixed, we will usually call δ -markings on C_0 simply *markings*. For a marking $\vec{\mu}$ we denote the underlying set by μ .

Proof of Theorem 6.1. — Let $\vec{\mu} = (p_1, \dots, p_\delta)$ be a marking. Consider the germ $U_{\vec{\mu}}$ of the decorated Severi variety $U_{d,\delta}$ at $[C_0; p_1, \dots, p_\delta]$. By Proposition 6.3, it is smooth of dimension $3d + g - 1$, and $U_{\vec{\mu}}$ is isomorphic to its image under the natural projection ϕ to $|\mathcal{O}_{\mathbf{P}^2}(d)|$. Thus, $\phi(U_{\vec{\mu}})$ is a smooth branch of $\overline{V}_{g,d}$. Furthermore, the germ of $\overline{V}_{g,d}$ is covered by such branches. Plainly, $\phi(U_{\vec{\mu}})$ depends only on the underlying set μ , and we set $\text{Br}(\mu) := \phi(U_{\vec{\mu}})$. On the other hand, the set μ is determined by the branch $\text{Br}(\mu)$. Indeed, a general curve $[C] \in \text{Br}(\mu)$ corresponds to a nodal curve C of geometric genus g by Corollary 5.3. Thus, it has precisely δ nodes, and the set μ is the specialization of the set of nodes of C to C_0 . We conclude that the germ of $\overline{V}_{g,d}$ at $[C_0]$ consists of smooth branches indexed by subsets $\mu \subseteq C_0^{\text{sing}}$ of cardinality δ .

To prove (1), let $\mu = \{p_1, \dots, p_\delta\}$ and $\mu' = \{p'_1, \dots, p'_\delta\}$ be two subsets of C_0^{sing} of cardinality δ . Set $r := |\mu \cap \mu'|$. Without loss of generality, we may assume that $p_i = p'_i$ for all $1 \leq i \leq r$. Then the intersection $\text{Br}(\mu) \cap \text{Br}(\mu')$ contains $\text{Br}(\mu \cup \mu')$, which is the image of the germ of $U_{d,2\delta-r}$ at $[C_0; p_1, \dots, p_\delta, p'_{r+1}, \dots, p'_\delta]$. Recall that $\text{Br}(\mu \cup \mu')$ is smooth of dimension $\binom{d+2}{2} - 1 - (2\delta - r)$ by Proposition 6.3, and let us show that the scheme-theoretic intersection $\text{Br}(\mu) \cap \text{Br}(\mu')$ coincides with $\text{Br}(\mu \cup \mu')$. To do so, we compare the tangent spaces.

Let $\mathcal{I} \subseteq \mathcal{O}_{C_0}$, $\mathcal{I}' \subseteq \mathcal{O}_{C_0}$, and $\mathcal{J} \subseteq \mathcal{O}_{C_0}$ be the ideal sheaves of μ , μ' , and $\mu \cup \mu'$ respectively. Then the tangent space to $\text{Br}(\mu)$ is given by

$$T_{[C_0]}(\text{Br}(\mu)) = H^0(C_0, \mathcal{I}(d)) \subseteq H^0(C_0, \mathcal{O}_{C_0}(d));$$

cf. the proof of Proposition 6.3. Similarly, the tangent space to $\text{Br}(\mu')$ is given by $T_{[C_0]}(\text{Br}(\mu')) = H^0(C_0, \mathcal{I}'(d))$, and to $\text{Br}(\mu \cup \mu')$ by $T_{[C_0]}(\text{Br}(\mu \cup \mu')) = H^0(C_0, \mathcal{J}(d))$. Thus,

$$\begin{aligned} T_{[C_0]}(\text{Br}(\mu)) \cap T_{[C_0]}(\text{Br}(\mu')) &= H^0(C_0, \mathcal{I}'(d)) \cap H^0(C_0, \mathcal{I}(d)) \\ &= H^0(C_0, \mathcal{J}(d)) \\ &= T_{[C_0]}(\text{Br}(\mu \cup \mu')), \end{aligned}$$

which implies $\text{Br}(\mu) \cap \text{Br}(\mu') = \text{Br}(\mu \cup \mu')$ scheme-theoretically. In particular, the intersection $\text{Br}(\mu) \cap \text{Br}(\mu')$ is smooth of dimension $\binom{d+2}{2} - 1 - (2\delta - r)$. However, by Proposition 6.3, $\text{Br}(\mu)$ is smooth of dimension $\binom{d+2}{2} - 1 - \delta$. Hence the codimension of $\text{Br}(\mu) \cap \text{Br}(\mu')$ in $\text{Br}(\mu)$ is $\delta - r = |\mu' \setminus \mu|$ as asserted.

To prove (2), notice that a general curve C in a given branch $\text{Br}(\mu)$ is irreducible if and only if its normalization C^ν is so. But C^ν specializes to the partial normalization of C_0 at μ . Hence C^ν is irreducible if and only if the partial normalization of C_0 at μ is connected. The latter is clearly equivalent to $C_0^0 \setminus \mu$ being connected. \square

7. The proof of the Main Theorem

The Main Theorem follows from the following stronger assertion about the irreducibility of decorated Severi varieties:

Theorem 7.1. — *The decorated Severi variety $U_{d,\delta}^{\text{irr}}$ is either empty or irreducible. Furthermore, it is empty if and only if $\delta > \binom{d-1}{2}$.*

Indeed, set $\delta = \binom{d-1}{2} - g$. Then $\delta \leq \binom{d-1}{2}$ since $g \geq 0$. By Proposition 6.3 (3), $\bar{V}_{g,d}^{\text{irr}}$ is the image of $\bar{U}_{d,\delta}^{\text{irr}}$ under the natural projection to $|\mathcal{O}_{\mathbf{P}^2}(d)|$. By Theorem 7.1, $U_{d,\delta}^{\text{irr}}$ is non-empty and irreducible, and hence so is $V_{g,d}^{\text{irr}}$. Finally, by Proposition 2.7, the Severi variety $V_{g,d}$ is either empty or equidimensional of dimension $3d + g - 1$. Thus, the dimension of $V_{g,d}^{\text{irr}}$ is $3d + g - 1$, which completes the proof of the Main Theorem. \square

The proof of Theorem 7.1. — If $\delta > \binom{d-1}{2}$, then $U_{d,\delta}^{\text{irr}} = \emptyset$ since the geometric genus of any plane curve of degree d with δ nodes is at most $\frac{d-1}{2} - \delta < 0$, and hence none of such curves is irreducible. Vice versa, assume that $\delta \leq \binom{d-1}{2}$. Then

$$\delta \leq \binom{d-1}{2} = \binom{d}{2} - (d-1) = |C_0^{\text{sing}}| - (d-1).$$

Pick an irreducible component L of C_0 . Then the number of nodes of C_0 that belong to L is $d-1$, and hence there exists a marking $\vec{\mu}$ disjoint from L . Since L intersects any other component of C_0 , it follows that $\vec{\mu}$ is an irreducible marking. Thus, by Theorem 6.1 (2), $V_{g,d}^{\text{irr}}$ is not empty, and hence so is $U_{d,\delta}^{\text{irr}}$ by Proposition 6.3 (3). We conclude that $U_{d,\delta}^{\text{irr}} = \emptyset$ if and only if $\delta > \binom{d-1}{2}$.

Assume that $U_{d,\delta}^{\text{irr}}$ is not empty, and let us show that it is irreducible. We start by defining an equivalence relation on the set of markings. Let \mathfrak{S} be the full group of symmetries of the set C_0^{sing} of $\binom{d}{2}$ nodes of C_0 . The group \mathfrak{S} acts naturally on the set of markings, and we use the notation $\vec{\mu}^\sigma$ to denote the image of a marking $\vec{\mu}$ under the action of $\sigma \in \mathfrak{S}$. Let $\vec{\mu}$ be a marking, and L, L', L'' three different components of C_0 . Set $p := L \cap L', q := L' \cap L'',$ and $r := L \cap L'',$ and let $\tau \in \mathfrak{S}$ be the transposition switching q and r . If $p \notin \mu$, then we say that $\vec{\mu}^\tau$ is *similar* to $\vec{\mu}$. Plainly, $\vec{\mu}$ is irreducible if and only if so is $\vec{\mu}^\tau$. We define the equivalence relation \sim on the set of markings to be the minimal equivalence relation for which any pair of similar markings is equivalent. Since similarity is a symmetric relation, the equivalence relation \sim is nothing but the transitive closure of the similarity relation. Theorem 7.1 now follows from the following two lemmata – the first is purely combinatorial, and the second is algebraic.

Lemma 7.2. — *Any two irreducible markings are equivalent.*

Lemma 7.3. — *If the markings (p_1, \dots, p_δ) and (p'_1, \dots, p'_δ) are equivalent, then the tuples $[C_0; p_1, \dots, p_\delta]$ and $[C_0; p'_1, \dots, p'_\delta]$ belong to the same irreducible component of $\mathcal{U}_{d,\delta}$.*

Indeed, by the lemmata, there exists a unique irreducible component $U \subseteq \mathcal{U}_{d,\delta}$ that contains $[C_0; p_1, \dots, p_\delta]$ for any irreducible marking $\bar{\mu} = (p_1, \dots, p_\delta)$. If $U' \subseteq \mathcal{U}_{d,\delta}^{\text{irr}}$ is any irreducible component, then the projection $\phi(U')$ of U' to $|\mathcal{O}_{\mathbf{P}^2}(d)|$ is dense in an irreducible component of $\bar{V}_{g,d}^{\text{irr}}$ by Proposition 6.3, and $C_0 \in \phi(\bar{U}')$ by Theorem 5.1. Thus, there exists a marking $\bar{\mu} = (p_1, \dots, p_\delta)$ such that $[C_0; p_1, \dots, p_\delta] \in U'$. Since $U' \subseteq \mathcal{U}_{d,\delta}^{\text{irr}}$, the marking $\bar{\mu}$ is irreducible by Theorem 6.1 (2) and Proposition 6.3 (3). Hence $U' = U$, which completes the proof. \square

Proof of Lemma 7.2. — Although the lemma is a particular case of [Ty07, Lemma 3.10], we include its proof for the reader's convenience. Denote the irreducible components of C_0 by L_1, \dots, L_d , and set $q_{i,j} := L_i \cap L_j$. Let $\bar{\mu}$ be an irreducible marking.

First, let us show that $\bar{\mu}$ is equivalent to a marking disjoint from L_1 . If $\mu \cap L_1 \neq \emptyset$, then the union C of the components L of C_0 for which $L \cap L_1 \in \mu$ is not empty. Let C' be the union of the remaining components excluding L_1 . Since $C_0 \setminus \mu$ is connected, $C' \neq \emptyset$, and there exists $r \in C \cap C'$ that does not belong to μ . Let $L_i \subseteq C$ and $L_j \subseteq C'$ be the components containing r . Then $q_{1,i} \in \mu$, $q_{1,j} \notin \mu$, and hence $\bar{\mu} \sim \bar{\mu}^\tau$, where τ is the transposition switching $q_{1,i}$ and r . Since $r \notin \mu$ and $q_{1,i} \in \mu$, it follows that the intersection $\mu^\tau \cap L_1$ contains fewer points than $\mu \cap L_1$. Thus, repeating this process finitely many times, we can find a marking disjoint from L_1 and equivalent to $\bar{\mu}$. It remains to show that if $\bar{\mu}$ and $\bar{\mu}'$ are disjoint from L_1 , then $\bar{\mu} \sim \bar{\mu}'$.

Assume that $\bar{\mu}$ is disjoint from L_1 . Let τ be a transposition switching a pair of distinct nodes in $C_0^{\text{sing}} \setminus L_1$. It is sufficient to prove that $\bar{\mu} \sim \bar{\mu}^\tau$. Indeed, such transpositions generate the full symmetric group of nodes $C_0^{\text{sing}} \setminus L_1$. Therefore, if $\bar{\mu}'$ is disjoint from L_1 , then $\bar{\mu}' = \bar{\mu}^\sigma$ for some permutation σ of $C_0^{\text{sing}} \setminus L_1$, and hence $\bar{\mu} \sim \bar{\mu}'$. Let us prove that $\bar{\mu} \sim \bar{\mu}^\tau$.

Assume first, that there exist indices $i, j, k > 1$ such that τ switches the nodes $q_{i,j}$ and $q_{i,k}$. Let τ_j be the transposition switching $q_{i,j}$ and $q_{i,1}$, and τ_k the transposition switching $q_{i,k}$ and $q_{i,1}$. Since $q_{1,j}, q_{1,k} \notin \mu$, it follows that $\bar{\mu}$ is equivalent to $\bar{\mu}^{\tau_j \tau_k \tau_j}$. However, $\tau_j \tau_k \tau_j = \tau$, which implies the claim. Assume that there are no i, j, k as above. Then there exist four distinct indices $i, j, k, l > 1$ such that τ switches $q_{i,j}$ and $q_{k,l}$. Let τ_i be the transposition switching $q_{i,j}$ and $q_{i,l}$, and τ_l be the transposition switching $q_{i,l}$ and $q_{k,l}$. Then $\bar{\mu} \sim \bar{\mu}^{\tau_i \tau_l \tau_i}$ by the first case considered above. However, $\tau = \tau_i \tau_l \tau_i$, and we are done. \square

Proof of Lemma 7.3. — As before, we denote the irreducible components of C_0 by L_1, \dots, L_d , and set $q_{i,j} := L_i \cap L_j$. It is sufficient to prove the lemma for similar markings. Set $\mu := (p_1, \dots, p_\delta)$, and let τ be the transposition for which $\mu^\tau = (p'_1, \dots, p'_\delta)$. Without loss of generality we may assume that $q_{1,2} \notin \mu$, and τ is the transposition switching $q_{1,3}$ and $q_{2,3}$. Set $\delta' := \binom{d}{2} - 1$, and pick a δ' -marking $\bar{\mu}' = (p_1, \dots, p_{\delta'})$ extending μ such that

$\mu' = C_0^{\text{sing}} \setminus \{q_{1,2}\}$. It is enough to show that there exists an irreducible component U of $U_{d,\delta'}$ containing both $[C_0; p_1, \dots, p_{\delta'}]$ and $[C_0; \tau(p_1), \dots, \tau(p_{\delta'})]$. Indeed, consider the natural forgetful map $\psi: U_{d,\delta'} \rightarrow U_{d,\delta}$. Then $\psi(U) \subseteq U_{d,\delta}$ is an irreducible subvariety that contains both $\mathbf{C}_0 := [C_0; p_1, \dots, p_\delta]$ and $\mathbf{C}_0^\tau := [C_0; \tau(p_1), \dots, \tau(p_\delta)]$, and hence \mathbf{C}_0 and \mathbf{C}_0^τ belong to the same irreducible component of $U_{d,\delta}$, as asserted.

Let us show that U as above exists. To simplify the notation, we assume that $\delta = \delta'$ and reorder the points so that

$$(p_1, \dots, p_\delta) = (q_{1,3}, \dots, q_{1,d}, q_{2,3}, \dots, q_{2,d}, q_{3,4}, \dots, q_{3,d}, \dots, q_{d-1,d}).$$

In particular, $q_{1,3} = p_1$ and $q_{2,3} = p_{d-1}$. Then τ is the transposition on C_0^{sing} switching p_1 and p_{d-1} . Notice that the action of τ on $\{\mathbf{C}_0, \mathbf{C}_0^\tau\}$ extends to action on $U_{d,\delta}$ that switches the first marked point with the $(d-1)$ -st, and hence induces a permutation on the set of the irreducible components of $U_{d,\delta}$. Let U be the irreducible component containing \mathbf{C}_0 , and let us show that $\mathbf{C}_0^\tau \in U$.

By Proposition 6.3, the variety $U_{d,\delta}$ has pure dimension $2d+1$. In particular, $\dim(U) = 2d+1$. Thus, by dimension count, a general $\mathbf{C} := [C; r_1, \dots, r_\delta] \in U$ is a union of a general conic Q and $d-2$ lines L'_3, \dots, L'_d in general position. Furthermore, $Q \cap L'_i = \{r_{i-2}, r_{d+i-4}\}$ for all $3 \leq i \leq d$, and $L'_3 \cap L'_i = \{r_{2d+i-7}\}$ for all $4 \leq i \leq d$. We fix general Q, L'_4, \dots, L'_d as above, and set $D := (\cup_{i \geq 4} L'_i) \cup Q$.

Consider the dense open subset $W \subset (\mathbf{P}^2)^*$ parameterizing the lines $L' \subset \mathbf{P}^2$ that intersect D transversally and such that $L' \cap D^{\text{sing}} = \emptyset$. Let \tilde{W} be the incidence variety parameterizing tuples $[L'; r'_1, r'_{d-1}, r'_{2d-3}, \dots, r'_{3d-7}]$, where r'_1, r'_{d-1} are the points of intersection of L' with Q , and r'_{2d+i-7} is the point of intersection of L' with L'_i for all $4 \leq i \leq d$. Then $\tilde{W} \rightarrow W$ is an étale covering of degree two. Furthermore, there is a natural morphism $\iota: \tilde{W} \rightarrow U_{d,\delta}$, that maps $[L'; r'_1, r'_{d-1}, r'_{2d-3}, \dots, r'_{3d-7}]$ to

$$[L' \cup D; r'_1, r_2, \dots, r_{d-2}, r'_{d-1}, r_d, \dots, r_{2d-4}, r'_{2d-3}, \dots, r'_{3d-7}, r_{3d-6}, \dots, r_\delta].$$

It is easy to see that \tilde{W} is irreducible. Indeed, consider the natural projection $\tilde{W} \rightarrow Q^2$ mapping $[L'; r'_1, r'_{d-1}, r'_{2d-3}, \dots, r'_{3d-7}]$ to (r'_1, r'_{d-1}) . Since the line L' is uniquely determined by the pair of points (r'_1, r'_{d-1}) , the variety \tilde{W} is mapped bijectively onto its image in Q^2 , and hence the image has dimension

$$\dim(\tilde{W}) = \dim(W) = \dim((\mathbf{P}^2)^*) = 2 = \dim(Q^2).$$

Thus, $\tilde{W} \rightarrow Q^2$ is dominant, and bijective onto its image. Hence \tilde{W} is irreducible since so is Q^2 . To finish the proof, it remains to notice that both \mathbf{C} and \mathbf{C}^τ belong to $\iota(\tilde{W})$, and hence belong to the same irreducible component of $U_{d,\delta}$. But $\mathbf{C} \in U$ and $\mathbf{C}^\tau \in U^\tau$, which implies $U = U^\tau$. Thus, $\mathbf{C}_0^\tau \in U^\tau = U$, and we are done. \square

Acknowledgements

We are grateful to Michael Temkin for helpful discussions and to an anonymous referee for many insightful comments that helped us to improve the presentation.

Publisher's Note

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

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/>.

REFERENCES

- [ACGS20] D. ABRAMOVICH, Q. CHEN, M. GROSS and B. SIEBERT, Decomposition of degenerate Gromov-Witten invariants, *Compos. Math.*, **156** (2020), 2020–2075.
- [ACP15] D. ABRAMOVICH, L. CAPORASO and S. PAYNE, The tropicalization of the moduli space of curves, *Ann. Sci. Éc. Norm. Supér. (4)*, **48** (2015), 765–809.
- [AC81a] E. ARBARELLO and M. CORNALBA, Footnotes to a paper of Beniamino Segre, *Math. Ann.*, **256** (1981), 341–362. <https://doi.org/10.1007/BF01679702>.
- [AC81b] E. ARBARELLO and M. CORNALBA, Su una congettura di Petri, *Comment. Math. Helv.*, **56** (1981), 1–37.
- [BIMS15] E. BRUGALLÉ, I. ITENBERG, G. MIKHALKIN and K. SHAW, Brief introduction to tropical geometry, in *Proceedings of the Gökova Geometry-Topology Conference 2014 (Gökova Geometry/Topology Conference (GGT), Gökova)*, pp. 1–75, 2015.
- [BM09] E. BRUGALLÉ and G. MIKHALKIN, Floor decompositions of tropical curves: the planar case, in *Proceedings of Gökova Geometry-Topology Conference 2008 (Gökova Geometry/Topology Conference (GGT), Gökova)*, pp. 64–90, 2009.
- [BPR13] M. BAKER, S. PAYNE and J. RABINOFF, On the structure of non-Archimedean analytic curves, in *Tropical and Non-Archimedean Geometry*, *Contemp. Math.*, vol. 605, pp. 93–121, Am. Math. Soc., Providence, 2013.
- [CC99] L. CHIANTINI and C. CILIBERTO, On the Severi varieties on surfaces in \mathbf{P}^3 , *J. Algebraic Geom.*, **8** (1999), 67–83.
- [CCUW20] R. CAVALIERI, M. CHAN, M. ULIRSCH and J. WISE, A moduli stack of tropical curves, *Forum Math. Sigma*, **8** (2020), e23.
- [CHL06] H.-J. CHIANG-HSIEH and J. LIPMAN, A numerical criterion for simultaneous normalization, *Duke Math. J.*, **133** (2006), 347–390.
- [CHT22a] K. CHRIST, X. HE and I. TYOMKIN, Degeneration of curves on some polarized toric surfaces, *J. Reine Angew. Math.*, **787** (2022), 197–240.

- [CHT22b] K. CHRIST, X. HE and I. TYOMKIN, Irreducibility of Severi varieties and Hurwitz schemes in small characteristic, 2022, in preparation.
- [Del85] P. DELIGNE, Le lemme de Gabber, *Astérisque*, **127** (1985), 131–150, Seminar on arithmetic bundles: the Mordell conjecture (Paris, 1983/84).
- [dJ96] A. J. de JONG, Smoothness, semi-stability and alterations, *Publ. Math. Inst. Hautes Études Sci.*, **83** (1996), 51–93.
- [dJHS11] A. J. de JONG, X. HE and J. M. STARR, Families of rationally simply connected varieties over surfaces and torsors for semisimple groups, *Publ. Math. Inst. Hautes Études Sci.*, **114** (2011), 1–85.
- [DM69] P. DELIGNE and D. MUMFORD, The irreducibility of the space of curves of given genus, *Publ. Math. Inst. Hautes Études Sci.*, **36** (1969), 75–109.
- [DPT80] M. DEMAZURE, H. C. PINKHAM and B. TEISSIER (eds.), *Séminaire sur les Singularités des Surfaces*, Lecture Notes in Mathematics, vol. 777, pp. 1976–1977, Springer, Berlin, 1980, held at the Centre de Mathématiques de l'École Polytechnique, Palaiseau.
- [EC85] F. ENRIQUES and O. CHISINI, *Lezioni sulla teoria geometrica delle equazioni e delle funzioni algebriche. 1. Vol. I, II*, Collana di Matematica [Mathematics Collection], vol. 5, Zanichelli Editore S.p.A., Bologna, 1985, Reprint of the 1915 and 1918 editions.
- [Ful69] W. FULTON, Hurwitz schemes and irreducibility of moduli of algebraic curves, *Ann. Math. (2)*, **90** (1969), 542–575.
- [GM07] A. GATHMANN and H. MARKWIG, The numbers of tropical plane curves through points in general position, *J. Reine Angew. Math.*, **602** (2007), 155–177.
- [Har86] J. HARRIS, On the Severi problem, *Invent. Math.*, **84** (1986), 445–461.
- [HM82] J. HARRIS and D. MUMFORD, On the Kodaira dimension of the moduli space of curves, *Invent. Math.*, **67** (1982), 23–88, with an appendix by William Fulton.
- [Hur91] A. HURWITZ, Über Riemann'sche Flächen mit gegebenen Verzweigungspunkten, *Math. Ann.*, **39** (1891), 1–60.
- [Ill71] L. ILLUSIE, *Complexe cotangent et déformations. I*, Lecture Notes in Mathematics, vol. 239, Springer, Berlin, 1971.
- [Kle82] F. KLEIN, *Über Riemann's Theorie der algebraischen Functionen und ihrer Integrale*, Teubner, Leipzig, 1882.
- [Knu83] F. F. KNUDSEN, The projectivity of the moduli space of stable curves. III. The line bundles on $M_{g,n}$, and a proof of the projectivity of $\overline{M}_{g,n}$ in characteristic 0, *Math. Scand.*, **52** (1983), 200–212.
- [Kon95] M. KONTSEVICH, Enumeration of rational curves via torus actions, in *The Moduli Space of Curves*, Progr. Math., vol. 129, *Texel Island, 1994*, pp. 335–368, Birkhäuser, Boston, 1995.
- [KS13] S. L. KLEIMAN and V. V. SHENDE, On the Göttsche threshold, in *A Celebration of Algebraic Geometry*, Clay Math. Proc., vol. 18, pp. 429–449, Am. Math. Soc., Providence, 2013, with an appendix by Ilya Tyomkin.
- [Lan20] L. LANG, Monodromy of rational curves on toric surfaces, *J. Topol.*, **13** (2020), 1658–1681.
- [Mik05] G. MIKHALKIN, Enumerative tropical algebraic geometry in \mathbf{R}^2 , *J. Am. Math. Soc.*, **18** (2005), 313–377.
- [Mor79] S. MORI, Projective manifolds with ample tangent bundles, *Ann. Math. (2)*, **110** (1979), 593–606.
- [Ran19] D. RANGANATHAN, Logarithmic Gromov-Witten theory with expansions, *Algebr. Geom.*, **9** (2022), 714–761. <https://doi.org/10.14231/AG-2022-022>.
- [Rau17] J. RAU, A first expedition to tropical geometry, 2017, <https://www.math.uni-tuebingen.de/user/jora/downloads/FirstExpedition.pdf>.
- [Ros52] M. ROSENBLIETH, Equivalence relations on algebraic curves, *Ann. Math. (2)*, **56** (1952), 169–191.
- [Sev21] F. SEVERI, *Vorlesungen über algebraische Geometrie, Anhang F*, Teubner, Leipzig, 1921.
- [Sta20] The Stacks project authors, The stacks project, 2020, <https://stacks.math.columbia.edu>.
- [Tyo07] I. TYOMKIN, On Severi varieties on Hirzebruch surfaces, *Int. Math. Res. Not.*, **23** (2007), rnm109.
- [Tyo12] I. TYOMKIN, Tropical geometry and correspondence theorems via toric stacks, *Math. Ann.*, **353** (2012), 945–995.
- [Tyo13] I. TYOMKIN, On Zariski's theorem in positive characteristic, *J. Eur. Math. Soc.*, **15** (2013), 1783–1803.
- [Tyo14] I. TYOMKIN, An example of a reducible Severi variety, in *Proceedings of the Gökova Geometry-Topology Conference 2013 (Gökova Geometry/Topology Conference (GGT), Gökova)*, pp. 33–40, 2014.
- [Zar82] O. ZARISKI, Dimension-theoretic characterization of maximal irreducible algebraic systems of plane nodal curves of a given order n and with a given number d of nodes, *Am. J. Math.*, **104** (1982), 209–226.

K. C.
Department of Mathematics,
Ben-Gurion University of the Negev,
P.O. Box 653,
Be'er Sheva, 84105, Israel
and
Institute of Algebraic Geometry,
Leibniz University Hannover,
Welfengarten 1,
30167, Hanover, Germany
kchrist@math.uni-hannover.de

X. H.
Yau Mathematical Sciences Center, Ningzhai,
Tsinghua University,
Haidian District,
Beijing, 100084, China
and
Einstein Institute of Mathematics,
The Hebrew University of Jerusalem,
Giv'at Ram,
Jerusalem, 91904, Israel
xianghe@mail.tsinghua.edu.cn

I. T.
Department of Mathematics,
Ben-Gurion University of the Negev,
P.O. Box 653,
Be'er Sheva, 84105, Israel
tyomkin@math.bgu.ac.il

Manuscrit reçu le 5 octobre 2021
Version révisée le 20 septembre 2022
Manuscrit accepté le 15 novembre 2022
publié en ligne le 1 décembre 2022.