

RESEARCH ARTICLE

Logarithmic Donaldson–Thomas theory

Davesh Maulik¹ and Dhruv Ranganathan²

¹Department of Mathematics Massachusetts Institute of Technology, Cambridge, MA, USA;
E-mail: maulik@mit.edu (corresponding author).

²Department of Pure Mathematics & Mathematical Statistics University of Cambridge, Cambridge, UK;
E-mail: dr508@cam.ac.uk (corresponding author).

Received: 19 September 2022; **Revised:** 14 December 2023; **Accepted:** 5 January 2024

2020 Mathematics Subject Classification: 14A21, 14N35, 14T90

Abstract

Let X be a smooth and projective threefold with a simple normal crossings divisor D . We construct the Donaldson–Thomas theory of the pair $(X|D)$ enumerating ideal sheaves on X relative to D . These moduli spaces are compactified by studying subschemes in expansions of the target geometry, and the moduli space carries a virtual fundamental class leading to numerical invariants with expected properties. We formulate punctual evaluation, rationality and wall-crossing conjectures, in parallel with the standard theory. Our formalism specializes to the Li–Wu theory of relative ideal sheaves when the divisor is smooth and is parallel to recent work on logarithmic Gromov–Witten theory with expansions.

Contents

1	Flavours of tropicalization	6
2	Tropical degenerations and algebraic transversality	13
3	Moduli of target expansions	25
4	Moduli of ideal sheaves	34
5	Logarithmic DT theory	45
6	First examples	50
7	Transverse limits and generically expanded targets	56
	References	62

Introduction

A basic technique in enumerative geometry is to degenerate from a smooth variety to a singular one, whose irreducible components may be easier to analyze. In Gromov–Witten theory, this is an essential tool in the subject, with both symplectic and algebraic incarnations [33, 43, 44, 45]. Associated to a smooth pair (X, D) , *relative* Gromov–Witten theory studies maps with fixed tangency along D and uses these to study degenerations of a smooth variety to two smooth components meeting transversely. The key geometric idea is that the target is dynamic, expanding along D to prevent components mapping to the divisor.

Logarithmic Gromov–Witten theory was developed [2, 17, 27] to handle more general logarithmically smooth pairs (X, D) , for instance, where D is a normal crossings boundary divisor; these occur as components of degenerations with more complicated combinatorics. In this approach, the target is no longer dynamic and, instead, logarithmic structures are used to maintain a form of transversality in limits.

When X is a threefold, there are sheaf-theoretic approaches to enumerating curves, developed by Donaldson–Thomas and Pandharipande–Thomas [21, 73, 64]. In the case of a smooth pair (X, D) , there is again a powerful relative theory, developed by Li–Wu [47], which is one of the central tools of the subject. The goal of this paper is to construct the Donaldson–Thomas and Pandharipande–Thomas theories of a pair (X, D) , where D is a simple normal crossings divisor; these will be referred to here as the logarithmic DT and PT theory of $(X|D)$.

One challenge in building such a theory is that there is no clear analogue of the unexpanded formalism of [27]. The moduli space of prestable curves and its universal family are themselves normal crossings pairs, so the logarithmic mapping space provides a starting point. The Hilbert scheme exhibits no such structure, and there is not, as yet, a good definition of a logarithmic structure on an ideal sheaf. Instead, our strategy is closer to the original expanded formalism of [47]. In an upcoming paper [56], we complete the parallel to logarithmic Gromov–Witten (GW) theory, with its associated degeneration formalism [3, 4, 68], and formulate the logarithmic version of the correspondence between Gromov–Witten and DT/PT theories.

0.1. Main results

We work over the complex numbers. Let X be a smooth and projective threefold equipped with a simple normal crossings divisor D . We further assume the intersection of any subset of irreducible components of D is connected. Our main results are described below. Precise definitions and statements can be found in the main text. Fix a curve class β in $H_2(X)$ and an integer n for the holomorphic Euler characteristic. The main construction of the paper is a moduli problem $DT_{\beta,n}(X|D)$ associated to these data.

Theorem A. *The moduli problem $DT_{\beta,n}(X|D)$ is representable by a proper Deligne–Mumford stack. It compactifies the moduli space of ideal sheaves on X relative to D with numerical invariants β and n and is equipped with universal diagram*

$$\begin{array}{ccc} \mathcal{Z} & \xrightarrow{\quad} & \mathcal{Y} \longrightarrow X \\ & \searrow & \swarrow \\ & DT_{\beta,n}(X|D) & \end{array}$$

The fibers of $DT_{\beta,n}(X|D)$ parameterize relative ideal sheaves on expansions of X along D . The space $DT_{\beta,n}(X|D)$ carries a perfect obstruction theory and a virtual fundamental class with expected properties.

The structure of our result mirrors the structure of relative DT theory. Namely, we first construct a moduli stack parametrizing allowable target expansions; the moduli space of stable relative ideal sheaves will then consist of ideal sheaves on fibers of the universal family, satisfying a certain stability condition. However, in our setting, both the classifying stack and the correct notion of stability are more subtle, and ideas from tropical geometry are crucial to finding the correct formulation. The approach is outlined in the final sections of this introduction.

Numerical invariants are obtained by integration against the virtual fundamental class. Primary and descendant fields are given by the Chern characters of the universal ideal sheaf, as in the standard theory [51]. The Hilbert scheme of points on the divisor components, relative to the induced boundary divisor, can be analogously compactified.

Theorem B. *Each DT moduli space $DT_{\beta,n}(X|D)$ is equipped with an evaluation morphism*

$$ev: DT_{\beta,n}(X|D) \rightarrow Ev(D).$$

The space $Ev(D)$ is a compactification of the Hilbert scheme of points on the smooth locus of D . DT invariants are defined by pairing the virtual class with cohomology classes pulled back from $Ev(D)$ and primary and descendant fields.

Once the numerical invariants are defined, we formulate the basic conjectures of the sheaf theory side of the subject in our general logarithmic setting. These conjectures include the precise evaluation of the punctual series and the rationality of the normalized generating function for primary DT invariants. These are collected in Section 5.4.

Formal similarities between the map and sheaf sides give a natural generalization of the relative GW/DT correspondence to normal crossings geometries and motivate a study of the correspondence via normal crossings degenerations [52, 53]. A prototype version of this theory plays a key role in the study of the GW/Pairs correspondence of Pandharipande and Pixton [63]. These ideas are a central motivation for our study and are developed in our follow-up [56].

The results complete an exact parallel to [68, Theorem A], where an expanded version of logarithmic Gromov–Witten theory is constructed. As in that paper, a notable feature of our moduli spaces is that they are not unique. Rather, there is a combinatorial choice in constructing the stack of expansions, which leads to an infinite collection of moduli spaces with the desired properties, naturally organized into an inverse system. There are natural compatibilities between virtual classes so that the numerical invariants are independent of this choice. This is very similar to the combinatorial choice in studying families of degenerating abelian varieties and constructing toroidal compactifications. While the geometry is a little different, the nonuniqueness arises in an identical fashion, via a system of polyhedral structures on a fanlike object.

However, there is a critical difference between this paper and [68], which creates an additional layer of complexity. In Gromov–Witten theory with expansions, we *start* with the existing mapping stack of logarithmic maps to a fixed target and then apply logarithmic modifications to construct the expanded moduli space. As a result, the output is automatically proper. In our case, there is no unexpanded space to start with, and we are forced to construct the entire system of modifications directly. Notably, properness of the moduli problem needs to be understood in a new way. In order to do this, we first provide a tropical algorithm for finding transverse limits for families of subschemes of X and then use this tropical algorithm to guess the correct moduli problem, after which we establish algebraicity, boundedness and somewhat tautologically, properness. In the following subsections, we give a detailed outline of our approach.

0.2. Transversality

Let $Z \subset X$ be a subscheme with ideal sheaf \mathcal{I}_Z . We are interested in subschemes that intersect D in its smooth locus, with the property no one-dimensional components or embedded points of Z lie in D . Algebraically, this is the condition that the map

$$\mathcal{I}_Z \otimes \mathcal{O}_D \rightarrow \mathcal{O}_D$$

is injective. We refer to subschemes satisfying this condition as *algebraically transverse subschemes*. The locus of algebraically transverse subschemes is a nonproper open subscheme of the Hilbert scheme of X . We aim to find a compactification of this moduli problem with the prescription that the universal subscheme continues to be algebraically transverse in an appropriate sense.

0.3. Limits from tropicalization, after Tevelev

In order to achieve transversality for limits of families, the scheme X must be allowed to break. The degenerations here are built from tropical geometry, using an elegant argument due to Tevelev, based on Kapranov’s *visible contours* [34, 71].

The cone over the dual complex of D is denoted Σ and can be identified with a union of faces inside $\mathbb{R}_{\geq 0}^k$. Let Σ^+ be product of the cone $\mathbb{R}_{\geq 0}$ with Σ . We view these as the fans associated to the toric stacks \mathbb{A}^k and $\mathbb{A}^k \times \mathbb{A}^1$, respectively. Given an injection of cone complexes $\Delta \hookrightarrow \Sigma^+$, the toric dictionary gives rise to a modification

$$\mathcal{Y} \rightarrow X \times \mathbb{A}^1 \rightarrow \mathbb{A}^1.$$

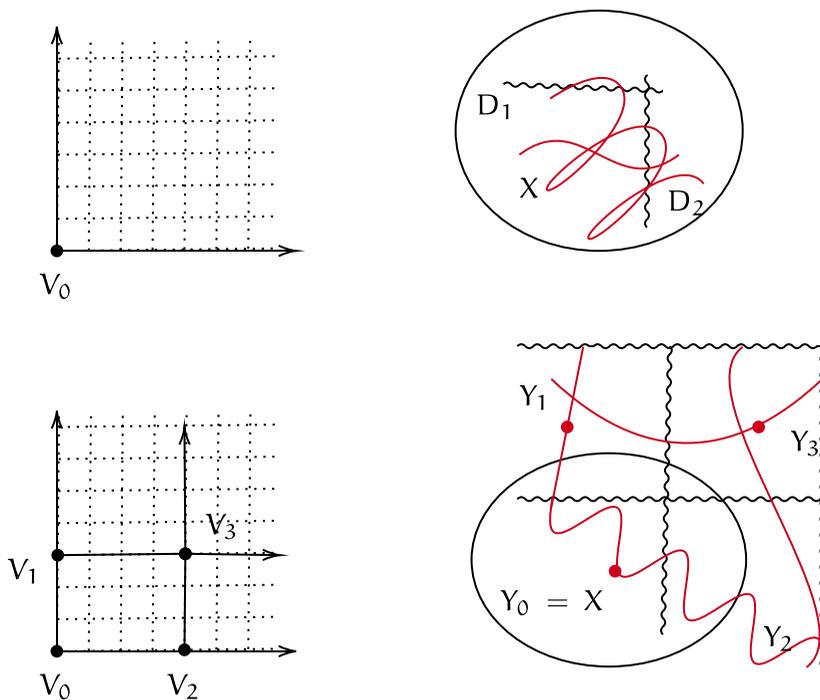


Figure 1. This captures how the authors visualize expansions and subschemes in them. The two pictures on the left are typical tropicalizations of subschemes. The top left is the tropicalization of a subscheme in the interior of the moduli problem, while the bottom left is a more complicated one. The pictures on the left are outputs of the Tevelev theorem and determine expansions of X – a gluing together of projective bundles over the strata of X . The reader can notice, for example in the second row, a bijection between components of the expansion and the vertices of the graph on the left. The circled component on the right, in both pictures, is the main component, that is, X itself. The subschemes drawn, indicated in red, can have singularities, or embedded components – three embedded points can be seen on the bottom right. However (i) the subscheme must be disjoint from the codimension 2 strata, that is, the corners, and (ii) the embedded points must lie in the interiors of irreducible components.

Geometrically, the expansion is obtained from the constant family by performing birational modifications to the strata of D in the special fiber and passing to an open subscheme. When D is a smooth divisor, this essentially recovers the class of expansions considered by Li and Li–Wu [44, 47], and we explain this in detail in Section 6.7. In the general case, they recover the class of targets in [68].

We can now apply Tevelev’s approach. Given an algebraically transverse family of subschemes \mathcal{Z}_η over a \mathbb{C}^* , the tropicalization of \mathcal{Z}_η is a subset of Σ^+ . This subset can be given the choice of a fan Δ contained in Σ^+ . For an appropriate choice of fan structure, this produces a degeneration \mathcal{Y} of X over \mathbb{A}^1 . The flat limit of \mathcal{Z}_η in this degeneration is algebraically transverse to the strata of \mathcal{Y} . This was proved by Tevelev when $(X|D)$ is toric, and we make the necessary extensions in the main text. These limits have strong uniqueness and functoriality properties, making them appropriate for constructing moduli. In Figure 1, we caricature a subscheme in $(X|D)$ and a potential expansion of $(X|D)$ with a limiting subscheme, together with the corresponding tropical data.

Precedent for building moduli spaces, particularly in contexts adjacent to toric geometry, via Tevelev’s work is provided by Hacking–Keel–Tevelev, and the method has been used in Gromov–Witten theory before; see [29, 38, 67].

0.4. The universal tropical expansion

By axiomatizing the output of Tevelev's argument, we propose a class of *ideal sheaves on X relative to D* for the DT moduli problem. These are subschemes of expansions of X along D that are transverse to the strata. In order to construct a global moduli problem, we identify an Artin stack that encodes the possible expansions of X that could arise from Tevelev's procedure.

The discussion above predicts the one-parameter degenerations of the target, but subtleties arise in extending them over higher-dimensional bases, having to do with flattening the universal degeneration. We tackle this by first studying an appropriate tropical moduli problem, using recent work that identifies a category of certain locally toric Artin stacks with purely combinatorial objects [16].

The outcome of the tropical study is a system of moduli stacks of universal expansions, related to each other by birational transformations, and organized into an inverse system. Each element in this system is 'good enough' to function as a stack of expansions for our moduli problem, but there is typically no distinguished choice. The system depends only on the combinatorics of the boundary divisor $D \subset X$. Both the moduli space and its universal family have this structure and compatible choices give rise to a universal degeneration.

0.5. Moduli space of stable relative ideal sheaves

After fixing a stack of expansions, we define a notion of DT stability for ideal sheaves on X relative to D. An important subtlety appears when considering *tube subschemes*, namely subschemes in a component of an expansion that are pulled back from a surface in its boundary. In relative DT theory, these are ruled out by stability; in our setting, they are forced on us by the combinatorial algorithms in an analogous fashion to how *trivial bubbles* arise in the stable maps geometry. In any fiber of the universal family of the stack of expansions, there are distinguished irreducible components, denoted tube components; the DT stability condition we impose is that these are precisely the components which host tube subschemes. We show this defines a moduli problem with the expected properties. The transversality hypotheses guarantee that the morphism to the moduli stack of expansions

$$\mathrm{DT}_{\beta,n}(X|D) \rightarrow \mathrm{Exp}(X|D)$$

has a perfect obstruction theory and consequently a virtual class. This establishes an appropriate DT theory for the pair $(X|D)$. In Section 6.7, we explain how, in hindsight, the tube geometry above can be artificially introduced into Li–Wu's theory and why it can be avoided in that case.

0.6. Pairs and so on

We have chosen to focus on the moduli theory of ideal sheaves in this paper, but the methods appear to be adaptable to other settings. In particular, one can define logarithmic stable pair invariants by replacing the Hilbert scheme with the stable pair moduli spaces of Pandharipande and Thomas. The stable pair adaptations are recorded in Remark 4.6.2. It seems reasonable to hope for further applications. For instance, it may be possible to rederive logarithmic Gromov–Witten theory, relying on target expansions from the very beginning. In another direction, the logarithmic theory of quasimaps has only been treated in the smooth pair case [11].

0.7. Toroidal embeddings

The expectation is that the theory set up in this paper can be extended to any logarithmically smooth target, without either the simple normal crossing or connectivity restrictions placed on D. The arguments in the present paper carry over with cosmetic changes to treat generalized Deligne–Faltings logarithmic structures [2]. This includes all singular toric varieties. A more delicate modification of the combinatorics can likely be used to treat divisor geometries with disconnected intersections. These two variants, together with the case where D has a self-intersecting component, will be addressed elsewhere. Logarithmically étale descent and virtual birational invariance techniques are likely to play a role; see [6, 9].

0.8. Further directions and recent progress

We mention briefly some natural directions to pursue with the theory constructed here. In a recent sequel [56] to this paper, we develop a degeneration formalism, generalizing that of [47], and parallel to [68]. We also develop a logarithmic version of the GW/Pairs correspondence, and show it is compatible with normal crossings degenerations. With this in place, our subsequent goal is to extend the inductive strategy of Pandharipande–Pixton and prove the GW/Pairs correspondence for a broader class of threefolds, that is, varieties which are not easily studied by double-point degenerations. For example, one can envisage a proof GW/Pairs for threefolds admitting an algebro-geometric SYZ fibration, that is, a normal crossings degeneration to a union of rational varieties. A natural class of examples comes from taking zero loci of sections of toric vector bundles.

In another direction, the formalism of relative DT theory (in cohomology and K-theory) interacts well with the representation-theoretic structure on the Hilbert scheme of points on a surface [54], and we expect that our logarithmic theory will extend this circle of ideas. In a similarly speculative vein, logarithmic Gromov–Witten invariants in genus 0 are related to the symplectic cohomology of the open variety $X \setminus D$, which is the natural replacement for quantum cohomology for open geometries. It would be interesting to examine whether logarithmic DT theory can be related to the symplectic cohomology of Hilbert schemes of points on open surfaces.

Our focus here is on logarithmic moduli spaces of subschemes of dimension at most 1, due to the applications for DT theory. The methods here make use of the simplicity of the transversality condition for one-dimensional subschemes, and the simpler combinatorics in this case. However, in the time since this paper first appeared on arXiv, further progress has been made. In recent work, Kennedy–Hunt [41] proposes a general logarithmic Quot scheme with no constraints on the dimension of the support, building on the techniques introduced here.

Outline of paper

We briefly outline the sections of this paper. In Section 1, we review some basic constructions and results from tropical geometry. In Section 2, we give the tropical algorithm for constructing algebraically transverse flat limits of subschemes. In Section 3, we construct the tropical moduli spaces of expansions as well as their geometric counterparts. In Section 4, we define stable relative ideal sheaves and show their moduli functor is represented by a proper Deligne–Mumford stack. In Section 5, we study the virtual structure on the moduli space, define logarithmic DT invariants and state the basic conjectures. In Section 6, we give a handful of simple examples, demonstrating the basic theory. In Section 7, we complete the proof of the valuative criterion of properness, initiated in Section 2, to deal with the case where the generic fiber is expanded.

Background and conventions

We have put some effort into minimizing the amount of logarithmic geometry that is explicitly used in this paper, and there is nothing that we use beyond [5, Sections 1–5]. We do use the combinatorics of cone complexes and cone spaces heavily and refer the reader to [16, Sections 2 & 6]. Logarithmic schemes and stacks that appear will be fine and saturated unless otherwise specified, locally of finite type and over the complex numbers.

1. Flavours of tropicalization

We require some elementary notions from logarithmic geometry, and a reference that is well suited to our point of view is [5, Sections 3–5]. There are a few different ways in which tropicalizations arise in logarithmic geometry, and we recall these for the reader.

1.1. Cone complexes and their morphisms

We start with the building blocks of toric geometry. A *polyhedral cone with integral structure* (σ, M) is a topological space σ equipped with a finitely generated abelian group M of continuous real-valued functions from σ to \mathbb{R} such that the evaluation

$$\sigma \rightarrow \text{Hom}(M, \mathbb{R})$$

is a homeomorphism onto a strongly convex polyhedral cone. If this cone is *rational* with respect to the dual lattice of M , then we say that (σ, M) is rational. We define the lattice of integral points of σ by taking the preimage of the dual lattice of M under the evaluation map. The set of elements of M that are nonnegative on σ form a monoid S_σ referred to as the *dual monoid*. The cone σ is recovered as the space of monoid homomorphisms $\text{Hom}(S_\sigma, \mathbb{R}_{\geq 0})$.

Henceforth, a *rational polyhedral cone with integral structure* will be referred to as a *cone*.

Definition 1.1.1 (Cone complexes). A *rational polyhedral cone complex* is a topological space that is presented as a colimit of a partially ordered set of cones, where all arrows are given by isomorphisms onto proper faces. A *morphism* of cone complexes is a continuous map

$$\Sigma' \rightarrow \Sigma$$

such that (i) the image of every cone in Σ' is contained in a cone of Σ and (ii) the restriction of the map to any cone of Σ' is given by an integer linear map.

Cone complexes are nearly identical to the *fans* considered in toric geometry [24]. The key differences are that they do not come equipped with a global embedding into a vector space, and two cones can intersect along a union of faces. However, we will typically restrict to the case where an intersection of two cones is a face of each, so the main thing to keep in mind is the lack of a global embedding. Concretely, the reader may keep in mind that toric fans of \mathbb{P}^1 and $\mathbb{A}^2 \setminus \{(0, 0)\}$ are isomorphic as cone complexes.

A cone complex is *smooth* if every cone is isomorphic to a standard orthant with its canonical integral structure. We record two combinatorial notions associated to morphisms between cone complexes.

Definition 1.1.2 (Flat maps and reduced fibers). Let Σ be a smooth cone complex, and let $\pi : \Sigma' \rightarrow \Sigma$ be a morphism of cone complexes. Then π is *flat* if the image of every cone of Σ' is a cone of Σ . A flat map is said to have *reduced fibers* if for every cone σ' of Σ' with image σ , the image of the lattice of σ' is equal to the lattice of σ .

The terminology is compatible with the identically named geometric notions, when applied to toric maps [7, Section 4 & 5].

1.2. Tropicalization for the target

A basic fact from the theory of logarithmic schemes is that every Zariski logarithmically smooth scheme gives rise to a cone complex. We take a moment to unpack this statement in concrete terms. Let $(X|D)$ be a smooth scheme of finite type over \mathbb{C} equipped with a simple normal crossings divisor D . Assume that the intersections of irreducible components of D are always connected. The presence of D gives X the structure of a logarithmically smooth scheme. We let X° be the complement of the divisor, where the logarithmic structure is trivial.

We unwind the definition of the logarithmic structure in this case for the benefit of the reader. The components of D give rise to a distinguished class of functions in the structure sheaf. For an open set $U \subset X$, we record the values of the logarithmic structure sheaf and the characteristic monoid sheaf:¹

¹These may be regarded as sheaves in either the Zariski or étale topology.

$$M_X(U) = \{f \in \mathcal{O}_X(U) : f \text{ is invertible on } U \setminus D\} \quad \text{and} \quad \overline{M}_X(U) = M_X(U)/\mathcal{O}_X^*(U).$$

Locally at each point $x \in X$, there are functions, canonical up to multiplication by a unit, cutting out the irreducible components of D passing through x . In particular, the stalk of the characteristic monoid sheaf at x is naturally identified with \mathbb{N}^e , where e is the number the such components.

The dual cones of the stalks of the characteristic monoids of X give rise to a collection of cones, one for each point of X . The generalization maps naturally give rise to gluing morphisms, and these cones form a finite type cone complex. We denote it by Σ_X , refer to it as the *tropicalization* of X or the *cone complex of X* and regard it as a cone complex equipped with an integral structure.

In our case, the tropicalization has a more practical description. Given the pair (X, D) , let d be the number of irreducible components of D and enumerate these components D_1, \dots, D_d . Each k -dimensional face of the cone $\mathbb{R}_{\geq 0}^d$ is spanned by rays v_{i_1}, \dots, v_{i_k} . Call such a face *relevant to (X, D)* if the corresponding intersection $\overline{D}_{i_1}, \dots, \overline{D}_{i_k}$ is nonempty. Then the tropicalization Σ_X is the union of cones in $\mathbb{R}_{\geq 0}^d$ that are relevant to (X, D) . In the case of a toric variety this construction recovers the fan, as an abstract cone complex. The construction extends to the case of logarithmically smooth schemes [1, 74]. Further details and generalizations may be found in the references [5, 16, 35, 39].

1.3. Subdivisions

Let $(X|D)$ be a simple normal crossings pair with tropicalization Σ_X . For simplicity, we assume that the nonempty intersections of components of D are connected. We will produce target expansions by using subdivisions of the tropicalization.

Definition 1.3.1. A *subdivision* is a cone complex Δ and a morphism of cone complexes

$$\Delta \hookrightarrow \Sigma_X$$

that is injective on the support of Δ and further such that the integral points of the image of each cone $\tau \in \Delta$ are exactly the intersection of the integral points of Σ_X with τ .

Note that this more flexible than the standard definition; the underlying map of sets need not be a bijection because we wish to have the flexibility to discard closed strata after blowing up. When the map on integral points is a bijection we call it a *complete subdivision*.

In toric geometry, subdivisions give rise to possibly nonproper birational models of X . The same is true for $(X|D)$. The cone complex Σ_X associated to D is smooth, and the intersection of any two cones is a face of each. If there are d rays in Σ_X , we may embed it via

$$\Sigma_X \hookrightarrow \mathbb{R}_{\geq 0}^d$$

by mapping each ray in Σ_X isomorphically onto the positive ray on the corresponding axis. The positive orthant in this vector space is the fan associated to the toric variety \mathbb{A}^d with dense torus \mathbb{G}_m^d . The subdivision

$$\Delta \hookrightarrow \Sigma_X \hookrightarrow \mathbb{R}_{\geq 0}^d$$

defines a noncomplete fan with associated toric variety \mathbb{A}_Δ , equipped with an \mathbb{G}_m^d -equivariant birational map $\mathbb{A}_\Delta \rightarrow \mathbb{A}^d$. By passing to quotients, we have a morphism of stacks

$$[\mathbb{A}_\Delta/\mathbb{G}_m^d] \rightarrow [\mathbb{A}^d/\mathbb{G}_m^d].$$

The presence of D gives rise to a tautological morphism $(X|D) \rightarrow [\mathbb{A}^d/\mathbb{G}_m^d]$.

Definition 1.3.2. The *birational model of X associated to the subdivision $\Delta \hookrightarrow \Sigma$* is given by

$$X_\Delta := X \times_{[\mathbb{A}^d/\mathbb{G}_m^d]} [\mathbb{A}_\Delta/\mathbb{G}_m^d].$$

Remark 1.3.3. In practice, we apply this construction to $X \times \mathbb{A}^1$ with the divisor $X \times \{0\} \cup D \times \mathbb{A}^1$. The deformation to the normal cone of a stratum of D is a special case of the construction.

Remark 1.3.4. It is natural to formulate subdivisions in terms of subfunctors of the functor on logarithmic schemes defined by X and by its cone complex. The birational model associated to a subdivision is defined by pulling back the subdivision along the tropicalization map; see [16, 35].

1.4. Tropicalization for subschemes via valuations

In this section, we assume that X is proper, with the exception of Remark 1.5.5 and situations where it becomes active.

Locally at each point on X , the components of the divisor D provide a distinguished set of functions – a subset of coordinates – at that point. The tropicalization of a subscheme is the image of its ‘coordinatewise valuation’, in these coordinates.

Let (K, ν) be a rank 1 valued field extending \mathbb{C} , and let ν be the valuation. Consider a morphism

$$\text{Spec } K \rightarrow X^\circ.$$

Since X is proper, this morphism extends to the valuation ring

$$\text{Spec } R \rightarrow X.$$

Let x be the image of the closed point. The smallest closed stratum of D containing x is the intersection of a (possibly empty) subset of irreducible divisor components D_{i_1}, \dots, D_{i_r} . The associated equations generate the stalk of the characteristic sheaf at the point x , which is abstractly isomorphic to \mathbb{N}^r . For each element $\bar{f} \in \mathbb{N}^r$, we may lift it to a function f on X in a neighborhood of the point x , pull back to $\text{Spec } R$ along the map above and compose with the valuation on R . The ratio of two such lifts is a unit, so this gives rise to a well-defined element

$$[\mathbb{N}^r \rightarrow \mathbb{R}_{\geq 0}] \in \Sigma_X \subset \text{Hom}(\mathbb{N}^r, \mathbb{R}_{\geq 0}).$$

For any valued field K , we have a well-defined morphism

$$\text{trop}: X^\circ(K) \rightarrow \Sigma_X.$$

Let K be a valued field whose associated valuation map $K^\times \rightarrow \mathbb{R}$ is surjective. Let $Z^\circ \subset X^\circ$ be a subscheme. Let $\text{trop}(Z^\circ)$ be the subset of Σ_X obtained by restricting trop to $Z^\circ(K)$; it is independent of the choice of valued field K , and its and its functoriality properties are outlined in [28, 75].²

1.5. Properties of the tropicalization

The shapes of tropicalizations are governed by the Bieri–Groves theorem [14, 74].

Theorem 1.5.1. *Let $Z^\circ \subset X^\circ$ be a closed subscheme. Then the set $\text{trop}(Z^\circ)$ is the support of a rational polyhedral cone complex of Σ_X . The topological dimension of $\text{trop}(Z^\circ)$ is bounded above by the algebraic dimension of Z° . If X° is a closed subvariety of an algebraic torus, the topological dimension of $\text{trop}(Z^\circ)$ is equal to the algebraic dimension of X .*

The set $\text{trop}(Z^\circ)$ has no distinguished polyhedral structure in general [48, Example 3.5.4]. It is simply a set, and this set *can* be given the structure of a cone complex.

²Comprehensive treatments of tropicalization regularly make use of Berkovich’s analytification. We omit this here to avoid additional technical machinery, although the proofs of these basic results use them. See [5] for a survey of the relationship between tropical, logarithmic, and non-Archimedean geometry.

The role of tropicalization in degeneration and compactification problems has its origin in the following two theorems, proved by Tevelev for toric varieties and Ulirsch for logarithmic schemes; see [71, Theorem 1.2] and [74, Theorem 1.2].

The first concerns the properness of closures of subschemes in partial compactifications of X° . Let Z° be a subscheme of X° and let X' be a simple normal crossings compactification. Let $X'' \subset X'$ be the complement of a union of closed strata of X' and let $\Sigma_{X''}$ be the subfan of $\Sigma_{X'}$ obtained by deleting the corresponding union of open cones.

Theorem 1.5.2. *The closure Z of Z° in the partial compactification X'' is proper if and only if $\text{trop}(Z^\circ)$ is set theoretically contained in $\Sigma_{X''}$.*

The second concerns transversality. We say the closure Z of Z° in X' intersects strata in the expected dimension if

$$\dim E \cap Z = \dim Z - \text{codim}_{X'} E.$$

Theorem 1.5.3. *The closure Z of Z° in the compactification X' of X° intersects strata in the expected dimension if and only if $\text{trop}(Z^\circ)$ is a union of cones in $\Sigma_{X'}$.*

Remark 1.5.4. When X is a toric variety, Tevelev has shown that there exists a toric blowup $X' \rightarrow X$ such that the closure has a stronger transversality property, hinted at in the introduction, called *algebraic transversality*. We will require and refine this result in the course of our main result [71, Theorem 1.2]; a simple proof is given in [48, Theorem 6.4.17].

Remark 1.5.5. If X is not proper, there is no longer a tropicalization map defined on the set $X^\circ(\mathbb{K})$ because limits need not exist. However, the above relationship with the cone complex persists. For each valued field \mathbb{K} with valuation ring R as above, let $X^{\square}(\mathbb{K})$ be the subset $X^\circ(\mathbb{K})$ consisting of those \mathbb{K} -points that extend to R -points. There is a morphism

$$X^{\square}(\mathbb{K}) \rightarrow \Sigma_X$$

defined exactly as defined above. Similarly, given a subscheme $Z \subset X^\circ$, we can define its tropicalization as the image of $Z^{\square}(\mathbb{K})$ in Σ_X ; see [75, Section 5.2].

1.6. Tropicalization via compactifications

In the previous section, tropicalizations were seen to select partial compactifications of X° , in which the closure of a subvariety meets each stratum in the expected codimension. There is a partial converse.

Let X' be a simple normal crossings compactification of X° . Let Z be a subscheme of X' whose intersection with the strata of X' have the expected codimension, with

$$Z^\circ = X^\circ \cap Z.$$

The following perspective is due to Hacking–Keel–Tevelev, extended by Ulirsch and is sometimes referred to as *geometric tropicalization*; see [29, 74].

Theorem 1.6.1. *The tropicalization of Z° is equal to the union of cones σ in the underlying set of $\Sigma_{X'}$ such Z nontrivially intersects the locally closed strata dual to σ . That is, there is an equality of subsets of $\Sigma_{X'}$ given by*

$$\text{trop}(Z^\circ) = \bigcup_{\sigma \in \Sigma_{X'}: V(\sigma) \cap Z \neq \emptyset} \sigma \quad \text{in } |\Sigma_{X'}|,$$

where $V(\sigma)$ is the locally closed stratum of X' corresponding to the cone σ .

The definition of $\text{trop}(Z^\circ)$ via coordinatewise valuation described in the previous subsection depends on a choice of compactification X of X° , but if X is replaced with a blowup along a stratum, the set $\text{trop}(Z^\circ)$ is unchanged. The result above can therefore be viewed as a computational tool; it describes the tropicalization using a single compactification in which Z is dimensionally transverse. We return to this in Section 7.

Remark 1.6.2 (Asymptotics and stars). Let $Z^\circ \subset X^\circ$ be a subscheme with tropicalization $\text{trop}(Z^\circ)$ in Σ_X . Assume that the tropicalization is a union of cones in Σ_X , and let Z denote the closure of Z° . If $D_i \subset X$ is an irreducible component, we can view it as a simple normal crossings pair in its own right with interior D_i° and divisor equal to the intersection of D_i with the remaining components of D . It contains the subscheme $Z_i^\circ = Z \cap D_i^\circ$. The tropicalization of $Z_i^\circ \subset D_i^\circ$ can be read off from the larger $\text{trop}(Z^\circ)$ as follows. The divisor D_i determines a ray ρ_i in Σ_X . The star of ρ_i is the union of cones that contain ρ_i and can be identified with Σ_{D_i} . The tropicalization of $\text{trop}(Z_i^\circ)$ is the union of the cones under this identification where $\text{trop}(Z^\circ)$ itself is supported. In practice, it is visible as the collection of asymptotic directions of $\text{trop}(Z^\circ)$ parallel to ρ_i .

1.7. Tropicalization for a family of subschemes

The constructions extend to flat one-parameter families of subschemes. We will consider subschemes in X that are defined over a valued field K that extends the trivially valued ground field \mathbb{C} .

In order to avoid foundational issues, we will assume that all subschemes are defined over the localization of a smooth algebraic curve of finite type. The assumption is made so we can appeal to the relationship between tropicalization via valuations and via logarithmic geometry, which has only received a definitive treatment under these hypotheses [74]. The valued field will arise for us in the study of the valuative criterion for properness. The relevant moduli spaces will be shown to be of finite type, so this is a harmless assumption.

1.7.1. Tropicalization over a valued field

Let $(X|D)$ be a simple normal crossings compactification as above, with interior X° . Let K be a valued field extending \mathbb{C} as above. Consider an algebraically transverse subscheme

$$Z_\eta \hookrightarrow X \times_{\mathbb{C}} K,$$

and let Z_η° be the open subscheme contained in X° . By the algebraic transversality hypothesis, this open subscheme is dense. After passing from K to a valued extension L with real surjective valuation, we may once again consider the tropicalization map

$$\text{trop}: Z_\eta^\circ(L) \rightarrow \Sigma_X.$$

The tropicalization is independent of the choice of L , provided the valuation is surjective onto the real numbers.³

The basic structure result extends to this nontrivially valued setting with a small twist.

Theorem 1.7.1. *The set $\text{trop}(Z_\eta^\circ)$ is the support of a rational polyhedral complex of Σ_X . The topological dimension of $\text{trop}(Z_\eta^\circ)$ is bounded above by the algebraic dimension of Z° . If X° is a closed subvariety of an algebraic torus, then the topological dimension of the tropicalization is equal to the algebraic dimension of Z° .*

The toric case of this result was established by Gubler without restriction on the base field [28]. The statement for simple normal crossings targets is a consequence of [15, Theorem 5.1].

³In the toric setting, the independence of L follows, for example, from nonemptiness results for the fibers of the tropicalization map [66], applied to the toric charts of X , following [74]. It can also be deduced from elementary properties of the Berkovich analytification and the formalism in [1, 74].

To emphasize the point, the tropicalization of a variety that is defined over a nontrivial valued extension of \mathbb{C} is polyhedral but *not necessarily conical*. In other words, it can have bounded cells. The target X itself is still defined over \mathbb{C} , not merely over K , so *its* tropicalization remains conical.

1.7.2. Families over a punctured curve

Let C be a smooth algebraic curve of finite type with a distinguished point $0 \in C$; let C° be the complement of this point. The pair $(C, 0)$ has a cone complex Σ_C , canonically identified with the positive real line $\mathbb{R}_{\geq 0}$.

Let $(X|D)$ be a pair as before and X° its interior. Consider a flat family of one-dimensional subschemes over C°

$$\mathcal{Z}^\circ \subset C^\circ \times X^\circ.$$

There are two tropicalizations associated to this family. For the first, we consider \mathcal{Z}° as a two-dimensional subscheme of $C^\circ \times X^\circ$. The partial compactification $C \times X$ determines a tropicalization, as explained in Remark 1.5.5:

$$\text{trop}(\mathcal{Z}^\circ) \subset \mathbb{R}_{\geq 0} \times \Sigma_X.$$

For the second, note that the function field of C is equipped with a discrete valuation arising from order of vanishing at 0 . Let K be an extension of this field with real surjective valuation, and note that there is a canonical inclusion of $\text{Spec } K$ to C° . Consider the tropicalization of the K -valued points of the base change

$$\mathcal{Z}_\eta = \mathcal{Z}^\circ \times_{C^\circ} K \subset X(K),$$

where the morphism $\text{Spec } K \rightarrow C^\circ$ is determined by the fraction field of the local ring of C at 0 . Denote the result by $\text{trop}(\mathcal{Z}_\eta)$ – note that the subscript η indicates that we are considering it as a generic fiber.

These two procedures are related. When considering the total space of \mathcal{Z}° as a surface over C° , there is a map

$$v: \text{trop}(\mathcal{Z}^\circ) \rightarrow \mathbb{R}_{\geq 0}.$$

The fiber $v^{-1}(1)$ coincides with the second tropicalization $\text{trop}(\mathcal{Z}_\eta)$. This follows from a tracing of definitions for the tropicalization of fibers of maps of Berkovich spaces and the functoriality results in [74].

1.8. Transversality for one-parameter families

The appropriate generalizations of the properness and transversality statements earlier in this section are as follows. Let R be the valuation ring of K , and equip $\text{Spec } R$ with the divisorial logarithmic structure at the closed point. The scheme $X \times \text{Spec } R$ is equipped with the simple normal crossings divisor given by $D \times \text{Spec } R \cup X \times 0$. Once equipped with this divisor, it has tropicalization $\Sigma_X \times \mathbb{R}_{\geq 0}$.

Let $\mathcal{Y}' \rightarrow X \times \text{Spec } R$ be a toroidal modification of the constant family, and let $\mathcal{Y}'' \subset \mathcal{Y}'$ be the complement of a union of closed strata. In practice, \mathcal{Y}'' will be the complement of all closed strata of codimension 2 in the special fiber. These have tropicalizations $\Sigma_{\mathcal{Y}'}$ and $\Sigma_{\mathcal{Y}''}$ that are, respectively, a complete subdivision and a subdivision of $\Sigma_X \times \mathbb{R}_{\geq 0}$. There is a morphism of cone complexes by composition:

$$\Sigma_{\mathcal{Y}''} \rightarrow \mathbb{R}_{\geq 0},$$

and we denote the fiber over 1 by $\Sigma_{\mathcal{Y}''}(1)$. We view this as a polyhedral complex.

Let \mathcal{Z}° be a subscheme of $X^\circ \times \text{Spec } K$. We examine the question of when the closure is proper.

Theorem 1.8.1. *The closure \mathcal{Z} of \mathcal{Z}° in the degeneration \mathcal{Y}'' of X is proper over $\text{Spec } R$ if and only if $\text{trop}(\mathcal{Z}^\circ)$ is set theoretically contained in $\Sigma_{\mathcal{Y}''}(1)$.*

Proof. This is well known to experts, but in the form stated, we have been unable to locate a suitable reference. We explain how it can be deduced from results that do appear in the literature.

First, we note that the closure of \mathcal{Z}° in the larger degeneration \mathcal{Y}' is certainly proper over $\text{Spec } R$ because X is proper and $\mathcal{Y}' \rightarrow X \times \text{Spec } R$ is a proper and birational morphism. We now use the hypothesis on K . Specifically, since K is a localization of the function field of a smooth curve C with the valuation associated to a closed point 0 with complement C° . Moreover, the given subscheme \mathcal{Z}° , as well as the degenerations \mathcal{Y}' and \mathcal{Y}'' , can be assumed to arise via base change from corresponding families over C along the inclusion

$$\text{Spec } R \rightarrow C$$

of the local ring at 0 . Rather than overburdening the notation, we replace the families over $\text{Spec } R$ with the corresponding families over C . We now view $C, X \times C, \mathcal{Y}'$ and \mathcal{Y}'' as logarithmic schemes over C . Note that, due to the potential nonproperness of C , Remark 1.5.5 is in effect. The reader is also advised to keep in mind the relationship between the two tropicalizations of \mathcal{Z}° described above.

We are now in a position to apply the Tevelev–Ulirsch lemma for logarithmic schemes as stated in [74, Lemma 4.1]. From it, we deduce that the closure of \mathcal{Z}° in \mathcal{Y}' coincides with the closure in \mathcal{Y}'' precisely under the hypotheses stated in the theorem. We conclude the result. \square

We keep the notation above and now deal with the corresponding transversality statement.

Theorem 1.8.2. *The closure \mathcal{Z} of \mathcal{Z}° in the degeneration \mathcal{Y}'' of X intersects the strata of \mathcal{Y}' in the expected dimension if and only if $\text{trop}(\mathcal{Z}^\circ)$ is a union of polyhedra in $\Sigma_{\mathcal{Y}''}(1)$.*

Proof. Proceed as in the proof of the theorem above and spread out the family until it is defined over a curve. To calculate the dimension of the intersections of \mathcal{Z} with the strata of \mathcal{Y} , the map to $\text{Spec } R$ is not relevant, so we directly apply Theorem 1.5.3 in the previous section and conclude. \square

2. Tropical degenerations and algebraic transversality

Let $(X|D)$ be a simple normal crossing pair such that the all intersections of irreducible components of D are connected. We further assume that X is proper. In this section, we first introduce the precise notion of expansion of X that we consider in this paper. We then examine how to use tropical data to construct algebraically transverse flat limits for families of subschemes of X , along the lines of work of Tevelev [71] and Ulirsch [74]. This will motivate our construction of the stack of expansions in the next section and provide the main ingredient in our proof of properness. Throughout this section, we specialize our discussion to subschemes of dimension at most 1.

2.1. The plan

Given a simple normal crossings pair $(X|D)$, we have recalled in the section above how a *conical subdivision* of its tropicalization Σ_X determines an open subset in a proper birational modification of X . A conical subdivision of $\Sigma_X \times \mathbb{R}_{\geq 0}$, by the same dictionary, will determine a *expansion*, a special type of degeneration, of X along its boundary D . It leads to a family defined over \mathbb{A}^1 .

A conical subdivision

$$\widetilde{\Sigma_X \times \mathbb{R}_{\geq 0}} \rightarrow \Sigma_X \times \mathbb{R}_{\geq 0}$$

can be visualized as a *polyhedral* subdivision of Σ_X : Take the height 1 slice of the subdivision under the projection to $\mathbb{R}_{\geq 0}$. The result is a union of polyhedra in Σ_X glued along faces – a polyhedral complex. The cone over this polyhedral subdivision recovers the original conical subdivision.

As we have seen in the previous section, the tropicalization of a one-dimensional subscheme of X , as defined in the previous section, determines a polyhedral subdivision and therefore an expansion.

We capture the possible tropicalizations of subschemes by the notion of a 1-complex and use them to prove flat limit algorithms which will eventually establish the properness and separatedness of our yet-to-be-proposed moduli spaces. In order to build these moduli, we will form the parameter space for such 1-complexes and use logarithmic geometry to turn this into a parameterizing stack for expansions. The logarithmic Hilbert scheme will be built on top of this stack of expansions.

2.2. Graphical preliminaries

Let \underline{G} be a finite graph, possibly disconnected but without loops or parallel edges. We enhance it with two additional pieces of data. The first is a finite set of rays, formally given by a finite set $R(\underline{G})$ equipped with a map to the vertex set

$$r: R(\underline{G}) \rightarrow V(\underline{G}).$$

The second is the *metrization* of the edge set, given by the edge length function

$$\ell: E(\underline{G}) \rightarrow \mathbb{R}_{>0}.$$

Definition 2.2.1. An *abstract 1-complex* is a triple $(\underline{G}, r: R(\underline{G}) \rightarrow V(\underline{G}), \ell: E(\underline{G}) \rightarrow \mathbb{R}_{>0})$ consisting of a finite graph, a collection of rays and an edge length on the edges.

These data give rise simultaneously to a metric space and a polyhedral complex, both enhancing the topological realization of \underline{G} . The topological realization of \underline{G} is a CW complex, and we endow an edge E with a metric by identifying it with an interval in \mathbb{R} of length $\ell(E)$. For each element $h \in R(\underline{G})$, we glue on a copy of the metric space $\mathbb{R}_{\geq 0}$ to the point $r(h)$. As we are now free to think of each edge or half edge as being either a polyhedron or a metric space, the result is a space G that is simultaneously a metric space and a polyhedral complex of dimension at most 1. In particular, it makes sense to talk about real-valued continuous piecewise affine functions on G .

Let Σ be a smooth cone complex such that the intersection of any two cones is a face of each. Let $|\Sigma|$ be the associated topological space. Note that Σ embeds canonically as a subcomplex of a standard orthant via $\Sigma \hookrightarrow \mathbb{R}_{\geq 0}^{\Sigma(1)}$.

A *piecewise affine map* $F: G \rightarrow \Sigma$ is the data of a continuous map on the underlying topological spaces

$$|G| \rightarrow |\Sigma|$$

such that every face of G maps to a cone of Σ and such that the map is integer affine upon restriction to each face.

We can measure the *slope along an edge* of G , well defined up to sign, as follows. Each noncontracted edge E maps to a cone σ and thus maps onto a line L_E in σ^{gp} . We refer to the expansion factor of the induced map $F: E \rightarrow L_E$ as the *slope of F along the edge E* .

Definition 2.2.2. An *embedded 1-complex* in Σ is a piecewise affine map, written,

$$G \rightarrow \Sigma$$

that is injective on underlying topological spaces subject to the following conditions: (i) the slope along all edges E in G is equal to 1, and (ii) the image of each ray of G is parallel to a one-dimensional face of the cone containing it.

This latter condition will be satisfied for all the embedded 1-complexes that occur for us since it will be forced by algebraic transversality of subschemes in expansions. The reader can also drop condition,

as these would correspond to components of the stack of expansions that are not relevant for our moduli spaces.

When it is clear from context that G has been embedded, we refer to it simply as a 1-complex and will use the notation G to refer to this embedded object.

2.3. Target geometry

Let Σ_X be the cone complex of $(X|D)$. An embedded 1-complex $G \subset \Sigma_X$ gives rise to a class of target geometries as follows. Place the embedded 1-complex G at height 1 inside $\Sigma_X \times \mathbb{R}_{\geq 0}$, and let $C(G)$ be the cone over it.

Lemma 2.3.1. *The cone $C(G)$ over G is a cone complex embedded in $\Sigma_X \times \mathbb{R}_{\geq 0}$.*

Proof. The cone over each face in G certainly forms a cone, so we check that this collection of cones meet along faces of each. Following the argument in [26, Theorem 3.4], this check is nontrivial in the fiber over 0 in $\Sigma_X \times \mathbb{R}_{\geq 0}$. Since G has dimension either 0 or 1, the cones in the 0 fiber are either rays starting from the origin or the origin itself. Since two such rays either coincide or intersect only at the origin, this implies the statement. \square

By applying the construction in Section 1.3 and Remark 1.3.3 to the subdivision

$$C(G) \hookrightarrow \Sigma_X \times \mathbb{R}_{\geq 0},$$

we obtain a target expansion

$$\mathcal{Y}_G \rightarrow X.$$

It is typically not proper.

Remark 2.3.2. The construction of such target expansions using subdivisions goes back at least to Mumford’s work on degenerations of abelian varieties [60]. Its first appearance in enumerative geometry is in work of Nishinou and Siebert [61].

Every affine toric variety carries a canonical logarithmic structure. Consider a toric monoid P with associated toric variety U_P and a closed point u in the closed torus orbit. By pulling back the toric logarithmic structure to u , we obtain the P -logarithmic point. It is denoted $\text{Spec } \mathbb{C}_P$. The standard logarithmic point is $\text{Spec } \mathbb{C}_{\mathbb{N}}$ and is the pullback to 0 of the toric logarithmic structure on \mathbb{A}^1 .

Definition 2.3.3 (Expansions and families). A rough expansion of X over $\text{Spec } \mathbb{C}_{\mathbb{N}}$ is a logarithmic scheme $\mathcal{Y} \rightarrow X \times \text{Spec } \mathbb{C}_{\mathbb{N}}$ which is the fiber over 0 in the modification of $(X|D) \times \mathbb{A}^1$ induced by a subdivision

$$\Delta \hookrightarrow \Sigma_X \times \mathbb{R}_{\geq 0}.$$

A rough expansion is called an expansion if, in addition, the following two conditions are satisfied

- (E1) The scheme \mathcal{Y} is reduced.
- (E2) The subdivision is given by the cone over an embedded 1-complex in Σ_X . Equivalently, at every point of \mathcal{Y} , the relative characteristic monoid of $\mathcal{Y} \rightarrow \text{Spec } \mathbb{C}_{\mathbb{N}}$ is free of rank at most 1.

If S is a fine and saturated logarithmic scheme, a morphism $\mathcal{Y}/S \rightarrow X \times S/S$ is called an expansion of X over S if all pullbacks $\text{Spec } \mathbb{C}_{\mathbb{N}} \rightarrow S$ are expansions and if the morphism $\mathcal{Y} \rightarrow S$ is flat and logarithmically smooth.

We make a few comments about the definition. The first concerns (E2). We deal exclusively with subschemes of dimension at most 1 in this paper, and the transversality condition that we will impose on these subschemes will ensure that they meet the strata of their ambient spaces in a dimensionally

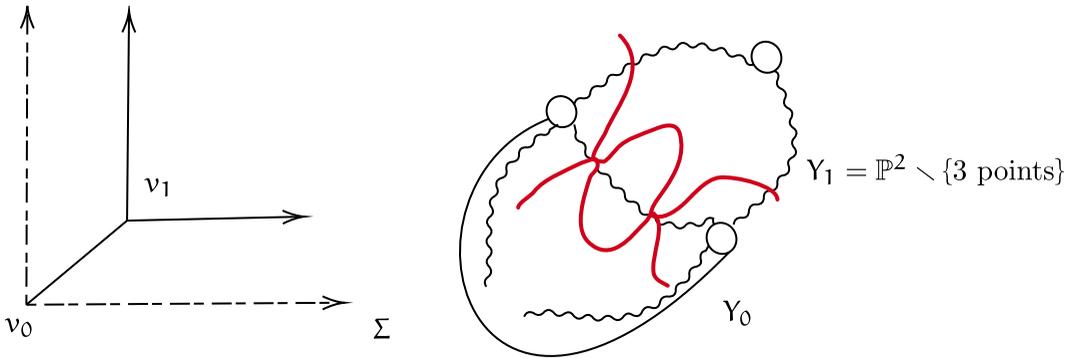


Figure 2. On the left, we sketch a 1-complex in the cone complex $\mathbb{R}_{\geq 0}^2$. The cone complex Σ_X is $\mathbb{R}_{\geq 0}^2$ drawn with a dashed arrow while the 1-complex is solid. On the right is a cartoon of the corresponding expansion. It is obtained by performing deformation to the normal cone of a codimension 2 stratum and then passing to an open. The wavy lines indicate the divisors where the logarithmic structure is nontrivial. The holes indicate codimension 2 strata that have been removed from the expansion. The red curve in the middle is a subscheme of the type that we will soon introduce.

transverse fashion. Once a rough expansion is constructed, the codimension 2 strata become irrelevant, and by removing these, one is led to more efficient moduli spaces. The second concerns the definition of an expansion over general bases S . For experts in logarithmic geometry, we note that one can replace the condition ‘flat and logarithmically smooth’ with ‘logarithmically smooth and integral’. Similarly, the reducedness of the fibers can be replaced with ‘saturated’. Finally, we note that although the expansions themselves tend to be nonproper, we always have a map from \mathcal{Y} to the original target X , which has been assumed to be projective throughout. Moreover, we will only ever consider sheaves on \mathcal{Y} with proper support, so for example, the support always defines a cycle in X by pushforward.

A 1-complex and the corresponding geometric expansion are shown in Figure 2.

Remark 2.3.4. The definition has a peculiar feature when we consider X itself. Specifically, X is a rough expansion of itself, but it is only an expansion of itself when it contains no codimension 2 strata. The complement of the codimension 2 strata in X is an expansion of X , as is the interior of X . The reason for imposing this condition is that we will later require our families of subschemes to have nonempty intersection with all logarithmic strata, which is not satisfied for a one-dimensional transverse subscheme unless there are no codimension 2 strata. This definition yields a cleaner universal property when we consider the valuative criterion of properness.

2.4. Flat limit algorithms: existence

We let X° be the complement of D in X . Let C be a smooth curve with a distinguished point 0 , whose complement is denoted C° .

Proposition 2.4.1. Consider a flat family of subschemes $\mathcal{Z}^\circ \subset X^\circ \times C^\circ$ over C° . There exists a rough expansion $\mathcal{Y} \rightarrow X \times C$ over C such that the closure \mathcal{Z} of \mathcal{Z}° in \mathcal{Y} has the following properties:

- (D1) The scheme \mathcal{Z} is proper and flat over C .
- (D2) The scheme \mathcal{Z} has nonempty intersection with each stratum of \mathcal{Y} .
- (D3) The scheme \mathcal{Z} intersects all the strata of \mathcal{Y} in the expected dimension.

Moreover, after replacing C with a ramified base change, the degeneration of X can be guaranteed to have reduced special fiber, that is, the rough expansion can be chosen to be an expansion.

We will refer to a family satisfying these properties as a *dimensionally transverse family*. A sketch picture of dimensional transversality can be seen on the right in Figure 2. Notice in particular that in

the main component Y_0 of the expansion shown there, the subscheme does not intersect the logarithmic divisors. Correspondingly, the 1-complex on the left of the image does not include the dashed axes.

Proof. Let

$$\text{Spec } \mathbb{C}[[t]] \rightarrow \mathbb{C}$$

be the spectrum of the completed local ring of \mathbb{C} at 0. Pull back the subscheme family \mathcal{Z}° to the generic point of this valuation ring to obtain a subscheme \mathcal{Z}_η° of X° over $\mathbb{C}((t))$. After passing to a valued field extension with value group \mathbb{R} , apply the tropicalization map to obtain

$$\text{trop}(\mathcal{Z}^\circ) \subset \Sigma_X.$$

By the structure result for tropicalizations recorded in Theorem 1.5.1, this is the support of a polyhedral complex of dimension at most 1 in Σ . Choose a polyhedral structure on this set, and call it G . In practice, this amounts to expressing the tropicalization as a union of vertices, closed bounded edges and unbounded rays. This results in a polyhedral complex embedded in Σ_X . Let $C(G)$ be the cone over G in $\Sigma_X \times \mathbb{R}_{\geq 0}$. We obtain an associated target family

$$\mathcal{Y}_G \rightarrow \mathbb{C}.$$

By construction, the tropicalization of \mathcal{Z}° when viewed as a single subscheme of the total space $X^\circ \times \mathbb{C}^\circ$ is equal to $C(G)$. The theorems in Section 1.4 imply that the closure is proper. Since the base is a smooth curve, the closure is also flat over \mathbb{C} .

We examine the strata intersections. Given a stratum W of \mathcal{Y}_G dual to a cone σ in $C(G)$, choose a point $v \in \sigma^\circ$. By definition, this is the image of a K -valued point of \mathcal{Z}° under the tropicalization map. Since the family of subschemes is proper, this extends to an \mathbb{R} -valued point where \mathbb{R} is the valuation ring of K . Since the tropicalization of this K -valued point lies in the interior of σ , the closed point maps to W , so it follows that \mathcal{Z} intersects all strata.

Finally, we show that we can obtain a degeneration with reduced special fiber after a base change. By using the toric dictionary, we observe that the special fiber of $\mathcal{Y}_G \rightarrow \mathbb{C}$ is reduced if and only if the vertices of G are lattice points in Σ_X . This can be engineered by using Kawamata’s cyclic covering trick, explained in [7]. The toroidal procedure is carried out by replacing the integral lattice in Σ_C with a finite index sublattice, thereby ensuring that the fiber over the new primitive generator of Σ_C is a dilation of G which has integral vertices. The main result of [7, Section 5] is that this produces the requisite ramified base change. The statements concerning strata intersections are unaffected by the base change, and we conclude the result. □

2.5. Flat limit algorithms: uniqueness

The tropical limit algorithm in the previous section inherits a uniqueness property that should be thought of as close to a universal closedness result for the forthcoming moduli problem, where by ‘close to’ we mean that we will need to eventually strengthen the notion of transversality.

Let K be the discretely valued field associated to the valuation at 0 in \mathbb{C} , and let \mathbb{R} be its valuation ring.

Proposition 2.5.1. *Let \mathcal{Z}_η° be a flat family of subschemes of X° over $\text{Spec } K$. Assume that the closure of \mathcal{Z}_η° in $X \times \text{Spec } K$ is dimensionally transverse. Then there exists a canonical triple $(R', \mathcal{Y}', \mathcal{Z}')$ comprising of a ramified base change $R \subset R'$ with fraction field K' and expansion of X*

$$\mathcal{Y}' \rightarrow \text{Spec } R'$$

such that the closure \mathcal{Z}' of $\mathcal{Z}_\eta^\circ \otimes_K K'$ in \mathcal{Y} is dimensionally transverse. Moreover, the triple satisfies the following uniqueness property:

(★) For any other choice $(R'', \mathcal{Y}'', \mathcal{Z}'')$ satisfying these requirements, there exists a unique toroidal birational morphism

$$\mathcal{Y}'' \rightarrow \mathcal{Y}'_1''$$

over $\text{Spec } R''$ with subscheme \mathcal{Z}''_1 and a ramified covering $\text{Spec } R'' \rightarrow \text{Spec } R$ such that $\mathcal{Z}''_1 \subset \mathcal{Y}'_1''$ is obtained from $\mathcal{Z}' \subset \mathcal{Y}'$ by base change.

We deduce the uniqueness results from two simple combinatorial observations.

Lemma 2.5.2. *Let $\iota : G \hookrightarrow \Sigma_X$ be an embedded 1-complex. The underlying set $|G|$ of G carries a unique minimal polyhedral structure \overline{G} such that the map ι descends to an embedding $\overline{G} \hookrightarrow \Sigma_X$.*

The term minimal here is used in the sense that any other polyhedral structure is obtained from the putative unique minimal one by subdividing along edges. Note that a morphism of polyhedral complexes is required to map vertices and edges of G to into cones of Σ_X .

Proof. Recall that G comes with a distinguished vertex set and $|G|$ has the structure of a metric space independent of the chosen vertex set. The nonbivalent vertices of G are necessarily contained in the vertex set of any polyhedral structure on $|G|$. Let $x \in G$ be a 2-valent vertex. Call v *inessential* if both of the following conditions hold.⁴

- (1) There exists an open neighborhood U_x of x that is completely contained in the relative interior σ° of a cone $\sigma \in \Sigma_X$.
- (2) Every point in this neighborhood lies on the same line in the vector space σ^{gp} .

A vertex that is not inessential is essential. Let \overline{G} be the polyhedral complex obtained from the metric space $|G|$ by declaring the vertices of \overline{G} to be the essential vertices of G . The morphism ι descends to an embedding $\overline{G} \rightarrow \Sigma_X$. Conversely, any polyhedral structure must contain the points of $|G|$ that are not inessential to ensure that ι is a morphism of polyhedral complexes. It follows that $\iota : \overline{G} \rightarrow \Sigma_X$ is minimal. □

Passing to the cone over \overline{G} , we obtain a uniqueness property concerning the dilations of \overline{G} obtained by uniformly scaling the edge lengths.

Lemma 2.5.3. *Let $G \hookrightarrow \Sigma_X$ be an embedded 1-complex. Let $\mathcal{C}(\overline{G}) \hookrightarrow \Sigma_X \times \mathbb{R}_{\geq 0}$ be the cone over the minimal polyhedral structure \overline{G} of G . There exists a minimum positive integer b such that all vertices in the fiber of $\mathcal{C}(\overline{G})$ over b lie in the lattice of Σ_X .*

Proof. We take b to be the least common multiple of the denominators appearing in the coordinates of all vertices of G , with respect to the coordinates on Σ_X . □

Proof of Proposition

The uniqueness is a translation of the combinatorial observations above, as we now explain. Let \mathcal{Z}°_η be a subscheme as in the proposition. The algorithm for finding limits has the following steps. We calculate the tropicalization of \mathcal{Z}°_η and obtain a canonical 1-complex $\overline{G} \hookrightarrow \Sigma_X$. We then pass to the cone over this complex $\mathcal{C}(\overline{G})$ in $\Sigma_X \times \mathbb{R}_{\geq 0}$. Finally, we find the minimum positive integer b in $\mathbb{R}_{\geq 0}$ over which the fiber in $\mathcal{C}(\overline{G})$ has integral vertices. We perform the order b cyclic ramified base change to obtain the requisite triple $(R', \mathcal{Y}', \mathcal{Z}')$.

Now, consider another set $(R'', \mathcal{Y}'', \mathcal{Z}'')$ satisfying these properties. Let $\mathcal{C}(G'') \rightarrow \mathbb{R}_{\geq 0}$ be the fan associated to this expansion. Tropicalization is invariant under taking valued field extensions; since the subscheme \mathcal{Z}'' nontrivially intersects all the strata of the expansion, we can conclude using Theorem 1.6.1 that the support of the tropicalization is equal to $\mathcal{C}(G'')$, where G'' is the height 1 slice of the cone. If we replace G'' with its minimal polyhedral structure \overline{G}'' , we obtain a refinement of cone

⁴Intuitively, x is inessential if the map ι doesn't bend at x ; thus, the map ι descends to a polyhedral map after coarsening the polyhedral structure to exclude x from the vertex set.

complexes $C(G'') \rightarrow C(\overline{G}'')$. The latter cone complex gives rise to an expansion of X in which the closure of Z_η° is still dimensionally transverse.

For the base change, by using the preceding lemma we let \overline{G}' be the minimal dilation of \overline{G} whose vertices are all integral. The cones over \overline{G}' and \overline{G}'' coincide, so \overline{G}'' must be a dilation of \overline{G}' . It follows that the morphism

$$C(\overline{G}'') \rightarrow \mathbb{R}_{\geq 0}$$

is obtained from

$$C(\overline{G}') \rightarrow \mathbb{R}_{\geq 0}$$

by passing to a finite index sublattice in the integral structure of the base. The result follows. \square

Before proceeding to algebraic transversality, we offer two examples of the procedure above. We begin with a zero-dimensional example.

Example 2.5.4. Let X be the \mathbb{A}^2 equipped with its toric logarithmic structure. Consider the subscheme Z given by the following union of two reduced points defined over the field $\mathbb{C}((t))$:

$$Z = \{(t^2, t^3)\} \cup \{(t^4, t^5)\} \subset \mathbb{A}^2.$$

The flat limit is a nonreduced point supported at the origin in \mathbb{A}^2 . The tropicalization of Z consists of two points

$$\text{trop}(Z) = \{(2, 3)\} \cup \{(4, 5)\} \subset \mathbb{R}_{\geq 0}^2.$$

According to the algorithm above, associated expansion is obtained by performing a toric modification of $\mathbb{A}^2 \times \mathbb{A}^1$, adding rays passing through points $(2, 3, 1)$ and $(4, 5, 1)$ in the fan $\mathbb{R}_{\geq 0}^2 \times \mathbb{R}_{\geq 0}$. The resulting expansion consists of two components, both isomorphic to \mathbb{G}_m^2 . The new flat limit in the expanded family is the union of two reduced points, with one in each of these two components.

We now consider a one-dimensional example.

Example 2.5.5. Let X be \mathbb{P}^2 equipped with its toric logarithmic structure and homogeneous coordinates X, Y and Z . Consider the subscheme Z given by a line defined over the field $\mathbb{C}((t))$:

$$Z = \mathbb{V}(t^{-3}X + t^{-2}Y + Z) \subset \mathbb{P}^2.$$

The tropicalization is easily computed by hand as the subset of \mathbb{R}^2 obtained by the union of 3 rays ρ_1, ρ_2, ρ_3 at the point $(-3, -2)$. The ray ρ_1 is parallel to the negative x -axis, the ray ρ_2 is parallel to the negative y -axis and ρ_3 is parallel to the first-quadrant diagonal. It is the translate of the 1-skeleton of the fan of \mathbb{P}^2 to the point $(-2, -3)$. The associated 1-complex G has two vertices: one vertex at the point $(-3, -2)$ which is trivalent and we call V_1 and another at the point $(-1, 0)$ where the set $|G|$ intersects the 1-skeleton of the fan of \mathbb{P}^2 , which we call V_2 . See Figure 3

The associated expansion consists of two components corresponding to the two vertices above. The component X_1 corresponding to V_1 is a copy of \mathbb{P}^2 minus its three fixed points. The component X_2 corresponding to V_2 is a copy of $\mathbb{P}^1 \times \mathbb{G}_m$. A straightforward calculation shows that the limiting subscheme is a line in X_1 union a fiber in X_2 .

2.6. Algebraic transversality from dimensional transversality

We have constructed limits of transverse subschemes of X inside expansions of X along D that are dimensionally transverse to the strata of the expansion. We require a stronger form of transversality.

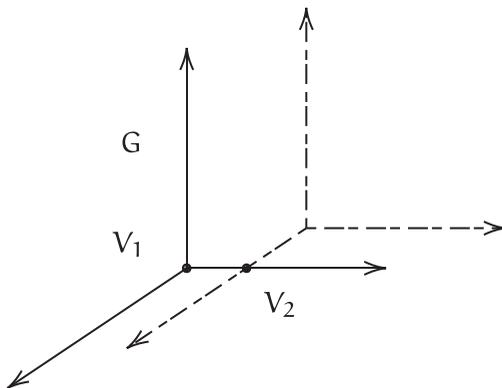


Figure 3. The 1-complex depicted in solid arrows with the fan of \mathbb{P}^2 in dashed arrows.

Definition 2.6.1 (Algebraic transversality). Let \mathcal{Y} be an expansion of X , and let $\mathcal{Z} \subset \mathcal{Y}$ be a subscheme with ideal sheaf $\mathcal{I}_{\mathcal{Z}}$. Then \mathcal{Z} is said to be *algebraically transverse* if it is dimensionally transverse and for every closed divisor stratum $\mathcal{S} \subset \mathcal{Y}$, the induced map

$$\mathcal{I}_{\mathcal{Z}} \otimes_{\mathcal{O}_{\mathcal{Y}}} \mathcal{O}_{\mathcal{S}} \rightarrow \mathcal{O}_{\mathcal{Y}} \otimes_{\mathcal{O}_{\mathcal{Y}}} \mathcal{O}_{\mathcal{S}}$$

is injective.

We call this *algebraic transversality* because it depends on more than just the dimension of the intersection with the strata, which is only set theoretic. On the other hand, it is weaker than genuine geometric transversality because it allows the intersection of the subscheme with a divisor to be nonreduced.⁵

Let us give a few different ways to think about this condition. We maintain the notation above.

Proposition 2.6.2. *The subscheme $\mathcal{Z} \subset \mathcal{Y}$ is algebraically transverse if and only all of the following conditions hold:*

- (i) \mathcal{Z} intersects every stratum of \mathcal{Y} ,
- (ii) there are no embedded points or components of \mathcal{Z} that are contained in the codimension 1 strata of \mathcal{Y} and
- (iii) if \mathcal{S} is any double divisor contained in irreducible components Y_i and Y_j of \mathcal{Y} , the restrictions of the subscheme to Y_i and Y_j have the same intersection with \mathcal{S} .

Proof. Assume $\mathcal{Z} \subset \mathcal{Y}$ is algebraically transverse. Since this implies dimensional transversality by definition, it is clear that (i) holds; see Proposition 2.4.1. Now, the injectivity condition in the definition of algebraic transversality can be checked formally local on \mathcal{Y} . If we look in a neighborhood U of a double divisor \mathcal{S} , then \mathcal{Y} can be identified with $\mathbb{A}^{\dim X - 1} \times \mathcal{N}$, where \mathcal{N} is a nodal curve, that is, given by $\{uv = 0\}$. The stated injectivity condition is equivalent to flatness of $\mathcal{Z}_U \rightarrow \mathcal{N}$. The condition (ii) now follows. Finally, condition (iii) follows by pulling back the flat morphism $\mathcal{Z}_U \rightarrow \mathcal{N}$ to each of the two branches $\{u = 0\}$ and $\{v = 0\}$ of the node and then to the node $\{u = v = 0\}$ itself. Since the result is independent of which branch we restrict to, (iii) follows.

Conversely, suppose the conditions (i)–(iii) hold. We can again restrict to a formal neighborhood $U = \mathbb{A}^{\dim X - 1} \times \mathcal{N}$, and examine the resulting subscheme to \mathcal{Z}_U . We claim that under these hypotheses, the map $\mathcal{Z}_U \rightarrow \mathcal{N}$ is flat. Since \mathcal{N} is reduced, we can deduce this from the valuative criterion for flatness [70, Corollaire 4.2.10], that is, we can check flatness after pulling back to the spectrum B of a DVR. Working over the base B , flatness of $\mathcal{Z}_B \rightarrow B$ is equivalent to the condition that the total space is the closure of the generic fiber. This is clearly implied by (i)–(iii) □

⁵In an earlier version of this paper, which was in circulation from 2020 until 2023, we had used the terminology ‘strongly transverse’ instead of ‘algebraically transverse’. We thank the referee, as well as numerous friends and colleagues of their objections to this.

Remark 2.6.3. The proposition allows one to check algebraic transversality in practice. Given a one-dimensional subscheme $\mathcal{Z} \hookrightarrow \mathcal{Y}$ in an expansion of X , our definition of dimensional transversality already includes condition (i). Therefore, in order to guarantee algebraic transversality, one needs to simply check that there are no embedded points on the divisorial strata and that the induced subscheme on the double divisor from of its branches is the same.

A more conceptual characterization comes from the proof of the proposition above. Recall that Σ_X is a union of faces in the standard orthant fan $\mathbb{R}_{\geq 0}^d$. Let A_{Σ} be the associated union of orbits in the Artin stack $[\mathbb{A}^d/\mathbb{G}_m^d]$. The expansion $\mathcal{Y} \rightarrow X$ is pulled back from an associated expansion $A_{\mathcal{Y}}$. We stress that this is a gluing together of Artin fans along divisors.

Proposition 2.6.4. *A subscheme $\mathcal{Z} \hookrightarrow \mathcal{Y}$ is algebraically transverse if and only if the composite map $\mathcal{Z} \rightarrow A_{\mathcal{Y}}$ is both flat and surjective.*

Proof. We can pass to an open neighborhood in \mathcal{Y} , and as in the previous proposition, algebraic transversality becomes equivalent to (i) flatness of \mathcal{Z}_U over the nodal curve in the local model $U = \mathbb{A}^{\dim X - 1} \times \mathcal{N}$, plus (ii) the condition that all strata have nonempty intersection with \mathcal{Z} . The map from U to the Artin fan $A_{\mathcal{Y}}$ is certainly flat and surjective, so the forward implication is clear. The converse follows immediately from the characterization of algebraic transversality in the previous proposition. \square

Via the proposition above, algebraic transversality is equivalent to the *logarithmic flatness* of \mathcal{Z} , when equipped with the pullback logarithmic structure from the expansion, together with the condition all strata have nonempty intersection with \mathcal{Z} . The perspective coming from logarithmic flatness play a key role in Kennedy-Hunt’s work on the general logarithmic Hilbert scheme [41].

Finally, we note that by basic properties of tensor products, algebraic transversality is equivalent to the condition that the higher Tor functors of $\mathcal{O}_{\mathcal{Z}}$ with $\mathcal{O}_{\mathcal{S}}$ vanish. This is how transversality is stated by Li–Wu, who refer to the condition as *normality to the divisor*; see [47].

Proposition 2.6.5. *Let $\mathcal{Z}_{\eta}^{\circ}$ be a flat family of subschemes of X° over $\text{Spec } K$ whose closure in X is algebraically transverse. Then there exists a canonical triple $(R', \mathcal{Y}', \mathcal{Z}')$ comprised of a ramified base change $R \subset R'$ with fraction field K' and expansion of X*

$$\mathcal{Y}' \rightarrow \text{Spec } R'$$

such that the closure \mathcal{Z}' of $\mathcal{Z}_{\eta}^{\circ} \otimes_K K'$ in \mathcal{Y} is algebraically transverse, satisfying the following uniqueness property:

(\star) *For any other choice $(R'', \mathcal{Y}'', \mathcal{Z}'')$ satisfying these requirements, there exists a unique toroidal birational morphism*

$$\mathcal{Y}'' \rightarrow \mathcal{Y}'$$

over $\text{Spec } R''$ with subscheme \mathcal{Z}'' , and a ramified covering $\text{Spec } R'' \rightarrow \text{Spec } R$ such that $\mathcal{Z}'' \subset \mathcal{Y}''$ is obtained from $\mathcal{Z}' \subset \mathcal{Y}'$ by base change.

Proof. The proposition is established in the following sequence of lemmas. \square

2.6.1. A proof via the work of Li–Wu

Given the dimensional transversality statement that we have already established, one can deduce the proposition above formally from the results of Li–Wu. We give a sketch of this in case the reader wants to simply skip the rest of this section without a loss of continuity. Using the arguments of the previous section on dimensional transversality, given a family of algebraically transverse subschemes over a valued field $\text{Spec } K$, a limit can be found over a ramified base change $\text{Spec } R'$ of its valuation ring, where the intersections of all locally closed strata with the flat limit of the subscheme are either empty or have the expected dimension. By removing the codimension 2 strata from the degeneration, we

obtain a degeneration with only double points such that the closure of the general fiber of subschemes forms a proper and flat family. In other words, we find a dimensionally transverse limit.

We can now appeal to the results of [47] concerning degenerating subschemes in double point degenerations. The results of Li–Wu are not explicitly stated for nonproper total spaces, but the properness is only used to establish *some* proper and flat limit of the generic fiber. This is already achieved by the first step. The methods of [47] apply with minor changes to give the result above. We leave the details of this to an interested reader.

2.6.2. A proof via Gröbner theory

We take a different approach, inspired by Gröbner theory considerations. The approach has a slightly more constructive nature.

The strategy is to reduce to a theorem of Tevelev on subvarieties of toric varieties [71]. Specifically, if X is a toric variety with dense torus T , Tevelev proves that, for any subscheme $Z \hookrightarrow X$, there exists a toric blowup $X' \rightarrow X$, with X' smooth such that the strict transform $Z' \hookrightarrow X'$ is what he calls *tropical* [71, Theorem 1.2]. This means that the induced map $Z \rightarrow [X'/T]$ is flat. When this theorem is applied to a flat family of subschemes in X of relative dimension 1, over \mathbb{A}^1 , the output is precisely the algebraically transverse family that we seek.

We will reduce our more general situation to the toric one. The reduction is essentially straightforward; once dimensional transversality has been achieved, the additional birational modifications required to produce an algebraically transverse family can be analyzed affine locally. The details follow.

We will use the following lemma in the course of our proof, which allows us to take an algebraically transverse limit and contract unnecessary blowups.

Lemma 2.6.6 (Contraction lemma). *Let $\mathcal{Y} \rightarrow \text{Spec } R$ be an expansion of X whose special fiber includes two irreducible components Y_1 and Y_2 that meet transversely along a smooth divisor D . Let $\mathcal{Y}' \rightarrow \mathcal{Y}$ be the blowup at D with exceptional E . Assume that the special fiber of \mathcal{Y}' is reduced.*

Suppose $\mathcal{Z}' \hookrightarrow \mathcal{Y}'$ is a flat and proper family of algebraically transverse subschemes of dimension 1, with generic fiber Z'_η , and such that $\mathcal{Z}' \cap E$ is the pullback of a subscheme along $E \rightarrow D$.

Let \mathcal{Z} be the closure of the generic fiber Z'_η in \mathcal{Y} . Then \mathcal{Z} is a flat family of algebraically transverse subschemes.

Proof. Throughout the proof, we add the subscript 0 to indicate special fibers. Suppose \mathcal{Z} were not algebraically transverse. By the characterization of algebraic transversality, this means that the special fiber $\mathcal{Z}_0 \hookrightarrow \mathcal{Y}_0$ has embedded points contained in D . Since \mathcal{Z} is the closure of its generic fiber, it is flat over $\text{Spec } R$. It follows that the special fibers of \mathcal{Z}'_0 and \mathcal{Z}_0 have the same holomorphic Euler characteristic. However, this now gives a numerical contradiction. Since \mathcal{Z}'_0 is algebraically transverse, its holomorphic Euler characteristic is given by

$$\chi_{\text{hol}}(\mathcal{Z}'_0) = \chi_{\text{hol}}(\mathcal{Z}'_0 \cap Y_1) + \chi_{\text{hol}}(\mathcal{Z}'_0 \cap Y_2) + \chi_{\text{hol}}(\mathcal{Z}'_0 \cap E) - \chi_{\text{hol}}(\mathcal{Z}'_0 \cap E \cap Y_1) - \chi_{\text{hol}}(\mathcal{Z}'_0 \cap E \cap Y_2).$$

But we also have

$$\chi_{\text{hol}}(\mathcal{Z}'_0 \cap E \cap Y_i) = \chi_{\text{hol}}(\mathcal{Z}'_0 \cap E)$$

since $\mathcal{Z}'_0 \cap E$ is a \mathbb{P}^1 -bundle over $\mathcal{Z}'_0 \cap E \cap Y_i$.

Now, the intersection of \mathcal{Z}' with $(Y_1 \cup Y_2) \setminus E$ maps isomorphically onto its image in the complement of D in \mathcal{Y}_0 . On the one hand, the closure of this image cannot contain the embedded points above and is a closed subscheme, properly contained in \mathcal{Z}_0 . On the other hand, the Euler characteristic of this closure is equal to that of \mathcal{Z}'_0 above. Since the embedded points contribute positively to the holomorphic Euler characteristic of \mathcal{Z}_0 , we arrive at a contradiction. □

We move on to the main proof. We first identify the subdivision that is required to produce an algebraically transverse limit and then prove that it has the properties required. We require the notion of

an *initial degeneration* from tropical geometry. In the notation of the proposition, let \mathcal{Z}_η° be a flat family of subschemes of X° over $\text{Spec } K$. Then we have

$$\text{trop}(\mathcal{Z}_\eta) \subset \Sigma_X$$

which is the support of a one-dimensional polyhedral complex embedded in Σ_X . Given a rational point w in Σ_X , we consider a nonproper degeneration of X as follows. Consider the linear map

$$\mathbb{R}_{\geq 0} \rightarrow \Sigma_X \times \mathbb{R}_{\geq 0}$$

mapping isomorphically onto the ray joining the origin to $(w, 1)$. This determines an associated flat degeneration

$$\mathcal{Y}_w \rightarrow \text{Spec}(\mathbb{R})$$

whose general fiber is the interior X° of X and whose special fiber is a torus torsor over a locally closed stratum of X and, precisely, the locally closed stratum corresponding to the cone that contains w in its interior. We denote the special fiber by X_w . We note, in particular, that, up to isomorphism, the scheme X_w is constant when w varies in the interior of each cone of Σ_X .

Definition 2.6.7. The *initial degeneration* of \mathcal{Z}_η at a point w of Σ_X is the special fiber of the closure of $\mathcal{Z}_\eta \cap X^\circ$ in \mathcal{Y}_w , viewed as a subscheme of X_w . It will be denoted $\text{in}_w(\mathcal{Z}_\eta)$.

Recall from the previous section that $\text{trop}(\mathcal{Z}_\eta)$ has a minimal polyhedral structure, and this determines a degeneration \mathcal{Y}_Γ . Let w be a point on an edge e of this minimal polyhedral structure. Then X_w is a \mathbb{G}_m -torsor over a codimension 1 stratum X_e in the special fiber of the associated degeneration. We will say that the initial degeneration $\text{in}_w(\mathcal{Z}_\eta)$ is *tubular* if it is the schematic preimage of a zero-dimensional subscheme under the projection $X_w \rightarrow X_e$.

The following basic finiteness theorem holds.

Proposition 2.6.8. *The subset of points w in $\text{trop}(\mathcal{Z}_\eta)$ with the property that $\text{in}_w(\mathcal{Z}_\eta)$ is not tubular is finite.*

Proof. In the notation of the preceding paragraph, if we are given a finite set of points in the set $\text{trop}(\mathcal{Z}_\eta)$, there is a unique coarsest polyhedral structure on the tropicalization that includes these points as vertices. It of course refines the minimal polyhedral structure on $\text{trop}(\mathcal{Z}_\eta)$. We refer to the vertices of this new polyhedral structure that do not appear in the minimal one as *extra* vertices. They should be viewed as lying in the interior of an edge in the minimal structure.

By the paragraph above and the dictionary explained in the background section, each such choice of finite set produces a model of X over the DVR. We assume it has reduced special fiber by performing an appropriate base change. The closure of \mathcal{Z}_η in this model is a flat family of subschemes.

Let w be an extra vertex. The component it is dual to is a \mathbb{P}^1 -bundle. Since the Hilbert polynomial of the subscheme is constant in the family, for all but finitely many choices of w in $\text{trop}(\mathcal{Z}_\eta)$, the curve class of the subscheme in this component is a fiber class. Similarly, since the holomorphic Euler characteristic is fixed, we find that $\text{in}_w(\mathcal{Z}_\eta)$ must be tubular for all but finitely many choices of w . \square

We are now in a position to construct an algebraically transverse limit.

Construction 2.6.9 (The Gröbner limit). Let \mathcal{S} denote the unique minimal polyhedral structure on $\text{trop}(\mathcal{Z}_\eta)$ whose vertices include all points w in $\text{trop}(\mathcal{Z}_\eta)$ whose initial degeneration is *not* tubular. Consider the associated rough expansion $\mathcal{V}_\mathcal{S} \rightarrow \text{Spec}(\mathbb{R})$. Perform a ramified base change $\mathbb{R} \subset \mathbb{R}'$ of order equal to the least common multiple of the denominators appearing in the coordinates of the vertices of \mathcal{S} . Let $\mathcal{V}_\mathcal{S}' \rightarrow \text{Spec}(\mathbb{R}')$ be the associated expansion. Let $\mathcal{Z} \rightarrow \text{Spec}(\mathbb{R}')$ be the new flat limit, obtained as the closure of \mathcal{Z}_η after extension of scalars.

We now show that this Gröbner limit is algebraically transverse.

Lemma 2.6.10. *The Gröbner limit family $\mathcal{Z} \rightarrow \text{Spec}(\mathbb{R}')$ is a family of algebraically transverse subschemes of $\mathcal{Y}_{\mathcal{G}}$.*

Proof. Consider the Gröbner family of subschemes $\mathcal{Z} \hookrightarrow \mathcal{Y}_{\mathcal{G}}$ over $\text{Spec}(\mathbb{R}')$ as above. The family is certainly dimensionally transverse, as we have already shown, and we assume that the generic fiber is algebraically transverse. Algebraic transversality is therefore equivalent to the condition that there are no embedded points on the double locus or the divisorial strata in the special fiber. We treat the case of the double locus; the other case is essentially identical.

The condition may be checked Zariski locally on the subscheme \mathcal{Z} and depends only on the map from \mathcal{Z} to the Artin fan of the target $\mathcal{Y}_{\mathcal{G}}$. By shrinking to a neighborhood U around the double divisor and denoting the subscheme by \mathcal{Z}_U , each divisor may be assumed to be principal. Choose generators for the coordinate ring of U that include the equations for these principal divisors, to obtain a maps

$$\mathcal{Z}_U \rightarrow \mathbb{A}_{\mathbb{R}}^n \rightarrow \mathbb{A}_{\mathbb{R}}^2,$$

where the second arrow is given by projection onto the first two coordinates. Note that with respect to the logarithmic structure on $\mathbb{A}_{\mathbb{R}}^2$ the tropicalization of \mathcal{Z}_U is a single edge, and we henceforth replace \mathcal{G} to be this edge. Moreover, the logarithmic structure induced on \mathcal{Z}_U by pullback from $\mathcal{Y}_{\mathcal{G}}$ is the pullback of the logarithmic structure coming from the coordinate axes of \mathbb{A}^2 . Our task is check flatness of the induced map

$$\mathcal{Z}_U \rightarrow [\mathbb{A}^2/\mathbb{G}_m^2].$$

We now reduce to the toric case. Choose $n - 2$ generic hyperplanes and add them to the logarithmic structure of \mathbb{A}^n above; this logarithmic structure is now toric. This gives rise to a factorization of the above map as

$$\mathcal{Z}_U \rightarrow [\mathbb{A}^n/\mathbb{G}_m^n] \rightarrow [\mathbb{A}^2/\mathbb{G}_m^2]. \tag{1}$$

The second map is certainly flat, so we need only check flatness for the first.

With the subscheme now embedded in a toric variety, with its toric logarithmic structure, we may now appeal to the methods of Gröbner theory and tropical geometry, in the sense of [28, Section 12] or [48, Section 6.4]. Since the affine hyperplanes are generic, the closures of the initial degenerations of \mathcal{Z}_U are constant on the faces of the tropicalization of \mathcal{Z}_U . In order to see this, we note that the tropicalization of \mathcal{Z}_U in the new embedding is obtained from \mathcal{G} by adding unbounded rays to each of its two vertices. Note that, since the affine hyperplanes are generic, the initial degenerations at the points of these unbounded rays are tubular and are constant along these new rays.

We apply the result of [28, Theorem 12.3]. This constructs (nonuniquely) a Gröbner complex \mathcal{G}' on the underlying set of \mathcal{G} , which is an explicit polyhedral structure on the tropicalization of \mathcal{Z}_U . The associated degeneration $\mathcal{Y}_{\mathcal{G}'}$ of the generic fiber $\mathcal{Y}_{\mathcal{G},\eta}$ has the property that (i) the initial degenerations are constant along faces, and (ii) the strict transform of \mathcal{Z}_U in the associated degeneration is algebraically transverse.⁶

The degenerations $\mathcal{Y}_{\mathcal{G}'}$ and $\mathcal{Y}_{\mathcal{G}}$ can only differ by further subdivision of edges in the tropicalization; call these new 2-valent vertices *extra vertices*. By algebraic transversality of the closure of \mathcal{Z}_U in $\mathcal{Y}_{\mathcal{G}'}$, if we restrict the subscheme to the components corresponding to the extra vertices, the resulting subscheme is pulled back along the blowup $\mathcal{Y}_{\mathcal{G}'}$ and $\mathcal{Y}_{\mathcal{G}}$. We now apply Lemma 2.6.10 to erase all the extra vertices and see that the closure in $\mathcal{Y}_{\mathcal{G}}$ is already algebraically transverse. This is equivalent to the flatness of the first arrow in Equation (1), and the proof of the lemma is complete. \square

We now establish the uniqueness property that we need.

⁶Note that in the tropical geometry literature, compactifications and degenerations with algebraic transversality are called ‘flat tropical compactifications’ in [48, Theorem 6.4.17] or merely ‘tropical compactifications’ in [71, 28].

Lemma 2.6.11. *The Gröbner limit family $\mathcal{Z} \rightarrow \text{Spec}(R')$ satisfies the uniqueness property (\star) of Proposition 2.6.5.*

Proof. Consider another family $\mathcal{Z}' \rightarrow \text{Spec}(R'')$ that is algebraically transverse obtained from a polyhedral structure on $\text{trop}(Z_\eta)$. The first claim is that if a point w on $\text{trop}(Z_\eta)$ has nontubular initial degeneration, then it must appear in the vertex set of any algebraically transverse family. To see this, first observe that if \mathcal{Z}' is an algebraically transversely family of subschemes in \mathcal{Y} , then if we blowup \mathcal{Y} further, then the strict and total transforms of \mathcal{Z}' must coincide. However, the strict transform is precisely the flat limit of \mathcal{Z}_η in the new family.

With this in mind, we consider a polyhedral structure Λ for which there is a point w not contained in the vertex set of such that the initial degeneration at w is not tubular. Form a new polyhedral structure Λ' by introducing this point w as part of the vertex set. Examine the total transform of \mathcal{Z} under the blowup $\mathcal{Y}_{\Lambda'} \rightarrow \mathcal{Y}_\Lambda$ restricted to the locally closed stratum corresponding to this vertex w . By algebraic transversality, it must coincide with the initial degeneration at w , but the total transform is necessarily tubular since it is the preimage of the intersection of \mathcal{Z} with a double divisor. We arrive at a contradiction.

The reducedness condition for an expansion $\mathcal{Y}_{\Lambda'}$ is that every vertex of Λ' must be integral. There is a unique minimal set of vertices – precisely those with nontubular initial degenerations – that must be included in every algebraically transverse model. It follows that the order of every base change for the Gröbner limit divides the order of the base change for every other algebraically transverse model. The uniqueness statement is a consequence. \square

Remark 2.6.12. We have the following important consequence of the algorithm for producing algebraically transverse limits. If $G \hookrightarrow \Sigma_X$ denotes the 1-complex associated to the minimal algebraically transverse limit, G may involve refining the minimal polyhedral structure \bar{G} by adding bivalent vertices to subdivide edges. However, if we look at the irreducible components Y_i corresponding to these bivalent vertices, the embedded subscheme $Z_i \subset Y_i$ is stable, that is, not fixed by the natural \mathbb{G}_m -action on Y_i .

With the results of this section as motivation, we define a *subscheme of X relative to D* .

Definition 2.6.13 (Relative subschemes). Let S be a logarithmic scheme. A *subscheme of X relative to D over S* is an expansion \mathcal{Y}/S of X over S together with a flat family of algebraically transverse subschemes $\mathcal{Z} \subset \mathcal{Y}$.

Proposition 2.6.5 suggests that algebraically transverse subschemes will satisfy the valuative criterion for properness. Our goal is now to construct moduli for these objects – first for the expansions themselves which happens in the next section and then for the algebraically transverse subschemes which happens in the one after.

3. Moduli of target expansions

The purpose of this section is to construct moduli spaces for the expansions in the previous section in analogy with the stack of expansions of a smooth pair; see [5, Section 6.1] for an exposition of the latter. Our approach is to use the dictionary between combinatorial moduli spaces and Artin fans, which we review in the first subsection. Using this, we then study the combinatorial moduli space parametrizing embedded 1-complexes inside Σ_X and show it can be given the structure of a cone space in the sense of [16].

Two important subtleties arise in this process. First, as discussed in the introduction, the choice of cone decomposition on the space of embedded 1-complexes is not unique and requires an auxiliary combinatorial choice. Second, in order to endow the universal family with the structure of a cone space, further subdivision is required, which translates into allowing additional codimension-1 bubbling in the geometric expansion. See Remarks 3.4.2 and 3.5.1.

3.1. Artin fans and cone stacks

The stack $\text{Exp}(X|D)$ is similar in nature to Olsson’s stack LOG of logarithmic structures and are instances of the Artin fans that are extracted from LOG; see [9, 62].

An *Artin cone* is a global toric quotient stack:

$$A_P = \text{Spec } \mathbb{C}[P] / \text{Spec } \mathbb{C}[P^{gp}],$$

where P is a toric monoid. These stacks are endowed with a logarithmic structure in the smooth topology from the toric logarithmic structure on $\text{Spec } \mathbb{C}[P]$. Observe that A_P recovers P from its logarithmic structure, so these data naturally determine each other.

Definition 3.1.1. An *Artin fan* is a logarithmic algebraic stack A that has a strict étale cover by a disjoint union of Artin cones.

In practice, this means that we are interested in algebraic stacks that arise by gluing stacks of the form A_P along open substacks corresponding to torus invariant subvarieties of $\text{Spec } \mathbb{C}[P]$. Note that we do not impose the faithful monodromy condition on Artin fans, but our constructions are explicit rather than conceptual, and in practice all of our Artin fans will be monodromy free.

Under mild assumptions, every logarithmic structure on a scheme arises from a morphism to an Artin fan [5, Proposition 3.2.1]. In practice, it is convenient pass to an equivalent but explicitly combinatorial 2-category, for which we follow [16, Sections 2.1–2.2].

Definition 3.1.2 (Cone spaces). A *cone space* is a collection $\mathcal{C} = \{\sigma_\alpha\}$ of rational polyhedral cones together with a collection of face morphisms \mathcal{F} such that the collection \mathcal{F} is closed under composition, the identity map of every cone σ_α lies in \mathcal{F} , and every face of a cone is the image of exactly one face morphism in \mathcal{F} .

A *cone complex* is obtained when the further assumption is placed that there is at most one morphism between any two cones. A *cone stack* can be defined similarly, as a fibered category over the category of rational polyhedral cones satisfying natural conditions. We will not need to work with cone stacks in a serious way, so we refer to the reader to [16, Section 2] for a definition. Their results allow us to systematically work with Artin fans using combinatorics. In fact, any cone space can be subdivided into a cone complex, so even cone spaces can be avoided for our purposes. However, it can be convenient to allow them to construct minimal choices of stacks of target expansions; for example, the stack of expansions in [47] arises from a cone space.

Theorem 3.1.3 [16, Theorem 3]. *The 2-categories of Artin fans and of cone stacks are equivalent.*

Under this equivalence, an Artin cone A_P is carried to the cone $\text{Hom}(P, \mathbb{R}_{\geq 0})$. If Σ is a fan embedded in a vector space and W_Σ is the resulting toric variety with dense torus \mathbb{G}_Σ , the equivalence carries the global quotient Artin fan $[W_\Sigma/\mathbb{G}_\Sigma]$ to the cone complex Σ .

3.2. Embedded complexes

Section 2 constructed expansions of X along D that accommodate limits of transverse subschemes. The construction was presented over valuation rings with logarithmic structure. In order to globalize this, we construct moduli for the tropical data over higher-dimensional cones and then glue these cones to form a cone space. By the categorical equivalence in Theorem 3.1.3, we will have produced $\text{Exp}(X|D)$ as an Artin stack.

Recall that Σ is a cone complex with smooth cones such that the intersection of any two cones is a face of each, and we view it as a subcomplex of the standard orthant by the embedding

$$\Sigma \hookrightarrow \mathbb{R}_{\geq 0}^{\Sigma(1)}.$$

We will make use of the Euclidean geometry of this ambient vector space.

In Section 2.2 we defined an *abstract 1-complex* to be a possibly disconnected metric graph G together with polyhedral complex structure. An *embedded 1-complex in Σ* , or *1-complex* for short, was an injective piecewise affine map of polyhedral complexes

$$G \hookrightarrow \Sigma$$

given by piecewise affine functions of slope 1. Let $|\mathcal{T}(\Sigma)|$ be the set of isomorphism classes of 1-complexes G equipped with an embedding

$$G \hookrightarrow \Sigma.$$

For each embedded 1-complex, we can dilate the embedded complex by any positive real scalar, giving rise to an copy of $\mathbb{R}_{\geq 0}$ in $|\mathcal{T}(\Sigma)|$. We will endow this set with the structure of a cone space, whose rays will be a subset of the rays obtained in this fashion.

3.3. Construction of cones of embedded graphs

By the conventions in this text, a *graph* refers to a possibly disconnected finite graph without loop or parallel edges, together with a finite collection of rays placed at the vertices. As is standard in Gromov–Witten theory, we visualize the rays as ‘legs’.

In order to put a cone space structure on $|\mathcal{T}(\Sigma)|$, we first study spaces of maps from graphs to Σ and then identify maps with the same image. For now, we suppress discussion of the integral structure and fix the correct integral lattice afterwards.

Definition 3.3.1. A *combinatorial 1-complex G in Σ* is a graph \underline{G} with vertex, ray and edge sets $V(\underline{G})$, $R(\underline{G})$ and $E(\underline{G})$, equipped with the following labeling data:

- (C1) For each vertex $v \in V(\underline{G})$, a cone σ_v in Σ .
- (C2) For each flag consisting of a vertex v incident to an edge or leg $e \in E(\underline{G}) \cup R(\underline{G})$, a choice of cone σ_e together with a nonzero primitive integral vector in the lattice of σ_e (referred to as a *edge direction*)

subject to the compatibility condition that if v is an incident vertex of an edge or ray e , the cone σ_v is a face of σ_e , including possibly σ_e itself.

We define a *linear path* in G to be a path (consisting of edges and/or rays) where every edge/ray has the same direction vector. Let $P(G)$ denote the set of linear paths in G .

An isomorphism of combinatorial 1-complexes is an isomorphism of the underlying graphs that is compatible with the labels. Fix a representative for each isomorphism class. We consider embeddings of a combinatorial 1-complex into Σ . Let $|\Sigma|$ denote the support of the cone complex. Define the set \mathbb{X}_G to be the set of functions

$$f : V(\underline{G}) \rightarrow |\Sigma|$$

such that

- (1) For each $v \in V(\underline{G})$, the image $f(v)$ lies in σ_v .
- (2) For each $e \in E(\underline{G})$ with adjacent vertices v and w , the line segment between $f(v)$ and $f(w)$ has direction vector equal to the one labeling e .

Each such function $f \in \mathbb{X}_G$ determines a polyhedral complex in Σ . It can be constructed as follows. Take the collection of points obtained as images of vertices of \underline{G} under f . Given an edge e between v and w introduce a segment between $f(v)$ and $f(w)$, whose edge direction is necessarily the one given by e . For each ray, glue an unbounded ray with the corresponding base point and edge direction dictated

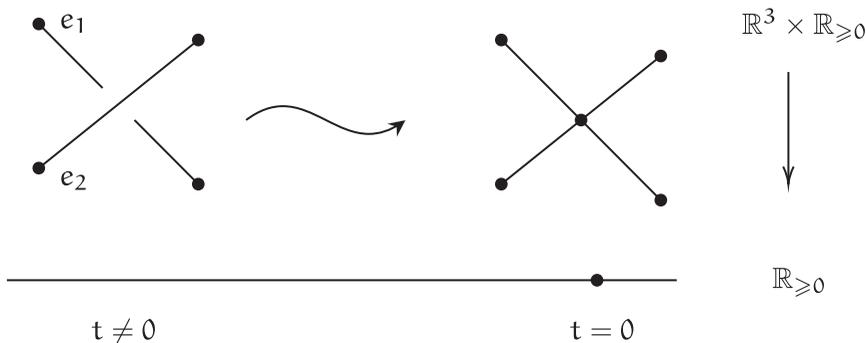


Figure 4. The figure shows a family of embedded 1-complexes in \mathbb{R}^3 over a one-dimensional base. The 1-complex on the left corresponds to a generic point, but at $t = 0$ the two skew edges cross forming a new vertex. As a consequence, the total space has two two-dimensional cones meeting at a point in their interior. In order to obtain a total space that, after taking cones, is a cone complex, vertices have to be added to e_1 and e_2 .

by G . This topological space has a minimal polyhedral complex structure whose vertices are of two types: (1) images of vertices of \underline{G} under f and (2) points where edges or rays intersect.

We refer to this embedded complex as the *image*,⁷ and denote it by $\text{im}(f)$. We restrict to those G such that there exists at least one function $f \in \mathbb{X}_G$ such that $\text{im}(f)$ is isomorphic to the complex $|\underline{G}|$, the latter with its canonical polyhedral structure.

Lemma 3.3.2. *The set \mathbb{X}_G has the structure of a cone.*

Proof. The set of all functions $f : V(\underline{G}) \rightarrow |\Sigma|$ satisfying only the first of the two conditions above is given by the product of σ_v ranging over all vertices v . Inside this product, we consider the locus where the second condition is imposed. Consider an edge E with incident vertices v and w . Consider the line containing $f(v)$ with edge direction given by e . The intersection point of this line with the cone σ_w is linear in the coordinates of $f(v)$. It follows that the locus where the second condition above is satisfied for a given edge e is linear and that \mathbb{X}_G is cut out from the product of all cones σ_v by a finite collection of linear equations. This gives us a cone structure on \mathbb{X}_G . \square

As mentioned above, each point $f \in \mathbb{X}_G$ determines an embedded 1-complex $\text{im}(f)$. Our primary interest is in embedded complexes, so we would like to replace the space \mathbb{X}_G with the corresponding set of images. However, the topological type of the image is not constant even for points in the relative interior of a cone in \mathbb{X}_G . A key example is provided by two skew edges in a three-dimensional cone that intersect in their interiors. We record the example explicitly to illustrate the behavior.

Example 3.3.3. Consider the family of embedded 1-complexes in $\mathbb{R}_{\geq 0}^3$ parameterized by a cone σ which we identify with $\mathbb{R}_{\geq 0}^2$ and furnish with coordinates (s, t) . At a point (s, t) , the associated 1-complex is a union of two components $\rho_1(s, t)$ and $\rho_2(s, t)$. At the point (s, t) , the ray $\rho_1(s, t)$ is the unbounded ray connecting $(0, 0, s)$ and $(1, 1, s)$. At the point (s, t) , the segment $\rho_2(s, t)$ is the line segment that connects $(t, 0, 0)$ and $(0, t, t)$. The two segments intersect along a codimension 1 subspace in the interior of the base cone σ . The total spaces of $\rho_1(s, t)$ and $\rho_2(s, t)$, taken individually, form two-dimensional cones in the four-dimensional space $\mathbb{R}_{\geq 0}^3 \times \sigma$. However, their union is not a cone complex, as their intersection is not a face of either. For each fixed value of (s, t) , there is a minimal polyhedral structure on the union $\rho_1(s, t) \cup \rho_2(s, t)$. The combinatorial type, and in particular the face poset, of this fiber can change in the interior of the cone σ . The figure below displays a slice of this picture; see Figure 4

⁷The terminology is mildly inaccurate since there may be no polyhedral map from the source of f to its image as a polyhedral complex without first subdividing the source. It is, however, the smallest polyhedral structure on the topological image with the property that images of vertices under f are vertices.

The topological changing in the interior of a cone, as in the above example, can be thought of as a nonflatness of the family of 1-complexes. The following definition is meant to capture this phenomenon of varying topological type. The topological type can also change due to specialization, for example, a vertex moving from a cone to a face. But it is the changes of topological type in the interior that is the exotic phenomenon, as compared to earlier moduli problems [1, 16, 27].

Definition 3.3.4. Let G and H be combinatorial 1-complexes in Σ_X . A *surjection* from G to H is given by the a pair of maps

$$\begin{aligned} \tau : V(\underline{G}) &\rightarrow V(\underline{H}) \\ \nu : E(\underline{G}) \sqcup R(\underline{G}) &\rightarrow P(\underline{H}) \end{aligned}$$

subject to the following conditions for:

- (1) If v and w are the vertices adjacent to e , then $\tau(v)$ and $\tau(w)$ are the endpoints of the path $\nu(E)$.
- (2) Let e be an edge of \underline{G} . Orient the edges of $\nu(e)$ to form a directed path and choose the orientation of E compatibly, that is, if e is directed from endpoint v to endpoint w then $\nu(e)$ is oriented from $\tau(v)$ to $\tau(w)$. Then,
 - If e is an edge of \underline{G} , then the edge directions for each edge in $\nu(e)$ coincide with that of e .
- (3) For each vertex v of \underline{G} , the cone $\sigma_{\tau(v)}$ is equal to a face of σ_v .
- (4) Each edge or ray in H is contained in a path which is in the image of ν .

Each point $f \in \mathbb{X}_G$ determines a surjection of G , as follows. Given a point $f \in \mathbb{X}_G$, we consider graph associated to the image $\text{im}(f)$, considered with its natural structure as a combinatorial 1-complex. A basic finiteness result holds.

Lemma 3.3.5. *Let G and \mathbb{X}_G be as above. The set of surjections of G associated to points of \mathbb{X}_G is finite.*

Proof. We may view the cone \mathbb{X}_G as a cone of tropical maps with type G to Σ in the sense of [27, Remark 1.21]. It comes equipped with a moduli diagram

$$\begin{array}{ccc} \mathbf{G} & \longrightarrow & \Sigma \times \mathbb{X}_G \\ \downarrow & & \\ \mathbb{X}_G & & \end{array}$$

The fibers of the vertical map are the metrizations of G given by endowing each edge e with length equal to the distance in Σ between the images of its endpoints. Since the horizontal map is linear on cones, the image of \mathbf{G} in $\Sigma \times \mathbb{X}_G$ can be given the structure of a cone complex. Choose any such cone complex structure $\text{im}(\mathbf{G})$. After subdividing this image further and replacing \mathbb{X}_G by a subdivision, the induced map

$$\tilde{\mathbf{G}} \rightarrow \tilde{\mathbb{X}}_G$$

is flat in the combinatorial sense – every cone of the source surjects onto a cone of the image. This follows from [7, Section 4]. After this subdivision, the combinatorial type of the 1-complex in the fibers in the relative interior of any cell is constant. It follows that there are only finitely many surjections that appear from taking image, as claimed. □

Let G be as above, and let H be a surjection of G associated to a point in \mathbb{X}_G . Note that we will only consider surjections associated to points of cones \mathbb{X}_G . Let \mathbb{X}_H be the cone associated to the combinatorial 1-complex H . There is an associated inclusion

$$\mathbb{X}_H \hookrightarrow \mathbb{X}_G.$$

Specifically, a point of \mathbb{X}_H in particular determines a function on $V(\underline{G})$ by using the map on vertex sets $V(\underline{G}) \rightarrow V(\underline{H})$.

Lemma 3.3.6. *Let $G \rightarrow H$ and $G \rightarrow K$ be two surjections of G obtained from points of \mathbb{X}_G by the image construction. Consider a point*

$$f \in \mathbb{X}_H \cap \mathbb{X}_K \subset \mathbb{X}_G.$$

The surjection $G \rightarrow J$ determined by f is a common surjection of H and K .

Proof. The lemma follows immediately from the definitions by taking J to be the image of f . □

Fix a combinatorial 1-complex G and let Aut_G be the automorphism group. There is a natural action of Aut_G on the cone \mathbb{X}_G . Our goal will be construct an automorphism-equivariant subdivision of \mathbb{X}_G such that the subsets \mathbb{X}_H for surjections H of G become subcomplexes. The following lemma will be of use to us.

Lemma 3.3.7. *Let C be a cone, and let $F \subset C$ be an embedded cone complex. Let Γ be a finite group acting on C and such that F is Γ -stable. There exists a Γ -equivariant complete subdivision \tilde{C} of C such that F is a union of faces of \tilde{C} . If F is a smooth cone complex, then \tilde{C} can be chosen to be smooth.*

Proof. Let $N_{\mathbb{R}}$ denote the associated vector space of the cone C . We consider F and C as subsets of this vector space. The existence of equivariant completions for toric varieties guarantees that there exists a complete fan Δ in $N_{\mathbb{R}}$ that contains F as a subcomplex [23]. Consider the common refinement of Δ and C :

$$C' = \{\sigma_1 \cap \sigma_2 : \sigma_1 \in \Delta \ \sigma_2 \in C\}.$$

This is a cone complex structure on the intersection of the supports of the two fans, that is on $|C|$; see [48, Section 2.3]. It clearly contains F as a subcomplex. The subdivision is not Γ equivariant. We fix this by averaging over the group. Given $\gamma \in \Gamma$, the set of translates of cones in C' form a cone complex with support $|C|$. The common refinement over all group elements gives rise to the subdivision \tilde{C} as claimed in the lemma. Since F is Γ -stable, this last common refinement step does not subdivide it. The final statement on smoothness follows from the fact that toric resolution of singularities for any fan can be performed without subdividing smooth cones [20, Theorem 11.1.9] and the fact that these subdivisions can be made Γ equivariant [8, Section 2]. □

Note that this procedure typically involves choices and is not canonical. Our application is to the cones associated to combinatorial one complexes.

Lemma 3.3.8. *For each combinatorial 1-complex G , there exists an Aut_G -equivariant subdivision $\mathbb{Y}_G \rightarrow \mathbb{X}_G$ such that for each surjection $G \rightarrow H$ induced by a point of \mathbb{X}_G , the induced map*

$$\mathbb{Y}_H \hookrightarrow \mathbb{Y}_G$$

is an inclusion of a subcomplex. If \mathbb{Y}_H is smooth, then \mathbb{Y}_G can be chosen to be smooth.

Proof. Given G , the set of all surjections of G forms a partially ordered set since the notion is stable under composition. Choose any total order extending this partial order. At the minimal elements of this partial order, there is no subdivision necessary, and the union of the cones associated to these types forms a cone complex in \mathbb{X}_G . Proceeding in the chosen order, we repeatedly apply Lemma 3.3.7. Given H and its surjections J_1, \dots, J_k , we observe that by Lemma 3.3.6, the union of the inductively chosen \mathbb{Y}_{J_i} is a cone complex. We may therefore apply the preceding lemma to obtain an Aut_H -equivariant subdivision \mathbb{Y}_H of \mathbb{X}_H such that each previously constructed \mathbb{Y}_{J_i} is a union of faces. The lemma follows. The smoothness statement follows from the corresponding statement in the preceding lemma. □

3.4. Tropical moduli of expansions

We now construct the moduli space of expansions via a colimit procedure. Fix a set of subdivisions \mathbb{Y}_G as guaranteed by the lemma above, one for each choice of combinatorial 1-complex G . Consider the cone complex given by the disjoint union $\sqcup\{\mathbb{Y}_G\}$ taken over all such G . We want to identify two cones σ and σ' if they are identified via an inclusion $\mathbb{Y}_H \hookrightarrow \mathbb{Y}_G$ induced by a surjection or identified via an automorphism of G . Consider the equivalence relation on cones generated by these identifications. In what follows, we will take the quotient by this equivalence relation. The quotient exists in the category of generalized cone complexes,⁸ defined by Abramovich–Caporaso–Payne [1, Section 2.6]. In showing that such quotients exist in their category, those authors also explain that the quotient also carries a natural cone space (in fact, even a cone complex) structure after subdivision, obtained as follows. We subdivide each \mathbb{Y}_G by taking barycentric subdivision of all the cones. Under the barycentric subdivision procedure above, if an element of Aut_G stabilizes a cone of \mathbb{Y}_G , then it fixes it pointwise. The cones of the quotient are precisely the equivalence classes considered above, generated by inclusions as G varies or by automorphisms of G . Note that the notion of a graph surjection captures both simple specialization phenomena, such as contracting edges of G , as well as the more exotic topological type changes in the interior.

Definition 3.4.1. A moduli space of tropical expansions T is the quotient of the disjoint union of cone complexes \mathbb{Y}_G over all combinatorial 1-complexes G , with cone structure as described in the preceding paragraph. The underlying set will be denoted $|T|$.

The quotient as a set is precisely the set of embedded 1-complexes in Σ . This also endows it, for example, with the structure of a topological space. We find this topological structure useful in thinking about the problem, but do not need it any technical sense here. This latter structure is examined more thoroughly in [41].

By toric resolution of singularities, we can guarantee that the cones of T are smooth. From this point forth, we only work with those T that are smooth. In parallel constructions in Gromov–Witten theory, it is natural to allow toric singularities in analogous spaces [9, 27]. We choose not to do this here; however, this is largely for convenience in the applications that we currently foresee [56].

Remark 3.4.2. Note that the cone spaces T constructed above are nonunique. The ambiguity arises from the inductive construction of the subdivision \mathbb{Y}_G as a subdivision of \mathbb{X}_G for each graph G . The construction is analogous to, and derived from, the result on equivariant completions of toric varieties and on completions of fans [23]. In this sense, the root cause of the nonuniqueness of our spaces is the same as the nonuniqueness seen in the compactification problems for tori. In particular, dropping the smoothness discussed above does *not* alleviate the nonuniqueness.

3.5. Universal family and integral structure

Given a moduli space T of tropical expansions, we want to equip it with a universal expansion. More precisely, we will construct a diagram of cone spaces

$$\begin{array}{ccc} \mathbb{Y} \subset \Sigma \times T & \longrightarrow & \Sigma \\ \downarrow & & \\ T, & & \end{array}$$

where $\mathbb{Y} \subset \Sigma \times T$ is a subdivision and the map $\mathbb{Y} \rightarrow T$ is flat in the combinatorial sense: That is, cones map surjectively onto cones.

⁸A generalized cone complex can be thought of as the coarse moduli space of a cone stack. We are interested here only in cone spaces, so the formalisms of [1] and [16] coincide.

In order to construct this diagram, we observe that each \mathbb{X}_G comes equipped with a universal graph G and a map to $\Sigma \times \mathbb{X}_G$. In order to construct Y we simply replicate every step in the construction of T above, modified as follows.

For each G , we consider the augmented graph G_* consisting of the disjoint union of G with a new distinguished vertex $*$ (with no edges connecting it to vertices of G). If we denote s a choice of vertex, edge or ray of G , we consider the set $\mathbb{U}_G(s)$ consisting of a function from G_* to Σ whose restriction to G is in \mathbb{X}_G and such that $f(*)$ lies on the image of $f(s)$.

Each element of $\mathbb{U}_G(s)$ determines a 1-complex in Σ equipped with distinguished *basepoint* lying on s ; this basepoint is the image $f(*)$. The set $\mathbb{U}_G(s)$ forms a cone since the choice for the basepoint is a cone, namely the cone formed by the edge or vertex s , as f varies in \mathbb{X}_G . As we consider all possible choices of s , these cones $\mathbb{U}_G(s)$ form a cone complex \mathbb{U}_G equipped with an action of Aut_G and a morphism

$$\iota_G : \mathbb{U}_G \rightarrow \mathbb{X}_G \times \Sigma.$$

Given a surjection $G \rightarrow H$, an edge s of G and an edge t of H which is contained in the path $\nu(s)$, we have

$$\mathbb{U}_H(t) \hookrightarrow \mathbb{U}_G(s).$$

We have similar maps when s or t are rays or vertices. Notice given a surjection, there could be many distinct embeddings of $\mathbb{U}_H(t)$ into \mathbb{U}_G .

Following our previous approach, Lemmas 3.3.5, 3.3.6, 3.3.7 and 3.3.8 each apply in this setting of pointed graphs. They give rise to another system of subdivisions $\widetilde{\mathbb{U}}_G$ of \mathbb{U}_G , which are compatible with relabeling, the pullback of the subdivision \mathbb{Y}_G and the inclusions of $\widetilde{\mathbb{U}}_H(t)$. If we apply the map ι_G to this subdivision, it is injective on every cone, and the images of any two cones intersect along a unique face of each. Therefore, there is a natural cone structure on its image $\mathbb{V}_G \subset \mathbb{Y}_G \times \Sigma$, along with natural maps $\mathbb{V}_H \rightarrow \mathbb{V}_G$ associated to surjections. We pass to the quotient as before and declare the colimit to be Y .

Finally, we fix the integral structure on T and Y as follows. Given a rational point P of T , we say it is *integral* if every vertex of the embedded 1-complex $Y_P \subset \Sigma$ lies on an integral lattice point of Σ . Note that the process is not circular: The lattice structure on the target Σ is unambiguous.⁹ Similarly, a rational point $Q \in Y$ is integral if its image in $T \times \Sigma$ is integral. One can check that this definition of integral points in each cone σ of T is compatible with the usual notion of integral structure, that is, it is the restriction of a lattice in the ambient rational vector space. It follows from our definition that the map $Y \rightarrow T$ is flat and reduced.

Remark 3.5.1. Given any point P of T , we obtain two 1-complexes associated to this point. First, since we can identify points of the topological realization of T with the points $|T|$, we have the 1-complex G_P associated to this point. Second, we can take the fiber Y_P of the map $Y \rightarrow T$. It follows from the construction that there is a map

$$Y_P \rightarrow G_P$$

which is a refinement of polyhedral sets, obtained by adding 2-valent vertices along edges. We will refer to the vertices of Y_P that are not present in G_P as *tube vertices*, and these will play an important role in our definition of stability.

As with our construction of T , a combinatorial choice is required in the construction of Y . Consequently, given P , while the 1-complex G_P is canonically given, the subdivision Y_P depends on this choice. Moreover, the integral structure on T also depends on this additional choice.

Remark 3.5.2 (Relative compactification). Fix a moduli space of tropical expansions T and a finite type subspace $T_{\beta,n}$. We abuse notation mildly, allowing Y to denote the universal embedded 1-complex

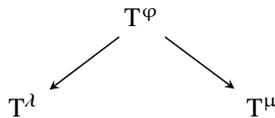
⁹Said differently, a 1-complex is integral if all its vertices are integral. A point of T is integral if the vertices of the universal 1-complex above it is integral

in $T_{\beta,n} \times \Sigma_X$. The subdivisions we have constructed for the universal target expansions are typically nonproper. For appropriate choices in the construction of T , it is possible to ensure that Y can be completed to complete, and in fact, projective, subdivision of $T_{\beta,n} \times \Sigma_X$. Indeed, after iterated barycentric subdivision $T_{\beta,n}$ can be embedded in a vector space [6, Section 4.6]. Results on equivariant completions for fans of toric varieties and toric Chow’s lemma give rise to the completion [23]. Toroidal semistable reduction guarantees flatness after a further sequence of blowups of $T_{\beta,n}$.

3.6. Common refinements

The different choices involved in constructing T above lead to different cone spaces, but they only differ by subdivision. We record this for future use. Let Λ denote the set of all conical structures on the set $|T|$ that arise from the construction above. Given $\lambda \in \Lambda$, let T^λ denote the associated moduli space of tropical expansions.

Proposition 3.6.1. *Any two conical structures $\lambda, \mu \in \Lambda$ share a common refinement. Precisely, there exists $\varphi \in \Lambda$ and morphisms*



such that each of the two arrows is a composition of a complete subdivision and passage to a finite index sublattice.

Proof. We reinspect the procedure described above. For each combinatorial 1-complex G and surjection $G \rightarrow H$, we choose a subdivision of \mathbb{X}_G such that the cone complex \mathbb{Y}_H is a union of faces of this subdivision. In turn, there are choices involved in constructing \mathbb{Y}_H from \mathbb{X}_H , given by studying surjections $H \rightarrow K$ and so on. The moduli spaces T are then formed by taking the colimit of the quotients of these cones by the appropriate automorphism groups. However, we note that two cone complexes on the same underlying space always have a common refinement, whose cones are given by intersections of the cones in the two complexes. Given cone complexes \mathbb{Y}_G^λ and \mathbb{Y}_G^μ , as their support is the cone \mathbb{X}_G we can construct their common refinement \mathbb{Y}_G^φ . The common refinement yields a compatible system of subdivisions for the different combinatorial 1-complexes in the following sense. Given a graph surjection $G \rightarrow H$ and an inclusion

$$\mathbb{X}_H \hookrightarrow \mathbb{X}_G,$$

we may take the common refinement of \mathbb{X}_H with each of these three cone complex structures on \mathbb{X}_G to obtain \mathbb{Y}_G^φ to obtain $\mathbb{Y}_H^\lambda, \mathbb{Y}_H^\mu$ and \mathbb{Y}_H^φ . A definition chase reveals that \mathbb{Y}_H^φ is the common refinement of \mathbb{Y}_H^λ and \mathbb{Y}_H^μ . Further, the common refinement of a Aut_G -equivariant subdivisions of \mathbb{X}_G remains equivariant. It follows that the preceding colimit construction can be run with the quotients $\mathbb{Y}_G^\varphi // \text{Aut}_G$. By construction, this cone space refines both T^λ and T^μ . \square

Remark 3.6.2. There is a broader context for the construction here that we have skirted around in order to keep the discussion concrete. Given a cone complex Σ , the integer points in its topological realization is canonically identified with the \mathbb{N} -valued points of the functor that Σ defines on monoids. Precisely, Σ defines the functor

$$\begin{aligned}
 F_\Sigma &: \text{Monoids} \rightarrow \text{Sets} \\
 P &\mapsto \text{Hom}(P^\vee, \Sigma).
 \end{aligned}$$

However, we notice that the \mathbb{N} -valued points of Σ and any complete subdivision of Σ coincide. More generally, the functors coincide on all *valuative monoids*. In this sense, although the construction of T above is not canonical, the associated functor on *valuative monoids* is canonical. One may view the discussion above as reverse engineering this. We take the integral points set $|T|$ to be the \mathbb{N} -valued points of a putative set of cones T^{\downarrow} on all monoids, which we construct. Functors on valutive monoids (resp. valutive logarithmic schemes) often arise as inverse limits of cone complexes (resp. logarithmic schemes) along subdivisions (resp. logarithmic modifications). However, a functor defined directly on the valutive category needs to satisfy additional properties in order to arise in this fashion. Once these are satisfied, the functor in question, in our case $|T|$ is determined as a polyhedral complex up to refinement. Rather than working with the inverse limit directly, we work with the entire inverse system, making compatible choices to produce morphisms where necessary. There are, at this point, relatively few genuinely valutive approaches to moduli problems. The logarithmic Picard group is a notable exception [59].

3.7. Geometric moduli

We now pass to Artin fans to produce our moduli stack of target expansions. Consider Σ_X the cone complex for $(X|D)$.

Given a moduli space of tropical expansions T for Σ_X and its universal family Y , we set $\text{Exp}(X|D)$ to be the Artin fan associated to the cone space T via Theorem 3.1.3. It carries a natural logarithmic structure and is smooth if we choose T to be smooth. The points of $\text{Exp}(X|D)$ correspond to cones of T .

To produce a universal expansion over $\text{Exp}(X|D)$, we apply the Artin fan construction to the inclusion $Y \hookrightarrow \Sigma \times T$ to obtain

$$A_Y \rightarrow A_{\Sigma} \times \text{Exp}(X|D).$$

Using the natural map $X \rightarrow A_{\Sigma}$, we set

$$\mathcal{Y} := A_Y \times_{A_{\Sigma} \times \text{Exp}(X|D)} (X \times \text{Exp}(X|D))$$

and consider the diagram

$$\begin{array}{ccc} \mathcal{Y} & \longrightarrow & X \\ \downarrow & & \\ \text{Exp}(X|D) & & \end{array}$$

Since we arranged for the morphism of cone complexes $Y \rightarrow T$ to be flat and reduced, \mathcal{Y} is a family of expansions of X over $\text{Exp}(X|D)$.

Recall that, for every point P of T , we have a distinguished set of *tube vertices* which are 2-valent vertices that appear in the subdivision $Y_P \rightarrow G_P$. Given a cone σ of T , the combinatorics of the universal complex does not change in the interior of σ , so the tube vertices of Y are well defined over σ by choosing any point P in its interior. If \mathcal{Y}_{σ} denotes the geometric expansion corresponding to this point of $\text{Exp}(X|D)$, the irreducible components corresponding to these tube vertices will be called *tube components*. These are components that are not needed for applying the valutive criterion but appear when producing families over the entire stack of degenerations. They play a significant role in the next section.

4. Moduli of ideal sheaves

In this section, we define stable ideal sheaves on X relative to D and show that, for fixed numerical invariants, their families are parameterized by a proper Deligne–Mumford stack $\text{DT}(X|D)$. For expository clarity, we will assume that X has dimension 3. Removing this assumption is straightforward but notationally cumbersome.

Let $\text{Exp}(X|D)$ be a stack of target expansions¹⁰ of X along D , and let $T(\Sigma_X)$ the associated tropical moduli stack. We will denote it by T when there is no chance of confusion.

4.1. Tube components and tube subschemes

The expansions constructed in the previous section contain certain distinguished components. In order to obtain well-behaved moduli spaces, we restrict subschemes inside such components.

Let $G \subset \Sigma_X$ be an embedded 1-complex. A *linear 2-valent vertex* is a 2-valent vertex v such that two conditions hold: (i) v and both of its outgoing edges lie in the interior of a single cone of Σ_X and (ii) the directions of the two flags are opposite, that is, a neighborhood of v lies on a line inside the cone σ_v .

Fix a moduli space of 1-complexes T equipped with a cone structure as in the previous section, as well as a cone complex on the associated family $Y \rightarrow T$. Recall that the underlying set $|T|$ is in natural bijection with the set of embedded 1-complexes in Σ_X .

Choose a point p of T . There are two 1-complexes that one can associate to this point. First, let G_p be the 1-complex obtained by examining the fiber of p in Y . Second, let G_p^* be the embedded 1-complex associated to the point $p \in |T|$. There is a natural refinement $c : G_p \rightarrow G_p^*$ which erases some – but not all – the linear 2-valent vertices. We maintain this notation for the next definition.

Definition 4.1.1 (Tube vertex). Let v be a linear 2-valent vertex of G_p . The vertex v is called a *tube vertex* if its image $c(v)$ under the refinement map above is not a vertex of G_p^* .

Remark 4.1.2. Tube vertices are a subset of linear 2-valent vertices, and we have defined a tube vertex by forcing the linear 2-valency condition for clarity. It should be noted that vertices that are not linear and 2-valent could never arise from a subdivision, so in fact the linear 2-valency condition could be removed from the preceding definition without change.

The notion geometrizes as follows. Given a point of $\text{Exp}(X|D)$, there is an associated expansion \mathcal{Y}_p there is an associated cone of T . Given any point p in the interior of this cone, we obtain a 1-complex G_p . A component of \mathcal{Y}_p is called a *tube component* if it corresponds to a tube vertex in G_p .

Note that each tube component is a \mathbb{P}^1 -bundle \mathbb{P} over a surface S .

Given a \mathbb{P}^1 - bundle \mathbb{P} over a surface S , a subscheme of \mathbb{P} is a *tube subscheme* if it is the schematic preimage of a zero-dimensional subscheme in S .¹¹ We say that a component of an expansion is *linear 2-valent* if the corresponding vertex in its 1-complex is linear 2-valent. Linear 2-valent components are all \mathbb{P}^1 -bundles over a surface S .

Tube terminology: a summary There are three different objects that have the adjective ‘tube’. A tube vertex is a distinguished linear 2-valent vertex of a 1-complex. A tube component is a component of an expansion that corresponds to such a tube vertex. And finally, a tube subscheme is a subscheme of a linear 2-valent component of an expansion (so in particular a \mathbb{P}^1 -bundle) that is pulled back along the bundle. A priori, there is nothing that stops a nontube, linear 2-valent component from containing a tube subscheme. We will explicitly stop this from happening in our moduli problem, which we now come to.

4.2. Stable ideal sheaves

We now state our main definition.

Definition 4.2.1. A *stable ideal sheaf on X relative D* over a closed point p is a morphism

$$p \rightarrow \text{Exp}(X|D)$$

with associated expansion \mathcal{Y}_p and a proper, algebraically transverse subscheme $\mathcal{Z}_p \subset \mathcal{Y}_p$, which satisfies the following stability condition:

¹⁰We suppress the choice from the notation for brevity.

¹¹These are the analogues of ‘trivial bubbles’ in relative and expanded Gromov–Witten theory.

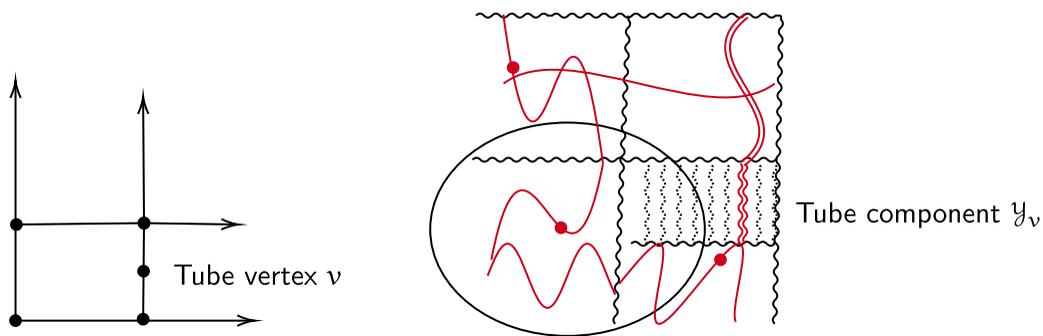


Figure 5. An 1-complex with a tube vertex on the left and the corresponding tube component on the right. The subscheme depicted in that tube component is a tube subscheme – it is attempting to depict the preimage of a zero-dimensional subscheme of length 2, along the projection. On the figure on the right, the circle in the bottom left indicates the main component of the expansion, while the rest of the components are bundles over strata. The tube component contains a tube subscheme, and the picture is meant to depict a subscheme satisfying DT stability.

DT stability: the subscheme Z_P restricted to a linear 2-valent component of \mathcal{Y}_P is a tube subscheme if and only if that component is a tube component.

A family of stable ideal sheaves on X relative to D over a scheme S is a morphism $S \rightarrow \text{Exp}(X|D)$ with associated expansion \mathcal{Y}_S and a flat family of algebraically transverse proper subschemes $Z_S \hookrightarrow \mathcal{Y}_S$ such that DT stability is satisfied in each geometric fiber.

Recall that algebraic transversality implies that Z meets every stratum of \mathcal{Y}_P . Also, notice that a family of stable ideal sheaves over a scheme S induces a logarithmic structure on S by pulling back the natural structure on the Artin fan $\text{Exp}(X|D)$.

Remark 4.2.2. The DT stability condition stems from the fact that while $|T|$ genuinely parameterizes 1-complexes, it is not a cone space or even a cone stack. On the other hand, given a cone space structure on T and its universal family Y , there may be points of T such that the universal family may be identical.

The DT stability condition itself can be understood as follows. Suppose we have two points P and Q in T such that the fiber in Y gives the same embedded 1-complex, and let us denote this 1-complex by Θ . However, P and Q are also points of $|T|$ which determine *distinct* 1-complexes. It follows that the fibers Y_P and Y_Q , which are both identified with Θ , are obtained from two distinct 1-complexes via subdivision. Necessarily then, the choices of points P and Q determine two distinct ways to coarsen Θ by erasing *certain* but *not all* 2-valent vertices. A vertex of Θ that is erased according to the point P will be a tube vertex for Y_P and will determine a tube component. However, such a vertex need not be erased according to the point Q . Passing to the corresponding geometric picture, the 1-complex Θ determines an expansion, and some of the components in this expansion may be \mathbb{P}^1 -bundle components. But a given such component might be a tube component in the expansion corresponding to P but not in the one corresponding to Q , and vice versa. If the component is a tube, the expansion can be blown down by erasing the vertex. The DT stability condition then insists that the subscheme is pulled back along this blow down. Figure 5 depicts the situation.

4.3. Geometry of the space of relative ideal sheaves

Let $\text{DT}(X|D)$ denote the fibered category over schemes of families of stable ideal sheaves on X relative to D . We prove that the moduli problem is algebraic with finite stabilizers.

Theorem 4.3.1. *For each fixed stack $\text{Exp}(X|D)$, the fibered category $\text{DT}(X|D)$ is an algebraic stack.*

The stack comes with a map to $\text{Exp}(X|D)$, and we equip it with the pullback logarithmic structure.

Proof. The proof is broken up into four steps; we prove algebraicity, first for the stack of expansions, then for the relative Hilbert scheme on the universal family of the stack of expansions, then for the substack of algebraically transverse maps and finally for the substack of transverse maps satisfying the DT stability condition.

First, the stack $\text{Exp}(X|D)$ is an Artin fan and therefore is algebraic. Second, since the map $\mathcal{Y} \rightarrow \text{Exp}(X|D)$ is separated, it follows from [72, Tag 0D01] that the Hilbert functor $\text{Hilb}(\mathcal{Y}/\text{Exp}(X|D))$ is representable by an algebraic space over $\text{Exp}(X|D)$.

We next consider algebraic transversality. Denote by $\text{Hilb}^\dagger(\mathcal{Y}/\text{Exp}(X|D))$ the open subspace parameterizing algebraically transverse subschemes without the DT stability condition.

Suppose we are given a family of expansions $\mathcal{Y}_S \rightarrow S$ and a family of subschemes $Z_S \subset Y_S$. Since algebraic transversality is characterized by the vanishing of higher Tors, it is a constructible condition. Therefore, it suffices to show it is stable under generization, so we assume S is a DVR. Recall that any expansion is naturally stratified since it arises as the special fiber of a logarithmically smooth degeneration. Given any expansion over S , the closure of any stratum lying over the generic fiber contains strata of the closed fiber. Therefore, by flatness, every stratum of the generic fiber must intersect Z nontrivially. Similarly, by upper semicontinuity, the Tor-vanishing condition for divisor strata of the generic fiber follows from the same condition on the closed fiber (by choosing a divisor stratum of the closed fiber in the closure).

Finally, we deduce openness of the DT stability condition from the following lemmas. □

Lemma 4.3.2. *The DT stability condition is open, and therefore $\text{DT}(X|D) \subset \text{Hilb}(\mathcal{Y}/\text{Exp}(X|D))$ is an open algebraic substack.*

Proof. As before, we need to check that the condition of DT stability is preserved by generization. Let B be the spectrum of a discrete valuation ring with generic point η and closed point b , and consider a morphism $B \rightarrow \text{Hilb}^\dagger(\mathcal{Y}/\text{Exp}(X|D))$, to the locus in the relative Hilbert scheme parameterizing algebraically transverse subschemes. Assume that DT stability is satisfied at 0 . We will prove that it is satisfied at the generic point.

Examine the map $B \rightarrow \text{Exp}(X|D)$. The generic and closed point each determine strata of the latter space. Let F_η and F_b be the corresponding cones in T ; note that F_η is a face of F_b . Each component determines a 1-complex embedded in Σ_X and by slight abuse of notation, we let G_η and G_b be the combinatorial types of the associated graphs and note that there is a specialization map¹² from G_b to G_η . To see this specialization, we may choose particular metric structures on G_b and G_η such that there is a morphism from a polyhedral complex $\mathcal{G} \rightarrow \mathbb{R}_{\geq 0}$, with $\mathcal{G} \subset \Sigma_X \times \mathbb{R}_{\geq 0}$ whose fibers over 0 and 1 are G_0 and G_1 , respectively. The combinatorial types of G_1 and G_b agree, and the combinatorial types of G_0 and G_η also agree. The specialization map is obtained by taking the cone over a face in G_1 and intersecting with the 0 fiber and then making the natural identifications with G_b and G_η .

Let \mathcal{Z} be the family of subschemes over B with generic and special fibers \mathcal{Z}_η and \mathcal{Z}_b . By the DT stability condition, we see that, at the closed point b , the subscheme \mathcal{Z}_b is a tube along the tube vertices of G_b . We wish to show that the same is true at η .

We claim that, if v is a tube vertex of G_η , then all vertices w of G_b that specialize to v are also tube vertices. Indeed, tube vertices arise by adding 2-valent vertices to subdivide the minimal graph structure on G_b and then specializing these vertices to obtain a subdivision of the minimal structure on G_η . It follows that the set of vertices specializing to a tube vertex must all be tube vertices.

Let v be a tube vertex of G_η . The star around v in G_η can be identified with the fan of \mathbb{P}^1 . Let w_1, \dots, w_k be the vertices of G_b specializing to v . Since tube vertices are always 2-valent, the star of v in the total space \mathcal{G} can be identified with the cone over a polyhedral subdivision of \mathbb{R} with k vertices.

¹²The reader may wish to have in mind that G_η typically simpler than G_b , in the sense of having fewer vertices.

The toric dictionary provides us with the following structure for the component Y_v dual to v over B . We have morphisms

$$Y_v \rightarrow S \rightarrow B,$$

where S is a smooth surface over B and $Y_v \rightarrow S$ is a proper family of semistable genus 0 curves that is generically a \mathbb{P}^1 bundle – that is, $X_{\eta} \rightarrow S_{\eta}$ is a \mathbb{P}^1 -bundle. The surface S is obtained as a locally closed stratum in a logarithmic modification of X and, in particular, is typically nonproper. Let Y_b be the union of the irreducible components corresponding to the vertices w_i .

After replacing our degeneration with this local model, we are left to check that if Z_b is a tube subscheme on every component of Y_b , then Z_{η} is also a tube subscheme. This is established by the lemma below, which then concludes the proof. \square

Lemma 4.3.3. *As above, let*

$$Y_v \rightarrow S \rightarrow B$$

be a generically smooth semistable genus 0 curve fibration over a family of smooth surfaces over B . Let Z_b and Z_{η} be the special and generic subschemes obtained by restriction to the local model as above. If $Z_b \subset Y_b$ is the preimage of a zero-dimensional subscheme $W_b \subset S_b$, then there exists a subscheme $W \subset S$ such that Z is the preimage of W .

Proof. By applying the criteria for properness and transversality from Theorems 1.5.2 and 1.5.3, we see that the family of subschemes obtained by the construction above to the component above is flat and proper. Properness and upper-semicontinuity of dimension for the family of subschemes guarantees that the image of Z_{η} in S_{η} is zero-dimensional. Since the special fiber Z_b is the preimage of a zero-dimensional subscheme along a nodal curve fibration, the special fiber has no embedded points. The Cohen–Macaulay condition is open, so the generic fiber also has no embedded points [72, Tag 0EOH].

We examine the Hilbert polynomial of the subscheme. By flatness of the family of subschemes, the fiber degree of the subscheme is constant. We therefore conclude that the subscheme Z_{η} is a Cohen–Macaulay thickening of a union of fibers in a \mathbb{P}^1 -bundle over S_{η} . For each such fiber, the thickening gives rise to a morphism from \mathbb{P}^1 to the Hilbert scheme of points $\text{Hilb}(S_{\eta}, d)$ on the surface S_{η} . The minimal possible holomorphic Euler characteristic of this thickening occurs if and only if the subscheme is a tube. Indeed, if the subscheme is not a tube, then the morphism $\mathbb{P}^1 \rightarrow \text{Hilb}(S_{\eta}, d)$ is nonconstant. The difference of the Euler characteristic and the degree of this moduli map is constant, from which the claim follows. \square

4.4. Boundedness

Fix a curve class β and holomorphic Euler characteristic n and consider the moduli space $\text{DT}_{\beta,n}(X|D)$ of ideal sheaves on X relative to D with these numerical data. Specifically, we consider subschemes for which the pushforward of the fundamental cycle to X is β . We establish boundedness of the moduli problem.

Theorem 4.4.1. *There exists a finite type scheme S and a map*

$$S \rightarrow \text{DT}_{\beta,n}(X|D)$$

which exhausts the geometric points of the target.

Proof. We prove this in two lemmas. First, in Lemma 4.4.2, we bound the number of tropical types that can arise by examining points of $\text{DT}_{\beta,n}(X|D)$. Second, in Lemma 4.4.4, we exhaust the locus of transverse subschemes in a fixed degenerate target that give a fixed tropicalization. \square

Let G and G' be embedded 1-complexes in Σ_X . We say that G and G' have the same type if they lie in the interior of the same cone of $T(\Sigma_X)$.

Lemma 4.4.2. *There are finitely many combinatorial types of embedded 1-complexes in Σ_X that arise from ideal sheaves in X relative to D with fixed curve class β and Euler characteristic n . That is, the moduli map*

$$DT_{\beta,n}(X|D) \rightarrow \text{Exp}(X|D)$$

factors through a finite type substack of the target.

Proof. The blueprint for this result on the Gromov–Witten side is [27, Section 3.1]. The proof bounds the possible slopes of edges of 1-complexes embedded in Σ_X that can arise from subschemes with given discrete data. This is possible by means of the balancing condition, which we formulate below, together with known finiteness results for balanced tropical curves. \square

STEP I. EMBEDDED TROPICAL CURVES

Let $\mathcal{Z} \subset \mathcal{Y}$ be an algebraically transverse subscheme of an expansion of X along D , defined over a point. This point is equipped with a logarithmic structure pulled back from $\text{Exp}(X|D)$. It has characteristic monoid P . After pulling back the family to a standard logarithmic point via a sharp map $P \rightarrow \mathbb{N}$, we obtain a 1-complex $\mathbb{G}_{\mathcal{Z}}$ embedded in Σ_X . We want to bound the combinatorial possibilities for these 1-complexes. The first step is to observe that 1-complexes coming from such $\mathcal{Z} \subset \mathcal{Y}$ can be decorated with additional information.

First, decorate each vertex v with the homology class of the associated subscheme Z_v in the homology of X . To be precise, we note that there is a proper cycle Z_v in the component dual to v . Its pushforward in X defines a class in $H_2(X)$. Denote this homology class at a vertex v by β_v . In particular, note that vertices such that the subscheme Z_v lies entirely in fiber of the projection $\mathcal{Y} \rightarrow X$ are decorated with 0.

Next, let v be a vertex of $\mathbb{G}_{\mathcal{Z}}$ and e an edge or ray incident to v . Note that e is dual to a codimension 1 stratum of \mathcal{Y} . The intersection of \mathcal{Z} with this stratum is a zero-dimensional subscheme. Define w_e to be the length of this subscheme. Since the subschemes on either side of a double divisor \mathcal{S} share the same intersection with \mathcal{S} , if e is an edge rather than a ray, the number w_e depends only on e and not on v . Range over all v to determine decorations for all edges and rays in $\mathbb{G}_{\mathcal{Z}}$.

As discussed previously, Σ_X can be embedded as a subcomplex of $\mathbb{R}_{\geq 0}^k$ – the latter with its standard cone structure – for k equal to the number of irreducible components in D . We obtain a 1-complex with vertex and edge decorations in $\mathbb{R}_{\geq 0}^k \subset \mathbb{R}^k$. We refer to this object as an *embedded tropical curve*. We abuse notation slightly and continue to denote it by $\mathbb{G}_{\mathcal{Z}}$, even after adding the decorations.

STEP II. PIECEWISE LINEAR FUNCTIONS

The group of integral linear functions on $\mathbb{R}_{\geq 0}^k$ is naturally identified with \mathbb{Z}^k . Indeed, given a set of integers (a_1, \dots, a_k) we can consider the function whose slope along the i^{th} coordinate axis is a_i . Every linear function in \mathbb{Z}^k restricts to a function on Σ_X that is piecewise linear and linear on each cone.¹³ A piecewise linear function, in this sense, gives rise to a Cartier divisor on X . If α is a piecewise linear function, we let $(\mathcal{O}_X(-\alpha), s_\alpha)$ be the associated Cartier divisor¹⁴.

We can make sense of this construction on any component \mathcal{Y}_v of the expansion as well. Let v be a vertex of $\mathbb{G}_{\mathcal{Z}}$. Restrict α to the star $\mathbb{G}_{\mathcal{Z}}(v)$ of v in $\mathbb{G}_{\mathcal{Z}}$. The result is a piecewise *affine* function at the star of v ; subtract the constant $\alpha(v)$ to obtain a piecewise linear function on the (zero or) one-dimensional

¹³Note we use piecewise linear, rather than linear, as the adjective for these functions. The reason is that complexes come with a natural notion of piecewise linear function but no natural notion of linear function, unlike fans in toric geometry. The embedding of Σ_X into $\mathbb{R}_{\geq 0}^k$ endows it with a notion of linear functions, but a different embedding would lead to a different notion of linear.

¹⁴The notation is chosen to be consistent with convention in Fulton’s text on toric varieties [24]. Specifically, on the toric variety \mathbb{P}^1 , the function that has slope 1 along a ray and slope 0 along the other ray is $\mathcal{O}_{\mathbb{P}^1}(-1)$

$\text{fan } \mathbb{G}_{\mathcal{Z}}(v)$. The construction above again gives a line bundle and a section on the associated component of the expansion.

STEP III. THE BALANCING CONDITION

Fix a linear function $\alpha \in \mathbb{Z}^k$. By the procedure above, this determines a Cartier divisor $(\mathcal{O}_{\mathcal{Y}_v}(-\alpha), s_{v,\alpha})$ on the component \mathcal{Y}_v of the expansion, together with a canonical section. Let Z_v be the subscheme in the component \mathcal{Y}_v . The divisor associated to this piecewise linear function can be intersected with Z_v to give a homomorphism:

$$\begin{aligned} \varphi_v : \mathbb{Z}^k &\rightarrow \mathbb{Z} \\ \alpha &\mapsto \text{deg}([Z_v] \cap c_1(\mathcal{O}_{\mathcal{Y}_v}(-\alpha))). \end{aligned}$$

The ambient space of the embedded tropical curve is canonically identified with the real positive dual of this copy of \mathbb{N}^k . In other words, we have

$$\mathbb{G}_{\mathcal{Z}} \hookrightarrow \mathbb{R}_{\geq 0}^k = \text{Hom}_{\text{Mon}}(\mathbb{N}^k, \mathbb{R}_{\geq 0}).$$

The homomorphism φ_v above is an integral element in the ambient cone $\mathbb{R}_{\geq 0}^k$; we use this to enhance $\mathbb{G}_{\mathcal{Z}}$ by adding yet more decoration. We may consider the negative of φ_v as a lattice point in the ambient vector space \mathbb{R}^k . To each vertex v , appended a ray e_v starting at v in the direction of $-\varphi_v$; we call these *appended rays*. The vector $-\varphi_v$ is a positive integral multiple of the primitive generator in the direction of e_v , and we equip these rays with a weight equal to this positive multiple.

Finally, we call the resulting object *the balanced tropical curve associated to \mathcal{Z}* . We view it as a new graph $\widehat{\mathbb{G}}_{\mathcal{Z}}$ equipped with an immersion into \mathbb{R}^k that restricts to an embedding on the 1-complex $\mathbb{G}_{\mathcal{Z}}$ into the standard quadrant.

The term *balanced* is justified as follows. When equipped with these new edges and their weight decorations, the sum of the outgoing vectors at each vertex, weighted by the edge decoration, is zero. Explicitly, we have at each vertex v :

$$\sum_{e \in E(\widehat{\mathbb{G}}_{\mathcal{Z}}) : v < e} w_e \cdot \vec{e} = \varphi_v,$$

where (i) the sum is taken over all edges e in the edge set of $\widehat{\mathbb{G}}_{\mathcal{Z}}$ that are incident to v , and (ii) the symbol \vec{e} denotes the primitive integral vector along e , directed away from the v .

The balancing condition follows from geometric considerations – we calculate the degree on Z_v of the line bundles associated to piecewise linear functions $\alpha \in \mathbb{Z}^k$ in two different ways, as follows. The intersections of Z_v with the divisors $\text{div}(s_{v,\alpha})$ are algebraically transverse, so we may calculate the numbers on the right-hand side of the balancing equation:

$$[Z_v] \cap c_1(\mathcal{O}_{\mathcal{Y}_v}(-\alpha))$$

by explicitly calculating lengths of scheme theoretic intersections. It suffices to check the balancing condition coordinate-by-coordinate in \mathbb{Z}^k . If we let φ_i denote the piecewise linear function corresponding to the i^{th} divisor, the sum of the i^{th} -coordinate components of the weighted edge directions incident to v . This gives the left-hand side of the equation above. Ranging over over all i gives the equation.

In particular, if β_v is equal to 0, then there is no appended ray e_v , and the balancing condition simply states that the weighted of edges incident to a vertex in $\mathbb{G}_{\mathcal{Z}}$ is equal to 0. More generally, by the formula above, the direction and weight of the appended ray is determine by the intersection numbers $\beta_v \cdot D_i$ over all i .

STEP IV. BOUNDING THE BALANCED CURVES

We now examine the asymptotic structure of $\widehat{\mathbb{G}}_{\mathcal{Z}}$. Since β is fixed and equal to an effective sum of the β_v , there are only finitely many vertices of $\mathbb{G}_{\mathcal{Z}}$ that can have nonzero curve class decoration. From the discussion above, for any such vertex v with β_v nonzero, there is a uniquely determined appended ray e_v with a weight decoration that makes the balancing condition hold. As a consequence, there are finitely many possibilities for the appended rays.

The unbounded rays in $\mathbb{G}_{\mathcal{Z}}$ itself, rather than the enhanced $\widehat{\mathbb{G}}_{\mathcal{Z}}$, are each parallel to one of the axes in $\mathbb{R}_{\geq 0}^k$. The sum of the weights of rays parallel to the i^{th} -ray is $\beta \cdot D_i$, so there are finitely many possibilities for the unbounded edges of $\widehat{\mathbb{G}}_{\mathcal{Z}}$. Any unbounded ray in $\widehat{\mathbb{G}}_{\mathcal{Z}}$ is either an appended ray or an unbounded in $\mathbb{G}_{\mathcal{Z}}$.

Putting the two preceding paragraphs together, there are only finitely many combinatorial possibilities for the asymptotic structure, meaning, outside of some large compact neighborhood of the origin, of the balanced tropical curve $\widehat{\mathbb{G}}_{\mathcal{Z}}$.

We now claim that with these data fixed there are only finitely many balanced tropical curves $\widehat{\mathbb{G}}_{\mathcal{Z}}$ that can appear for fixed discrete data. Indeed, Nishinou–Siebert present a Newton polygon argument to deduce this in the course of the proof of [61, Proposition 2.1] which assures this.¹⁵

The key consequence of the Nishinou–Siebert argument for us is that the underlying set of every balanced tropical curve in \mathbb{R}^n can be obtained as a subset of a complete intersection of $n - 1$ tropical hypersurfaces. Moreover, each such tropical hypersurface is the preimage of a balanced tropical curve in \mathbb{R}^2 under an integer linear projection. The number of combinatorial types tropical curves in \mathbb{R}^2 of given asymptotics is bounded by the number of coherent subdivisions of the dual lattice polytope of this projected tropical curve. The lattice polytope is determined by the images of the asymptotic directions under the projection, of which there are certainly a finite number of possibilities. See [57, Section 3] for a detailed discussion of the plane curve case.

The argument above immediately implies that, if we range over all $\widehat{\mathbb{G}}_{\mathcal{Z}}$, and ignore 2-valent vertices and all isolated vertices v , there are finitely many possibilities for the combinatorial type of $\widehat{\mathbb{G}}_{\mathcal{Z}}$. The number of isolated points in the tropicalization is bounded because we have fixed the the curve class and holomorphic Euler characteristic. By the DT stability condition, nontube 2-valent vertices cannot contain a tube subscheme, so contribute to either the curve class or increase the Euler characteristic. Therefore, there are only finitely many such 2-valent vertices. The number of tube vertices only depends on the combinatorial type determined in the previous paragraph. We conclude that the number of combinatorial types of embedded 1-complexes arising from ideal sheaves is bounded after fixing the numerical data.

The following remark should give some intuition for what an expansion looks like, which might be valuable in navigating the arguments that follow.

Remark 4.4.3 (The shape of a component in an expansion). Let us work in dimension 3 for concreteness. Let \mathcal{Y} be an expansion of a threefold X . The top-dimensional components of \mathcal{Y} have the following four forms (i) toroidal modification of X ; there is at most one such component, (ii) a compactified \mathbb{G}_m -bundle over a locally closed boundary surface in X , (iii) a compactified \mathbb{G}_m^2 -bundle over a boundary curve in X or (iv) a toric threefold.

To conclude boundedness we claim that, after fixing the Hilbert polynomial with respect to the pullback of a polarization from X , the space of subschemes whose expansion has a given combinatorial type with numerical data fixed (i.e., subschemes with a fixed tropicalization) is parameterized by a quasi-projective scheme of finite type.

Let us briefly explain the nature of this claim. The polarization is pulled back from X , it is typically not ample on an expansion, and we may have subschemes in these components that only contribute to

¹⁵Alternatively, a balanced tropical curve in \mathbb{R}^n is a polyhedral subcomplex of the dual tropical hypersurface of a subdivision of the associated Chow polytope; see [48, End of Section 4.4]. The types are therefore bounded by the number of coherent subdivisions of the Chow polytope. See Bertrand–Brugalle–Lopez de Medrano [13] and Yu [76, Proposition 4.1] for related boundedness results.

the Hilbert polynomial via their holomorphic Euler characteristic. A basic principle is that the higher the codimension of the stratum of X over which a component lives, the less control the Hilbert polynomial provides in the choice of subscheme in that component. The extreme case occurs where the stratum in question is a toric variety Y that is collapsed to a point on projection to X . The problem here reduces to examining the moduli of subschemes of Y with fixed contact orders along the toric boundary. But while the Hilbert polynomial is not a strong constraint, the balancing condition compensates for this. It constrains the contact orders of the subscheme and thereby constrains the homology class of the subscheme.¹⁶

Lemma 4.4.4. *The stack of ideal sheaves on X relative to D with fixed combinatorial type and fixed numerical data is bounded.*

Proof. Fix an expansion \mathcal{Y} and examine the irreducible components. We study the subschemes in an irreducible component of this expansion. As noted above, the components are all compactified torus bundles over locally closed strata of X . By the algebraic transversality hypothesis, the boundedness is insensitive to further logarithmic modifications or partial compactifications of \mathcal{Y} . In particular, after replacing \mathcal{Y} with a further logarithmic modification, any component V of \mathcal{Y} can be presented as

$$V \rightarrow W \rightarrow B$$

with $V \rightarrow W$ logarithmically étale and with B a closed stratum in X and W a \mathbb{P}^r -bundle over B , with $r = \text{codim}_X V$. The logarithmically étale morphism can be taken to be a composition of an open immersion and sequence of blowups.

Let $Z \subset V$ be a subscheme with the given combinatorial data. We claim that there are finitely many choices for the homology class of Z in V . We argue as follows. Since the Hilbert polynomial in X is fixed, there are finitely many choices for the homology class of the pushforward of $[Z]$ to B . On the other hand, since we have fixed the contact orders of Z , there are finitely many choices for the fiber degree of the class $[Z]$ after pushforward to the toric bundle W . Moreover, the cohomology of V can be obtained as a module over that of W by the blowup formula and is generated by classes of exceptional divisors and exceptional curves in the blowup. The homology class $[Z]$ is determined by its intersection multiplicities with these exceptionals, which are in turn determined by the contact order.

Choose a very ample line bundle on V and observe that the degree of the subscheme Z with respect to this very ample line bundle is now bounded. Since the holomorphic Euler characteristic of Z is also fixed, there are finitely many possible values of the Hilbert polynomial for the subscheme Z . The possible choices for the ideal sheaf are parameterized by an open subset in a finite disjoint union of connected components of the Hilbert scheme of V . The result follows. \square

4.5. Finiteness of automorphisms

We continue with the notation from the previous sections. Let $\text{DT}_{\beta,n}(X|D)$ be the moduli space of ideal sheaves on X relative to D , with a fixed choice of target expansion moduli. There is a structure map

$$\text{DT}_{\beta,n}(X|D) \rightarrow \text{Exp}(X|D).$$

We use this to deduce finiteness of automorphisms the space of relative ideal sheaves.

Theorem 4.5.1. *The isotropy groups of the moduli space $\text{DT}_{\beta,n}(X|D)$ are finite.*

Proof. We show that the objects of $\text{DT}_{\beta,n}(X|D)$ have finite automorphism group. Recall that a point of this stack is given by a morphism

$$\text{Spec } \mathbb{C} \rightarrow \text{Exp}(X|D)$$

¹⁶This problem, with slightly different hypotheses, was considered by Katz–Payne, and their [37, Theorem 1.2] can be seen as a manifestation of the same phenomenon.

together with an ideal sheaf $\mathcal{I}_{\mathcal{Z}}$ on the expansion \mathcal{Y} of $(X|D)$; here, \mathcal{Z} denotes the associated subscheme. Moreover, we have imposed the condition that the ideal sheaf gives rise to a subscheme that is a tube precisely along the tube components. In order to analyze the stabilizer group of the ideal sheaf $\mathcal{I}_{\mathcal{Z}}$, we first note that the expansion \mathcal{Y} maps to a substack $B\mathbb{G}_m^r$ of $\text{Exp}(X|D)$. Specifically, the map above determines a locally closed stratum of $\text{Exp}(X|D)$ corresponding to a cone σ of the associated cone space. The isotropy group at this point is naturally identified with the torus whose lattice of one-parameter subgroups is $\sigma^{\text{gp}}(\mathbb{N})$. We claim that $\mathcal{I}_{\mathcal{Z}}$ is not stabilized by any positive dimensional subgroup of this isotropy group.

To see this, we unwind how the isotropy group acts on \mathcal{Y} . The components of the expansion are either modifications of X or an equivariantly compactified \mathbb{G}_m^k bundle over S° , where S° is a locally closed stratum in X of codimension k . Each of these bundles therefore possesses a translation action by the torus. Upon restriction to each component, the isotropy group acts as by scaling the torus directions described above (however, not all such translation actions come from the action of isotropy).

We wish to show that, in the setup above, no positive-dimensional subtorus in \mathbb{G}_m^r stabilizes the ideal sheaf $\mathcal{I}_{\mathcal{Z}}$. If there was such a subgroup, then in particular, $\mathcal{I}_{\mathcal{Z}}$ would be preserved by any one-parameter subgroup of that subgroup. It therefore suffices to show that there is no one-parameter subgroup of \mathbb{G}_m^r that stabilizes $\mathcal{I}_{\mathcal{Z}}$.

Assume to the contrary, that there were a one-parameter subgroup of \mathbb{G}_m^r that stabilizes $\mathcal{I}_{\mathcal{Z}}$. We may pullback the expansion \mathcal{Y} along the inclusion $B\mathbb{G}_m \rightarrow B\mathbb{G}_m^r$. This choice of $B\mathbb{G}_m$ point corresponds to a ray in a cone inside T . On the interior of this ray, there may be tube vertices on the associated 1-complex. Erasing these vertices by passing to the associated 1-complex in $|T|$ gives rise to a contraction morphism

$$\begin{array}{ccc} \mathcal{Y} & \xrightarrow{\quad} & \overline{\mathcal{Y}} \\ & \searrow & \swarrow \\ & B\mathbb{G}_m & \end{array}$$

which collapses the \mathbb{P}^1 -bundles corresponding to these vertices to their base. Furthermore, the subscheme \mathcal{Z} is the inverse image, along this contraction, of a subscheme $\overline{\mathcal{Z}} \subset \overline{\mathcal{Y}}$. By the definition of algebraic transversality, every component of \mathcal{Y} has nonempty intersection with the subscheme. The only possibility for a \mathbb{G}_m to stabilize a subscheme is by scaling the fibers of a \mathbb{P}^1 -bundle component in $\overline{\mathcal{Y}}$. But in order for \mathbb{G}_m to stabilize such a subscheme, that subscheme, restricted to this component, has to be pulled back along the projection, and therefore it has to be a tube subscheme. But we have blown down all the tube components of \mathcal{Y} , and thus by DT stability, the remaining \mathbb{P}^1 -bundle components cannot contain tube subschemes. It follows that the stabilizer must be finite, as required. \square

4.6. Properness

We now verify properness of the moduli problem. We only prove the properness here under the simplifying assumption that the generic fiber of a family of expansions is unexpanded. This conveys the essential geometric arguments in the proof. The removal of the assumption adds some technical subtleties, and we delay this to Section 7.

Theorem 4.6.1. *The moduli space $DT_{\beta,n}(X|D)$ is proper.*

Proof. The stack $\text{Exp}(X|D)$ is locally of finite presentation; see [16, Corollary 6.7]. The moduli space $DT_{\beta,n}(X|D)$ is a relative Hilbert scheme over the stack of expansions and is therefore also quasiseparated; the arguments in the previous section show that it is of finite type. Therefore, we need only check the weak valuative criterion for properness. Let K be a valued field and consider a morphism

$$\text{Spec } K \rightarrow DT_{\beta,n}(X|D).$$

Since $\mathrm{DT}_{\beta,n}(X|D)$ is of finite type, we may assume that K is the function field of a smooth curve, and the valuation is given by order of vanishing at a closed point. After pulling back the universal expansion and universal subscheme, we must check that, after replacing K with an extension, the resulting family extends to yield a map

$$\mathrm{Spec} R \rightarrow \mathrm{DT}_{\beta,n}(X|D).$$

We prove the result here in the case where the fiber over K is unexpanded. In the unexpanded case, this follows directly from Proposition 2.6.5; indeed, our construction has been rigged to arrange this, and we spell out the details as follows. The general case is handled by a nearly identical argument in Section 7, but a more complicated formalism is required to carry it out.

Suppose we have a cone space T with Artin fan A_T . The data of a morphism $\mathrm{Spec} R \rightarrow A_T$ which sends the generic point to the open stratum are equivalent to giving an integral point in T . Indeed, such a statement holds for an Artin cone by a direct analysis, and the general statement follows. If we pass to a base change of R , the corresponding integral point of the composed morphism is obtained by scaling the original point.

To prove the existence of a stable limit, apply Proposition 2.6.5; possibly after a base change, there exists a 1-complex $G \subset \Sigma_X$ such that the corresponding expansion $\mathcal{Y}_G \rightarrow \mathrm{Spec} R$ has an algebraically transverse limit. Furthermore, by Remark 2.6.12, there are no tube subschemes contained in the components associated to bivalent vertices of G .

The 1-complex G defines a rational point of T . After a possible scaling to ensure integrality, along with the corresponding base change of R , we have a point $P \in T$ and an associated morphism $\mathrm{Spec} R \rightarrow \mathrm{Exp}(X|D)$. By construction, the corresponding expansion $\mathcal{Y}_R \rightarrow \mathrm{Spec} R$ factors through \mathcal{Y}_G , and the contracted components are precisely the tube vertices of \mathcal{Y}_P . As a result the flat limit of \mathcal{Z}_K in \mathcal{Y}_R will be algebraically transverse and stable.

For uniqueness of the limit, given an expansion $\mathcal{Y} \rightarrow \mathrm{Spec} R$ associated to an integral point P' of T , the universal property of G implies that there is a subdivision $\mathcal{Y}_{P'} \rightarrow G$ and corresponding modification $\mathcal{Y} \rightarrow \mathcal{Y}_G$. There are no tube subschemes for \mathcal{Y}_G , so the new vertices of the subdivision are precisely the tube vertices of this family, so P' and P coincide and the stable limit is unique. \square

Remark 4.6.2 (Stable pair moduli). In this remark, we sketch how to modify the definitions and constructions to handle stable pairs instead of ideal sheaves. A *stable pair* consists of a pair (F, s) consisting of a sheaf F of fixed Hilbert polynomial with purely one-dimensional support and a section $\mathcal{O}_X \xrightarrow{s} F$ whose cokernel is zero-dimensional [64]. A stable pair (\mathcal{F}, s) on an expansion \mathcal{Y} of X is algebraically transverse if the scheme theoretic support of \mathcal{F} is algebraically transverse and the support of the cokernel is disjoint from the divisorial strata of \mathcal{Y} .

The existence and uniqueness of algebraically transverse limits is proven analogously to the case of subschemes. Specifically, given a family of algebraically transverse stable pairs over a valuation ring, we find a dimensionally transverse limit for the schematic support using the tropicalization. We can then apply the stable pairs arguments used by Li and Wu [47] to produce an algebraically transverse limit. As before with the Gröbner limit, one can tropically characterize the algebraically transverse limit. A *tube stable pair* is a pair that is pulled back from a zero-dimensional subscheme on a surface along a \mathbb{P}^1 -bundle, with the condition that the cokernel of the section vanishes. The bivalent vertices added to the dimensionally transverse limit correspond to those components where the limit stable pair is not a tube stable pair (and one can use this to give an alternate proof of existence as well).

The moduli space $\mathrm{PT}(X|D)$ can be constructed as a relative stable pair space on the universal expansion over $\mathrm{Exp}(X|D)$, with the condition that only tube components host tube stable pairs. Properness follows from the tropical limit algorithms, and the virtual class follows from [55, Proposition 10] as in the ideal sheaf case.

5. Logarithmic DT theory

We construct logarithmic DT invariants for a simple normal crossings threefold pair $(X|D)$, by showing that the moduli spaces constructed in the previous section come equipped with virtual classes and evaluation morphisms to the relative divisors. We also show that the virtual class, and hence the numerical theory, does not depend on the choices that have been in constructing the moduli stack of target expansions. We define DT generating functions and state the basic rationality conjecture, analogous to the nonrelative setting.

5.1. Virtual classes

Each moduli space of ideal sheaves comes equipped with a structure map

$$t : \text{DT}_{\beta,n}(X|D) \rightarrow \text{Exp}(X|D).$$

Let \mathcal{I} be an ideal sheaf. Recall that given a coherent sheaf \mathcal{E} on X , there is a trace map:

$$\text{Ext}^i(\mathcal{E}, \mathcal{E}) \rightarrow H^i(X, \mathcal{O}_X),$$

see [31, Chapter 10]. The kernel $\text{Ext}_0(\mathcal{E}, \mathcal{E})$ of this map is called the *traceless* Ext group.

The groups

$$\text{Ext}^1(\mathcal{I}, \mathcal{I})_0 \quad \text{Ext}^2(\mathcal{I}, \mathcal{I})_0$$

govern the deformation and obstruction spaces of the ideal sheaf \mathcal{I} respectively. Additional details concerning the fact that the traceless Ext groups govern the obstructions to deformations of sheaves can be found in [32, Section 4] and [73, Section 3].

Proposition 5.1.1. *The structure map*

$$t : \text{DT}_{\beta,n}(X|D) \rightarrow \text{Exp}(X|D).$$

is equipped with a perfect obstruction theory, with deformation and obstruction spaces as above. If β is the curve class, then there is an associated virtual fundamental class $[\text{DT}_{\beta,n}(X|D)]^{\text{vir}}$ in Chow homology of degree

$$\text{vdim } \text{DT}_{\beta,n}(X|D) = \int_{\beta} c_1(T_X).$$

The virtual class in Borel–Moore homology is defined by taking the image under the cycle class map.

Proof. The existence of the relative obstruction theory as described is established in the requisite generality in [55, Proposition 10]. Since the base $\text{Exp}(X|D)$ of the obstruction theory is a connected Artin fan, it is irreducible and equidimensional. Virtual pullback of the fundamental class gives rise to the virtual fundamental class in the usual fashion [49]. □

5.2. Virtual birational models

The spaces constructed in the paper to this point have all required a choice of tropical target moduli. We show that different choices lead to compatible theories.

We supplement the notation in this section to keep track of the different theories arising from the tropical choices. Recall that the construction of T in Section 3 produces an infinite collection of spaces, depending on the choices made in the algorithmic steps of Section 3.3. Moreover, there are additional choices that are made in the construction of the universal family.

We denote by Λ the set of outputs of that construction. More precisely, each element in Λ is a morphism of polyhedral complexes

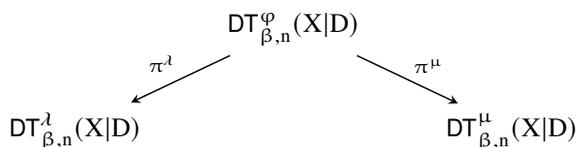
$$\pi^\lambda : \Upsilon^\lambda \rightarrow \mathbb{T}^\lambda,$$

where \mathbb{T}^λ is a moduli space of tropical expansions and π^λ is the projection from the universal family. Let $\text{Exp}(X|D)^\lambda$ be the associated moduli space of expansions, and let

$$\mathcal{Y}^\lambda \rightarrow \text{Exp}(X|D)^\lambda$$

be the projection from the universal family. We fix the discrete data and establish the following birational invariance statement, parallel to [68, Section 3.6], which is itself modeled on [9, Section 6].

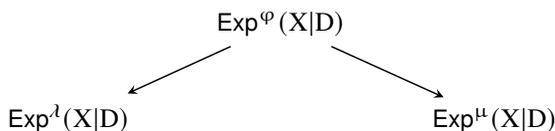
Theorem 5.2.1 (Independence of choices). *For any $\lambda, \mu \in \Lambda$ there exists $\varphi \in \Lambda$ and logarithmic modifications*



such that

$$\pi_{\star}^\lambda [\text{DT}_{\beta,n}^\varphi(X|D)]^{\text{vir}} = [\text{DT}_{\beta,n}^\lambda(X|D)]^{\text{vir}} \quad \text{and} \quad \pi_{\star}^\mu [\text{DT}_{\beta,n}^\varphi(X|D)]^{\text{vir}} = [\text{DT}_{\beta,n}^\mu(X|D)]^{\text{vir}}.$$

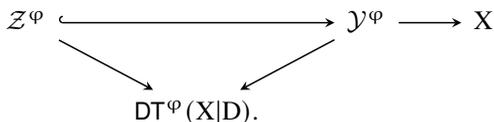
Proof. We drop the discrete data to avoid overcrowding the notation. By applying Proposition 3.6.1 in conjunction with the categorical equivalence between Artin fans and cone stacks, we can find a common refinement



where both vertical arrows are birational. The definitions furnish a forgetful morphism

$$\text{DT}^\varphi(X|D) \rightarrow \text{DT}^\lambda(X|D).$$

We argue that pushforward identifies virtual classes. We have a universal diagram



The analogous diagram exists for λ . We claim that the induced morphism

$$p : \mathcal{Y}^\varphi \rightarrow \mathcal{Y}^\lambda \times_{\text{Exp}^\lambda(X|D)} \text{Exp}^\varphi(X|D)$$

is obtained by adding tube components to the target geometry. Indeed, this is a tautology – by definition, tropicalizations of the expansions of this family of targets induce the same moduli map to $|\mathbb{T}|$. Since the DT stability condition guarantees the tube components only contain tube subschemes, the universal

subscheme \mathcal{Z}^φ is simply obtained by pulling back the family \mathcal{Z}^λ along the modification. It follows that we have a commutative square

$$\begin{CD} DT^\lambda(X|D) @>>> DT^\mu(X|D) \\ @VVV @VVV \\ Exp^\lambda(X|D) @>>> Exp^\mu(X|D) \end{CD}$$

As we have argued above, each of the two vertical arrows have perfect obstruction theories. The introduction of tube components has no effect on the obstruction theory. Indeed, the universal ideal sheaf \mathcal{I}^φ is pullback of the universal ideal sheaf \mathcal{I}^λ along the induced map of targets, so the right vertical obstruction pulls back to the *same* obstruction theory as the one on the left. Since the lower horizontal is a subdivision, it is proper and birational; equality of virtual classes follows from the Costello–Herr–Wise theorem [19, 30] \square

Remark 5.2.2. There is a further birational invariance statement that one might expect for the *target*. Specifically, if $X' \rightarrow X$ is the blowup of X along a stratum of D with total transform D' , one may hope that there is an associated map $DT(X'|D') \rightarrow DT(X|D)$ that identifies virtual classes, in direct analogy with [9]. In fact, when phrased without the DT stability condition, such a result is obvious since every expansion of X' is also an expansion of X . However, a nontube component in a subscheme of an expansion of X' can become a tube component with respect to X . An appropriately corrected analysis does yield a birational invariance statement. We delay this to future work.

5.3. Evaluations at the relative divisors

Let $(X|D)$ be as above. Let E_1, \dots, E_k be the irreducible components of D . These components are smooth by hypothesis. Each intersection of distinct divisors $E_i \cap E_j$ is either empty or a connected divisor, again by hypothesis. Thus, each variety E_i is equipped with a simple normal crossings divisor of its own. We abuse notation slightly and denote these by $(E_i|D)$.

By applying the constructions of the preceding sections to E_i , we obtain a stack of expansions $Exp(E_i|D)$. For our purposes, it will suffice to examine with the substack parameterizing those expansions whose associated 1-complex is a finite set of points. We pass to the subscheme of the universal expansion where the irreducible components of the expansion are disjoint.

For each positive integer d_i , we consider the Hilbert scheme of points in the universal expansion

$$\mathcal{E} \rightarrow Exp^\lambda(E_i|D)$$

of zero dimensional subschemes of length d_i . As before, the symbol λ indicates that we have chosen some stack of target expansions. By studying the locus comprising subschemes of expansions such that no component of the expansion is nonempty, we obtain a proper and finite type stack $Hilb_\lambda^{d_i}(E_i|D)$ as a special case of the construction above.

As a special case of the virtual birational invariance of the DT moduli spaces, we obtain the parallel statement for the Hilbert schemes of points on a curve or surface. However, since the Hilbert schemes of points on smooth curves and surfaces are unobstructed moduli schemes, the same argument produces an honest birationality statement – different choices of λ give rise to birational models of the logarithmic Hilbert scheme of points on logarithmically smooth (curves and) surfaces.

Given a curve class β , define d_i to be the intersection number $D_i \cdot \beta$. We therefore have an evaluation space for relative insertions

$$Ev_\beta(X|D) = Hilb^{d_1}(E_1|D) \times \dots \times Hilb^{d_k}(E_k|D).$$

Lemma 5.3.1. *There exists compatible stacks $\text{Exp}(X|D)$ and $\text{Exp}(E_i|D)$ and an evaluation morphism*

$$\text{ev} : \text{DT}_{\beta,n}(X|D) \rightarrow \text{Ev}_{\beta}(X|D)$$

of the associated moduli spaces of ideal sheaves and points.

Proof. We first explain how the evaluation looks at the tropical level. The reader may first wish to reconsult with Remark 1.6.2. Fix a space of expansions for $\text{Exp}(X|D)$ with associated tropical moduli space $T(X|D)$. The latter is a particular conical structure on the space of embedded 1-complexes. Examine such a 1-complex $G \hookrightarrow \Sigma$. Each divisor component E_i determines a ray ρ_i . The transversality hypothesis implies that outside a bounded set in Σ , the 1-complex is made up of a finite set of unbounded rays, each parallel to one of the rays ρ_i . The star of ρ_i in Σ is a cone complex of dimension one lower than that of Σ . Following the toric dictionary, rays parallel to ρ_i determine points in the star $\Sigma(\rho_i)$. This star is canonically embedded in the tropicalization of $(E_i|D)$. For each cone σ in the cone space of $T(X|D)$, we therefore obtain a family over σ of points in $\Sigma(\rho_i)$.

We now apply the construction in Section 3.3 to the cone complex $\Sigma(\rho_i)$ to produce a tropical moduli space $T(E_i|D)$. By the description above, given a cone of $T(X|D)$, we produce a map to $T(E_i|D)$. The image of σ need not be contained in a single cone of $T(E_i|D)$, but after replacing both $T(E_i|D)$ and σ with subdivisions $T'(E_i|D)$ and σ' , we obtain a morphism

$$\sigma' \rightarrow T'(E_i|D)$$

of cone spaces. For any choice of conical model, the morphism from $\text{DT}_{\beta,n}(X|D) \rightarrow \text{Exp}(X|D)$ factors through a finite type substack of the image, given by a finite collection of cones in $T(X|D)$. It follows that after a finite collection of subdivision operations, we obtain a morphism

$$\text{DT}_{\beta,n}(X|D) \rightarrow \text{Exp}'(E_i|D)$$

again for an appropriate system of choices. In order to lift this to an evaluation morphism, we pullback the universal expansion \mathcal{E} over $\text{Exp}'(E_i|D)$ to $\text{DT}_{\beta,n}(X|D)$. There is an inclusion of expansions

$$\mathcal{E} \times_{\text{Exp}'(E_i|D)} \text{Exp}'(X|D) \hookrightarrow \mathcal{Y}.$$

Intersecting the universal subscheme \mathcal{Z} with the expansion, the algebraic transversality condition implies that we obtain a subscheme of length d_i , and thereby a point of $\text{Hilb}^{d_i}(E_i|D)$. Performing the construction for each divisor, we have constructed the claimed morphism. □

Remark 5.3.2. When combined with the projection formula, the birationality statements for the moduli spaces of sheaves and of points arising from Theorem 5.2.1 guarantee that numerical invariants are independent of the choices made in the construction of the spaces of expansions.

Remark 5.3.3. This evaluation morphism imposes schematic tangency conditions via the Hilbert scheme of points on the boundary surfaces. A more refined evaluation is to the relative Hilbert scheme of points on the universal surfaces over $\text{Exp}(X|D)$. This refined object is more natural from the point of view of the degeneration formula, and for the GW/DT correspondence. In particular, this is the way in which the gluing formula is proved on the Gromov–Witten side [68]. We will examine this refined evaluation in future work, in its appropriate context.

5.4. Logarithmic DT invariants

We define descendent insertions for logarithmic DT invariants as follows. Let

$$(\pi_{\text{DT}}, \pi_X) : \mathcal{Y} \rightarrow \text{DT}_{\beta,n}(X|D) \times X$$

denote the universal expansion over the DT moduli space. While the map π_{DT} is not proper, it is proper when restricted to the universal subscheme $\mathcal{Z} \hookrightarrow \mathcal{Y}$. We claim the universal ideal sheaf $\mathcal{I}_{\mathcal{Z}}$ is a perfect

complex. Indeed, it suffices to show it is Tor-finite, and for this we can take the derived tensor product with the skyscraper sheaf k_p at any closed point p of \mathcal{Y} . Since $\mathcal{I}_{\mathcal{Z}}$ is flat over $\text{DT}_{\beta,n}(X|D)$, this is equivalent to showing Tor-finiteness of $\mathcal{I}_{\mathcal{Z}}$ for the point

$$[Z] = \pi_{\text{DT}}(p) \in \text{DT}_{\beta,n}(X|D).$$

However, this follows from the algebraic transversality of Z , by Proposition 3.5 of [47].

Since $\mathcal{I}_{\mathcal{Z}}$ is perfect, it admits Chern classes; we denote by $\text{ch}_k(\mathcal{I}) \in H^*(\mathcal{Y})$ the cohomological degree k part of the Chern character. These classes are supported on \mathcal{Z} for $k \geq 1$.

Given a cohomology class $\gamma \in H^*(X)$, we define descendent operators for $k \geq 0$

$$\tau_k(\gamma) : H_{\star}(\text{DT}_{\beta,n}(X|D), \mathbb{Q}) \rightarrow H_{\star}(\text{DT}_{\beta,n}(X|D), \mathbb{Q})$$

via the formula

$$\tau_k(\gamma) := (-1)^{k+1} \pi_{\text{DT},\star}(\text{ch}_{k+2}(\mathcal{I}) \cdot \pi_X^{\star}(\gamma)) \cap \pi_{\text{DT}}^{\star}.$$

Here, the pullback on homology is well-defined since π_{DT} is flat by [22], and the pushforward is well defined since π_{DT} is proper on \mathcal{Z} .

Given a cohomology class $\mu \in H^*(\text{Ev}_{\beta}(X|D))$, and cohomology classes $\gamma_1, \dots, \gamma_r \in H^*(X)$ and indices k_1, \dots, k_r , we can define descendent invariants

$$\langle \tau_{k_1}(\gamma_1) \dots \tau_{k_r}(\gamma_r) | \mu \rangle_{\beta,n} = \text{deg} \left(\prod_{i=1}^r \tau_{k_i}(\gamma_i) \text{ev}^{\star} \mu \cap [\text{DT}]^{\text{vir}} \right),$$

where deg denotes the degree of the pushforward to a point. If we sum over n , we have the partition function

$$Z_{\text{DT}} \left(X, D; q | \prod_{i=1}^r \tau_{k_i}(\gamma_i) | \mu \right)_{\beta} = \sum_{\chi} \langle \prod_{i=1}^r \tau_{k_i}(\gamma_i) | \mu \rangle_{\beta,n} q^{\chi}$$

which is a Laurent series since for each fixed curve class β , the logarithmic DT spaces are empty for $\chi \ll 0$.

In analogy with absolute and relative DT invariants, we have the following basic conjecture regarding this series.

Conjecture 5.4.1.

(i) *The DT series for zero-dimensional subschemes is given by*

$$Z_{\text{DT}}(X, D; q)_{\beta=0} = M(-q) \int_X c_3(T_X^{\text{log}} \otimes K_X^{\text{log}}),$$

where

$$M(q) = \prod_{n \geq 1} \frac{1}{(1 - q^n)^n}$$

is the McMahon function.

(ii) *For any curve class β , and insertions $\gamma_1, \dots, \gamma_r$ of degree ≥ 2 and relative insertion μ , the normalized DT series*

$$Z'_{\text{DT}} \left(X, D; q | \prod_{i=1}^r \tau_{k_i}(\gamma_i) | \mu \right)_{\beta} := \frac{Z_{\text{DT}}(X, D; q | \prod_{i=1}^r \tau_{k_i}(\gamma_i) | \mu)_{\beta}}{Z_{\text{DT}}(X, D; q)_{\beta=0}}$$

is the Laurent expansion of a rational function in q .

We expect the first part of this conjecture can be proven once the degeneration formalism is in place.

Remark 5.4.2. The conjecture above is natural in the context of the Gromov–Witten/DT conjectures, which were studied for Calabi–Yau threefolds in [51] and generalized to more general targets, including smooth pair relative geometries, in the sequel [52]. The rationality of the reduced DT partition function is required in order for a precise formulation of the GW/DT equivalence, as the latter involves an exponential change of variables. The DT series itself is *not* a rational function, as evidenced by the first part of the conjecture concerning the degree 0 part, which must first be removed. The degree 0 evaluation is known for empty and smooth boundary divisors [12, 42, 46]. Rationality of the DT generating series is known for all toric targets [53].

The normalized DT series that removed the degree 0 contribution of the MacMahon function has a geometric interpretation in terms of the stable pairs theory of Pandharipande and Thomas [64]. By replacing the DT invariants in the definition of Z_{DT} with the corresponding stable pairs invariants, we obtain a PT generating series denoted $Z_{PT}(X, D; q | \prod_{i=1}^r \tau_{k_i}(\gamma_i) | \mu)_\beta$. We expect rationality and wall-crossing statements for this generating series.

Conjecture 5.4.3. *For any curve class β , relative insertion μ , and insertions $\gamma_1, \dots, \gamma_r$, the PT series*

$$Z_{PT}\left(X, D; q \mid \prod_{i=1}^r \tau_{k_i}(\gamma_i) \mid \mu\right)_\beta$$

is a rational function in q . If the degree of each γ_i is at least 2, this series equals the normalized DT series Z'_{DT} from the previous conjecture.

The GW/DT/PT correspondences, including the degree 0 evaluations and their compatibility with logarithmic degeneration formulas, will be addressed more systematically in forthcoming work [56].

6. First examples

We sketch a handful of basic examples, focusing on the space of targets, where the main new complexity lies. The goal of this section is to show that in a reasonable range of test cases, the logarithmic DT moduli spaces can be worked with fairly explicitly. The final part of this section explains how to think about the tube vertices in the traditional Li–Wu context.

6.1. Target yoga

There is a well-known procedure by which the strata of the expanded target geometry may be recognized from the tropical pictures, as a consequence of the toric dictionary concerning closed torus orbits and stars around cones [24, Section 3.1]. Let $G \hookrightarrow \Sigma$ be an embedded 1-complex. Let \mathcal{Y}_G be the associated expansion of X . Let v be a vertex of G that maps to a cone σ of Σ of dimension k . There is an associated irreducible component

$$Y_v \hookrightarrow \mathcal{Y}_G \rightarrow X.$$

This map factors through the inclusion of the closed stratum $X_\sigma \hookrightarrow X$. The induced morphism

$$Y_v \rightarrow X_\sigma$$

is a partially compactified torus bundle of rank k . For instance, if X is a threefold and σ has dimension 3, then Y_v is simply a toric threefold. The torus bundles are obtained from the line bundles associated to the divisorial logarithmic structure on X and its blowups. The partial compactification of Y_v is obtained (fiberwise) as the toric variety associated to the star of v in G , which gives rise to a fan embedded in the star of σ .

6.2. Dual plane

Let $(X|D)$ be the pair consisting of \mathbb{P}^2 and its toric boundary divisor. The most basic example is the moduli space of lines in \mathbb{P}^2 , with transverse contact orders. There is a canonical space of ideal sheaves $DT(\mathbb{P}^2|D)$, and it is identified with the blowup of the dual \mathbb{P}^2 at its three torus fixed points. This example is worked out carefully in [67, Section 4.1] for stable maps, and the situation is identical to the sheaf theory setup. The figures in op. cit. may be of particular use.

The fibers of the expanded target family $\mathcal{Y} \rightarrow DT(\mathbb{P}^2|D)$ can be described as follows. The target is either isomorphic to $\mathbb{P}^2 \setminus \{p_1, p_2, p_3\}$ or the complement of the codimension 2 strata in the deformation to the normal cone of a line in \mathbb{P}^2 . The latter arises as a transverse replacement in two ways, either when the subscheme limits to a coordinate line, or when the intersection point of the subscheme with the coordinate line limits to a coordinate point.

6.3. Plane curves

The case of degree d curves in \mathbb{P}^2 makes contact with the study of discriminants and, in particular, the secondary fan of a toric fan, constructed by Gelfand–Kapranov–Zelevinsky [25]. A detailed construction of the stable pair moduli spaces in these cases has been undertaken Kennedy–Hunt [40]. The main result of this paper shows that the PT moduli spaces can be expressed explicitly as the vanishing locus of a section of a tautological vector bundle on smooth space that lies over a toric variety. The discussion that follows collapses to the dual plane for $d = 1$.

Let \mathcal{P}_d be the convex hull of the lattice points $(0, 0)$, $(d, 0)$ and $(0, d)$. Let \mathbb{P}_d be the linear system of degree d curves in \mathbb{P}^2 , which coincides with the Hilbert scheme of degree d curves. An open subset of this projective space parameterizes curves that are algebraically transverse to the coordinate lines in \mathbb{P}^2 . This subset is stable under the \mathbb{G}_m^2 action on \mathbb{P}^2 . Limits under the action of this torus capture the relevant degenerations to construct the DT moduli space, as we now sketch.

Any expansion \mathbb{P}^2 along its toric strata arising from our construction is necessarily toric. Let $G \subset \Sigma_{\mathbb{P}^2}$ be a 1-complex. Given an algebraically transverse subscheme

$$Z \subset \mathcal{Y}_G,$$

of the associated expansion, we may proceed as in Section 4.4 and label each edge E of G with the length of the intersection of Z with the divisor attached to E . Since the target is toric, the graph G together with this weighting furnishes a balanced tropical curve. We abuse notation slightly and continue to denote this decorated object by the symbol G .

The secondary fan T_d^\sim of the polytope \mathcal{P}_d enters as a moduli space of tropical plane curves up to translation by the natural additive action of \mathbb{R}^2 . Recall that the secondary fan of \mathcal{P}_d parameterizes subdivisions of \mathcal{P}_d that are *regular* or *coherent*, that is, they are induced by a strictly convex piecewise linear function; see the chapter in the text [25]. The dual of such a subdivision gives rise to a polyhedral subdivision of \mathbb{R}^2 . The edges and rays of this subdivision carry a natural weight – they correspond to edges in the coherent subdivision of \mathcal{P}_d , and the weight is defined to be the lattice length of the edge. This weighted one-dimensional polyhedral complex is a balanced tropical curve in \mathbb{R}^2 [48, Remark 2.3.10]. The top-dimensional cones correspond to maximal triangulations of \mathcal{P}_d .

Geometrically, toric variety of the secondary fan is isomorphic to the Chow quotient $\mathbb{P}_d//\mathbb{G}_m^2$, of the Hilbert scheme by the dilating action of the torus in \mathbb{P}^2 . The Chow quotient carries a universal flat family of broken toric surfaces, arising as limits of \mathbb{G}_m^2 -orbits in \mathbb{P}_d ; see [10, Section 3], and the total space of the universal family is also a toric variety. The total space of this family is birational to \mathbb{P}_d .¹⁷ We explain how some combinatorial adjustments give rise to an instance of the spaces constructed here.

¹⁷Notice that the Chow quotient itself has dimension 2 lower than \mathbb{P}_d . It can be interpreted as a ‘rubber variant’ of the logarithmic DT space.

Let T_d be the fan of the universal family of the Chow quotient. At the combinatorial level, these data are captured by the maps of polyhedral complexes

$$\begin{array}{ccc} \Sigma_d & \longrightarrow & \mathbb{R}^2 \\ \downarrow & & \\ T_d & & \end{array}$$

The toric variety associated to Σ_d is the universal family of the Chow quotient, interpreted a moduli space of pairs of a tropical curve together with a point in \mathbb{R}^2 . See [40] for further details. In order to make a connection to the DT moduli space, a further modification is required. The expansions we propose are always equipped with a map to \mathbb{P}^2 , while the degenerations arising from regular triangulations need not be modifications of the constant family. The adaptation is as follows. The secondary fan comes equipped with a universal family of polyhedral subdivisions of \mathbb{R}^2 :

$$\begin{array}{ccc} \Sigma_d & \longrightarrow & \mathbb{R}^2 \\ \downarrow & & \\ T_d & & \end{array}$$

Let $\tilde{\Sigma}_d$ be the fiber product $\Sigma_d \times_{\mathbb{R}^2} \Sigma_{\mathbb{P}^2}$, which is the common refinement of Σ_d with the constant family $T_d \times \Sigma_{\mathbb{P}^2}$. The toric map induced by $\tilde{\Sigma}_d \rightarrow T_d$ need not be flat with reduced fibers, but by universal weak semistable reduction, there is a canonical such family over a new fan T'_d ; see [40, 58].

Finally, let T''_d denote the quotient of T'_d by the equivalence relation that identifies two embedded tropical curves if they have the same underlying 1-complex (i.e., we forget the weights along the edges described above). It is straightforward to check that this is a cone space, and is a union of cones in a moduli space of tropical expansions $T(\Sigma)$ constructed in Section 3.3. We refer the reader to [40] for further details.

Let $\text{Exp}(\mathbb{P}^2|D)$ and A_d be the Artin fans associated to T''_d and T'_d respectively. The constructions in the previous section give rise to a DT moduli space; given a subscheme in an expansion, we can recover the weights on the edges using the schematic intersection numbers with the corresponding divisors. As a result, the structure morphism to $\text{Exp}(\mathbb{P}^2|D)$ factors through

$$\text{DT}(\mathbb{P}^2|D) \rightarrow A_d.$$

The subdivision $\tilde{\Sigma}$ gives rise to a (nonrepresentable) birational modification \mathbb{P}'_d of the universal orbit of the Chow quotient $\mathbb{P}_d//G_m^2$. This toric stack is naturally identified with $\text{DT}(\mathbb{P}^2|D)$. A rigorous treatment of the identification is given in [40]. Analogous statements hold for any toric surface.

We conclude with a series of remarks.

Remark 6.3.1. The stable pair moduli spaces in degree d on \mathbb{P}^2 but with arbitrary holomorphic Euler characteristic admit an explicit and elegant description. A complete picture, including the basic geometry and virtual class, has been given in [40]. The upshot of Kennedy-Hunt's results are that the logarithmic pair spaces are obtained as the relative Hilbert schemes of points on the universal subscheme over the toric variety constructed above, parallel to [65, Proposition B.10]. At the combinatorial level, this amounts to allowing a fixed additional number of vertices on the balanced 1-complexes considered above. The additional vertices allow for expansions along the divisors to accommodate support for the cokernel of the section of the stable pair.

Remark 6.3.2. The subdivision step above can be avoided by working directly with the logarithmic multiplicative group G_{\log}^2 . The tropicalization is the \mathbb{R}^2 appearing above, with 'fan' structure given by the single nonstrictly convex cone \mathbb{R}^2 . We avoid the detour through the details of the logarithmic

multiplicative group, but a reader interested in the details of this perspective may extract them using the arguments in [69].

Remark 6.3.3. The moduli space above is naturally identified with the universal family over the Chow quotient. The Chow quotient itself can also be interpreted as a DT moduli space for the rubber moduli space \mathbb{P}^2 ; see [50]. The rubber moduli will be discussed elsewhere.

Remark 6.3.4. The relationship between the secondary fan and enumerative geometry should be credited to Katz [36, Section 9]. Katz suggested that the universal family of the Hilbert quotient of a toric variety should function as a moduli space of polarized target expansions for Gromov–Witten theory relative to the toric boundary. In retrospect, it seems more natural to realize this connection within the context of ideal sheaves or stable pairs. A detailed study of the role of the secondary fan in logarithmic curve counting theories seems worthwhile.

6.4. Subgroups in toric threefolds

Let X be a smooth projective toric threefold with D the toric boundary divisor. Let Σ_X be the fan in the cocharacter space $N_{\mathbb{R}}$. We make the further assumption that there exists a nonzero primitive vector v such that the ray generated by v and by $-v$ are both rays in Σ_X . This guarantees that the associated compactified one-parameter subgroup

$$\varphi_v : \mathbb{P}^1 \rightarrow X$$

is transverse to D .

The DT moduli space in the class $[\text{im}(\varphi)]$ is again related to the geometry of Chow quotients. Consider the line ℓ_v spanned by v . The relevant embedded 1-complexes in this instance are precisely the parallel translates of ℓ in $N_{\mathbb{R}}$. Given such a translate ℓ of ℓ_v , we equip it with a canonical polyhedral structure by taking the common refinement with Σ_X inside $N_{\mathbb{R}}$.

The moduli space of such embedded 1-complexes is canonically identified with the fan of the Chow quotient $X//\mathbb{G}_m$ by results of Chen–Satriano [18]. The identification is somewhat indirect, so we spell it out. In loc. cit., the authors identify the Chow quotient with the moduli space of logarithmic stable maps to X in the class of φ_v . The tropical moduli space of maps is precisely the set of parallel translates ℓ_v with the polyhedral structure described above, for instance from the arguments in [67]. However, it is straightforward to see that all tropical maps in this case are forced to be embedded complexes, so we may identify the tropical moduli space with the set of embedded 1-complexes. The DT moduli space for these data is then equal to the Chow quotient. The identification of coarse moduli spaces occurs can be upgraded using the Chow quotient stack.

Remark 6.4.1. Stable map, ideal sheaf and stable pair spaces coincide in this case.

6.5. An embedded point and a subgroup

We maintain the notation from the previous example and consider subschemes in the curve class φ_v , but we increase the Euler characteristic by 1. Geometrically, we consider the class of a one-parameter subgroup together with an embedded point. As the additional point may wander the target, the space of relative ideal sheaves includes expansions caused by the presence of the additional point on the boundary D . The relevant 1-complexes are therefore translates of the line ℓ together with an additional point, which can lie anywhere in $N_{\mathbb{R}}$, and whose presence introduces a vertex.

The space of expansions can be constructed as follows. Let S_v be the toric fan associated to the Chow quotient $X//\mathbb{G}_m$ as above, and let U be the universal orbit of the Chow quotient. The latter is a refinement of Σ_X . The product fan $\Sigma_U \times S_v$ can be viewed as a moduli space of pairs (p, ℓ) where, as above, ℓ is a parallel translate of ℓ_v and p is a point in Σ_U . Consider the incidence locus:

$$\text{IL} := \{(p, [\ell]) : p \in \ell\} \subset \Sigma_U \times S_v.$$

This incidence locus has a fan structure by viewing it as the preimage of the diagonal in the projection

$$\Sigma_U \times S_v \rightarrow S_v \times S_v$$

where the morphisms are the projection of the universal family and the identity map on the two factors. Note that since the vertical maps are both combinatorially flat with reduced fibers, the fiber product is canonically a cone complex.

Consider any subdivision T of $\Sigma_U \times S_v$ such that the locus IL is a union of faces. Over T , we obtain two families of subcomplexes – the first is the family of parallel translates of ℓ_v while the second is simply the universal family U over T . Each of these map to $T \times \Sigma_U$, and we choose a conical structure on the union of their images. The image gives the required conical structure.

The Artin fan associated to T gives rise to a moduli space $\text{Exp}(X|D)$ of expansions of X along D . Note that the fibers of the morphism

$$DT(X|D) \rightarrow \text{Exp}(X|D)$$

also have a straightforward description. Fixing the target expansion also fixes a 1-complex in Σ_X . For each vertex of Σ_X that lies on the line ℓ , we choose a torsor for the one-parameter subgroup v , in the corresponding component. The closure provides a subscheme Z with the required homology class. Therefore, additional moduli is provided by the choice of point in the expansion. In order to account for the moduli of tangent directions when the point lies on Z , consider the moduli space of points in the expansion and blowup the incidence locus with Z .

6.6. Lines in \mathbb{P}^3

Let $(X|D)$ be the toric pair $(\mathbb{P}^3|H)$ where H is the union of the coordinate planes. We take the fan of \mathbb{P}^3 as formed by the standard basis vectors and the negative of their sum in \mathbb{R}^3 . The class of a generic line in \mathbb{P}^3 can be described using similar consideration as the one above. First we consider the moduli spaces of lines in \mathbb{P}^3 . The generic line in \mathbb{P}^3 that meets only the codimension 1 strata of the toric boundary has tropicalization a single vertex with 4 outgoing infinite rays, in each of the three coordinate directions and $(-1, -1, -1)$. The relevant embedded 1-complexes required to construct the moduli space of expansions are balanced tropical curves whose asymptotic rays are the ones described above.

Once again the space of 1-complexes can naturally be identified with the elements of the moduli space of tropical maps because all tropical maps for this moduli problem are forced to be embeddings. The moduli space coincides with an explicit blowup of $\overline{\mathcal{M}}_{0,4} \times \mathbb{P}^3$, following [67, Theorem B].

If we consider a line together with an embedded point, thereby raising the holomorphic Euler characteristic, the tropical space of expansions is obtained from the one above by adding an additional point, and subdividing the incidence locus where this point lies on the 1-complex. As before, the DT moduli space has a map to the space of targets, whose fibers can be explicitly described as the blowup of an algebraic incidence locus, stratum-wise on the space of target expansions. We leave the details to the reader.

6.7. Revisiting the Li–Wu setup

The notions of tube vertices, tube components, and tube subschemes are not present in Li–Wu’s construction of the DT moduli problem for smooth pairs $(X|D)$, and our formalism can be made to specialize to theirs by choosing the right cone structure on the cone space T of 1-complexes. However, tube components can be *forced* to appear in the Li–Wu setup if one makes less efficient choices for T . We explain this now, since it may help readers who are familiar with relative DT and GW theory already.¹⁸

¹⁸The content of the present subsection came out of a email discussion between R. Thomas and the second author. We take the opportunity to thank him for his questions.

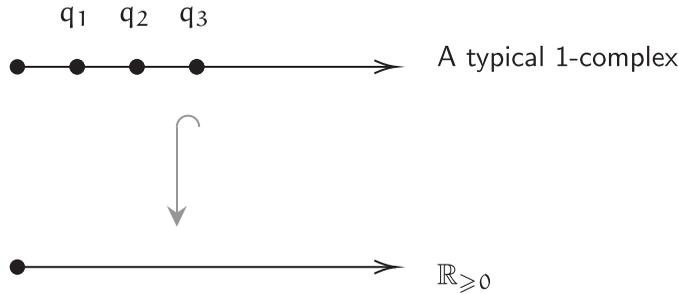


Figure 6. A typical 1-complex arising from the Li–Wu theory of expansions.

6.7.1. Li–Wu via 1-complexes

Consider a pair $(X|D)$ with D smooth and connected. The cone complex Σ_X of this space is just $\mathbb{R}_{\geq 0}$. A particular type of 1-complex in Σ_X can be obtained by choosing a finite collection of distinct points $q_1 < q_2 < \dots < q_L$ in $\mathbb{R}_{\geq 0}$. Given this choice, we can subdivide $\mathbb{R}_{\geq 0}$ by adding these points as vertices, taking all edges between them, as well as the unbounded ray starting at q_L . This subdivision can be viewed as a 1-complex embedded in $\mathbb{R}_{\geq 0}$. See Figure 6.

Note that there are other 1-complexes, which come from disconnected graphs and the discussion that follows can be carried over to that setup, but we ignore these both for simplicity, and because they can be easily avoided in the smooth pair case.

Denote the set of 1-complexes above by $|T^c(\Sigma_X)|$. It is a subset of $|T(\Sigma_X)|$ discussed in Section 3.4. There is a canonical cone space structure on $|T^c(\Sigma_X)|$, where two points lie in the same cone if the associated 1-complexes have the same number of vertices. This cone space is not of finite type: it has one cone of dimension r for each natural number. In the corresponding universal 1-complex, there are no tube vertices.

Correspondingly, there is a canonical stack of expansions which we denote $\text{Exp}^c(X|D)$, and it is equipped with a universal family

$$\mathcal{Y} \rightarrow \text{Exp}^c(X|D).$$

The stack and its universal family are precisely those found in [47], as well as Li’s earlier work in GW theory [44]. Using this stack, the constructions in Section 4 can be adapted to reproduce the Li–Wu moduli space.

6.7.2. Artificially introducing tube complexes

To introduce tube phenomena, perform any sequence of blowup $\mathcal{Y}' \rightarrow \mathcal{Y}$ along strata. Examine the composition map

$$\mathcal{Y}^\sim \rightarrow \text{Exp}^c(X|D).$$

This map is likely not flat, but by performing universal toroidal semistable reduction [7, 58], we obtain a map between blowups:

$$\mathcal{Y}' \rightarrow \text{Exp}^b(X|D)$$

that is now flat with reduced fibers.

To orient ourselves, let us choose a point q in $\text{Exp}^b(X|D)$. It maps to a point of $\text{Exp}^c(X|D)$ that parameterizes expansions with k expanded components. By pulling back to q , we now have two expansions over q :

$$\pi : \mathcal{Y}'_q \rightarrow \mathcal{Y} \rightarrow q.$$

The space \mathcal{Y}'_q is still an expansion of X along D since any subdivision of an expansion is still an expansion. It may have, and can always be arranged to have, more than k expanded components. Let us assume that \mathcal{Y}'_q has $k + m$ expanded components. The map π contracts these new components. Let us call these m components the *extra components*.

The construction above can be performed using the formalism developed in Section 3.4. Precisely, since $\text{Exp}^b(X|D)$ and \mathcal{Y}' are both subdivisions, they can be obtained from appropriate cone structures on $|\mathbb{T}^c(\Sigma_X)|$ and on its universal 1-complex. The tube vertices, as defined earlier in the paper, give tube components in this context. These correspond exactly to the extra components.

We can now pull back the DT moduli space $\text{DT}_{\beta,n}^b(X|D)$ along the blowup map

$$\text{Exp}^c(X|D) \rightarrow \text{Exp}^b(X|D).$$

We can pullback the universal subscheme in \mathcal{Y} to the blowup \mathcal{Y}' . Since in the usual Li–Wu space, a component of the expansion can never contain a tube subscheme, we can now see that the only subschemes that can be contained in a tube/extra components of a fiber of \mathcal{Y}' are tube subschemes.

The moral of the story is that in the smooth pair setting, all the possible outputs from different choices in our construction can be obtained by the procedure described above. Since we have the Li–Wu space already in this setting, the formalism of Section 3.4 can be avoided – the paper does not produce anything beyond the system of logarithmic modifications of the Li–Wu space. But in the general setting of simple normal crossings pairs, there does not seem to exist a minimal model for the stack $\text{Exp}(X|D)$. As such, the system of models must be constructed independently of the minimal model. The situation is similar to the work of Molcho–Wise on the Picard group [59] and is explored in detail in forthcoming work of Kennedy–Hunt on general logarithmic Hilbert and Quot schemes [41].

7. Transverse limits and generically expanded targets

We verify the details of the valuative criterion for Theorem 4.6.1 that were postponed, that is, for families with a generic target expansion. An analysis via normalization requires a thorough treatment of the rubber geometry. Instead, we handle it with some additional formalism concerning logarithmic structures and tropicalizations in these circumstances; the overall structure is similar to the special case dealt with previously, so parallel reading may be advisable.

7.1. Preliminaries on valuation rings

Let K be a complete discretely valued field with algebraically closed residue field and valuation ring R . Denote the spectra of K and R by S° and S respectively, and the inclusion by $j : S^\circ \rightarrow S$. Equip the scheme S° with the standard logarithmic structure with monoid \mathbb{N} . The tropicalization of S° is canonically identified with the real dual $\text{Hom}(\mathbb{N}, \mathbb{R}_{\geq 0})$ of \mathbb{N} , where the Hom is taken in the category of monoids. The tropicalization is therefore a ray and we denote it by ρ .

In the situation typically considered, and in particular the one handled in Section 2.6, one starts with a trivial logarithmic structure on the generic point of the spectrum of a DVR $\text{Spec } R$ and endows the full spectrum with the divisorial logarithmic structure. The characteristic monoid at the closed point can be understood as being the image of R under the valuation. We are now in the case where there is an existing logarithmic structure on the generic point, which we can and will choose to have characteristic monoid \mathbb{N} , that must be extended to the closed point. We now explain how to do this. A geometric perspective is recorded in the remark below.

Consider the rank 2 lattice obtained as the direct sum of $\overline{M}_{S^\circ}^{\text{gp}}$ and the group of exponents of a uniformizer for R ; as presented, there is a natural identification with \mathbb{Z}^2 , by sending the generator of \overline{M}_{S° to $(1, 0)$ and the positive generator of the value group to $(0, 1)$. Let N_S be the dual vector space of this group. There is a generalization map on characteristic monoids from the closed point to the open point. Dualizing this map, the ray ρ above, that is, the dual cone of the characteristic monoid at S° , embeds in

N_S . Finally, there is a canonical quotient map $N_S \rightarrow \mathbb{R}$, defined by taking the quotient of the span of ρ . The image of the lattice N_S is canonically identified with the value group of K .

We examine fine and saturated extensions of the logarithmic structure from S° to S ; denote the inclusion by $j: S^\circ \hookrightarrow S$. A sufficient class is obtained as follows. Begin with the direct image logarithmic structure $j_*M_{S^\circ}$.¹⁹ The characteristic monoid of this structure at the closed point is isomorphic to the set of lexicographically nonnegative elements in \mathbb{Z}^2 . The dual of the characteristic group at the closed point is the vector space N_S defined above and the ray ρ is equipped with an embedding in N_S .

Collect the set of all two-dimensional rational polyhedral cones $\sigma \subset N_S$ containing ρ as a face and whose image in the value group is the nonnegative elements. There is an ordering among elements in this set by inclusion. The dual monoids P_σ determine fine and saturated submonoids of the characteristic monoid of direct image. We obtain, for each choice of cone σ , a fine and saturated logarithmic structure extending the one on S° whose characteristic monoid at the closed point is P_σ . Denote the extension by S_σ .

Terminology

The logarithmic structures extending the given one on S° above will be called *logarithmic extensions*. The set of such cones σ form a filtered system under inclusion. In the discussion that follows will say that a statement holds for *sufficiently small* σ to indicate that it is true after replacing σ with any element in a lower interval in the filtered system. We will correspondingly refer to the extensions as *sufficiently small (fine and saturated) extensions*.

Remark 7.1.1 (Geometric perspective on the extensions). Geometrically, the reader may visualize S as the germ of a curve whose generic point lies in the interior of a boundary divisor of a toric variety and specializes to a codimension 2 orbit, meeting this orbit with multiplicity 1. One obtains a natural logarithmic structure on S by pulling back the toric logarithmic structure. However, we can also blow up this toric variety, lift S to this blowup by taking strict transform, and perform the same construction. Iterating this construction, we obtain a sequence of logarithmic structures on S . These are all fine and saturated, but not naturally isomorphic. Further blowups lead to smaller cones in the description above. The reader may wish to contrast this with the situation where S specializes from the interior of a toric variety to the interior of a divisor, meeting the divisor with multiplicity 1. In this case as well, pullback gives rise to a natural logarithmic structure but performing blowups to the ambient toric variety do not affect the pullback logarithmic structure on S .

For a finite extension K' of K equipped with the natural valuation, the corresponding morphism $\text{Spec } R' \rightarrow \text{Spec } R$ induces a map from the value group of K' to the value group of K . It is ramified cover, étale away from from the closed points. We can equip $\text{Spec } R'$ with the a logarithmic structure by pulling back the logarithmic structure from $\text{Spec } R$ and then saturating the result. The map on value groups is an inclusion of a finite index sublattice, and in particular, if we denote by the spectrum by S' , then the associated vector spaces N_S and $N_{S'}$ above are naturally identified. The ray ρ is canonically identified the dual cone of the characteristic monoid at the generic point in $\text{Spec } R$, that is, $\text{Spec } K$. Since the map $\text{Spec } K' \rightarrow \text{Spec } K$ is étale, the morphism above is strict with the above logarithmic structures. In particular, the characteristic monoids at the generic points are naturally identified. As a result, the integral structure on the ray ρ induced from the lattice in N_S and in $N_{S'}$ are the same.

Before proceeding, we remind the reader that a map from the logarithmic scheme S_σ to an Artin cone is equivalent to the data of a map from σ to the corresponding cone [16, Section 6.2].

¹⁹The direct image logarithmic structure is typically not fine. It is possible to work with it nonetheless, but as we prefer to use standard toric machinery, we will avoid the direct use of this logarithmic structure.

7.2. Expansions and tropicalizations

Let $\mathcal{Y}_\rho \rightarrow S^\circ$ be an expansion of X , and let $\Sigma_\rho \rightarrow \rho$ be the associated tropical family, noting that Σ_ρ is embedded in $\Sigma_X \times \rho$. Let Z_ρ be an algebraically transverse subscheme of \mathcal{Y}_ρ .²⁰

Lemma 7.2.1. *There exists a rough expansion \mathcal{Y}_σ of X over a sufficiently small extension S_σ of S° , extending the family \mathcal{Y}_ρ , such that the closure of Z_ρ in \mathcal{Y}_σ is dimensionally transverse. Moreover, the family $\mathcal{Y}_\sigma \rightarrow S_\sigma$ can be chosen to have relative logarithmic rank at most 1, that is, the rank of the relative characteristic monoid of the expansion over the base at most 1.*

Proof. Let σ_0 be any smooth cone giving rise to a fine and saturated extension and consider $\Sigma_X \times \sigma_0$. The given family Σ_ρ is embedded in $\Sigma_X \times \sigma_0$. Since $\Sigma_X \times \sigma_0$ is smooth it can be embedded in a vector space, and by the completion theorem for fans, we obtain a complete fan refining $\Sigma_X \times \sigma_0$ and extending Σ_ρ ; this was referenced in Remark 3.5.2. For any sufficiently small replacement σ of σ_0 , every cone of this subdivision surjects onto a cone of σ , guaranteeing flatness. We obtain a complete subdivision Σ_σ^0 of $\Sigma_X \times \sigma$ and therefore a rough expansion of X over S_σ .

Consider the set of all subdivisions Σ_σ of Σ_σ^0 that do not change Σ_ρ . We claim any sufficiently refined subdivision has the property that the closure Z_ρ is dimensionally transverse. To see this, consider the closure of a component Y of the generic fiber of \mathcal{Y}_ρ . By the tropical compactification results for subschemes of logarithmic schemes, there exists a sequence of blowups of strata in the closure of this component in \mathcal{Y}_σ , such that the closure of $Z_\rho \cap Y$ is dimensionally transverse [74, Theorems 1.1 & 1.2]. The generic fiber subscheme has been supposed to be algebraically transverse and has dimension 1, so any stratum that it intersects has codimension at most 1 and dominates the base S_σ . Therefore, the generic fiber is disjoint from the centers of these blowups. Repeat this for each component Y of \mathcal{Y}_ρ to obtain a new total space \mathcal{Y}_σ for the target with cone complex Σ_σ and the closure of the subscheme Z_ρ is dimensionally transverse. After replacing σ with a sufficiently small member in the filtered system and pulling back the target family, every cone of Σ_σ maps surjectively onto a cone in σ , guaranteeing flatness. Pass to the union of strata that meet the closure of Z_ρ nontrivially and obtain a rough expansion of relative logarithmic rank equal to 1. □

Given a subscheme $Z_\rho \hookrightarrow \mathcal{Y}_\rho$ as above, choose an extension \mathcal{Y}_σ over S_σ such that the closure of Z_ρ is dimensionally transverse. Let Σ_σ be the corresponding embedded complex in $\Sigma_X \times \sigma$; we make the following definition.

Definition 7.2.2 (Tropicalization with expanded target). The tropicalization of Z_ρ with respect to the family \mathcal{Y}_σ over S_σ above is the union of cones in Σ_σ such that the closure of Z_ρ has nonempty intersection with the corresponding logarithmic stratum of \mathcal{Y}_σ . The tropicalization is denoted $\text{trop}(Z_\rho \hookrightarrow \mathcal{Y}_\sigma)$ viewed as a subset of $\Sigma_X \times \sigma$.

The tropicalization of a subscheme of the interior of a logarithmically smooth scheme is the image under a valuation map and is manifestly insensitive to subdivisions.

Lemma 7.2.3 (Birational invariance). *Let $\mathcal{Y}'_{\sigma'}$ and \mathcal{Y}_σ be two rough expansions over $S_{\sigma'}$ and S_σ , both extending \mathcal{Y}_ρ . After restricting both families to a sufficient small cone σ'' , the two tropicalizations*

$$\text{trop}(Z_\rho \hookrightarrow \mathcal{Y}'_{\sigma'}) \text{ and } \text{trop}(Z_\rho \hookrightarrow \mathcal{Y}_\sigma)$$

give the same subsets of $\Sigma_X \times \sigma''$.

Proof. Consider two families as in the statement, and let $\Sigma'_{\sigma'}$ and $\Sigma_{\sigma'}$ be the associated cone complexes. As we have done previously, we can complete both complexes to complete subdivisions and guarantee them to be flat after shrinking. By shrinking further, we can assume that σ and σ' coincide, and examine them as cone complexes embedded in $\Sigma_X \times \sigma$. Let Σ''_σ be the common refinement. After shrinking again, this common refinement is flat over σ . Since any subdivision of a cone complex is dominated

²⁰As before, we assume that these schemes are defined over a finitely generated subfield of K to avoid foundational issues.

by an iterated stellar subdivision, we reduce further to the case where Σ_σ is replaced by a single stellar subdivision, that is, a weighted blowup of a stratum, necessarily in the special fiber. Since the closure Z_σ of Z_ρ is dimensionally transverse to the strata of Y_σ , the set of strata that the proper and total transforms intersect nontrivially are the same. The strict transform is equal to the closure of Z_ρ in a weighted blowup \mathcal{Y}'_σ so the tropicalizations coincide. \square

Lemma 7.2.4 (Algebraic transversality). *After replacing S_σ with a ramified base change S'_σ , there exists an expansion \mathcal{Y}_σ of X over a sufficiently small extension S'_σ of S_ρ , extending the family \mathcal{Y}_ρ , such that the closure of Z_ρ in \mathcal{Y}_σ is algebraically transverse.*

Proof. After obtaining a rough expansion whose logarithmic structure has relative rank 1 with the requisite dimensional transversality property, we perform the ramified base change to obtain an expansion of X with reduced fibers [7, Section 5]. We handle the passage from dimensional transversality to algebraic transversality similarly to Section 2.6 as we now explain. However, the reader is informed, however, that after establishing dimensional transversality as we have, the problem is formally identical to the one considered by Li–Wu, whose arguments can now be plugged in [47].

Let \mathcal{Y}'_σ denote the expansion guaranteeing dimensional transversality. Any failure of algebraic transversality necessarily occurs at the double intersection of two components of the special fiber of \mathcal{Y}'_σ . If there is a Zariski open neighborhood of this double intersection inside the total space that is irreducible, then the Zariski local situation is identical to Section 2.6. In this case, the equation at this double intersection is $xy = f$, where f is an element of the valuation ring R . The tropicalization of this local neighborhood is therefore an edge of length equal to the valuation of f . We associate, as Section 2.6, an initial degeneration to each rational point on this edge, and again all but finitely many of these points have tubular initial degeneration. We may blowup by including these nontubular points to the polyhedral structure. After a base change to make these components reduced, algebraic transversality follows as in the case already treated.

The remaining case is similar. We may pass to a Zariski open neighborhood U , in the total space of the degeneration, of the double intersection. The neighborhood is necessarily reducible and its normalization consists of components U_1 and U_2 glued along a smooth divisor D . Let Z_1 and Z_2 be the restriction of Z_ρ to U_1 and U_2 , respectively, and note that Z_1 and Z_2 are both necessarily algebraically transverse. Treating the U_i separately, we may apply the results of Section 2.6. Specifically, the tropicalization of each U_i is a ray $\mathbb{R}_{\geq 0}$ and every rational point on this ray gives rise to an initial degeneration, finitely many of which are nontubular. These determine an expansion and a base change that give rise to a family of expansions of U_i , such that the closure of Z_i is algebraically transverse. We obtain expansions of both U_i along D such that the closure of Z_i is algebraically transverse.

These expansions can be naturally glued; indeed, any expansion of U_i along D can be obtained by first performing a subdivision of the family \mathcal{Y}_σ and then pulling back along the inclusion of U_i . Since algebraic transversality can be checked after normalizing, it follows that there exists a family \mathcal{Y}_σ over S'_σ of expansions such that the closure of the generic fiber is algebraically transverse. Finally, blow down any components in the special fiber whose subschemes are tubular using Lemma 2.6.10. We obtain an algebraically transverse expansion without tubular components. \square

7.3. Verification of the valuative criterion

We fix a conical structure on the moduli space T of tropical expansions and on its universal family. This determines a moduli space $DT(X|D)$ of relative ideal sheaves equipped with a structure map to the stack $\text{Exp}(X|D)$ of target expansions. We check that for each map

$$\text{Spec } K \rightarrow DT(X|D)$$

from the spectrum of a discretely valued field, there exists a canonical base change $K \subset K'$ of a well-defined smallest order, and an extension of the map to the spectrum of the corresponding valuation ring.

We have shown existence in the case where $\text{Spec } K$ maps to the locus in $\text{Exp}(X|D)$ where the logarithmic structure is trivial and therefore assume that it maps to the boundary of $\text{Exp}(X|D)$. By replacing T with a subdivision, that is, by performing a weighted blowing up of stratum to which it maps, we may choose a lift of the morphism from $\text{Spec } K$ so that it maps to the interior of a logarithmic divisor in $\text{Exp}(X|D)$. Note that the blowup is proper, so it is sufficient to check the valuative criterion in the case where the generic point maps to a divisor. Indeed, this follows from applying the valuative criterion to the blowup map itself: If we find a completion of the map to $\text{Spec } R$ on the blowup, we can compose with the blow down to obtain an extension.

We have fixed $S^\circ \subset S$. Pull back the logarithmic structure, universal expansion and subscheme to form a family

$$\begin{array}{ccc} \mathcal{Z}_\rho & \xrightarrow{\quad} & \mathcal{Y}_\rho \\ & \searrow & \swarrow \\ & S^\circ & \end{array}$$

In keeping with the notation of the preceding discussion, ρ denotes the ray corresponding to the exceptional divisor that contains the image of $\text{Spec } K$. We fix these data for the following.

Proposition 7.3.1. *There exists a sufficiently small extension S_σ of S° and its logarithmic structure, a ramified base change $S'_\sigma \rightarrow S_\sigma$, and an expansion $\mathcal{Y}_{\sigma'}$ over $S_{\sigma'}$ such that, after pullback, the closure of \mathcal{Z}_ρ is algebraically transverse and satisfies the DT stability condition.*

Proof. Apply Lemma 7.2.4 to obtain a finite extension of K and a sufficiently small extension S'_σ and an expansion $\mathcal{Y}_\sigma \rightarrow S'_\sigma$ such that the closure of \mathcal{Z}_ρ is algebraically transverse. The tropicalization gives rise to a family of 1-complexes $\Sigma_\sigma \rightarrow \sigma$. Choose any point P in the relative interior of Σ , and let \mathbb{G}_P be the associated embedded 1-complex. Each vertex v of the 1-complex \mathbb{G}_P corresponds to a component of the special fiber of the expansion family. Let $Z_v \hookrightarrow Y_v$ be the associated subscheme. A 2-valent vertex that lies on a line in Σ_X determines a component of the expansion that is a \mathbb{P}^1 -bundle over a codimension 2 stratum in X ; we examine whether the associated subscheme Z_v is a tube, that is, it is pulled back from a zero-dimensional subscheme along the bundle projection. Let $\underline{\mathbb{G}}_P$ be the 1-complex obtained by erasing all such tube vertices. Formally, if v is such a vertex that is incident to two edges, delete both v and these edges and connect the two vertices that were incident to v by a single edge. The result is an embedded 1-complex. This gives rise to a map from the ray generated by P to the complex $|\mathbb{T}|$. After shrinking σ further, for example until P is extremal, this defines a map

$$\sigma \rightarrow T.$$

The union of the 1-complexes $\underline{\mathbb{G}}_P$ do not necessarily form a cone complex, though their support is certainly the same as that of a cone complex.

By pulling back the universal tropical expansion over T to σ we obtain a subdivision Υ_σ of $\Sigma_X \times \sigma$ whose support equal to that of Σ_σ . Pass to a common refinement Δ_σ , which again, after shrinking σ is flat over σ and therefore gives rise to a rough expansion over S'_σ . After performing another ramified base change, we obtain a new expansion and a morphism of expansions associated to the map $\Delta_\sigma \rightarrow \Upsilon_\sigma$

$$\mathcal{Y}_\sigma^\Delta \rightarrow \mathcal{Y}_\sigma^\Upsilon, \text{ over } S''_\sigma.$$

As algebraic transversality is stable under additional subdivision, the closure of \mathcal{Z}_ρ inside $\mathcal{Y}_\sigma^\Delta$ is algebraically transverse. By construction, any component of the special fiber that is contracted by this map carries a tube subscheme. The resulting family is therefore also algebraically transverse and satisfies the DT stability condition. We conclude existence of limits. □

The order of the base change required for an extension can be controlled. We produced a map $\sigma \rightarrow T$ and momentarily view it as a map only at the level of the \mathbb{Q} -structures. The cone σ can be equipped with

an integral structure coming from the value group of R or any of finite extensions. When equipped with the integral structure coming from R the map to T above may not be given by an integer linear map. The restriction to ρ is an inclusion. There is a minimal base change in these circumstances, obtained by passage to a finite index sublattice of the standard lattice in N_S , such that the map can be extended, giving rise to a family of stable ideal sheaves. Algebraic transversality and DT stability are unaffected by further base change, we can conclude *a posteriori* that the family produced by this minimal extension is transverse and stable.

Call the minimal extension constructed above as the *distinguished* extension. We conclude the proof of properness by verifying the distinguished extension is the only one.

Proposition 7.3.2. *Given a morphism $\text{Spec } K \rightarrow \text{DT}(X|D)$ and an extension $\text{Spec } R'' \rightarrow \text{DT}(X|D)$ from a ramified base change, it is obtained by pulling back the distinguished extension above.*

Proof. Given the morphism from $S^\circ \rightarrow \text{DT}(X|D)$, we have obtained the distinguished extension to a logarithmic morphism $S'_\circ \rightarrow \text{DT}(X|D)$ as in the previous proposition. Let $R \subset R''$ be a finite base change, and consider an extension

$$\text{Spec } R'' \rightarrow \text{DT}(X|D).$$

By the universal property of the direct image, this lifts to a logarithmic map from the direct image logarithmic structure on $\text{Spec } R''$ to $\text{DT}(X|D)$. The target is fine and saturated so for any sufficiently small fine and saturated extension, the schematic map is enhanced to a logarithmic map $S''_\tau \rightarrow \text{DT}(X|D)$. This can be seen by direct analysis of the monoids; see also [59, Section 2.2.5].

We view τ and σ as two cones in the same vector space N_S , but that span possibly different rank 2 lattices. By shrinking, we may and do assume that τ and σ coincide set theoretically on N_S , that is, their real points coincide, and we obtain two families of embedded 1-complexes in $\Sigma_X \times \sigma$. Denote these by Y and Y' respectively, and note that they are different polyhedral structures on the same set in $\Sigma_X \times \sigma$. We may therefore pass to the common refinement of these two complexes and choose a sublattice of σ such that the associated map has reduced fibers.

Let P be a point in the relative interior of σ , and let \mathbb{G}_P be the fiber of this common refinement over P . As in the previous proposition, we may examine the closure of Z_ρ for vertices in \mathbb{G}_P that carry tube subschemes and erasing the corresponding tube vertices. This gives rise to a new graph and map from σ to $|T|$ as discussed in the previous proposition. Now note that the outcome of erasing the vertices with tube subschemes for \mathbb{G}_P is the same as erasing such vertices for the fibers over p in Y and Y' . The map from the base cone σ to T that is induced by the map above therefore coincides in both cases, at the level of the real points.

For the base change, observe that as discussed immediately before the proposition statement, there is a minimum order of base change required to extend the given map from $\text{Spec } K$. It follows that up to a further ramified base change, the two extensions coincide, establishing separatedness. \square

Acknowledgements. We are grateful to Dan Abramovich, Dori Bejleri, Qile Chen, Johan de Jong, Mark Gross, Eric Katz, Max Lieblich, Hannah Markwig, Navid Nabijou, Andrei Okounkov, Sam Payne, Mattia Talpo, Richard Thomas, Martin Ulirsch, Jeremy Usatine and Jonathan Wise for numerous conversations over the years on related topics. We extend a special thanks to Patrick Kennedy-Hunt for numerous questions and comments on earlier versions of this paper that helped improve both the presentation and the mathematics. The first author would like to thank Daniel Tiger for providing extra childcare during the final stages of this project. The second author was a Moore Instructor at MIT and a visitor at the Chennai Mathematical Institute during work on this project and is grateful to both institutions for excellent working conditions. Two anonymous referees provided detailed feedback which led to various improvements to both exposition and mathematics.

Competing interest. The authors have no competing interest to declare.

Financial support. D.R. is partially supported by EPSRC grant EP/V051830/1.

References

- [1] D. Abramovich, L. Caporaso and S. Payne, ‘The tropicalization of the moduli space of curves’, *Ann. Sci. Éc. Norm. Supér.* **48** (2015), 765–809.
- [2] D. Abramovich and Q. Chen, ‘Stable logarithmic maps to Deligne–Faltings pairs II’, *Asian J. Math.* **18** (2014), 465–488.
- [3] D. Abramovich, Q. Chen, M. Gross and B. Siebert, ‘Decomposition of degenerate Gromov–Witten invariants’, *Compos. Math.* **156** (2020), 2020–2075.
- [4] D. Abramovich, Q. Chen, M. Gross and B. Siebert, ‘Punctured logarithmic maps’, Preprint, 2020, [arXiv:2009.07720](https://arxiv.org/abs/2009.07720)
- [5] D. Abramovich, Q. Chen, S. Marcus, M. Ulirsch and J. Wise, ‘Skeletons and fans of logarithmic structures’, in *Nonarchimedean and Tropical Geometry*, M. Baker and S. Payne, eds., Simons Symposia, (Springer, 2016), 287–336.
- [6] D. Abramovich, Q. Chen, S. Marcus and J. Wise, ‘Boundedness of the space of stable logarithmic maps’, *J. Eur. Math. Soc.* **19** (2017), 2783–2809.
- [7] D. Abramovich and K. Karu, ‘Weak semistable reduction in characteristic 0’, *Invent. Math.* **139** (2000), 241–273.
- [8] D. Abramovich and J. Wang, ‘Equivariant resolution of singularities in characteristic 0’, *Math. Res. Lett.* **4** (1997), 427–433.
- [9] D. Abramovich and J. Wise, ‘Birational invariance in logarithmic Gromov–Witten theory’, *Comp. Math.* **154** (2018), 595–620.
- [10] K. B. Ascher and S. Molcho, ‘Logarithmic stable toric varieties and their moduli’, *Algebraic Geometry* **3** (2016), 296–319.
- [11] L. Battistella and N. Nabijou, ‘Relative quasimaps and mirror formulae’, *Int. Math. Res. Not. IMRN* (2021), 7885–7931.
- [12] K. Behrend and B. Fantechi, ‘Symmetric obstruction theories and Hilbert schemes of points on threefolds’, *Algebra & Number Theory* **2** (2008), 313–345.
- [13] B. Bertrand, E. Brugallé and L. López de Medrano, ‘Planar tropical cubic curves of any genus, and higher dimensional generalisations’, *Enseign. Math.* **64** (2018), 415–457.
- [14] R. Bieri and J. R. Groves, ‘The geometry of the set of characters induced by valuations’, *J. Reine Angew. Math. (Crelle’s Journal)* **347** (1984), 168–195.
- [15] M. Brandt and M. Ulirsch, ‘Symmetric powers of algebraic and tropical curves: a non-Archimedean perspective’, *Trans. Amer. Math. Soc. Ser. B* **9** (2022), 586–618.
- [16] R. Cavalieri, M. Chan, M. Ulirsch and J. Wise, ‘A moduli stack of tropical curves’, *Forum Math. Sigma* **8** (2020), 1–93.
- [17] Q. Chen, ‘Stable logarithmic maps to Deligne–Faltings pairs I’, *Ann. of Math.* **180** (2014), 341–392.
- [18] Q. Chen and M. Satriano, ‘Chow quotients of toric varieties as moduli of stable log maps’, *Algebra and Number Theory Journal* **7** (2013), 2313–2329.
- [19] K. Costello, ‘Higher genus Gromov–Witten invariants as genus zero invariants of symmetric products’, *Ann. of Math.* (2006), 561–601.
- [20] D. A. Cox, J. B. Little and H. K. Schenck, *Toric Varieties*, Graduate Studies in Mathematics, vol. 124 (American Mathematical Society, Providence, RI, 2011).
- [21] S. K. Donaldson and R. P. Thomas, ‘Gauge theory in higher dimensions’, in *The Geometric Universe (Oxford, 1996)* (Oxford Univ. Press, Oxford, 1998), 31–47.
- [22] A. Douady and J.-L. Verdier, ‘Séminaire de géométrie analytique de l’École normale supérieure 1974/1975’, *Asterisque exposes VI–IX* (1976).
- [23] G. Ewald and M.-N. Ishida, ‘Completion of real fans and Zariski–Riemann spaces’, *Tohoku Math. J.* **58** (2006), 189–218.
- [24] W. Fulton, *Introduction to Toric Varieties* (Princeton University Press, Princeton, NJ, 1993).
- [25] I. M. Gelfand, M. M. Kapranov and A. V. Zelevinsky, ‘Triangulations and secondary polytopes’, in *Discriminants, Resultants, and Multidimensional Determinants* (Springer, 1994), 214–251.
- [26] J. I. B. Gil and M. Sombra, ‘When do the recession cones of a polyhedral complex form a fan?’ *Discrete & Computational Geometry* **46** (2011), 789–798.
- [27] M. Gross and B. Siebert, ‘Logarithmic Gromov–Witten invariants’, *J. Amer. Math. Soc.* **26** (2013), 451–510.
- [28] W. Gubler, ‘A guide to tropicalizations’, in *Algebraic and Combinatorial Aspects of Tropical Geometry*, Contemp. Math., vol. 589 (Amer. Math. Soc., Providence, RI, 2013), 125–189.
- [29] P. Hacking, S. Keel and J. Tevelev, ‘Stable pair, tropical, and log canonical compactifications of moduli spaces of del Pezzo surfaces’, *Invent. Math.* **178** (2009), 173–227.
- [30] L. Herr and J. Wise, ‘Costello’s pushforward formula: errata and generalization’, *manuscripta mathematica* **171** (2023), 621–642.
- [31] D. Huybrechts and M. Lehn, *The Geometry of Moduli Spaces of Sheaves* (Cambridge University Press, 2010).
- [32] D. Huybrechts and R. P. Thomas, ‘Deformation-obstruction theory for complexes via Atiyah and Kodaira–Spencer classes’, *Math. Ann.* **346** (2010), 545–569.
- [33] E.-N. Ionel and T. H. Parker, ‘Relative Gromov–Witten invariants’, *Ann. of Math.* (2003), 45–96.
- [34] M. Kapranov, *Chow Quotients of Grassmannians I*, *IM Gelfand Seminar*, vol. **16**, 1993, 29–110.
- [35] K. Kato, *Logarithmic Structures of Fontaine–Illusie*, Algebraic Analysis, Geometry, and Number Theory (Baltimore, MD, 1988), (1989), 191–224.
- [36] E. Katz, ‘Tropical invariants from the secondary fan’, *Adv. Geom.* **9** (2009), 153–180.
- [37] E. Katz and S. Payne, ‘Realization spaces for tropical fans’, in *Combinatorial Aspects of Commutative Algebra and Algebraic Geometry* (Springer, 2011), 73–88.

- [38] S. Keel and J. Tevelev, ‘Geometry of Chow quotients of Grassmannians’, *Duke Math. J.* **134** (2006), 259–311.
- [39] G. Kempf, F. Knudsen, D. Mumford and B. Saint-Donat, ‘Toroidal embeddings I’, *Lecture Notes in Mathematics* **339** (1973).
- [40] P. Kennedy-Hunt, ‘Logarithmic Pandharipande–Thomas spaces and the secondary polytope’, Preprint, 2021, [arXiv:2112.00809](https://arxiv.org/abs/2112.00809)
- [41] P. Kennedy-Hunt, ‘The logarithmic quot space: foundations and tropicalisation’, Preprint, 2023, [arXiv:2308.14470](https://arxiv.org/abs/2308.14470).
- [42] M. Levine and R. Pandharipande, ‘Algebraic cobordism revisited’, *Invent. Math.* **176** (2009), 63–130.
- [43] A.-M. Li and Y. Ruan, ‘Symplectic surgery and Gromov–Witten invariants of Calabi-Yau 3-folds’, *Invent. Math.* **145** (2001), 151–218.
- [44] J. Li, ‘Stable morphisms to singular schemes and relative stable morphisms’, *J. Diff. Geom.* **57** (2001), 509–578.
- [45] J. Li, A degeneration formula of GW-invariants, *J. Diff. Geom.* **60** (2002), 199–293.
- [46] J. Li, ‘Zero dimensional Donaldson–Thomas invariants of threefolds’, *Geometry & Topology* **10** (2006), 2117–2171.
- [47] J. Li and B. Wu, ‘Good degeneration of Quot-schemes and coherent systems’, *Communications in Analysis and Geometry* **23** (2015), 841–921.
- [48] D. Maclagan and B. Sturmfels, *Introduction to Tropical Geometry*, Graduate Studies in Mathematics, vol. 161 (American Mathematical Society, Providence, RI, 2015).
- [49] C. Manolache, ‘Virtual pull-backs’, *J. Algebr. Geom.* **21** (2012), 201–245.
- [50] S. Marcus and J. Wise, ‘Logarithmic compactification of the Abel–Jacobi section’, Preprint, 2017, [arXiv:1708.04471](https://arxiv.org/abs/1708.04471)
- [51] D. Maulik, N. Nekrasov, A. Okounkov and R. Pandharipande, ‘Gromov–Witten theory and Donaldson–Thomas theory, I’, *Compos. Math.* **142** (2006), 1263–1285.
- [52] D. Maulik, N. Nekrasov, A. Okounkov and R. Pandharipande, ‘Gromov–Witten theory and Donaldson–Thomas theory, II’, *Compos. Math.* **142** (2006), 1286–1304.
- [53] D. Maulik, A. Oblomkov, A. Okounkov and R. Pandharipande, ‘Gromov-Witten/Donaldson-Thomas correspondence for toric 3-folds’, *Invent. Math.* **186** (2011), 435–479.
- [54] D. Maulik and A. Okounkov, Quantum groups and quantum cohomology, *Astérisque* (2019), ix+209.
- [55] D. Maulik, R. Pandharipande and R. P. Thomas, ‘Curves on K3 surfaces and modular forms’, *J. Topol.* **3** (2010), 937–996.
- [56] D. Maulik and D. Ranganathan, ‘Logarithmic enumerative geometry for curves and sheaves’, Preprint, 2023, [arXiv:2311.14150](https://arxiv.org/abs/2311.14150).
- [57] G. Mikhalkin, ‘Enumerative tropical geometry in \mathbb{R}^2 ’, *J. Amer. Math. Soc.* **18** (2005), 313–377.
- [58] S. Molcho, ‘Universal stacky semistable reduction’, *Israel J. Math.* **242** (2021), 55–82.
- [59] S. Molcho and J. Wise, ‘The logarithmic Picard group and its tropicalization’, *Compos. Math.* **158** (2022), 1477–1562.
- [60] D. Mumford, ‘An analytic construction of degenerating abelian varieties over complete rings’, *Compos. Math.* **24** (1972), 239–272.
- [61] T. Nishinou and B. Siebert, ‘Toric degenerations of toric varieties and tropical curves’, *Duke Math. J.* **135** (2006), 1–51.
- [62] M. C. Olsson, ‘Logarithmic geometry and algebraic stacks’, *Ann. Sci. Éc. Norm. Supér.* **36** (2003), 747–791.
- [63] R. Pandharipande and A. Pixton, ‘Gromov-Witten/pairs correspondence for the quintic 3-fold’, *J. Amer. Math. Soc.* **30** (2017), 389–449.
- [64] R. Pandharipande and R. P. Thomas, ‘Curve counting via stable pairs in the derived category’, *Invent. Math.* **178** (2009), 407–447.
- [65] R. Pandharipande and R. P. Thomas, ‘Stable pairs and BPS invariants’, *J. Amer. Math. Soc.* **23** (2010), 267–297.
- [66] S. Payne, ‘Fibers of tropicalization’, *Math. Z.* **262** (2009), 301–311.
- [67] D. Ranganathan, ‘Skeletons of stable maps I: Rational curves in toric varieties’, *J. Lond. Math. Soc.* **95** (2017), 804–832.
- [68] D. Ranganathan, ‘Logarithmic Gromov–Witten theory with expansions’, *Algebr. Geom.* **9** (2022), 714–761.
- [69] D. Ranganathan and J. Wise, ‘Rational curves in the logarithmic multiplicative group’, *Proc. Amer. Math. Soc.* **148** (2020), 103–110.
- [70] M. Raynaud and L. Gruson, ‘Criteres de platitude et de projectivité: Techniques de platisation d’un module’, *Invent. Math.* **13** (1971), 1–89.
- [71] J. Tevelev, ‘Compactifications of subvarieties of tori’, *Amer. J. Math.* **129** (2007), 1087–1104.
- [72] The Stacks Project Authors, *stacks project*. <http://stacks.math.columbia.edu>, 2019.
- [73] R. P. Thomas, ‘A holomorphic Casson invariant for Calabi-Yau 3-folds, and bundles on K3 fibrations’, *J. Differential Geom.* **54** (2000), 367–438.
- [74] M. Ulirsch, ‘Tropical compactification in log-regular varieties’, *Math. Z.* **280** (2015), 195–210.
- [75] M. Ulirsch, ‘Functorial tropicalization of logarithmic schemes: the case of constant coefficients’, *Proc. Lond. Math. Soc.* **114** (2017), 1081–1113.
- [76] T. Y. Yu, ‘The number of vertices of a tropical curve is bounded by its area’, *L’Enseignement Mathématique* **60** (2015), 257–271.