

Tropical geometry and mirror symmetry

Mark Gross

UCSD MATHEMATICS, 9500 GILMAN DRIVE, LA JOLLA, CA 92093-0112
E-mail address: mgross@math.ucsd.edu

Key words and phrases. Mirror symmetry, Tropical geometry

ABSTRACT. This monograph, based on lectures given at the NSF-CBMS conference on Tropical Geometry and Mirror Symmetry at Kansas State University, aims to present a snapshot of ideas being developed by Gross and Siebert to understand mirror symmetry via tropical geometry.

In this program, there are three worlds. The first part of the book presents these three linked realms: tropical geometry, in which geometric objects are piecewise linear objects; the A- and B-models of mirror symmetry (Gromov-Witten theory and period integrals respectively); and log geometry. Log geometry is used to go between the world of tropical geometry, on the one hand, and the world of the A- and B-models, on the other.

Next, one complete example is given in depth, namely mirror symmetry for \mathbb{P}^2 . Following Siebert and Nishinou, a complete proof of Mikhalkin's tropical curve counting theorem is given for toric surfaces. Gross's mirror result for \mathbb{P}^2 showing how period integrals can be computed directly in terms of tropical geometry is then given.

Finally, the book ends with a survey of the Gross-Siebert program in the Calabi-Yau case. A complete proof of the correspondence theorem between affine K3 surfaces and degenerations of K3 surfaces is given, a special case of the correspondence theorem for Calabi-Yau varieties proved by Gross and Siebert in general and proved for K3 surfaces by Kontsevich and Soibelman.

“Maybe the tropics,” somebody, probably the General said, “but never the Polar Region, it’s too white, too mathematical up there.”

—Thomas Pynchon, *Against the Day*.

Contents

Preface	ix
Introduction	xi
Part 1. The three worlds	1
Chapter 1. The tropics	3
1.1. Tropical hypersurfaces	3
1.2. Some background on fans	11
1.3. Parameterized tropical curves	13
1.4. Affine manifolds with singularities	19
1.5. The discrete Legendre transform	27
1.6. Tropical curves on tropical surfaces	30
1.7. References and further reading	32
Chapter 2. The A- and B-models	33
2.1. The A-model	33
2.2. The B-model	67
2.3. References and further reading	89
Chapter 3. Log geometry	91
3.1. A brief review of toric geometry	92
3.2. Log schemes	98
3.3. Log derivations and differentials	112
3.4. Log deformation theory	117
3.5. The twisted de Rham complex revisited	126
3.6. References and further reading	129
Part 2. Example: \mathbb{P}^2.	131
Chapter 4. Mikhalkin's curve counting formula	133
4.1. The statement and outline of the proof	133
4.2. Log world \rightarrow tropical world	138
4.3. Tropical world \rightarrow log world	143
4.4. Classical world \rightarrow log world	157
4.5. Log world \rightarrow classical world	165
4.6. The end of the proof	169
4.7. References and further reading	171
Chapter 5. Period integrals	173
5.1. The perturbed Landau-Ginzburg potential	173

5.2. Tropical descendent invariants	179
5.3. The main B-model statement	187
5.4. Deforming Q and P_1, \dots, P_k	192
5.5. Evaluation of the period integrals	222
5.6. References and further reading	244
Part 3. The Gross-Siebert program	245
Chapter 6. The program and two-dimensional results	247
6.1. The program	247
6.2. From integral tropical manifolds to degenerations in dimension two	255
6.3. Achieving compatibility: The tropical vertex group	290
6.4. Remarks and generalizations	305
6.5. References and further reading	306
Bibliography	307
Index of Symbols	313
General Index	315

Preface

The NSF-CBMS conference on Tropical Geometry and Mirror Symmetry was held at Kansas State University during the period December 13–17, 2008. It was organized by Ricardo Castaño-Bernard, Yan Soibelman, and Ilia Zharkov. During this time, I gave ten hours of lectures. In addition, talks were given by M. Abouzaid, K.-W. Chan, C. Doran, K. Fukaya, I. Itenberg, L. Katzarkov, A. Mavlyutov, D. Morrison, Y.-G. Oh, T. Pantev, B. Siebert, and B. Young.

My talks were meant to give a snapshot of a long-term program currently being carried out with Bernd Siebert aimed at achieving a fundamental conceptual understanding of mirror symmetry. Tropical geometry emerges naturally in this program, so in the lectures I took a rather ahistorical point of view. Starting with the tropical semi-ring, I developed tropical geometry and explained Mikhalkin's tropical curve-counting formulas, outlining the proof given by Nishinou and Siebert. I then explained my recent work in connecting this to the mirror side. Finally, I sketched the ideas behind recent work by myself and Siebert on constructing degenerations of Calabi-Yau manifolds from affine manifolds with singularities.

This monograph follows the structure of the lectures closely, filling in many details which were not given there. Like the lectures, this monograph only represents a snapshot of an evolving program, but I hope it will be useful to those who may wish to become involved in this program.

NSF grant DMS-0735319 provided support both for the conference and for this book.

Mark Gross, La Jolla, 2010

Introduction

The early history of mirror symmetry has been told many times; we will only summarize it briefly here. The story begins with the introduction of Calabi-Yau compactifications in string theory in 1985 [11]. The idea is that, since superstring theory requires a ten-dimensional space-time, one reconciles this with the observed universe by requiring (at least locally) that space-time take the form

$$\mathbb{R}^{1,3} \times X,$$

where $\mathbb{R}^{1,3}$ is usual Minkowski space-time and X is a very small six-dimensional Riemannian manifold. The desire for the theory to preserve the supersymmetry of superstring theory then leads to the requirement that X have $SU(3)$ holonomy, i.e., be a Calabi-Yau manifold. Thus string theory entered the realm of algebraic geometry, as any non-singular projective threefold with trivial canonical bundle carries a metric with $SU(3)$ holonomy, thanks to Yau's proof of the Calabi conjecture [113].

This generated an industry in the string theory community devoted to producing large lists of examples of Calabi-Yau threefolds and computing their invariants, the most basic of which are the Hodge numbers $h^{1,1}$ and $h^{1,2}$.

In 1989, a rather surprising observation came out of this work. Candelas, Lynker and Schimmrigk [12] provided a list of Calabi-Yau hypersurfaces in weighted projective space which exhibited an obvious symmetry: if there was a Calabi-Yau threefold with Hodge numbers given by a pair $(h^{1,1}, h^{1,2})$, then there was often also one with Hodge numbers given by the pair $(h^{1,2}, h^{1,1})$. Independently, guided by certain observations in conformal field theory, Greene and Plesser [36] studied the quintic threefold and its mirror partner. If we let X_ψ be the solution set in \mathbb{P}^4 of the equation

$$x_0^5 + \cdots + x_4^5 - \psi x_0 x_1 x_2 x_3 x_4 = 0$$

for $\psi \in \mathbb{C}$, then for most ψ , X_ψ is a non-singular quintic threefold, and as such, has Hodge numbers

$$h^{1,1}(X_\psi) = 1, \quad h^{1,2}(X_\psi) = 101.$$

On the other hand, the group

$$G = \frac{\{(a_0, \dots, a_4) | a_i \in \mu_5, \prod_{i=0}^4 a_i = 1\}}{\{(a, a, a, a, a) | a \in \mu_5\}}$$

acts diagonally on \mathbb{P}^4 , via

$$(x_0, \dots, x_4) \mapsto (a_0 x_0, \dots, a_4 x_4).$$

Here μ_5 is the group of fifth roots of unity. This action restricts to an action on X_ψ , and the quotient X_ψ/G is highly singular. However, these singularities can be

resolved via a proper birational morphism $\check{X}_\psi \rightarrow X_\psi/G$ with \check{X}_ψ a new Calabi-Yau threefold with Hodge numbers

$$h^{1,1}(\check{X}_\psi) = 101, \quad h^{1,2}(\check{X}_\psi) = 1.$$

These examples were already a surprise to mathematicians, since at the time very few examples of Calabi-Yau threefolds with positive Euler characteristic were known (the Euler characteristic coinciding with $2(h^{1,1} - h^{1,2})$).

Much more spectacular were the results of Candelas, de la Ossa, Green and Parkes [10]. Guided by string theory and path integral calculations, Candelas et al. conjectured that certain period calculations on the family \check{X}_ψ parameterized by ψ would yield predictions for numbers of rational curves on the quintic threefold. They carried out these calculations, finding agreement with the known numbers of rational curves up to degree 3. We omit any details of these calculations here, as they have been expounded in many places, see e.g., [43]. This agreement was very surprising to the mathematical community, as these numbers become increasingly difficult to compute as the degree increases. The number of lines, 2875, was known in the 19th century, the number of conics, 609250, was computed only in 1986 by Sheldon Katz [66], and the number of twisted cubics, 317206375, was only computed in 1990 by Ellingsrud and Strømme [22].

Throughout the history of mathematics, physics has been an important source of interesting problems and mathematical phenomena. Some of the interesting mathematics that arises from physics tends to be a one-off — an interesting and unexpected formula, say, which once verified mathematically loses interest. Other contributions from physics have led to powerful new structures and theories which continue to provide interesting and exciting new results. I like to believe that mirror symmetry is one of the latter types of subjects.

The conjecture raised by Candelas et al., along with related work, led to the study of Gromov-Witten invariants (defining precisely what we mean by “the number of rational curves”) and quantum cohomology, a way of deforming the usual cup product on cohomology using Gromov-Witten invariants. This remains an active field of research, and by 1996, the theory was sufficiently developed to allow proofs of the mirror symmetry formula for the quintic by Givental [34], Lian, Liu and Yau [75] and subsequently others, with the proofs getting simpler over time.

Concerning mirror symmetry, Batyrev [6] and Batyrev-Borisov [7] gave very general constructions of mirror pairs of Calabi-Yau manifolds occurring as complete intersections in toric varieties. In 1994, Maxim Kontsevich [68] made his fundamental Homological Mirror Symmetry conjecture, a profound effort to explain the relationship between a Calabi-Yau manifold and its mirror in terms of category theory.

In 1996, Strominger, Yau and Zaslow proposed a conjecture, [108], now referred to as the SYZ conjecture, suggesting a much more concrete geometric relationship between mirror pairs; namely, mirror pairs should carry dual special Lagrangian fibrations. This suggested a very explicit relationship between a Calabi-Yau manifold and its mirror, and initial work in this direction by myself [37, 38, 39] and Wei-Dong Ruan [97, 98, 99] indicates the conjecture works at a topological level. However, to date, the analytic problems involved in proving a full-strength version of the SYZ conjecture remain insurmountable. Furthermore, while a proof of the SYZ conjecture would be of great interest, a proof alone will not explain the finer aspects of mirror symmetry. Nevertheless, the SYZ conjecture has motivated

several points of view which appear to be yielding new insights into mirror symmetry: notably, the rigid analytic program initiated by Kontsevich and Soibelman in [69, 70] and the program developed by Siebert and myself using log geometry, [47, 48, 51, 49].

These ideas which grew out of the SYZ conjecture focus on the base of the SYZ fibration; even though we do not know an SYZ fibration exists, we have a good guess as to what these bases look like. In particular, they should be *affine manifolds*, i.e., real manifolds with an atlas whose transition maps are affine linear transformations. In general, these manifolds have a singular locus, a subset not carrying such an affine structure. It is not difficult to write down examples of such manifolds which we expect to correspond, say, to hypersurfaces in toric varieties.

More precisely,

DEFINITION 0.1. An *affine manifold* B is a real manifold with an atlas of coordinate charts

$$\{\psi_i : U_i \rightarrow \mathbb{R}^n\}$$

with $\psi_i \circ \psi_j^{-1} \in \text{Aff}(\mathbb{R}^n)$, the affine linear group of \mathbb{R}^n . We say B is *tropical* (respectively *integral*) if $\psi_i \circ \psi_j^{-1} \in \mathbb{R}^n \rtimes \text{GL}_n(\mathbb{Z}) \subseteq \text{Aff}(\mathbb{R}^n)$ (respectively $\psi_i \circ \psi_j^{-1} \in \text{Aff}(\mathbb{Z}^n)$, the affine linear group of \mathbb{Z}^n).

In the tropical case, the linear part of each coordinate transformation is integral, and in the integral case, both the translational and linear parts are integral.

Given a tropical manifold B , we have a family of lattices $\Lambda \subseteq \mathcal{T}_B$ generated locally by $\partial/\partial y_1, \dots, \partial/\partial y_n$, where y_1, \dots, y_n are affine coordinates. The condition on transition maps guarantees that this is well-defined. Dually, we have a family of lattices $\check{\Lambda} \subseteq \mathcal{T}_B^*$ generated by dy_1, \dots, dy_n , and then we get two torus bundles

$$\begin{aligned} f : X(B) &\rightarrow B \\ \check{f} : \check{X}(B) &\rightarrow B \end{aligned}$$

with

$$X(B) = \mathcal{T}_B/\Lambda, \quad \check{X}(B) = \mathcal{T}_B^*/\check{\Lambda}.$$

Now $X(B)$ carries a natural complex structure. Sections of Λ are flat sections of a connection on \mathcal{T}_B , and the horizontal and vertical tangent spaces of this connection are canonically isomorphic. Thus we can write down an almost complex structure J which interchanges these two spaces, with an appropriate sign-change so that $J^2 = -\text{id}$. It is easy to see that this almost complex structure on \mathcal{T}_B is integrable and descends to $X(B)$.

On the other hand, \mathcal{T}_B^* carries a canonical symplectic form which descends to $\check{X}(B)$, so $\check{X}(B)$ is canonically a symplectic manifold.

We can think of $X(B)$ and $\check{X}(B)$ as forming a mirror pair; this is a simple version of the SYZ conjecture. In this simple situation, however, there are few interesting compact examples, in the Kähler case being limited to the possibility that $B = \mathbb{R}^n/\Gamma$ for a lattice Γ (shown in [15]). Nevertheless, we can take this simple case as motivation, and ask some basic questions:


- (1) What geometric structures on B correspond to geometric structures of interest on $X(B)$ and $\check{X}(B)$?
- (2) If we want more interesting examples, we need to allow B to have singularities, i.e., have a tropical affine structure on an open set $B_0 \subseteq B$ with

$B \setminus B_0$ relatively small (e.g., codimension at least two). How do we deal with this?

By 2000, it was certainly clear to many of the researchers in the field that holomorphic curves in $X(B)$ should correspond to certain sorts of piecewise linear graphs in B . Kontsevich suggested the possibility that one might be able to actually carry out a curve count by counting these graphs. In 2002, Mikhalkin [79, 80] announced that this was indeed possible, introducing and proving curve-counting formulas for toric surfaces. This was the first evidence that one could really compute invariants using these piecewise linear graphs. For historical reasons which will be explained in Chapter 1, Mikhalkin called these piecewise linear graphs “tropical curves,” introducing the word “tropical” into the field.


This brings us to the following picture. Mirror symmetry involves a relationship between two different types of geometry, usually called the A-model and the B-model. The A-model involves symplectic geometry, which is the natural category in which to discuss such things as Gromov-Witten invariants, while the B-model involves complex geometry, where one can discuss such things as period integrals.

This leads us to the following conceptual framework for mirror symmetry:



Here, we wish to explain mirror symmetry by identifying what we shall refer to as tropical structures in B which can be interpreted as geometric structures in the A- and B-models. However, the interpretations in the A- and B-models should be different, i.e., mirror, so that the fact that these structures are given by the same tropical structures then gives a conceptual explanation for mirror symmetry. For the most well-known aspect of mirror symmetry, namely the enumeration of rational curves, the hope should be that tropical curves on B correspond to (pseudo)-holomorphic curves in the A-model and corrections to period calculations in the B-model.

The main idea of my program with Siebert is to try to understand how to go between the tropical world and the A- and B-models by passing through another world, the world of log geometry. One can view log geometry as half-way between tropical geometry and classical geometry:



As this program with Siebert is ongoing, with much work still to be done, my lectures at the CBMS regional conference in Manhattan, Kansas were intended to give a snapshot of the current state of this program. This monograph closely follows the outline of those lectures. The basic goal is threefold.

First, I wish to explain explicitly, at least in special cases, all the worlds suggested in the above diagram: the tropical world, the “classical” world of the A- and B-model, and log geometry.

Second, I would like to explain one very concrete case where the full picture has been worked out for both the A- and B-models. This is the case of \mathbb{P}^2 . For the A-model, curve counting is the result of Mikhalkin, and here I will give a proof of his result adapted from a more general result of Nishinou and Siebert [86], as that approach is more in keeping with the philosophy of the program. For the B-model, I will explain my own recent work [42] which shows how period integrals extract tropical information.

Third, I wish to survey some of the results obtained by Siebert and myself in the Calabi-Yau case, outlining how this approach can be expected to yield a proof of mirror symmetry. While for \mathbb{P}^2 I give complete details, this third part is intended to be more of a guide for reading the original papers, which unfortunately are quite long and technical. I hope to at least convey an intuition for this approach.

I will take a very ahistorical approach to all of this, starting with the basics of tropical geometry and working backwards, showing how a study of tropical geometry can lead naturally to other concepts which first arose in the study of mirror symmetry. In a way, this may be natural. To paraphrase Witten’s statement about string theory, mirror symmetry often seems like a piece of twenty-first century mathematics which fell into the twentieth century. Its initial discovery in string theory represents some of the more difficult aspects of the theory. Even an explanation of the calculations carried out by Candelas et al. can occupy a significant portion of a course, and the theory built up to define and compute Gromov-Witten invariants is even more involved. On the other hand, the geometry that now seems to underpin mirror symmetry, namely tropical geometry, is very simple and requires no particular background to understand. So it makes sense to develop the discussion from the simplest starting point.

The prerequisites of this volume include a familiarity with algebraic geometry at the level of Hartshorne’s text [57] as well as some basic differential geometry. In addition, familiarity with toric geometry will be very helpful; the text will recall many of the basic necessary facts about toric geometry, but at least some previous experience will be useful. For a more in-depth treatment of toric geometry, I

recommend Fulton's lecture notes [27]. We shall also, in Chapter 3, make use of sheaves in the étale topology, which can be reviewed in [83], Chapter II. However, this use is not vital to most of the discussion here.

I would like to thank many people. Foremost, I would like to thank Ricardo Castaño-Bernard, Yan Soibelman, and Ilia Zharkov for organizing the NSF-CBMS conference at Kansas State University. Second, I would like to thank Bernd Siebert; the approach in this book grew out of our joint collaboration. I also thank the many people who answered questions and commented on the manuscript, including Sean Keel, M. Brandon Meredith, Rahul Pandharipande, D. Peter Overholser, Daniel Schultheis, and Katharine Shultis.

Parts of this book were written during a visit to Oxford; I thank Philip Candelas for his hospitality during this visit. The final parts of the book were written during the fall of 2009 at MSRI; I thank MSRI for its financial support via a Simons Professorship.

I would like to thank Lori Lejeune for providing the files for Figures 17, 18 and 19 of Chapter 1 and Figures 7 and 8 of Chapter 6. Finally, and definitely not least, I thank Arthur Greenspoon, who generously offered to proofread this volume.

Convention. Throughout this book \mathbb{k} denotes an algebraically closed field of characteristic zero. \mathbb{N} denotes the set of natural numbers $\{0, 1, 2, \dots\}$.

Part 1

The three worlds

CHAPTER 1

The tropics

We start with the simplest of the three worlds, the tropical world. Tropical geometry is a kind of piecewise linear combinatorial geometry which arises when one starts to think about algebraic geometry over the so-called tropical semi-ring.

This chapter will give a rather shallow introduction to the subject. We will start with the definition of the tropical semi-ring and some elementary algebraic geometry over the tropical semi-ring. We move on to the notion of parameterized tropical curve, which features in Mikhalkin's curve counting results. Next, we introduce the type of tropical objects which arise in the Gross-Siebert program: affine manifolds with singularities. These arise naturally if one wants to think about curve counting in Calabi-Yau manifolds. We end with a duality between such objects given by the Legendre transform.

1.1. Tropical hypersurfaces

We begin with the *tropical semi-ring*,

$$\mathbb{R}^{\text{trop}} = (\mathbb{R}, \oplus, \odot).$$

Here \mathbb{R} is the set of real numbers, but with addition and multiplication defined by

$$a \oplus b := \min(a, b)$$

$$a \odot b := a + b.$$

Of course there is no additive inverse. This semi-ring became known as the tropical semi-ring in honour of the Brazilian mathematician Imre Simon. The word *tropical* has now spread rapidly.

We would like to do algebraic geometry over the tropical semi-ring instead of over a field. Of course, since there is no additive identity in this semi-ring, it is not immediately obvious what the zero-locus of a polynomial should be. The correct, or rather, useful, interpretation is as follows. Let

$$\mathbb{R}^{\text{trop}}[x_1, \dots, x_n]$$


denote the space of functions $f : \mathbb{R}^n \rightarrow \mathbb{R}$ given by tropical polynomials

$$f(x_1, \dots, x_n) = \sum_{(i_1, \dots, i_n) \in S} a_{i_1, \dots, i_n} x_1^{i_1} \cdots x_n^{i_n}$$

where $S \subseteq \mathbb{Z}^n$ is a finite index set. Here all operations are in \mathbb{R}^{trop} , so this is really the function

$$f(x_1, \dots, x_n) = \min\{a_{i_1, \dots, i_n} + \sum_{k=1}^n i_k x_k \mid (i_1, \dots, i_n) \in S\}$$

This is a piecewise linear function, and the *tropical hypersurface* defined by f , $V(f) \subseteq \mathbb{R}^n$, as a set, is the locus where f is not linear.

FIGURE 1. $0 \oplus (0 \odot x_1)$

In order to write these formulas in a more invariant way, in what follows we shall often make use of the notation

$$M = \mathbb{Z}^n, \quad M_{\mathbb{R}} = M \otimes_{\mathbb{Z}} \mathbb{R}, \quad N = \text{Hom}_{\mathbb{Z}}(M, \mathbb{Z}), \quad N_{\mathbb{R}} = N \otimes_{\mathbb{Z}} \mathbb{R}.$$

We denote evaluation of $n \in N$ on $m \in M$ by $\langle n, m \rangle$. We shall often use the notion of *index* of an element $m \in M \setminus \{0\}$; this is the largest positive integer r such that there exists $m' \in M$ with $rm' = m$. If the index of m is 1, we say m is *primitive*.

With this notation, we can view a tropical function as a map $f : M_{\mathbb{R}} \rightarrow \mathbb{R}$ written, for $S \subseteq N$ a finite set, as

$$f(z) = \sum_{n \in S} a_n z^n := \min\{a_n + \langle n, z \rangle \mid n \in S\}.$$

Now $V(f)$ will be a union of codimension one polyhedra in \mathbb{R}^n . Here, by a *polyhedron*, we mean:

DEFINITION 1.1. A *polyhedron* σ in $M_{\mathbb{R}}$ is a finite intersection of closed half-spaces. A *face* of a polyhedron is a subset given by the intersection of σ with a hyperplane H such that σ is contained in a half-space with boundary H .

The *boundary* $\partial\sigma$ of σ is the union of all proper faces of σ , and the *interior* $\text{Int}(\sigma)$ of σ is $\sigma \setminus \partial\sigma$.

The polyhedron σ is a *lattice polyhedron* if it is an intersection of half-spaces defined over \mathbb{Q} and all vertices of σ lie in M .

A *polytope* is a compact polyhedron.

Returning to $V(f)$, each codimension one polyhedron making up $V(f)$ separates two domains of linearity of f , in one of which f is given by a monomial with exponent $n \in N$ and in the other by a monomial with exponent $n' \in N$. Then the *weight* of this polyhedron in $V(f)$ is the index of $n' - n$. We then view $V(f)$ as a weighted polyhedral complex.

EXAMPLES 1.2. Figures 1 through 5 give examples of two-variable tropical polynomials and their corresponding “zero loci.” All edges have weight 1 unless otherwise indicated. We also indicate the monomial determining the function on each domain of linearity and the precise position of the vertices.

We now explain a simple way to see what $V(f)$ looks like. Given

$$f = \sum_{n \in S} a_n z^n,$$




FIGURE 2. $1 \oplus (0 \odot x_1) \oplus (0 \odot x_2)$




FIGURE 3. $0 \oplus (0 \odot x_1) \oplus (0 \odot x_2) \oplus (1 \odot x_1 \odot x_2) \oplus (5 \odot x_1 \odot x_1) \oplus (5 \odot x_2 \odot x_2)$




FIGURE 4. $0 \oplus (0 \odot x_1) \oplus (0 \odot x_2) \oplus (3 \odot x_1 \odot x_2) \oplus (9 \odot x_1 \odot x_1) \oplus (1 \odot x_2 \odot x_2)$

we consider the *Newton polytope* of S ,

$$\Delta_S := \text{Conv}(S) \subseteq N_{\mathbb{R}},$$



FIGURE 5. $0 \oplus (0 \odot x_1) \oplus (1 \odot x_2) \oplus (0 \odot x_1 \odot x_1 \odot x_2 \odot x_2)$ 

FIGURE 6

the convex hull of S in $N_{\mathbb{R}}$. The coefficients a_n then define a function

$$\varphi : \Delta_S \rightarrow \mathbb{R}$$

as follows. We consider *the upper convex hull* $\tilde{\Delta}_S$ of the set

$$\tilde{S} = \{(n, a_n) \mid n \in S\} \subseteq N_{\mathbb{R}} \times \mathbb{R},$$

namely

$$\tilde{\Delta}_S = \{(n, a) \in N_{\mathbb{R}} \times \mathbb{R} \mid \text{there exists } (n, a') \in \text{Conv}(\tilde{S}) \text{ with } a \geq a'\}.$$

We then define

$$\varphi(n) = \min\{a \in \mathbb{R} \mid (n, a) \in \tilde{\Delta}_S\}.$$

For example, considering the univariate tropical polynomial

$$f = 1 \oplus (0 \odot x) \oplus (0 \odot x^2) \oplus (2 \odot x^3),$$

we get Δ_S and $\tilde{\Delta}_S$ as depicted in Figure 6, with the lower boundary of $\tilde{\Delta}_S$ being the graph of φ .

This picture yields a polyhedral decomposition of Δ_S :

DEFINITION 1.3. A *(lattice) polyhedral decomposition* of a (lattice) polyhedron $\Delta \subseteq N_{\mathbb{R}}$ is a set \mathcal{P} of (lattice) polyhedra in $N_{\mathbb{R}}$ called *cells* such that




FIGURE 7. The left-hand figure is $\tilde{\Delta}_S$. The right-hand picture shows $\check{\mathcal{P}}$ on the x -axis and the graph of $\check{\varphi}$.

- (1) $\Delta = \bigcup_{\sigma \in \mathcal{P}} \sigma$.
- (2) If $\sigma \in \mathcal{P}$ and $\tau \subseteq \sigma$ is a face, then $\tau \in \mathcal{P}$.
- (3) If $\sigma_1, \sigma_2 \in \mathcal{P}$, then $\sigma_1 \cap \sigma_2$ is a face of both σ_1 and σ_2 .

For a polyhedral decomposition \mathcal{P} , denote by \mathcal{P}_{\max} the subset of maximal cells of \mathcal{P} . We denote by $\mathcal{P}^{[k]}$ the set of k -dimensional cells of \mathcal{P} .

Indeed, to get a polyhedral decomposition \mathcal{P} of Δ_S , we just take \mathcal{P} to be the set of images under the projection $N_{\mathbb{R}} \times \mathbb{R} \rightarrow N_{\mathbb{R}}$ of proper faces of $\tilde{\Delta}_S$. A polyhedral decomposition of Δ_S obtained in this way from the graph of a convex piecewise linear function is called a *regular decomposition* and these decompositions play an important role in the combinatorics of convex polyhedra, see e.g., [32].

We can now define the *discrete Legendre transform* of the triple $(\Delta_S, \mathcal{P}, \varphi)$:

DEFINITION 1.4. The *discrete Legendre transform* of $(\Delta_S, \mathcal{P}, \varphi)$ is the triple $(M_{\mathbb{R}}, \check{\mathcal{P}}, \check{\varphi})$ where:

(1)

$$\check{\mathcal{P}} = \{\check{\tau} \mid \tau \in \mathcal{P}\}$$

with

$$\check{\tau} = \left\{ m \in M_{\mathbb{R}} \mid \begin{array}{l} \exists a \in \mathbb{R} \text{ such that } \langle -m, n \rangle + a \leq \varphi(n) \\ \text{for all } n \in \Delta_S, \text{ with equality for } n \in \tau \end{array} \right\}.$$

(2) $\check{\varphi}(m) = \max\{a \mid \langle -m, n \rangle + a \leq \varphi(n) \text{ for all } n \in \Delta_S\}$.

Let us explain this in a bit more detail. First, if $\sigma \in \mathcal{P}_{\max}$, let $m_{\sigma} \in M$ be the slope of $\varphi|_{\sigma}$. Then in fact

$$\check{\sigma} = \{-m_{\sigma}\},$$

as follows from the convexity of φ . Second, the formula in (2) is a fairly standard way of describing the Legendre transformed function $\check{\varphi}$. We think of $\check{\varphi}(m)$ as obtained by taking the graph in $N_{\mathbb{R}} \times \mathbb{R}$ of a linear function on $N_{\mathbb{R}}$ with slope $-m$ and moving it up or down until it becomes a supporting hyperplane for $\tilde{\Delta}_S$. The value of this affine linear function at 0 is then $\check{\varphi}(m)$; see Figure 7. Note that if $m \in \text{Int}(\check{\tau})$, then the graph of $\langle -m, \cdot \rangle + \check{\varphi}(m)$ is then a supporting hyperplane for the face of $\tilde{\Delta}_S$ projecting isomorphically to τ .



FIGURE 8. The Newton polytope and subdivision for Figure 1.

In fact, $\check{\varphi}$ can be described in a more familiar way. Note that

$$\check{\varphi}(m) = \min\{\varphi(n) + \langle m, n \rangle \mid n \in \Delta_S\}.$$

From this, it is clear that $\check{\mathcal{P}}_{\max}$ consists of the maximal domains of linearity of $\check{\varphi}$, with $\check{\varphi}|_{\check{v}}$ having slope v for v a vertex (element of $\mathcal{P}^{[0]}$) of \mathcal{P} . Indeed, the minimum is always achieved at some vertex, and if this vertex is v , then $\check{\varphi}(m) = \varphi(v) + \langle m, v \rangle$. Thus $\check{\varphi}$ is linear on \check{v} with slope v . Furthermore, as necessarily $\varphi(v) + \langle m, v \rangle \leq \varphi(v') + \langle m, v' \rangle$ whenever $m \in \check{v}$, one sees that $\check{\varphi}$ is in fact given by the tropical polynomial

$$\sum_{n \in \mathcal{P}^{[0]}} \varphi(n) z^n.$$

This is not necessarily the original polynomial defining the function f . However, clearly the vertices of $\tilde{\Delta}_S$ are of the form (n, a_n) for $n \in \mathcal{P}^{[0]} \subseteq S$, so $\varphi(n) = a_n$ for $n \in \mathcal{P}^{[0]}$, and the tropical polynomial defining $\check{\varphi}$ is simply missing some of the terms of the original defining polynomial f . These missing terms are precisely ones of the form $a_n z^n$ with (n, a_n) not a vertex of $\tilde{\Delta}_S$. We can see that such terms are irrelevant for calculating f . Indeed, if (n, a_n) is not a vertex of $\tilde{\Delta}_S$ for some $n \in N \cap \Delta_S$, and $f(m) = \langle m, n \rangle + a_n$ for some $m \in M_{\mathbb{R}}$, then

$$\langle m, n \rangle + a_n \leq \langle m, n' \rangle + a_{n'}$$

for all $n' \in S$. But then the hyperplane in $N_{\mathbb{R}} \times \mathbb{R}$ given by

$$\{(n', r) \in N_{\mathbb{R}} \times \mathbb{R} \mid \langle m, n' \rangle + r = \langle m, n \rangle + a_n\}$$

is a supporting hyperplane for $\tilde{\Delta}_S$ which contains (n, a_n) , and hence must also contain a vertex $(n', a_{n'})$ of $\tilde{\Delta}_S$. Then $f(m)$ coincides with $\langle m, n' \rangle + a_{n'}$, and hence the term $a_n z^n$ was irrelevant for calculating f . Thus we see

$$\check{\varphi} = f.$$

Since the domains of linearity of $\check{\varphi}$ are the polyhedra of $\check{\mathcal{P}}_{\max}$, we see that

$$V(f) = \bigcup_{\tau \in \mathcal{P}^{[1]}} \check{\tau}.$$

Since the 0-cells of $\check{\mathcal{P}}$ are the cells $\check{\sigma} = \{-m_{\sigma}\}$ for $\sigma \in \mathcal{P}_{\max}$, it is usually easy to draw $V(f)$ using this description. Additionally, the weights are easily determined: for $\tau \in \mathcal{P}^{[1]}$, the weight of $\check{\tau}$ is just the affine length of τ , i.e., the index of the difference of the endpoints of τ .

To summarize, the function φ determines the dual decomposition $\check{\mathcal{P}}$, whose vertices are given by slopes of φ , and $V(f)$ is the codimension one skeleton of $\check{\mathcal{P}}$.

EXAMPLES 1.5. For the examples of Figures 1 through 5, the Newton polytopes along with their regular decomposition and values of a_n are given in Figures 8 through 12.




FIGURE 9. The Newton polytope and subdivision for Figure 2.




FIGURE 10. The Newton polytope and subdivision for Figure 3.




FIGURE 11. The Newton polytope and subdivision for Figure 4.

This description of $V(f)$ leads to an important condition known as the *balancing condition*. Specifically, for each $\tilde{\omega} \in \check{\mathcal{P}}^{[n-2]}$, a codimension two cell, let $\check{\tau}_1, \dots, \check{\tau}_k \in \check{\mathcal{P}}^{[n-1]}$ be the cells containing it in $V(f)$, with weights w_1, \dots, w_k . Note that ω is a two-dimensional cell of \mathcal{P} and τ_1, \dots, τ_k are the edges of ω . Let $n_1, \dots, n_k \in N$ be primitive tangent vectors to τ_1, \dots, τ_k , pointing in directions consistent with the orientations on τ_1, \dots, τ_k induced by some chosen orientation on ω . The vectors n_1, \dots, n_k are primitive normal vectors to $\check{\tau}_1, \dots, \check{\tau}_k$. Indeed, the endpoints of τ_i




FIGURE 12. The Newton polytope and subdivision for Figure 5.

give the slopes of $\check{\varphi} = f$ on the two domains of linearity of f on either side of $\check{\tau}_i$, so n_i must be constant on $\check{\tau}_i$. Obviously, we have

$$(1.1) \quad \sum_{i=1}^k w_i n_i = 0.$$

We call this the *balancing condition*.

In the case when $\dim M_{\mathbb{R}} = 2$, so that $V(f)$ is a curve, it is useful to rewrite this as follows. Let $V \in \check{\mathcal{P}}^{[0]}$ be a vertex of $V(f)$, contained in edges $E_1, \dots, E_k \in \check{\mathcal{P}}^{[1]}$ of $V(f)$, and let $m_1, \dots, m_k \in M$ be primitive tangent vectors to E_1, \dots, E_k pointing away from V . Suppose E_i has weight w_i . Then (1.1) is equivalent to

$$(1.2) \quad \sum_{i=1}^k w_i m_i = 0.$$

EXAMPLE 1.6. *The tropical Bézout theorem.* Suppose $\dim M_{\mathbb{R}} = 2$, and let e_1, e_2 be a basis for M . Let Δ_d be the polytope which is the convex hull of $0, de_1$, and de_2 . If $f = \sum_{n \in \Delta_d} a_n z^n$, then $V(f)$ is a tropical curve in $M_{\mathbb{R}}$, which we call a *degree d curve in the tropical projective plane*. For example, Figure 2 depicts a degree 1 curve, i.e., a tropical line, and Figures 3 and 4 depict degree 2 curves, i.e., tropical conics. These should be thought of as tropical analogues of ordinary lines and conics in \mathbb{P}^2 . These tropical versions often share surprising properties in common with the usual algebraic versions. We give one example here.

Let $C, D \subseteq M_{\mathbb{R}}$ be two tropical curves in the tropical projective plane of degree d and e respectively. Suppose that C and D intersect at only a finite number of points; this can always be achieved by translating C or D . In fact, we can similarly assume that none of these intersection points are vertices of C or D . We can define a notion of multiplicity of an intersection point of these two curves. Suppose that a point $P \in C \cap D$ is contained in an edge E of C and an edge F of D , of weights $w(E)$ and $w(F)$ respectively. Let m_1 be a primitive tangent vector to E and m_2 be a primitive tangent vector to F . Then we define the *intersection multiplicity of C and D at P* to be the positive integer

$$i_P(C, D) := w(E)w(F)|m_1 \wedge m_2|.$$




FIGURE 13. Two tropical conics meeting at four points.

Here $m_1, m_2 \in M \cong \mathbb{Z}^2$, and $\bigwedge^2 M \cong \mathbb{Z}$, so $|m_1 \wedge m_2|$ makes sense as a positive number no matter which isomorphism is chosen. We then have the tropical Bézout theorem, which states that

$$\sum_{P \in C \cap D} i_P(C, D) = d \cdot e.$$

This is exactly the expected result for ordinary algebraic curves in \mathbb{P}^2 , of course. For a proof, see [96], §4. See Figure 13 for an example.

1.2. Some background on fans

We will collect here a number of standard notions concerning fans. We send the reader to [27] for more details.

DEFINITION 1.7. A *strictly convex rational polyhedral cone* in $M_{\mathbb{R}}$ is a lattice polyhedron in $M_{\mathbb{R}}$ with exactly one vertex, which is $0 \in M_{\mathbb{R}}$.

A *fan* Σ in $M_{\mathbb{R}}$ is a set of strictly convex rational polyhedral cones such that

- (1) If $\sigma \in \Sigma$, and $\tau \subseteq \sigma$ is a face, then $\tau \in \Sigma$.
- (2) If $\sigma_1, \sigma_2 \in \Sigma$, then $\sigma_1 \cap \sigma_2$ is a face of σ_1 and σ_2 .

In other words, a fan Σ is a polyhedral decomposition of a set $|\Sigma| \subseteq M_{\mathbb{R}}$, called the *support* of Σ , with all elements of the polyhedral decomposition being strictly convex rational polyhedral cones.

A fan is *complete* if $|\Sigma| = M_{\mathbb{R}}$.

DEFINITION 1.8. Let Σ be a fan in $M_{\mathbb{R}}$. A PL (piecewise linear) function on Σ is a continuous function $\varphi : |\Sigma| \rightarrow \mathbb{R}$ which is linear when restricted to each cone of Σ .

The function φ is *strictly convex* if

- (1) $|\Sigma|$ is a convex set in $M_{\mathbb{R}}$;
- (2) For $m, m' \in |\Sigma|$, $\varphi(m) + \varphi(m') \geq \varphi(m + m')$, with equality holding if and only if m, m' lie in the same cone of Σ .

The function φ is *integral* if for each $\sigma \in \Sigma_{\max}$ there exists an $n_\sigma \in N$ such that n_σ and φ agree on σ .

The *Newton polyhedron* of a strictly convex PL function $\varphi : |\Sigma| \rightarrow \mathbb{R}$ is

$$\Delta_\varphi := \{n \in N_{\mathbb{R}} \mid \varphi(m) + \langle n, m \rangle \geq 0 \text{ for all } m \in |\Sigma|\}.$$

The Newton polyhedron of a function φ is unbounded if and only if Σ is not a complete fan. If Σ is complete, it is easy to see that

$$\Delta_\varphi = \text{Conv}(\{-n_\sigma \mid \sigma \in \Sigma_{\max}\})$$

where $n_\sigma \in N_{\mathbb{R}}$ is the linear function $\varphi|_\sigma$. Note there is a one-to-one inclusion reversing correspondence between cones in Σ and faces of Δ_φ , with $\sigma \in \Sigma$ corresponding to

$$\{n \in \Delta_\varphi \mid \varphi(m) + \langle n, m \rangle = 0 \text{ for all } m \in \sigma\}.$$

DEFINITION 1.9. If $\Delta \subseteq N_{\mathbb{R}}$ is a polyhedron, $\sigma \subseteq \Delta$ a face, the *normal cone to Δ along σ* is

$$N_\Delta(\sigma) = \{m \in M \mid m|_\sigma = \text{constant}, \langle m, n \rangle \geq \langle m, n' \rangle \text{ for all } n \in \Delta, n' \in \sigma\}.$$

If $\tau \subseteq \sigma$ is a subset, then $T_\tau\sigma$ denotes the *tangent wedge* to σ along τ , defined by

$$T_\tau\sigma = \{r(m - m') \mid m \in \sigma, m' \in \tau, r \geq 0\}.$$

The *normal fan* of Δ is

$$\check{\Sigma}_\Delta := \{N_\Delta(\sigma) \mid \sigma \text{ is a face of } \Delta\}.$$

One checks easily that

$$(1.3) \quad T_\sigma\Delta = (N_\Delta(\sigma))^\vee := \{n \in N \mid \langle n, m \rangle \geq 0 \quad \forall m \in N_\Delta(\sigma)\}.$$

The normal fan $\check{\Sigma}_\Delta$ to Δ carries a PL function $\varphi_\Delta : |\check{\Sigma}_\Delta| \rightarrow \mathbb{R}$ defined by

$$\varphi_\Delta(m) = -\inf\{\langle n, m \rangle \mid n \in \Delta\}.$$

It is easy to see that if $\varphi : |\Sigma| \rightarrow \mathbb{R}$ is strictly convex, then Σ is the normal fan to Δ_φ . This in fact gives a one-to-one correspondence between strictly convex PL functions φ on a fan Σ and polyhedra Δ with normal fan Σ . Note given Δ , $\Delta_{\varphi_\Delta} = \Delta$, and given $\varphi : |\Sigma| \rightarrow \mathbb{R}$, $\varphi_{\Delta_\varphi} = \varphi$.

DEFINITION 1.10. If Σ is a fan in $M_{\mathbb{R}}$, $\tau \in \Sigma$, we define the *quotient fan $\Sigma(\tau)$ of Σ along τ* to be the fan

$$\Sigma(\tau) := \{(\sigma + \mathbb{R}\tau)/\mathbb{R}\tau \mid \sigma \in \Sigma, \tau \subseteq \sigma\}$$

in $M_{\mathbb{R}}/\mathbb{R}\tau$, where $\mathbb{R}\tau$ is the linear space spanned by τ .

If $\varphi : |\Sigma| \rightarrow \mathbb{R}$ is a PL function on Σ , then φ induces a function $\varphi(\tau) : \Sigma(\tau) \rightarrow \mathbb{R}$, well-defined up to linear functions, as follows. Choose $n \in N_{\mathbb{R}}$ such that $\varphi(m) = \langle n, m \rangle$ for $m \in \tau$. Then for any σ containing τ , $\varphi|_\sigma - n$ is zero on τ , hence descends to a linear function on $(\sigma + \mathbb{R}\tau)/\mathbb{R}\tau$. These piece together to give a PL function $\varphi(\tau)$ on $\Sigma(\tau)$, well-defined up to a linear function (determined by the choice of n).

Note that if φ is strictly convex, then so is $\varphi(\tau)$. If $M'_{\mathbb{R}} = M_{\mathbb{R}}/\mathbb{R}\tau$ and $N'_{\mathbb{R}} = \text{Hom}(M'_{\mathbb{R}}, \mathbb{R})$, then $N'_{\mathbb{R}} = (\mathbb{R}\tau)^\perp$. It is then easy to see that $\Delta_{\varphi(\tau)}$ is just the translate by n of the face of Δ_φ corresponding to τ .

1.3. Parameterized tropical curves

We shall now use the discussion of the balancing condition in §1.1 to define tropical curves in a more abstract setting. In theory, similar definitions could be given for tropical varieties of higher dimension, but we will not do so here.

Let $\bar{\Gamma}$ be a connected graph with no bivalent vertices. Such a graph can be viewed in two different ways. First, it can be viewed as a purely combinatorial object, i.e., a set $\bar{\Gamma}^{[0]}$ of vertices and a set $\bar{\Gamma}^{[1]}$ of edges consisting of unordered pairs of elements of $\bar{\Gamma}^{[0]}$, indicating the endpoints of an edge.

We can also view $\bar{\Gamma}$ as the topological realization of the graph, i.e., a topological space which is the union of line segments corresponding to the edges. We shall confuse these two viewpoints at will, hopefully without any confusion.

Let $\bar{\Gamma}_\infty^{[0]}$ be the set of univalent vertices of $\bar{\Gamma}$, and write

$$\Gamma = \bar{\Gamma} \setminus \bar{\Gamma}_\infty^{[0]}.$$

Let $\Gamma^{[0]}, \Gamma^{[1]}$ denote the set of vertices and edges of Γ . Here we are thinking of Γ and $\bar{\Gamma}$ as topological spaces, so Γ now has some non-compact edges. Let $\Gamma_\infty^{[1]}$ be the set of non-compact edges of Γ . A *flag* of Γ is a pair (V, E) with $V \in \Gamma^{[0]}$ and $E \in \Gamma^{[1]}$ with $V \in E$.

In addition, all graphs will be weighted graphs, i.e., $\bar{\Gamma}$ comes along with a weight function

$$w : \bar{\Gamma}^{[1]} \rightarrow \mathbb{N} = \{0, 1, 2, \dots\}.$$

We will often consider *marked graphs*, $(\Gamma, x_1, \dots, x_k)$, where Γ is as above and x_1, \dots, x_k are labels assigned to non-compact edges of weight 0, i.e., we are given an inclusion

$$\begin{aligned} \{x_1, \dots, x_k\} &\hookrightarrow \Gamma_\infty^{[1]} \\ x_i &\mapsto E_{x_i} \end{aligned}$$

with $w(E_{x_i}) = 0$. We will use the convention in this book which is not actually quite standard in the tropical literature that $w(E) \neq 0$ unless $E = E_{x_i}$ for some x_i .

We can now define a marked parameterized tropical curve in $M_{\mathbb{R}}$, where as usual, $M = \mathbb{Z}^n$, $M_{\mathbb{R}} = M \otimes_{\mathbb{Z}} \mathbb{R}$ and $N = \text{Hom}_{\mathbb{Z}}(M, \mathbb{Z})$.

DEFINITION 1.11. A *marked parameterized tropical curve*

$$h : (\Gamma, x_1, \dots, x_k) \rightarrow M_{\mathbb{R}}$$

is a continuous map h satisfying the following two properties:

- (1) If $E \in \Gamma^{[1]}$ and $w(E) = 0$, then $h|_E$ is constant; otherwise, $h|_E$ is a proper embedding of E into a line of rational slope in $M_{\mathbb{R}}$.
- (2) *The balancing condition.* Let $V \in \Gamma^{[0]}$, and let $E_1, \dots, E_\ell \in \Gamma^{[1]}$ be the edges adjacent to V . Let $m_i \in M$ be a primitive tangent vector to $h(E_i)$ pointing away from $h(V)$. Then

$$\sum_{i=1}^{\ell} w(E_i) m_i = 0.$$

If $h : (\Gamma, x_1, \dots, x_n) \rightarrow M_{\mathbb{R}}$ is a marked parameterized tropical curve, we write $h(x_i)$ for $h(E_{x_i})$.

We will call two marked parameterized tropical curves $h : (\Gamma, x_1, \dots, x_k) \rightarrow M_{\mathbb{R}}$ and $h' : (\Gamma', x'_1, \dots, x'_k) \rightarrow M_{\mathbb{R}}$ *equivalent* if there is a homeomorphism $\varphi : \Gamma \rightarrow \Gamma'$ with $\varphi(E_{x_i}) = E_{x'_i}$ and $h = h' \circ \varphi$. We will define a marked tropical curve to be an equivalence class of parameterized marked tropical curves.

The *genus* of h is $b_1(\Gamma)$. □

We wish to talk about the *degree* of a tropical curve, and to do so, we need to fix a fan Σ . In fact, for the moment, we will only make use of the set of one-dimensional cones in Σ , $\Sigma^{[1]}$. Denote by T_{Σ} the free abelian group generated by $\Sigma^{[1]}$. For $\rho \in \Sigma^{[1]}$, denote by $t_{\rho} \in T_{\Sigma}$ the corresponding generator. We have a map

$$\begin{aligned} r : T_{\Sigma} &\rightarrow M \\ t_{\rho} &\mapsto m_{\rho} \end{aligned}$$

where m_{ρ} is the primitive generator of the ray ρ .

DEFINITION 1.12. A marked tropical curve h is *in* X_{Σ} if for each $E \in \Gamma_{\infty}^{[1]}$ which is not a marked edge, $h(E)$ is a translate of some $\rho \in \Sigma^{[1]}$.

If h is a curve in X_{Σ} , the *degree* of h is $\Delta(h) \in T_{\Sigma}$ defined by

$$\Delta(h) = \sum_{\rho \in \Sigma^{[1]}} d_{\rho} t_{\rho}$$

where d_{ρ} is the number of edges $E \in \Gamma_{\infty}^{[1]}$ with $h(E)$ a translate of ρ , counted with weight.

For $\Delta \in T_{\Sigma}$, $\Delta = \sum_{\rho \in \Sigma^{[1]}} d_{\rho} t_{\rho}$, define


$$|\Delta| := \sum_{\rho \in \Sigma^{[1]}} d_{\rho}.$$

The following lemma is a straightforward application of the balancing condition, obtained by summing the balancing conditions over all vertices of Γ :

LEMMA 1.13. $r(\Delta(h)) = 0$.

EXAMPLE 1.14. Let Σ be the fan for \mathbb{P}^2 . This is the complete fan in $M_{\mathbb{R}} = \mathbb{R}^2$ whose one-dimensional rays are ρ_0, ρ_1, ρ_2 generated by $m_0 = (-1, -1)$, $m_1 = (1, 0)$ and $m_2 = (0, 1)$; see Figure 14. The two-dimensional cones are $\sigma_{i, i+1}$, with indices taken modulo 3 and where $\sigma_{i, i+1}$ is generated by m_i and m_{i+1} . We shall see in Chapter 3 that this fan defines \mathbb{P}^2 as a toric variety (Example 3.2). In particular, we shall give meaning to the symbol “ X_{Σ} ”, which is actually a variety, and in the case of this particular Σ , $X_{\Sigma} = \mathbb{P}^2$. Then the examples of Figures 2, 3 and 4 are tropical curves in $X_{\Sigma} = \mathbb{P}^2$. The degree of Figure 2 is $t_{\rho_0} + t_{\rho_1} + t_{\rho_2}$, while the degree of Figures 3 and 4 is $2(t_{\rho_0} + t_{\rho_1} + t_{\rho_2})$. In general, a tropical curve in \mathbb{P}^2 will be, by the above lemma, of degree $d(t_{\rho_0} + t_{\rho_1} + t_{\rho_2})$, in which case we say the curve is *degree d in \mathbb{P}^2* (compare with Example 1.6). So in particular, Figure 2 is a *tropical line*, and all degree one curves in \mathbb{P}^2 are just translates of this example. Figures 3 and 4 are tropical conics.

It is reasonable to ask what the relationship is between this new definition of tropical curve and the earlier notion of a tropical hypersurface in $M_{\mathbb{R}}$ with $\dim M_{\mathbb{R}} = 2$. In particular, one can ask whether or not $h(\Gamma)$ is a tropical hypersurface in $M_{\mathbb{R}}$. Of course, to pose this question, one must first define weights on $h(\Gamma)$, as h is in general not an embedding. Viewing $h(\Gamma)$ as a one-dimensional polyhedral complex,

FIGURE 14. The fan for \mathbb{P}^2 .

we need to assign a weight $w(E)$ to each edge E of $h(\Gamma)$. We define this as follows. Pick a point $m \in E$ which is not a vertex of $h(\Gamma)$ and is not the image of any vertex of Γ , and define

$$w(E) = \sum_{\substack{E' \in \Gamma^{[1]} \\ E' \cap h^{-1}(m) \neq \emptyset}} w(E'),$$

i.e., the weight of E is the sum of weights of edges of Γ whose image under h contains m . It is easy to check that the balancing condition on h implies firstly that this weight is well-defined, i.e., doesn't depend on the choice of m , and secondly that $h(\Gamma)$ satisfies the balancing condition.

PROPOSITION 1.15. *If $h : \Gamma \rightarrow M_{\mathbb{R}}$ is a tropical curve with $\dim M_{\mathbb{R}} = 2$, then there exists a tropical polynomial f such that $h(\Gamma) = V(f)$, as weighted one-dimensional polyhedral complexes.*

PROOF. We define f as follows. $h(\Gamma)$ yields a polyhedral decomposition $\tilde{\mathcal{P}}$ of $M_{\mathbb{R}}$ whose maximal cells are closures of connected components of $M_{\mathbb{R}} \setminus h(\Gamma)$.

Choose some cell $\sigma_0 \in \tilde{\mathcal{P}}_{\max}$ and define $f|_{\sigma_0} \equiv 0$. We then define f inductively. Suppose f is defined on $\sigma \in \tilde{\mathcal{P}}_{\max}$. If $\sigma' \in \tilde{\mathcal{P}}_{\max}$ and $E = \sigma \cap \sigma'$ satisfies $\dim E = 1$, then we can define f on σ' as follows. Extend $f|_{\sigma}$ to an affine linear function $f_{\sigma} : M_{\mathbb{R}} \rightarrow \mathbb{R}$. Let $n_E \in N$ be a primitive normal vector to E which takes a constant value on E and takes larger values on σ than on σ' . Denote by $\langle n_E, E \rangle$ the value n_E takes on E . Then we define f to be $f_{\sigma} + w(E)(n_E - \langle n_E, E \rangle)$ on σ' . It is an immediate consequence of the balancing condition that this is well-defined. Indeed, if we define f on a sequence of polygons with a common vertex V , starting at σ and passing successively through edges E_1, \dots, E_n , then when we return to σ , we have constructed the function $f_{\sigma} + \sum_{i=1}^n w(E_i)(n_{E_i} - \langle n_{E_i}, E_i \rangle) = f_{\sigma}$ by the balancing condition.

Finally, we note that f is convex, i.e., given by a tropical polynomial. In addition, clearly $h(\Gamma) = V(f)$. \square

We are ready to talk about moduli spaces of such curves. For this, we need to talk about the *combinatorial type* of a marked tropical curve $h : (\Gamma, x_1, \dots, x_n) \rightarrow M_{\mathbb{R}}$. This is the data of the labelled graph $(\Gamma, x_1, \dots, x_n)$, the weight function w , along with, for each flag (V, E) of Γ , the primitive tangent vector $m_{(V,E)} \in M$ to $h(E)$ pointing away from $h(V)$. A *combinatorial equivalence class* is the set of all tropical curves of the same combinatorial type. We denote by $[h]$ the combinatorial equivalence class of a curve h .

DEFINITION 1.16. For $g, k \geq 0$, Σ a fan in $M_{\mathbb{R}}$, $\Delta \in T_{\Sigma}$ with $r(\Delta) = 0$, denote by $\mathcal{M}_{g,k}(\Sigma, \Delta)$ the set of tropical curves in X_{Σ} of genus g , degree Δ and with k marked points.

If $[h]$ is a combinatorial equivalence class of curves of genus g with k marked points of degree Δ in X_{Σ} , we denote by $\mathcal{M}_{g,k}^{[h]}(\Sigma, \Delta) \subseteq \mathcal{M}_{g,k}(\Sigma, \Delta)$ the set of all curves of combinatorial equivalence class $[h]$.

PROPOSITION 1.17.

$$\mathcal{M}_{g,k}(\Sigma, \Delta) = \coprod \mathcal{M}_{g,k}^{[h]}(\Sigma, \Delta),$$

where the disjoint union is over all combinatorial equivalence classes of curves of degree Δ and genus g with k marked points. For a given combinatorial equivalence class $[h]$ of a curve $h : (\Gamma, x_1, \dots, x_k) \rightarrow M_{\mathbb{R}}$, $\mathcal{M}_{g,k}^{[h]}(\Sigma, \Delta)$ is the interior of a polyhedron of dimension

$$\geq e + k + (3 - \dim M_{\mathbb{R}})(g - 1) - \text{ov}(\Gamma),$$

where

$$\text{ov}(\Gamma) = \sum_{V \in \Gamma^{[0]}} (\text{Valency}(V) - 3)$$

is the overvalence of Γ and e is the number of non-compact unmarked edges of Γ .

PROOF. First note the topological Euler characteristic $\chi(\bar{\Gamma}) = 1 - g$ satisfies

$$\begin{aligned} \chi(\bar{\Gamma}) &= \#\bar{\Gamma}^{[0]} - \#\bar{\Gamma}^{[1]} \\ &= \#\Gamma^{[0]} - \#(\Gamma^{[1]} \setminus \Gamma_{\infty}^{[1]}). \end{aligned}$$

On the other hand,

$$3(\#\Gamma^{[0]}) + \text{ov}(\Gamma) = \sum_{V \in \Gamma^{[0]}} \text{Valency}(V) = 2(\#\Gamma^{[1]} \setminus \Gamma_{\infty}^{[1]}) + \#\Gamma_{\infty}^{[1]},$$

from which we conclude that

$$\begin{aligned} \# \text{ of compact edges of } \Gamma &= \# \text{ non-compact edges of } \Gamma + 3g - 3 - \text{ov}(\Gamma) \\ &= e + k + 3g - 3 - \text{ov}(\Gamma). \end{aligned}$$

Now to describe all possible tropical curves with the given topological type, we choose a reference vertex $V \in \Gamma^{[0]}$, and we need to choose $h(V) \in M_{\mathbb{R}}$ and affine lengths¹ ℓ_E of each bounded edge $h(E)$. However, these lengths cannot be chosen independently. Indeed, suppose we have a cycle E_1, \dots, E_m of edges in Γ , $\partial E_i =$

¹The affine length of a line segment of rational slope in $M_{\mathbb{R}}$ with endpoints m_1, m_2 is the number $\ell \in \mathbb{R}_{>0}$ such that $m_1 - m_2 = \ell m_{\text{prim}}$ for some primitive $m_{\text{prim}} \in M$.

$\{V_{i-1}, V_i\}$ with $V_m = V_0$. We of course have by definition of $m_{(V_{i-1}, E_i)}$ that $V_i = V_{i-1} + \ell_{E_i} m_{(V_{i-1}, E_i)}$. Thus $V_0 = V_m = V_0 + \sum_{i=1}^m \ell_{E_i} m_{(V_{i-1}, E_i)}$, or

$$\sum_{i=1}^m \ell_{E_i} m_{(V_{i-1}, E_i)} = 0$$

in $M_{\mathbb{R}}$. So, for each cycle, we obtain the above linear equation, which imposes $\dim M_{\mathbb{R}}$ linear conditions on the ℓ_E 's. Thus, given that there exists a tropical curve of the given combinatorial type, the set of all curves of this combinatorial type is

$$M_{\mathbb{R}} \times (\mathbb{R}_{>0}^{e+k+3g-3-\text{ov}(\Gamma)} \cap L)$$

where $L \subseteq \mathbb{R}^{e+k+3g-3-\text{ov}(\Gamma)}$ is a linear subspace of codimension $\leq g \cdot \dim M_{\mathbb{R}}$ and hence the whole cell is of dimension

$$\geq e + k + (3 - \dim M_{\mathbb{R}})(g - 1) - \text{ov}(\Gamma).$$

□

REMARK 1.18. One should view the case where all vertices of Γ are trivalent as a generic situation. However, there are tropical curves of genus $g \geq 1$ in $M_{\mathbb{R}}$ for $\dim M_{\mathbb{R}} \geq 3$ which are not trivalent and cannot be viewed as limits of trivalent curves.

Of course, for $g = 0$, equality always holds for the dimension, but for $g \geq 1$ equality need not hold. A curve of a given combinatorial type is said to be *superabundant* if the moduli space of curves of that type is larger than

$$e + k + (3 - \dim M_{\mathbb{R}})(g - 1) - \text{ov}(\Gamma).$$

Otherwise a curve is called *regular*. Superabundant curves cause a great deal of difficulty for tropical geometry, and we shall handle this by restricting further to plane curves, i.e., $\dim M_{\mathbb{R}} = 2$. Furthermore, as we shall only need the genus zero case for our discussion, we shall often restrict our attention to this case also.

Restricting to the case that $\dim M_{\mathbb{R}} = 2$, we define

DEFINITION 1.19. A marked tropical curve $h : \Gamma \rightarrow M_{\mathbb{R}}$ for $\dim M_{\mathbb{R}} = 2$ is *simple* if

- (1) Γ is trivalent;
- (2) h is injective on the set of vertices and there are no disjoint edges E_1, E_2 with a common vertex V for which $h|_{E_1}$ and $h|_{E_2}$ are non-constant and $h(E_1) \subseteq h(E_2)$;
- (3) Each unbounded unmarked edge of Γ has weight one.

By our discussion above, simple curves move in a family of dimension at least $|\Delta| + k + g - 1$, as now $e = |\Delta|$. However, one can show that simple curves in dimension two are always regular, see [80], Proposition 2.21, so in fact simple curves move in $(|\Delta| + k + g - 1)$ -dimensional families. We know this for $g = 0$ already, but since we shall not be focussing on higher genus curves, we omit a proof of this fact.

LEMMA 1.20. *Fix Σ a fan in $M_{\mathbb{R}}$, $\dim M_{\mathbb{R}} = 2$, and a degree $\Delta \in T_{\Sigma}$. Let $P_1, \dots, P_{|\Delta|-1} \in M_{\mathbb{R}}$ be general points.² Then there are a finite number of marked*

²By general, we mean that $(P_1, \dots, P_{|\Delta|-1}) \in M_{\mathbb{R}}^{|\Delta|-1}$ lies in some dense open subset of $M_{\mathbb{R}}^{|\Delta|-1}$.

genus zero tropical curves $h : (\Gamma, x_1, \dots, x_{|\Delta|-1}) \rightarrow M_{\mathbb{R}}$ in X_{Σ} with $h(x_i) = P_i$ for all i . Furthermore, these curves are simple, and there is at most one such curve of any given combinatorial type.

PROOF. First note there are only a finite number of combinatorial types of curves of degree Δ in X_{Σ} . Indeed, given a curve $h : \Gamma \rightarrow M_{\mathbb{R}}$, we know by Proposition 1.15 that $h(\Gamma)$ is in fact a tropical hypersurface in $M_{\mathbb{R}}$. The degree Δ in fact determines the Newton polytope (up to translation) of a defining equation for $h(\Gamma)$. Furthermore, specifying a regular subdivision of the Newton polytope is equivalent to specifying the combinatorial type of $h(\Gamma)$. For each possible combinatorial type of $h(\Gamma)$, there are only a finite number of ways of parameterizing such a curve. Since there are a finite number of lattice subdivisions of the Newton polytope, this implies there are only a finite number of combinatorial types.

So in fact we can prove the result just by fixing one combinatorial type of curve, $[h]$. This gives the description as in the proof of Proposition 1.17,

$$\mathcal{M}_{0,|\Delta|-1}^{[h]}(\Sigma, \Delta) \cong M_{\mathbb{R}} \times \mathbb{R}_{>0}^{e+|\Delta|-4-\text{ov}(\Gamma)},$$

obtained after choosing a reference vertex $V \in \Gamma^{[0]}$. We have an evaluation map

$$\text{ev} : \mathcal{M}_{0,|\Delta|-1}^{[h]}(\Sigma, \Delta) \rightarrow (M_{\mathbb{R}})^{|\Delta|-1}$$

sending $h : (\Gamma, x_1, \dots, x_{|\Delta|-1}) \rightarrow M_{\mathbb{R}}$ to

$$\text{ev}(h) = (h(x_1), \dots, h(x_{|\Delta|-1})).$$

Note that in fact ev is an affine linear map. Indeed, to compute $h(x_i)$ given h corresponding to a point in $M_{\mathbb{R}} \times \mathbb{R}_{>0}^{e+|\Delta|-4-\text{ov}(\Gamma)}$, let E_1, \dots, E_n be the sequence of edges traversed from the reference vertex V to the vertex adjacent to E_{x_i} , with $\partial E_i = \{V_{i-1}, V_i\}$, $V_0 = V$. Then

$$h(x_i) = h(V) + \sum_{i=1}^n \ell_{E_i} m_{(V_{i-1}, E_i)}$$

where ℓ_{E_i} is the affine length of E_i . This shows that $h(x_i)$ depends affine linearly on $h(V)$ and the length of the edges.

Thus, unless $\dim \mathcal{M}_{0,|\Delta|-1}^{[h]}(\Sigma, \Delta) \geq \dim((M_{\mathbb{R}})^{|\Delta|-1})$, there is no curve of combinatorial type $[h]$ through general $(P_1, \dots, P_{|\Delta|-1}) \in (M_{\mathbb{R}})^{|\Delta|-1}$. This inequality of dimensions only holds if

$$e + |\Delta| - 2 - \text{ov}(\Gamma) \geq 2(|\Delta| - 1),$$

or

$$e - \text{ov}(\Gamma) \geq |\Delta|.$$

Since $e \leq |\Delta|$ and $\text{ov}(\Gamma) \geq 0$, strict inequality never holds and equality only holds if $e = |\Delta|$ and $\text{ov}(\Gamma) = 0$, i.e., all unbounded edges of Γ are weight 1 and Γ is trivalent. If this equality holds, then either the image of ev is codimension ≥ 1 , in which case again there are no curves of combinatorial type $[h]$ through general $(P_1, \dots, P_{|\Delta|-1})$, or else ev is a local isomorphism and then there is at most one curve of combinatorial type $[h]$ passing through any $P_1, \dots, P_{|\Delta|-1} \in M_{\mathbb{R}}$.

Finally, it is easy to see that the general curve in $\mathcal{M}_{0,|\Delta|-1}^{[h]}(\Sigma, \Delta)$ is injective on the set of vertices. Also, if there is a vertex V with attached edges E_1, E_2 with $h|_{E_1}, h|_{E_2}$ non-constant and $h(E_1) \subseteq h(E_2)$, then we are free to move $h(V)$ along

the affine line containing $h(E_i)$, violating the fact that there is only one such curve. So a curve of type $[h]$ passing through general points $P_1, \dots, P_{|\Delta|-1}$ is simple, as desired. \square

This result allows us to count tropical curves passing through a general set of points. However, as Mikhalkin showed, to get a meaningful result these must be counted with a suitable multiplicity, which we now define.

DEFINITION 1.21. Let $h : \Gamma \rightarrow M_{\mathbb{R}}$ be a simple tropical curve, with $\dim M_{\mathbb{R}} = 2$. We define for $V \in \Gamma^{[0]}$ with adjacent edges E_1, E_2 and E_3 ,

$$\begin{aligned} \text{Mult}_V(h) &= w_{\Gamma}(E_1)w_{\Gamma}(E_2)|m_{(V,E_1)} \wedge m_{(V,E_2)}| \\ &= w_{\Gamma}(E_2)w_{\Gamma}(E_3)|m_{(V,E_2)} \wedge m_{(V,E_3)}| \\ &= w_{\Gamma}(E_3)w_{\Gamma}(E_1)|m_{(V,E_3)} \wedge m_{(V,E_1)}| \end{aligned}$$

if none of E_1, E_2, E_3 are marked, and otherwise $\text{Mult}_V(h) = 1$. Here for $m_1, m_2 \in M$, we identify $\bigwedge^2 M$ with \mathbb{Z} so that $|m_1 \wedge m_2|$ makes sense. The equalities follow from the balancing condition.

We then define the (*Mikhalkin*) *multiplicity* of h to be

$$\text{Mult}(h) = \prod_{V \in \Gamma^{[0]}} \text{Mult}_V(h).$$

Finally, for a given fan Σ and degree Δ , we write

$$N_{\Delta, \Sigma}^{0, \text{trop}} = \sum_h \text{Mult}(h)$$

where the sum is over all $h \in \mathcal{M}_{0, |\Delta|-1}(\Sigma, \Delta)$ passing through $|\Delta|-1$ general points in $M_{\mathbb{R}}$.

While the generality of these points guarantees that the sum makes sense, it is not obvious that $N_{\Delta, \Sigma}^{0, \text{trop}}$ doesn't depend on the choice of these points. This will be shown later, twice, once in Chapter 4 and once in Chapter 5.

In fact, the same definition can be made for curves of genus g . Indeed, one can show that, for a choice of $|\Delta|+g-1$ general points in $M_{\mathbb{R}}$, there are a finite number of simple genus g curves passing through these points (see [80], Proposition 2.23). Using the same definition of multiplicity, one obtains numbers $N_{\Delta, \Sigma}^{g, \text{trop}}$.

If $\dim M_{\mathbb{R}} > 2$, there are similar definitions for counting formulas for genus zero curves: see [86] for precise statements. However, because of superabundant families of $g > 0$ curves, there are more serious issues in higher genus.

Mikhalkin's main result in [80] relates the numbers $N_{\Delta, \Sigma}^{g, \text{trop}}$, which are of course purely combinatorial, to counts of holomorphic curves in the toric variety X_{Σ} . Stating this result rather imprecisely here, he shows that the numbers $N_{\Delta, \Sigma}^{g, \text{trop}}$ coincide with the numbers $N_{\Delta, \Sigma}^{g, \text{hol}}$ of holomorphic curves of genus g in X_{Σ} passing through $|\Delta|+g-1$ points in general position. The fact that this count can be computed in this purely combinatorial fashion was the first significant result in tropical geometry. We shall give a proof of this result for genus zero in Chapter 4.

1.4. Affine manifolds with singularities

We shall now discuss possible generalizations of the discussion of tropical curves. The question we want to pose here is: all our curves have lived in $M_{\mathbb{R}}$, a real affine space. Are there interesting choices for more general target manifolds?

One could, for example, study tropical curves inside tropical hypersurfaces, as has been done, say, in [111]. However, this is not the point of view we want to take here. Instead, we want to look at target spaces which “locally look like $M_{\mathbb{R}}$,” in such a way that we can still talk about tropical curves. The main point is that to talk about tropical curves, one needs the structure of $M_{\mathbb{R}}$ as an affine space, but one also needs to know about the integral structure $M \subseteq M_{\mathbb{R}}$.

In what follows, we consider the group $\text{Aff}(M_{\mathbb{R}}) = M_{\mathbb{R}} \rtimes \text{GL}(M_{\mathbb{R}})$ of affine linear automorphisms of $M_{\mathbb{R}}$, given by $m \mapsto Am + b$, where $A \in \text{GL}_n(\mathbb{R})$ and $b \in M_{\mathbb{R}}$, and its subgroups

$$M_{\mathbb{R}} \rtimes \text{GL}(M) \subseteq \text{Aff}(M_{\mathbb{R}})$$

and

$$\text{Aff}(M) := M \rtimes \text{GL}(M) \subseteq \text{Aff}(M_{\mathbb{R}}).$$

DEFINITION 1.22. A *tropical affine manifold* is a real topological manifold B (possibly with boundary) with an atlas of coordinate charts $\psi_i : U_i \rightarrow M_{\mathbb{R}}$ with transition functions $\psi_i \circ \psi_j^{-1} \in M_{\mathbb{R}} \rtimes \text{GL}(M) \subseteq \text{Aff}(M_{\mathbb{R}})$.

An *integral affine manifold* is a tropical manifold with transition functions in $\text{Aff}(M)$.

We will often make use of two local systems on a tropical affine manifold, defining $\Lambda \subseteq \mathcal{T}_B$ to be the family of lattices locally generated by $\partial/\partial y_1, \dots, \partial/\partial y_n$ for y_1, \dots, y_n local affine coordinates. The sheaf $\check{\Lambda} \subseteq \mathcal{T}_B^*$ is the dual local system locally generated by dy_1, \dots, dy_n . The point is these families of lattices are well-defined on tropical manifolds because of the restriction on the transition maps. Note that \mathcal{T}_B carries a natural flat connection, ∇_B , with flat sections being \mathbb{R} -linear combinations of $\partial/\partial y_1, \dots, \partial/\partial y_n$.

It is easy to generalize the notion of parameterized tropical curve with target a tropical affine manifold, as locally the notion of a line segment with rational slope and the balancing condition make sense.

EXAMPLE 1.23. $B = M_{\mathbb{R}}/\Gamma$ for a lattice $\Gamma \subseteq M_{\mathbb{R}}$ gives an example of a compact tropical affine manifold. In [82], such spaces arise naturally as tropical Jacobians of tropical curves. Figure 15 gives an example of a two-dimensional torus containing a genus 2 tropical curve.

Unfortunately, tori are not particularly useful for the applications we have in mind, so we shall generalize the notion of tropical affine manifold as follows.

DEFINITION 1.24. A *tropical affine manifold with singularities* is a topological manifold B along with data

- a subset $\Delta \subseteq B$ which is a locally finite union of codimension ≥ 2 locally closed submanifolds of B ;
- a tropical affine structure on $B_0 := B \setminus \Delta$.

We say B is an *integral affine manifold with singularities* if the affine structure on B_0 is integral. The set Δ is called the *singular locus* or *discriminant locus* of B .

Note that we are assuming that B is still a topological manifold even at the singular points: the singularities lie in the affine structure in the sense that, in general, the affine structure cannot be extended across Δ . We shall see some examples later, but for the moment one can imagine a two-dimensional cone as an




FIGURE 15. A tropical curve of genus two in \mathbb{R}^2/Γ , where Γ is generated by (a, b) and (b, c) . The dotted lines give a fundamental domain, and the three external vertices of the curve are in fact identified, so these vertices represent a single trivalent vertex.

example, obtained by cutting an angular sector out of a piece of paper and then gluing together the two edges of the sector. (However, in fact we will not ultimately allow this particular example.)

A priori, the singularities of the affine structure can be arbitrarily complicated, and there are many reasonable examples which we shall not wish to consider. In addition, it is often convenient to consider restrictions on the nature of the boundary of B . To control the singularities and the boundary, we introduce a refined notion which arises from the following construction.

CONSTRUCTION 1.25. Let B be a topological manifold (possibly with boundary) equipped with a polyhedral decomposition \mathcal{P} , i.e.,

$$B = \bigcup_{\sigma \in \mathcal{P}} \sigma$$

where

- $\sigma \in \mathcal{P}$ is a subset of B equipped with a homeomorphism to a (not necessarily compact) polyhedron in $M_{\mathbb{R}}$ with faces of rational slope and at least one vertex. Thus in particular any $\sigma \in \mathcal{P}$ has a set of faces: these faces are inverse images of faces of the polyhedron in $M_{\mathbb{R}}$.
- If $\sigma \in \mathcal{P}$ and $\tau \subseteq \sigma$ is a face, then $\tau \in \mathcal{P}$.
- If $\sigma_1, \sigma_2 \in \mathcal{P}$, $\sigma_1 \cap \sigma_2 \neq \emptyset$, then $\sigma_1 \cap \sigma_2$ is a face of σ_1 and σ_2 .

For each $\sigma \in \mathcal{P}$, viewing $\sigma \subseteq M_{\mathbb{R}}$ yields a tangent space $\Lambda_{\sigma, \mathbb{R}} \subseteq M_{\mathbb{R}}$ to σ , and we can set

$$\Lambda_{\sigma} := \Lambda_{\sigma, \mathbb{R}} \cap M.$$

The assumption that σ has faces of rational slope implies in particular that if σ is of codimension at least one in $M_{\mathbb{R}}$, then the affine space spanned by σ has rational slope. Thus Λ_{σ} generates $\Lambda_{\sigma, \mathbb{R}}$ as an \mathbb{R} -vector space.

Now the interior of each $\sigma \in \mathcal{P}$ carries a natural affine structure. Indeed, σ is equipped with a homeomorphism with a polyhedron in $M_{\mathbb{R}}$, which is embedded in the affine space it spans in $M_{\mathbb{R}}$. This gives an affine coordinate chart on $\text{Int}(\sigma)$. However, this doesn't define an affine structure on B , but only an affine structure




FIGURE 16. A fan structure given by the map S_τ in a neighbourhood of the one-dimensional cell τ . This can be viewed as describing an affine structure in a direction transverse to τ .

on the subset of B given by

$$\bigcup_{\sigma \in \mathcal{P}_{\max}} \text{Int}(\sigma),$$

where \mathcal{P}_{\max} is the set of maximal cells in \mathcal{P} . This is insufficient for giving a structure of tropical affine manifold with singularities to B , so we need to extend this affine structure. To do so requires the choice of some extra structure, known as a fan structure.

DEFINITION 1.26. Let $\tau \in \mathcal{P}$. The *open star* of τ is

$$U_\tau := \bigcup_{\sigma \in \mathcal{P} \text{ s.t. } \tau \subseteq \sigma} \text{Int}(\sigma).$$

A *fan structure* along $\tau \in \mathcal{P}$ is a continuous map $S_\tau : U_\tau \rightarrow \mathbb{R}^k$ where $k = \dim B - \dim \tau$, satisfying

- (1) $S_\tau^{-1}(0) = \text{Int}(\tau)$.
- (2) If $\tau \subseteq \sigma$, then $S_\tau|_{\text{Int}(\sigma)}$ is an integral affine submersion onto its image, with $\dim S_\tau(\sigma) = \dim \sigma - \dim \tau$. By integral affine submersion we mean the following. We can think of σ as a lattice polytope in $\Lambda_{\sigma, \mathbb{R}}$. Then the map $S_\tau|_\sigma$ is induced by a surjective affine map $\Lambda_\sigma \rightarrow W \cap \mathbb{Z}^k$, for some vector subspace $W \subseteq \mathbb{R}^k$ of codimension equal to the codimension of σ in B .
- (3) For $\tau \subseteq \sigma$, define $K_{\tau, \sigma}$ to be the cone generated by $S_\tau(\sigma \cap U_\tau)$. Then

$$\Sigma_\tau := \{K_{\tau, \sigma} \mid \tau \subseteq \sigma \in \mathcal{P}\}$$

is a fan with $|\Sigma_\tau|$ convex.

Two fan structures S_τ, S'_τ are considered *equivalent* if $S_\tau = \alpha \circ S'_\tau$ for some $\alpha \in \text{GL}_k(\mathbb{Z})$.

If $S_\tau : U_\tau \rightarrow \mathbb{R}^k$ is a fan structure along $\tau \in \mathcal{P}$ and $\sigma \supseteq \tau$, then $U_\sigma \subseteq U_\tau$. We then obtain a fan structure along σ induced by S_τ given by the composition

$$U_\sigma \hookrightarrow U_\tau \xrightarrow{S_\tau} \mathbb{R}^k \rightarrow \mathbb{R}^k / L_\sigma \cong \mathbb{R}^l$$

where $L_\sigma \subseteq \mathbb{R}^k$ is the linear span of $K_{\tau,\sigma}$. This is well-defined up to equivalence. It is easy to see that with the induced fan structure on σ , $\Sigma_\sigma = \Sigma_\tau(K_{\tau,\sigma})$ in the notation of Definition 1.10.

See Figure 16 for a picture of a fan structure. The most important case is when $\tau = v$ a vertex of \mathcal{P} . Then a fan structure is an identification of a neighbourhood of v in B with a neighbourhood of the origin in \mathbb{R}^n ($n = \dim B$). This identification locally near v identifies \mathcal{P} with a fan Σ_v in \mathbb{R}^n .

Given a fan structure S_v at each vertex $v \in \mathcal{P}$, we can construct a tropical structure on B as follows. We first need to choose a discriminant locus $\Delta \subseteq B$. The precise details of this choice of discriminant locus in fact turn out not to be so important, and it can be chosen fairly arbitrarily, subject to certain constraints:

- (1) Δ does not contain any vertex of \mathcal{P} .
- (2) Δ is disjoint from the interior of any maximal cell of \mathcal{P} .
- (3) For any $\rho \in \mathcal{P}$ which is a codimension one cell not contained in ∂B , the connected components of $\rho \setminus \Delta$ are in one-to-one correspondence with vertices of ρ , with each vertex contained in the corresponding connected component.

For example, if $\dim B = 2$, we simply choose one point in the interior of each compact edge of \mathcal{P} not contained in ∂B , and take Δ to be the set of these chosen points. If B is compact without boundary of any dimension, we can take Δ to be the union of all simplices in the first barycentric subdivision of \mathcal{P} which neither contain a vertex of \mathcal{P} nor intersect the interior of a maximal cell of \mathcal{P} . For the general case, see [49], §1.1. We should also note that, for us, this is a maximal choice of discriminant locus, and if there is a subset $\Delta' \subseteq \Delta$ such that the affine structure on $B \setminus \Delta$ extends to an affine structure across $B \setminus \Delta'$, we will replace Δ with Δ' without comment.

For a vertex v of \mathcal{P} , let W_v denote a choice of open neighbourhood of v with $W_v \subseteq U_v$ satisfying the condition that if $v \in \rho$ with ρ a codimension one cell, then $W_v \cap \rho$ is the connected component of $\rho \setminus \Delta$ containing v . Then

$$\{\text{Int}(\sigma) \mid \sigma \in \mathcal{P}_{\max}\} \cup \{W_v \mid v \in \mathcal{P}^{[0]}\}$$

form an open cover of $B_0 := B \setminus \Delta$. We define an affine structure on B_0 via the already given affine structure on $\text{Int}(\sigma)$ for $\sigma \in \mathcal{P}_{\max}$,

$$\psi_\sigma : \text{Int}(\sigma) \hookrightarrow M_{\mathbb{R}}$$

and the composed maps

$$\psi_v : W_v \hookrightarrow U_v \xrightarrow{S_v} \mathbb{R}^{\dim B}$$

where the first map is the inclusion.

It is easy to see that this produces the structure of a tropical affine manifold with singularities on B . Indeed, the crucial point is that the affine charts ψ_v induced by the choice of fan structure are compatible with the charts ψ_σ on the interior of maximal cells of \mathcal{P} , but this follows precisely from item (2) in the definition of a fan structure.

If furthermore all polyhedra in \mathcal{P} are lattice polytopes, then in fact the affine structure is integral.

This construction provides a wide class of examples. However, these examples are still too general. We will impose one additional condition.

We say a collection of fan structures $\{S_v \mid v \in \mathcal{P}^{[0]}\}$ is *compatible* if, for any two vertices v, w of $\tau \in \mathcal{P}$, the fan structures induced on τ by S_v and S_w are equivalent. Note that given such a compatible set of fan structures, we obtain well-defined fan structures along every $\tau \in \mathcal{P}$.³

DEFINITION 1.27. A *tropical manifold* is a pair (B, \mathcal{P}) where B is a tropical affine manifold with singularities obtained from the polyhedral decomposition \mathcal{P} of B and a compatible collection $\{S_v \mid v \in \mathcal{P}^{[0]}\}$ of fan structures.

(B, \mathcal{P}) is an *integral tropical manifold* if in addition all polyhedra in \mathcal{P} are lattice polyhedra.

EXAMPLES 1.28. (1) Any lattice polyhedron σ with at least one vertex supplies an example of an integral tropical manifold, either bounded or unbounded, with $B = \sigma$ and \mathcal{P} the set of faces of σ . In this case the affine structure on $\text{Int}(\sigma)$ extends to give the structure of an affine manifold (with boundary) on σ . Here $\Delta = \emptyset$.

(2) The polyhedral decomposition of $B = M_{\mathbb{R}}$ given in Definition 1.4 is also an example of a tropical manifold (provided the tropical hypersurface in question has at least one vertex).

(3) Let $\Xi \subseteq M_{\mathbb{R}}$ be a reflexive lattice polytope, i.e., $0 \in \text{Int}(\Xi)$ and the polytope

$$\Xi^* := \{n \in N_{\mathbb{R}} \mid \langle n, m \rangle \geq -1 \quad \forall m \in \Xi\}$$

is also a lattice polytope.

Then $B = \partial\Xi$ carries the obvious polyhedral decomposition \mathcal{P} consisting of the proper faces of Ξ . These faces are lattice polytopes. So, to specify an integral tropical manifold structure on B , we need only specify a fan structure at each vertex v of Ξ . This is done via the projection $S_v : U_v \rightarrow M_{\mathbb{R}}/\mathbb{R}v$. Compatibility is easily checked, as the induced fan structure on a cell $\omega \in \mathcal{P}$ containing v is the projection $U_{\omega} \rightarrow M_{\mathbb{R}}/\mathbb{R}\omega$, where $\mathbb{R}\omega$ now denotes the vector subspace of $M_{\mathbb{R}}$ spanned by ω .

There are a number of refinements of this construction. For example, if we take a refinement \mathcal{P}' of \mathcal{P} by integral lattice polytopes, we can use the same prescription above for the fan structure at the vertices.

Note that, in this case, the description of $\Delta' \subseteq B$ of the singular locus as determined by the refinement \mathcal{P}' may give a much bigger discriminant locus, with $\Delta' \cap \text{Int}(\sigma) \neq \emptyset$ for σ a maximal proper face of Ξ . However, the affine structure induced by \mathcal{P}' on $\text{Int}(\sigma)$ is compatible with the obvious affine structure on $\text{Int}(\sigma)$, so in fact the affine structure extends across points of $\Delta' \cap \text{Int}(\sigma)$. Thus we can replace Δ' with

$$\Delta' \cap \bigcup_{\substack{\tau \in \mathcal{P}' \\ \dim \tau = \dim \Xi - 2}} \tau.$$

For example, let

$$\Xi_1 := \text{Conv}\{(-1, -1, -1), (3, -1, -1), (-1, 3, -1), (-1, -1, 3)\}.$$

³In the case we will focus on in this book, $\dim B = 2$, this compatibility condition is in fact trivial, since, provided $v \neq w$, τ is dimension one or two and there are not many choices for zero or one-dimensional fans! So, for the most part, the reader can ignore this condition.




FIGURE 17




FIGURE 18




FIGURE 19

Then choose \mathcal{P}' so that each edge of \mathcal{P} is subdivided into four line segments of affine length 1, as in Figure 17. Then Δ' consists of 24 points, one in the interior of each of these line segments.

We can repeat this in higher dimension, say taking

$$\Xi_2 := \text{Conv}\{(-1, -1, -1, -1), (4, -1, -1, -1), \dots, (-1, -1, -1, 4)\}.$$

Suppose \mathcal{P}' is chosen so that each 2-face of Ξ is triangulated by \mathcal{P}' as in Figure 18. Then Δ' restricted to such a two-face is depicted in Figure 19.

So far, these examples are all integral tropical manifolds. To obtain examples which are not integral, one can deform one of the above examples continuously.




FIGURE 20. (B, \mathcal{P}) is shown on the left, with the embedding in \mathbb{R}^2 shown giving the correct affine structures on σ_1 and σ_2 . The fan structures S_v and $S_{v'}$ depicted then turn (B, \mathcal{P}) into a tropical manifold.

Suppose $\Xi' \subseteq M_{\mathbb{R}}$ is a deformation of Ξ which has the same normal fan. For example, Ξ_1 and Ξ_2 can just be rescaled, but

$$\Xi_3 := \{(x, y, z) \mid -1 \leq x, y, z \leq 1\}$$

can be deformed to

$$\Xi'_3 := \{(x, y, z) \mid -a \leq x \leq a, -b \leq y \leq b, -c \leq z \leq c\}.$$

As before, take the polyhedral decomposition \mathcal{P}' of $B' = \partial\Xi'$ to consist of all proper faces of Ξ' . Then, for each vertex $v' \in \mathcal{P}'$ corresponding to a vertex v of Ξ , define a fan structure $S_{v'} : W_{v'} \rightarrow M_{\mathbb{R}}/\mathbb{R}v$ via

$$S_{v'}(m) = m - v' \pmod{\mathbb{R}v}.$$

One checks easily again that this defines a fan structure, and hence a tropical manifold (B', \mathcal{P}') which is not, in general, integral, as the elements of \mathcal{P}' are not lattice polytopes.

These are all special cases of much more general constructions given in [40, 54, 55]. The Batyrev-Borisov [7] construction, in particular, for complete intersection Calabi-Yau varieties in toric varieties yields vast numbers of examples, as discussed in [40, 55].

(4) In Chapter 6, we will focus largely on the two-dimensional case. In this case, the classification of singularities which may occur in tropical manifolds is particularly straightforward. We now give a simple model for the singularities which can occur.

Consider (B, \mathcal{P}) as depicted in Figure 20, with two maximal cells σ_1, σ_2 and one one-dimensional cell $\tau = \sigma_1 \cap \sigma_2$, and four vertices, including $v = (0, 0)$ and $v' = (0, 1)$, as depicted. The picture gives the affine structures on σ_1 and σ_2 , and to specify the full affine structure we need to specify fan structures at the vertices v and v' , which are shown on the right in Figure 20. We note that we are violating the condition that $|\Sigma_{v'}|$ be convex, if $d > 2$, but this can be rectified by embedding




FIGURE 21

this local example in a larger B . One can think of the induced affine structure as coming from two charts on a cover of $B_0 = B \setminus \Delta$ consisting of the open sets

$$\begin{aligned} U &= B \setminus \{(0, x) \in B \mid x \geq 1/2\} \\ U' &= B \setminus \{(0, x) \in B \mid x \leq 1/2\}. \end{aligned}$$

Figure 21 then shows the embeddings of U and U' into \mathbb{R}^2 . The best way to describe the singularity is to describe the monodromy of the local system Λ around, say, a counterclockwise loop $\gamma : [0, 1] \rightarrow B_0$ about Δ . Suppose $\gamma(0)$ is a point right below Δ . Using the chart on U , we identify the integral tangent vectors at $\gamma(0)$ with $\mathbb{Z}^2 \subseteq \mathbb{R}^2$, the latter being the tangent space of the chart on U . We then move into σ_2 along γ , and to cross back into σ_1 above Δ we need to switch to the chart on U' , which requires applying on σ_2 the linear transformation $\begin{pmatrix} 1 & 0 \\ d & 1 \end{pmatrix}$; this acts the same way on tangent vectors. We can then complete our loop γ through U' , and switching back to U requires no further change of coordinates. Hence the monodromy of Λ around γ is given in the standard basis by $\begin{pmatrix} 1 & 0 \\ d & 1 \end{pmatrix}$.

In fact, this example describes the local structures of all possible singularities which can occur in two-dimensional tropical manifolds.

An important feature of these singularities is that tangent vectors to the edge containing the singularity are left invariant by monodromy.

1.5. The discrete Legendre transform

We now wish to generalize the discrete Legendre transform that we saw in §1.1 to tropical manifolds. As we shall describe in Chapter 6, the motivation for this is that the discrete Legendre transform in fact describes a form of mirror symmetry. This section may be skipped on a first reading.

To define the discrete Legendre transform, we first need to generalize the notion of a convex piecewise linear function. Even if there are no singularities, a tropical manifold B may carry no convex functions, e.g., $B = M_{\mathbb{R}}/\Gamma$ a torus. Instead, we use multi-valued functions.

To begin, we need

DEFINITION 1.29. A function $f : M_{\mathbb{R}} \rightarrow \mathbb{R}$ is *affine linear* if it is of the form $m \mapsto \langle n, m \rangle + b$ for some $n \in N_{\mathbb{R}}, b \in \mathbb{R}$. It is *integral affine linear* if $n \in N, b \in \mathbb{Z}$.

DEFINITION 1.30. If B is an affine manifold, an *affine linear function* $f : B \rightarrow \mathbb{R}$ is a continuous function given on a coordinate chart $\psi_i : U_i \rightarrow M_{\mathbb{R}}$ by the composition of ψ_i with an affine linear function $M_{\mathbb{R}} \rightarrow \mathbb{R}$. Furthermore, if B is integral, we say $f : B \rightarrow \mathbb{R}$ is *integral affine* if it is locally given by $M_{\mathbb{R}} \rightarrow \mathbb{R}$ integral affine linear.

If (B, \mathcal{P}) is an (integral) tropical manifold, an (integral) affine function on an open set $U \subseteq B$ is a continuous map $\varphi : U \rightarrow \mathbb{R}$ that is (integral) affine on $U \setminus \Delta$.

An (integral) PL function on U is a continuous map $\varphi : U \rightarrow \mathbb{R}$ such that if $S_{\tau} : U_{\tau} \rightarrow \mathbb{R}^k$ is the fan structure along $\tau \in \mathcal{P}$, then

$$(1.4) \quad \varphi|_{U \cap U_{\tau}} = \lambda + \varphi_{\tau} \circ S_{\tau},$$

where

$$\varphi_{\tau} : |\Sigma_{\tau}| \rightarrow \mathbb{R}$$

is an (integral) PL function on the fan Σ_{τ} determined by the fan structure S_{τ} along τ , and λ is (integral) affine linear.

Denote by $\mathcal{A}ff(B, \mathbb{R})$ the sheaf of affine linear functions on B (and denote by $\mathcal{A}ff(B, \mathbb{Z})$ the sheaf of integral affine linear functions if B is integral), and similarly by $\mathcal{P}\mathcal{L}_{\mathcal{P}}(B, \mathbb{R})$ (or $\mathcal{P}\mathcal{L}_{\mathcal{P}}(B, \mathbb{Z})$) the sheaf of (integral) PL functions.

A *multi-valued (integral) PL function* is a section of $\mathcal{P}\mathcal{L}_{\mathcal{P}}(B, \mathbb{R})/\mathcal{A}ff(B, \mathbb{R})$ (respectively $\mathcal{P}\mathcal{L}_{\mathcal{P}}(B, \mathbb{Z})/\mathcal{A}ff(B, \mathbb{Z})$). In other words, it is a collection $\{(U_i, \varphi_i)\}$ of PL functions with $\varphi_i - \varphi_j$ affine linear on $U_i \cap U_j$.

We say a multi-valued PL function on U is *convex* if it is locally represented on $U \cap U_{\tau}$ by $\varphi_{\tau} \circ S_{\tau}$ with φ_{τ} a strictly convex function.

We can now generalize the discrete Legendre transform given in §1.1. Given a triple $(B, \mathcal{P}, \varphi)$, where (B, \mathcal{P}) is a tropical manifold and φ is a multi-valued strictly convex PL function on B , we will define the discrete Legendre transform of $(B, \mathcal{P}, \varphi)$, denoted by $(\check{B}, \check{\mathcal{P}}, \check{\varphi})$.

For any $\tau \in \mathcal{P}$, we have a representative $\varphi_{\tau} : |\Sigma_{\tau}| \rightarrow \mathbb{R}$ of φ as in (1.4), and set $\check{\tau} := \Delta_{\varphi_{\tau}}$, the Newton polyhedron of φ_{τ} . If $\tau \subseteq \sigma$ for $\tau, \sigma \in \mathcal{P}$, then Σ_{σ} is the quotient fan of Σ_{τ} by the cone in Σ_{τ} corresponding to σ and, up to a linear map, $\varphi_{\sigma} : |\Sigma_{\sigma}| \rightarrow \mathbb{R}$ is induced by $\varphi_{\tau} : |\Sigma_{\tau}| \rightarrow \mathbb{R}$ as in Definition 1.10. Hence $\check{\sigma}$ can be identified with a face of $\check{\tau}$.

This gives us the pair $(\check{B}, \check{\mathcal{P}})$, where $\check{\mathcal{P}} = \{\check{\sigma} \mid \sigma \in \mathcal{P}\}$, and \check{B} is obtained by identifying $\check{\tau}_1$ and $\check{\tau}_2$ along the common face $\check{\sigma}$ if $\sigma \in \mathcal{P}$ is the smallest cell of \mathcal{P} containing τ_1 and τ_2 .

It is easy to see \check{B} constructed in this way is a topological manifold, with $\check{\mathcal{P}}$ the dual polyhedral complex to \mathcal{P} , and in fact $B \setminus \partial B$ and $\check{B} \setminus \partial \check{B}$ are homeomorphic.

To complete the description of $(\check{B}, \check{\mathcal{P}}, \check{\varphi})$, we need to define the fan structure on $(\check{B}, \check{\mathcal{P}})$ and the function $\check{\varphi}$.

For $\sigma \in \mathcal{P}_{\max}$, $\check{\sigma}$ is a vertex of $\check{\mathcal{P}}$, and let $\check{\Sigma}_{\sigma}$ denote the normal fan (Definition 1.9) to σ . The cones of $\check{\Sigma}_{\sigma}$ are in one-to-one inclusion reversing correspondence with faces of σ , and for $\tau \subseteq \sigma$ the tangent wedge (Definition 1.9) to $\check{\tau}$ at $\check{\sigma}$ is $N_{\sigma}(\tau)$. Thus there is a natural fan structure

$$\check{S}_{\check{\sigma}} : U_{\check{\sigma}} \rightarrow N_{\mathbb{R}}.$$

Indeed, locally $U_{\check{\sigma}} \cap \check{\tau}$ looks like the tangent wedge to $\check{\tau}$ at $\check{\sigma}$, which is identified via $\check{S}_{\check{\sigma}}$ with the normal cone $N_{\sigma}(\tau)$. Finally, we define $\check{\varphi}$ by taking $\check{\varphi}_{\tau}$ on $\Sigma_{\check{\tau}} = \check{\Sigma}_{\tau}$ to be the function induced by τ as in Definition 1.9.

EXAMPLES 1.31. (1) Let $B = \sigma \subseteq M_{\mathbb{R}}$ be a strictly convex rational polyhedral cone, \mathcal{P} the set of faces of σ , and $\varphi \equiv 0$. Then \check{B} is simply the dual cone

$$\sigma^{\vee} := \{n \in N_{\mathbb{R}} \mid \langle n, m \rangle \geq 0 \text{ for all } m \in \sigma\}$$

and $\check{\varphi} \equiv 0$.

(2) Let $B \subseteq M_{\mathbb{R}}$ be a polyhedron, \mathcal{P} its set of faces, $\varphi \equiv 0$. Then $\check{\mathcal{P}}$ is the normal fan to B in $N_{\mathbb{R}}$, \check{B} is the support of the normal fan, and

$$\check{\varphi}(n) := -\inf\{\langle n, m \rangle \mid m \in B\}$$

is the PL function on the normal fan induced by B , as in Definition 1.9.

(3) Let $B \subseteq N_{\mathbb{R}}$ be a (compact) polytope, \mathcal{P} a polyhedral decomposition of B , φ a strictly convex PL function on B . Then $\check{B} = M_{\mathbb{R}}$, and $(\check{B}, \check{\mathcal{P}}, \check{\varphi})$ agrees with the discrete Legendre transform defined in §1.1, up to a change of sign of $\check{\varphi}$.

(4) Let $\Xi \subseteq M_{\mathbb{R}}$ be a reflexive polytope, $B = \partial\Xi$, \mathcal{P} the set of proper faces of B , so the construction of Example 1.28, (3), yields an integral tropical manifold (B, \mathcal{P}) . Define φ as follows. First, let Σ be the fan defined by $\Sigma = \{\mathbb{R}_{\geq 0}\sigma \mid \sigma \in \mathcal{P}\}$. The fact that Ξ is reflexive implies there is a strictly convex integral PL function $\psi : |\Sigma| \rightarrow \mathbb{R}$ which takes the value 1 on each vertex of Ξ . For $\tau \in \mathcal{P}$, choose $n_{\tau} \in N$ such that $n_{\tau} = \psi|_{\tau}$ on τ , and define φ_{τ} on the quotient fan $\Sigma(\mathbb{R}_{\geq 0}\tau)$ to be defined by the function induced by $\psi - n_{\tau}$, as in Definition 1.10. From the way the fan structure was defined on B , $\Sigma(\mathbb{R}_{\geq 0}\tau)$ is precisely the fan Σ_{τ} associated to $\tau \in \mathcal{P}$. Hence the collection of PL functions $\{\varphi_{\tau} \mid \tau \in \mathcal{P}\}$ defines a convex multi-valued integral PL function on B .

As an exercise in the definitions, check that $(\check{B}, \check{\mathcal{P}}, \check{\varphi})$ is the integral tropical manifold defined by the same construction using the dual reflexive polytope Ξ^* .

(5) Consider the surface depicted in Figure 22. This figure depicts a surface B which is the union of three two-dimensional simplices, each isomorphic to the standard two-simplex. Furthermore, there is one point of the discriminant locus along each interior edge. The figure shows an open subset of B obtained by removing the dotted lines along the edges. This open set is embedded in \mathbb{R}^2 as depicted by an affine coordinate chart. In this embedding, the vertices are $v_0 = (0, 0)$, $v_1 = (1, 0)$, $v_2 = (0, 1)$ and $v_3 = (-1, -1)$. This embedding defines a fan structure at the internal vertex, and we give fan structures at the three boundary vertices as depicted. As a result, the boundary is actually a straight line in the affine structure. This gives B and \mathcal{P} , and in addition, we can choose a strictly convex PL function φ . On the open set depicted in the figure, φ will be given by $(0, 0)$ on the upper right-hand triangle, $(-1, 0)$ on the left-hand triangle, and $(0, -1)$ on the lower left-hand triangle. In other words, φ is completely determined by the fact that it is zero at all vertices, except for the vertex v_3 , where φ is 1.

What is the discrete Legendre transform of $(B, \mathcal{P}, \varphi)$? First, the Newton polytope of φ at v_0 is the standard simplex, and the Newton polyhedra of φ at v_1, v_2 and v_3 are all isomorphic to $[0, 1] \times [0, \infty)$. This gives a picture roughly as in Figure 23. That figure is misleading though: all three unbounded edges are in fact parallel! Furthermore, the fan structure at each vertex is the normal fan to the standard simplex, i.e., the fan for \mathbb{P}^2 . This gives us \check{B} and $\check{\mathcal{P}}$. Finally, $\check{\varphi}$ can




FIGURE 22




FIGURE 23

still be taken to be single-valued, taking the value 0 on the compact maximal cell, and on each unbounded cell $[0, 1] \times [0, +\infty)$, $\tilde{\varphi}$ is just the second coordinate in this representation.

One can in fact compute the monodromy of Λ around the three singular points of \tilde{B} , and one finds at each point that it is given by $\begin{pmatrix} 1 & 0 \\ 3 & 1 \end{pmatrix}$ in a suitable basis. As a result, it is actually possible to pull apart each singular point into three singular points, each with monodromy $\begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix}$, as in Figure 24. We will use this deformation as an example in the next section.

1.6. Tropical curves on tropical surfaces

Briefly we show how to define tropical curves on two-dimensional tropical manifolds. The reason for restricting to dimension two is that it is not yet clear what the correct definitions in higher dimensions should be.

So fix a tropical manifold (B, \mathcal{P}) with $\dim B = 2$. Let $\bar{\Gamma}$ be a connected graph with no bivalent vertices. Let $\bar{\Gamma}_\infty^{[0]}$ be a *subset* of the set of univalent vertices (in




FIGURE 24

distinction from §1.3, where $\bar{\Gamma}_\infty^{[0]}$ denoted the set of *all* univalent vertices). Let $\Gamma = \bar{\Gamma} \setminus \bar{\Gamma}_\infty^{[0]}$. In addition, assume $\bar{\Gamma}$ has a weighting $w : \bar{\Gamma}^{[1]} \rightarrow \mathbb{N}$.

We will not bother with marked tropical curves here; one can fill in the details for this if the reader wishes. So we assume all weights are positive.

In what follows, recall that with $B_0 := B \setminus \Delta$, $B_0 \setminus \partial B_0$ is an honest tropical affine manifold, and hence carries a local system Λ . Let $i : B_0 \setminus \partial B_0 \hookrightarrow B$ be the inclusion. We will use below the sheaf $i_*\Lambda$. For U a contractible open set in B_0 , $\Gamma(U, i_*\Lambda) \cong \mathbb{Z}^2$, i.e., $i_*\Lambda$ is locally constant on B_0 . But if U is a small neighbourhood of a point of Δ , and the affine structure can't be extended across this point, then

$$(1.5) \quad \Gamma(U, i_*\Lambda) \cong \mathbb{Z},$$

the monodromy invariant part of the local system on U . Here, recall that the monodromy around Δ takes the form $\begin{pmatrix} 1 & 0 \\ d & 1 \end{pmatrix}$ for some $d \geq 0$ in a suitable basis.

DEFINITION 1.32. A continuous map $h : \Gamma \rightarrow B$ is a *parameterized tropical curve* if it is proper and satisfies the following two conditions:

- (1) For each edge E of Γ , $h|_E$ is an immersion (the image can self-intersect). Furthermore, there is a section $u \in \Gamma(E, h^{-1}(i_*\Lambda))$ which is tangent to every point of $h(E)$.
- (2) For every vertex V of Γ , let $E_1, \dots, E_m \in \Gamma^{[1]}$ be the edges adjacent to V . If $h(V) \in \Delta$, there is no further condition. Otherwise, let u_1, \dots, u_m be integral tangent vectors at $h(V)$, i.e., elements of the stalk $(i_*\Lambda)_{h(V)}$, with u_i primitive, tangent to $h(E_i)$, and pointing away from $h(V)$. Then

$$\sum_{j=1}^m w(E_j)u_j = 0.$$

Let's clarify what these conditions mean. (1) tells us that locally $h(E)$ is a line of rational slope; this is a well-defined notion in a tropical affine manifold. However, if $h(E)$ contains a point of Δ with non-trivial monodromy, the tangent direction to $h(E)$ near this point is completely determined, by (1.5). In other words, there is a




FIGURE 25. The darker lines give a tropical curve.

unique invariant direction at each point of Δ , and $h(E)$ must be tangent to that direction. Note that, in the analysis of two-dimensional singularities given in §1.4, this invariant direction is precisely the direction of the edge of \mathcal{P} passing through the point of Δ .

The second condition tells us that, besides the usual balancing condition, we can have edges terminating at singular points. Even if an edge terminates at a singular point, however, it still must be tangent to the invariant direction at the singular point.

EXAMPLE 1.33. Consider the surface \check{B} given in Example 1.31, (5), or rather, the variant with 9 singular points rather than 3. Figure 25 shows a tropical curve. In fact, there are 27 such tropical curves: three choices of unbounded edge, and 3^2 choices for the endpoints. In Example 6.2, we shall see that this surface corresponds to the cubic surface in \mathbb{P}^3 . Morally, these tropical curves correspond to the 27 lines on the cubic surface.

1.7. References and further reading

The material of §1.1 largely follows the more in depth introductory paper [96]. Standard references for fans are the books on toric varieties by Fulton [27] and Oda [87]. The material on parameterized tropical curves originates in Mikhalkin's work [80]. See also the book of Itenberg, Mikhalkin and Shustin, [62]. The material on affine manifolds with singularities and the discrete Legendre transform forms part of the Gross-Siebert program: see [48], [47], [49] and [41].

The A- and B-models

Our goal in this chapter is to explain the A- and B-models involved in the mirror symmetry story as needed for this book. The topics covered here are rather selective. The A-model is concerned with Gromov-Witten invariants. We shall be extremely brief about the definition of Gromov-Witten invariants, which have been examined thoroughly in other texts. On the other hand, we shall discuss in detail some of the structures produced by Gromov-Witten invariants, such as the notion of semi-infinite variations of Hodge structure. The B-model is concerned with period integrals. For the B-model, we shall ignore completely the Calabi-Yau case, which the reader can find discussed in many sources, including [43]. Instead, we focus on the case we shall be largely interested in here, the case of Fano manifolds. Again, in this case, the relevant structures are semi-infinite variations of Hodge structure.

2.1. The A-model

2.1.1. Stable maps and Gromov-Witten invariants. Let's begin with the basic definitions:

DEFINITION 2.1. A *stable n -pointed curve* is data (C, x_1, \dots, x_n) where C is a (possibly reducible) proper reduced connected algebraic curve over an algebraically closed field with only nodes as singularities, and $x_1, \dots, x_n \in C$ are distinct points which do not coincide with any of the nodes. Furthermore, the automorphism group of (C, x_1, \dots, x_n) must be finite. Here, an automorphism of (C, x_1, \dots, x_n) means an automorphism $\varphi : C \rightarrow C$ satisfying $\varphi(x_i) = x_i$. The *genus* of C is the arithmetic genus of C . An isomorphism of n -pointed curves (C, x_1, \dots, x_n) and (C', x'_1, \dots, x'_n) is an isomorphism $\varphi : C \rightarrow C'$ with $\varphi(x_i) = x'_i$.

We note that (C, x_1, \dots, x_n) having a finite automorphism group can be characterized as follows. Let $\nu : \tilde{C} \rightarrow C$ be the normalization of C , and call a point of \tilde{C} *distinguished* if it either maps via ν to a node or to one of the marked points x_1, \dots, x_n . Then the condition that (C, x_1, \dots, x_n) has a finite automorphism group is equivalent to every component of \tilde{C} of genus zero having at least 3 distinguished points, and every component of \tilde{C} of genus one having at least 1 distinguished point.

A famous result of Deligne and Mumford [19] is that the moduli space $\overline{\mathcal{M}}_{g,n}$ of n -pointed stable curves is an *irreducible proper smooth Deligne-Mumford stack*.

This phrase requires some discussion. We will take the attitude towards stacks exhibited in [56], pages 139–149. In other words, we won't give the definition of a stack, and simply will try not to worry too much about this gap. However, the basic point is as follows. In an ideal world, $\overline{\mathcal{M}}_{g,n}$ would be a fine moduli space, i.e., would be a scheme representing the functor from the category of schemes to

the category of sets given by

$$S \mapsto \left\{ \begin{array}{l} \text{isomorphism classes of flat families } \mathcal{C} \rightarrow S \text{ with sections} \\ \sigma_1, \dots, \sigma_n : S \rightarrow \mathcal{C} \text{ such that } (\mathcal{C}_{\bar{s}}, \sigma_1(\bar{s}), \dots, \sigma_n(\bar{s})) \text{ is a stable} \\ n\text{-pointed genus } g \text{ curve for every geometric point } \bar{s} \text{ of } S \end{array} \right\}.$$

Thus $\text{Hom}(S, \overline{\mathcal{M}}_{g,n})$ would coincide with the set of isomorphism classes of such flat families over S . In particular, there would be a universal family of stable n -pointed curves over $\overline{\mathcal{M}}_{g,n}$, corresponding to the identity in $\text{Hom}(\overline{\mathcal{M}}_{g,n}, \overline{\mathcal{M}}_{g,n})$. However, it is well-known that no such universal family exists due to the existence of automorphisms of curves: the automorphism group of a stable curve is finite, but need not be trivial. As a result, one needs to enlarge the category of schemes so that the moduli functor becomes representable. This is done using the category of algebraic stacks. A smooth Deligne-Mumford stack can be viewed as an algebro-geometric object which is an orbifold, i.e., a scheme which is locally the quotient of a smooth scheme by a finite group, but we remember this local description. We can still talk about such things as Chow groups (or cohomology groups if we are working over \mathbb{C} with the usual topology) and intersection products, but these are at best defined over \mathbb{Q} .

In particular, it is worth emphasizing the point that intersection numbers are only defined over \mathbb{Q} . Indeed, if locally we can describe a smooth Deligne-Mumford stack as X/G , with X smooth, then informally, if $Z_1, Z_2 \subseteq X/G$ are two subschemes of complementary dimension meeting in a zero-dimensional subscheme, we define the intersection number of Z_1 and Z_2 at $P \in Z_1 \cap Z_2$ as $\frac{1}{|G|} \sum_{Q \in \pi^{-1}(P)} i_Q$, where $\pi : X \rightarrow X/G$ is the projection and i_Q denotes the intersection multiplicity of $\pi^{-1}(Z_1)$ and $\pi^{-1}(Z_2)$ at Q .

For example, consider $X = \mathbb{P}^2$, $G = \mathbb{Z}_2$, with \mathbb{Z}_2 acting by

$$(x, y, z) \mapsto (-x, -y, z).$$

We can view X/G as a smooth Deligne-Mumford stack. Let D_1 and D_2 be the images of the lines L_1 and L_2 given by $x = 0$ and $y = 0$ in X/G . If $\pi : X \rightarrow X/G$ denotes the projection, we then have

$$D_1 \cdot D_2 = \frac{1}{2} L_1 \cdot L_2 = \frac{1}{2}.$$

With this in mind, given a smooth Deligne-Mumford stack X , we obtain $A_{\mathbb{Q}}^i(X)$, the codimension i Chow group with rational coefficients, along with intersection products

$$A_{\mathbb{Q}}^i(X) \times A_{\mathbb{Q}}^j(X) \rightarrow A_{\mathbb{Q}}^{i+j}(X),$$

giving a ring structure on $A_{\mathbb{Q}}^*(X)$.

Similarly, if we work over \mathbb{C} , in the usual topology, the rational cohomology ring $H^*(X, \mathbb{Q})$ makes sense, along with cup products

$$H^i(X, \mathbb{Q}) \times H^j(X, \mathbb{Q}) \rightarrow H^{i+j}(X, \mathbb{Q}).$$

We shall usually work in this latter context.

EXAMPLES 2.2. $\overline{\mathcal{M}}_{0,n}$ is empty for $n \leq 2$, and for $n = 3$ is just a point. In fact, $\overline{\mathcal{M}}_{0,n}$ is always a scheme, because n -pointed rational curves with $n \geq 3$ have no automorphisms.

$\overline{\mathcal{M}}_{0,4}$ is isomorphic to \mathbb{P}^1 , which we see as follows. Given $(\mathbb{P}^1, x_1, x_2, x_3, x_4)$, there is a unique element of $\text{PGL}(2)$ taking x_1 to 0, x_2 to 1, and x_3 to ∞ . Then




FIGURE 1

$x_4 \in \mathbb{P}^1 \setminus \{0, 1, \infty\}$ determines the isomorphism class of $(\mathbb{P}^1, x_1, x_2, x_3, x_4)$. (This is just the classical cross-ratio). Thus $\mathcal{M}_{0,4}$, the open subset of $\overline{\mathcal{M}}_{0,4}$ corresponding to non-singular curves, is isomorphic to $\mathbb{P}^1 \setminus \{0, 1, \infty\}$. The points $0, 1, \infty$ correspond to stable but singular curves depicted in Figure 1.

Note that the universal family over $\overline{\mathcal{M}}_{0,4} = \mathbb{P}^1$ is obtained by taking the family $\mathbb{P}^1 \times \mathbb{P}^1 \rightarrow \mathbb{P}^1$ given by projection onto the second component, and taking the sections $\{0\} \times \mathbb{P}^1$, $\{1\} \times \mathbb{P}^1$ and $\{\infty\} \times \mathbb{P}^1$ and the diagonal section as $\sigma_1, \dots, \sigma_4$ respectively. We then blow up the three points where σ_4 intersects one of the other three sections.

Traditionally, one views this as saying that as x_4 approaches one of the other three marked points, we have to *bubble off* a copy of \mathbb{P}^1 to maintain the distinctness of the marked points.

$\overline{\mathcal{M}}_{1,1}$ can be identified with the j -line, but not as a scheme. Every pointed elliptic curve has an automorphism given by negation, and elliptic curves with j -invariant 0 and 1 have larger automorphism groups. Thus even points other than $j = 0$ and $j = 1$ have to be viewed as stacky points, locally of the form U/\mathbb{Z}_2 , where \mathbb{Z}_2 acts trivially on U !

In general, $\dim \overline{\mathcal{M}}_{g,n} = 3g - 3 + n$.

Next, we consider stable maps.

DEFINITION 2.3. Let X be a variety. A stable n -pointed map to X is a map

$$f : (C, x_1, \dots, x_n) \rightarrow X$$

such that C is a proper connected reduced nodal algebraic curve, $x_1, \dots, x_n \in C$ are distinct points not coinciding with any nodes of C , such that f has a finite automorphism group. Here an automorphism of f is an automorphism φ of (C, x_1, \dots, x_n) such that $f \circ \varphi = f$.

Note that the finiteness of the automorphism group of f is equivalent to the statement that every component of the normalization of C on which f is constant has at least three distinguished points if it is genus zero and at least one distinguished point if it is genus one.

If $\beta \in H_2(X, \mathbb{Z})$, we say f represents β if $f_*([C]) = \beta$, where $[C] \in H_2(C, \mathbb{Z})$ is the fundamental class. (If we are not working over the field \mathbb{C} , we can replace H_2 with the dimension one Chow group).

Two stable maps $f : (C, x_1, \dots, x_n) \rightarrow X$, $f' : (C', x'_1, \dots, x'_n) \rightarrow X$ are viewed as being isomorphic if there is an isomorphism $\varphi : (C, x_1, \dots, x_n) \rightarrow (C', x'_1, \dots, x'_n)$ with $\varphi(x_i) = x'_i$, and $f' \circ \varphi = f$.

As before, we can ask whether the moduli functor

$$S \mapsto \left\{ \begin{array}{l} \text{isomorphism classes of flat families } \mathcal{C} \rightarrow S \text{ with sections} \\ \sigma_1, \dots, \sigma_n : S \rightarrow \mathcal{C} \text{ and a morphism } f : \mathcal{C} \rightarrow X \\ \text{such that } f : (\mathcal{C}_{\bar{s}}, \sigma_1(\bar{s}), \dots, \sigma_n(\bar{s})) \rightarrow X \text{ is a stable map} \\ \text{of genus } g \text{ representing } \beta \text{ for every geometric point } \bar{s} \text{ of } S \end{array} \right\}$$

is representable. It is not in general representable by a scheme, nor even by a *smooth* Deligne-Mumford stack. However, it is a proper Deligne-Mumford stack [9]. This will make life more difficult for us. We write the stack of n -pointed stable maps of genus g representing a class $\beta \in H_2(X, \mathbb{Z})$ as $\overline{\mathcal{M}}_{g,n}(X, \beta)$.

Let us also remark here that properness can be tested using the valuative criterion, which boils down to the following basic fact, known as *stable reduction*. See [28], §4.2 for a proof.

PROPOSITION 2.4. *Let S be a non-singular curve over \mathbb{k} , with $s \in S$ a point. Let $U = S \setminus \{s\}$. Suppose we have a family of stable maps $f : (C_U, \sigma_1, \dots, \sigma_n) \rightarrow X$ over U . Then there is an open neighbourhood V of s , a finite map $\pi : V' \rightarrow V$ with V' a non-singular curve and a point $s' \in V'$ such that*

- (1) *if $U' = V' \setminus \{s'\}$, $\pi|_{U'}$ is étale;*
- (2) *Let $C_{U'} := C_U \times_U U'$. This gives a pull-back family of stable maps $f' : (C_{U'}, \sigma'_1, \dots, \sigma'_n) \rightarrow X$ given by the composition $C_{U'} \rightarrow C_U \xrightarrow{f} X$. Then f' extends to a family of stable maps $(C_{V'}, \sigma'_1, \dots, \sigma'_n) \rightarrow X$ over V' .*

EXAMPLES 2.5. Consider the target space $X = \mathbb{P}^2$, $[\ell] \in H_2(X, \mathbb{Z})$ the homology class of a line. Then $\overline{\mathcal{M}}_{0,0}(X, [\ell])$ is of course $(\mathbb{P}^2)^*$, the dual of \mathbb{P}^2 . On the other hand, $\overline{\mathcal{M}}_{0,1}(X, [\ell])$ is the incidence correspondence $I \subseteq \mathbb{P}^2 \times (\mathbb{P}^2)^*$, with

$$I = \{(x, \ell) \mid x \in \ell\}.$$

$\overline{\mathcal{M}}_{0,0}(X, 2[\ell])$ is a bit more complicated. There are four types of stable maps in this moduli space. First, the domain C of f may be irreducible, with $f(C)$ a conic, or C may be a union of two lines, with $f(C)$ a reducible conic. Next, f could be a double cover of a line $\ell \subseteq \mathbb{P}^2$, with the domain either irreducible or reducible. In these two double cover cases, f has a non-trivial automorphism of order 2, and so these points are stacky points in $\overline{\mathcal{M}}_{0,0}(\mathbb{P}^2, 2[\ell])$.

Next consider $\overline{\mathcal{M}}_{1,0}(\mathbb{P}^2, [\ell])$. There are no maps from an elliptic curve to \mathbb{P}^2 representing the class $[\ell]$, so at first glance this moduli space would appear to be empty. However, we have maps f with domain $C = C_1 \cup C_2$ where C_1 is a \mathbb{P}^1 and C_2 is either an elliptic curve or a nodal rational curve attached to C_1 at one point. Furthermore f embeds C_1 as a line and is constant on C_2 . Since this requires both the choice of the image of C_1 in \mathbb{P}^2 and the choice of the points on C_1 and C_2 where C_1 and C_2 meet, we see that

$$\overline{\mathcal{M}}_{1,0}(\mathbb{P}^2, [\ell]) = \overline{\mathcal{M}}_{0,1}(\mathbb{P}^2, [\ell]) \times \overline{\mathcal{M}}_{1,1}.$$

□

We will not go into much detail here, but a standard consequence of deformation theory tells us that these moduli spaces have an *expected dimension*. This expected

dimension arises because given a stable map f with domain an n -pointed curve of genus g , one can compute two spaces, T_f^1 and T_f^2 , where T_f^1 is the tangent space to $[f] \in \overline{\mathcal{M}}_{g,n}(X, \beta)$ and T_f^2 is an obstruction space, which roughly means that locally $\overline{\mathcal{M}}_{g,n}(X, \beta)$ is cut out by $\dim T_f^2$ equations in T_f^1 . This means that the dimension of a small open neighbourhood of $[f]$ in $\overline{\mathcal{M}}_{g,n}(X, \beta)$ lies between

$$\dim T_f^1 - \dim T_f^2 \text{ and } \dim T_f^1.$$

The number $\dim T_f^1 - \dim T_f^2$ is the *expected*, or *virtual*, dimension of $\overline{\mathcal{M}}_{g,n}(X, \beta)$, and can be computed using Riemann-Roch, giving the formula

$$\dim T_f^1 - \dim T_f^2 = n + (\dim_{\mathbb{C}} X - 3)(1 - g) + \int_{\beta} c_1(\mathcal{T}_X).$$

For example, the virtual dimension of $\overline{\mathcal{M}}_{0,0}(\mathbb{P}^2, d[\ell])$ is $3d-1$, as $c_1(\mathcal{T}_{\mathbb{P}^2}) = 3[\ell]$. This is indeed the correct dimension, and in this case $\overline{\mathcal{M}}_{0,0}(\mathbb{P}^2, d[\ell])$ is a smooth Deligne-Mumford stack. On the other hand, the expected dimension of $\overline{\mathcal{M}}_{1,0}(\mathbb{P}^2, [\ell])$ is 3, while we saw that in fact $\dim \overline{\mathcal{M}}_{1,0}(\mathbb{P}^2, [\ell]) = 4$. So in fact, these moduli spaces need not be the expected dimension. More generally, they need not be smooth.

This presents a serious problem for what we want to do. For example, if the goal is to count the number of curves of genus g representing the class β in a Calabi-Yau threefold, (so $c_1(\mathcal{T}_X) = 0$), the expected dimension is zero, and so if the moduli space $\overline{\mathcal{M}}_{g,0}(X, \beta)$ is non-singular of dimension zero, we can just count the number of points in this moduli space. However, we can't expect this to happen all the time.

Historically, there are two approaches to solving this problem. One is to deform the complex structure on the tangent space X to a generic almost complex structure, for which these moduli spaces tend to be better behaved. This approach was pioneered by Ruan [100] and Ruan-Tian [101]. However, this approach takes us into the realm of symplectic geometry. The alternative approach involves the construction of a *virtual fundamental class*, pioneered by Behrend-Fantechi [8] and Li-Tian [74]. Since our goal is to summarize results here, we only give the outcome of the construction, namely a so-called *virtual fundamental class*

$$[\overline{\mathcal{M}}_{g,n}(X, \beta)]^{\text{vir}} \in A_d(\overline{\mathcal{M}}_{g,n}(X, \beta)) \otimes \mathbb{Q}$$

(or in $H_{2d}(\overline{\mathcal{M}}_{g,n}(X), \mathbb{Q})$), where d is the virtual dimension of $\overline{\mathcal{M}}_{g,n}(X, \beta)$.

We can use this class to define *Gromov-Witten invariants*. There are natural maps

$$\text{ev}_i : \overline{\mathcal{M}}_{g,n}(X, \beta) \rightarrow X$$

with, for $f : (C, x_1, \dots, x_n) \rightarrow X$ a stable map, $\text{ev}_i([f]) = f(x_i)$. Putting these together gives

$$\text{ev} = \text{ev}_1 \times \dots \times \text{ev}_n : \overline{\mathcal{M}}_{g,n}(X, \beta) \rightarrow X^n.$$

For classes $\alpha_1, \dots, \alpha_n \in H^*(X, \mathbb{Q})$, $\beta \in H_2(X, \mathbb{Z})$, we define the *Gromov-Witten invariant*

$$\langle \alpha_1, \dots, \alpha_n \rangle_{g,\beta} = \int_{[\overline{\mathcal{M}}_{g,n}(X, \beta)]^{\text{vir}}} \text{ev}^*(\alpha_1 \times \dots \times \alpha_n) \in \mathbb{Q}.$$

Here by the integral we mean we are evaluating a cohomology class on a homology class. This can also be performed with Chow groups if preferred. Note that if

$\alpha_i \in H^{d_i}(X, \mathbb{Q})$ (we write this by saying $\deg \alpha_i = d_i$), then this integral is zero unless twice the virtual dimension of $\overline{\mathcal{M}}_{g,n}(X, \beta)$ is equal to $\sum_{i=1}^n d_i$, i.e.,

$$(2.1) \quad 2(n + (\dim_{\mathbb{C}} X - 3)(1 - g) + \int_{\beta} c_1(\mathcal{T}_X)) = \sum_{i=1}^n d_i.$$

Intuitively, Gromov-Witten invariants can be thought of as follows. Fixing submanifolds Y_1, \dots, Y_n of X whose Poincaré dual cohomology classes are $\alpha_1, \dots, \alpha_n$, this Gromov-Witten invariant should be thought of as the number of curves of homology class β and genus g in X which pass through Y_1, \dots, Y_n . This is subject to the interpretation given above, i.e., by a curve we mean a stable map, and as regards number, we need to use the virtual fundamental class to get an actual number. Finally, in general, these numbers are only rational, and need not be positive.

There is one variant of Gromov-Witten invariants which we shall need in the sequel, namely the so-called *gravitational descendent invariants*. On $\overline{\mathcal{M}}_{g,n}(X, \beta)$, there are natural line bundles \mathcal{L}_i , $i = 1, \dots, n$, whose fibre at a point $[(C, x_1, \dots, x_n)]$ is the cotangent line $\mathfrak{m}_{x_i}/\mathfrak{m}_{x_i}^2$, where $\mathfrak{m}_{x_i} \subseteq \mathcal{O}_{C, x_i}$ is the maximal ideal. Let

$$\psi_i = c_1(\mathcal{L}_i) \in H^2(\overline{\mathcal{M}}_{g,n}(X, \beta), \mathbb{Q}).$$

Then for classes $\alpha_1, \dots, \alpha_n \in H^*(X, \mathbb{Q})$, $\beta \in H_2(X, \mathbb{Z})$, descendent Gromov-Witten invariants are defined by

$$\langle \psi^{p_1} \alpha_1, \dots, \psi^{p_n} \alpha_n \rangle_{g, \beta} = \int_{[\overline{\mathcal{M}}_{g,n}(X, \beta)]^{\text{vir}}} \psi_1^{p_1} \cup \dots \cup \psi_n^{p_n} \cup \text{ev}^*(\alpha_1 \times \dots \times \alpha_n) \in \mathbb{Q}.$$

Note that if $\deg \alpha_i = d_i$, this is zero unless twice the virtual dimension of $\overline{\mathcal{M}}_{g,n}(X, \beta)$ is equal to $\sum_{i=1}^n d_i + \sum_{i=1}^n 2p_i$, i.e.,

$$(2.2) \quad 2(n + (\dim_{\mathbb{C}} X - 3)(1 - g) + \int_{\beta} c_1(\mathcal{T}_X)) = \sum_{i=1}^n d_i + \sum_{i=1}^n 2p_i.$$

EXAMPLE 2.6. The example we shall focus on in this book is $X = \mathbb{P}^2$ and $g = 0$, and in this case $\overline{\mathcal{M}}_{0,n}(\mathbb{P}^2, d[\ell])$ is a smooth Deligne-Mumford stack for $d \geq 0$ (and $n \geq 3$ if $d = 0$) of the expected dimension. One can in fact show that if we take $\alpha_1, \dots, \alpha_n \in H^4(X, \mathbb{Q})$ to be the Poincaré dual class $[\text{pt}]$ of a point, then $\langle \alpha_1, \dots, \alpha_n \rangle_{0, d[\ell]}$ in fact coincides with the number of rational curves of degree d passing through general points $P_1, \dots, P_n \in \mathbb{P}^2$ when this number is expected to be finite, which is precisely the case, using (2.1), if $n - 1 + 3d = 2n$, i.e., $n = 3d - 1$. So, for example,

$$\langle [\text{pt}], [\text{pt}] \rangle_{0, [\ell]} = 1,$$

since there is one line through two points, and

$$\langle [\text{pt}]^5 \rangle_{0, 2[\ell]} := \langle [\text{pt}], [\text{pt}], [\text{pt}], [\text{pt}], [\text{pt}] \rangle_{0, 2[\ell]} = 1,$$

as there is again one conic passing through 5 points in \mathbb{P}^2 . Next

$$\langle [\text{pt}]^8 \rangle_{0, 3[\ell]} = 12,$$

using the well-known fact that 8 general points in \mathbb{P}^2 are contained in a pencil of cubic curves, and a general pencil of plane cubics contains 12 singular, hence rational, cubics.

In this case, we always obtain integers; there are no stacky phenomena here. Essentially, the point is that, in this case, only curves involving multiple covers contribute to stacky phenomena, and there are no such curves passing through a general set of points. However, this is not the case with descendent invariants; for example, it can be shown that

$$\langle \psi^4[\text{pt}] \rangle_{0,2[\ell]} = \frac{1}{8}.$$

(See Example 5.17.)

We will list here the most important properties of Gromov-Witten invariants and descendent Gromov-Witten invariants here, without proof. See for example [28] and [18] for more details.

For Gromov-Witten invariants, we have:

The Fundamental Class Axiom. If $n + 2g \geq 4$ or $\beta \neq 0$ and $n \geq 1$, and $[X] \in H^0(X, \mathbb{Q})$ is the fundamental class of X , then

$$\langle \alpha_1, \dots, \alpha_{n-1}, [X] \rangle_{g,\beta} = 0.$$

The Divisor Axiom. If $n + 2g \geq 4$ or $\beta \neq 0$ and $n \geq 1$, and $\alpha_n \in H^2(X, \mathbb{Q})$, then

$$\langle \alpha_1, \dots, \alpha_n \rangle_{g,\beta} = \left(\int_{\beta} \alpha_n \right) \langle \alpha_1, \dots, \alpha_{n-1} \rangle_{g,\beta}.$$

The Point Mapping Axiom. For $g = 0$, $\beta = 0$,

$$\langle \alpha_1, \dots, \alpha_n \rangle_{0,0} = \begin{cases} \int_X \alpha_1 \cup \alpha_2 \cup \alpha_3 & \text{if } n = 3, \\ 0 & \text{otherwise.} \end{cases}$$

Descendent Gromov-Witten invariants satisfy generalizations of these axioms, and some new ones. In particular, we have:

The Fundamental Class Axiom. If $n + 2g \geq 4$ or $\beta \neq 0$ and $n \geq 1$, then

$$\langle \psi^{p_1} \alpha_1, \dots, \psi^{p_{n-1}} \alpha_{n-1}, [X] \rangle_{g,\beta} = \sum_{i=1}^{n-1} \langle \psi^{p_1} \alpha_1, \dots, \psi^{p_i-1} \alpha_i, \dots, \psi^{p_{n-1}} \alpha_{n-1} \rangle_{g,\beta},$$

where the invariant on the right is taken to be zero if ψ appears with a negative power.

The Divisor Axiom. If $n + 2g \geq 4$ or $\beta \neq 0$ and $n \geq 1$, and $\alpha_n \in H^2(X, \mathbb{Q})$, then

$$\begin{aligned} & \langle \psi^{p_1} \alpha_1, \dots, \psi^{p_{n-1}} \alpha_{n-1}, \alpha_n \rangle_{g,\beta} \\ &= \left(\int_{\beta} \alpha_n \right) \langle \psi^{p_1} \alpha_1, \dots, \psi^{p_{n-1}} \alpha_{n-1} \rangle_{g,\beta} \\ & \quad + \sum_{i=1}^{n-1} \langle \psi^{p_1} \alpha_1, \dots, \psi^{p_i-1} (\alpha_n \cup \alpha_i), \dots, \psi^{p_{n-1}} \alpha_{n-1} \rangle_{g,\beta}, \end{aligned}$$

with the same caveat if a power of ψ is negative.

The Point Mapping Axiom. If $n \leq 3$, then

$$\langle \psi^{\nu_1} \alpha_1, \dots, \psi^{\nu_n} \alpha_n \rangle_{0,0} = 0$$

unless $n = 3$ and $\nu_1 = \dots = \nu_n = 0$.

The Dilaton Axiom.

$$\langle \psi[X], \psi^{p_1} \alpha_1, \dots, \psi^{p_n} \alpha_n \rangle_{g,\beta} = (2g - 2 + n) \langle \psi^{p_1} \alpha_1, \dots, \psi^{p_n} \alpha_n \rangle_{g,\beta}.$$

2.1.2. Quantum cohomology. One significant insight which came out of string theory (see [112], [110]) is that Gromov-Witten invariants can be used to perturb the usual cup product in cohomology. The fact that this new product, called the quantum product, remains associative then gives strong restrictions on Gromov-Witten invariants, in particular giving a complete determination of Gromov-Witten invariants for \mathbb{P}^2 .

We first define the *Gromov-Witten potential* of X as follows. Choose a basis T_0, \dots, T_m of $H^*(X, \mathbb{C})$. We will always use the convention that $T_0 \in H^0(X, \mathbb{C})$ is the fundamental class of X and T_1, \dots, T_p generate $H^2(X, \mathbb{C})$. Let $\gamma = \sum_{i=0}^m y_i T_i$. Then we define the Gromov-Witten potential as

$$\Phi = \sum_{n=0}^{\infty} \sum_{\beta \in H_2(X, \mathbb{Z})} \frac{1}{n!} \langle \gamma^n \rangle_{0, \beta}.$$

Here $\langle \gamma^n \rangle_{0, \beta} = \langle \gamma, \dots, \gamma \rangle_{0, \beta}$, with γ taken n times. So for a given n and β , we obtain a term which is a sum of monomials in y_0, \dots, y_m of degree n . If, for a given n , there are only a finite number of β such that $\langle \gamma^n \rangle_{0, \beta}$ is non-zero, then $\Phi \in \mathbb{C}[[y_0, \dots, y_m]]$, the ring of formal power series in y_0, \dots, y_m . This holds, for example, in the Fano case, i.e. $c_1(\mathcal{T}_X)$ is ample, as follows from (2.1). In the non-Fano case, it is sometimes necessary instead to work over a Novikov ring, but we shall not do so here. As a result, in this case we can consider Φ as a function on a formal neighbourhood of $0 \in H^*(X, \mathbb{C})$.

In the case that X has non-trivial cohomology in odd degree, there is some subtlety in the definition of Φ , since $T_i \cup T_j = (-1)^{\deg T_i + \deg T_j} T_j \cup T_i$. This essentially forces us to view $H^*(X, \mathbb{C})$ as a supermanifold. Assume that each T_i has some definite degree, i.e., $T_i \in H^{d_i}(X, \mathbb{C})$, with $\deg T_i = d_i$. Then from the definition of Gromov-Witten invariants, we have

$$\langle \dots, T_i, T_j, \dots \rangle_{g, \beta} = (-1)^{\deg T_i + \deg T_j} \langle \dots, T_j, T_i, \dots \rangle_{g, \beta}.$$

If, with $\gamma = \sum_{i=0}^m y_i T_i$, we take the coordinates y_i to be *supercommuting*, i.e.,

$$y_i y_j = (-1)^{\deg T_i + \deg T_j} y_j y_i, \quad y_i T_j = (-1)^{\deg T_i + \deg T_j} T_j y_i,$$

then $y_i T_i$ and $y_j T_j$ commute for all i, j . So Φ should be viewed as a function of supercommuting variables, and this commutation rule has to be applied uniformly. Since our main interest here is $X = \mathbb{P}^2$ where there is no odd cohomology, we will often restrict to the case of X having only even cohomology to avoid excessive worries about signs.

Using Φ , we define the (*big*) *quantum cohomology* of X as the ring¹

$$H^*(X, \mathbb{C}[[y_0, \dots, y_m]]),$$

with product given on generators by

$$T_i * T_j = \sum_k (\partial_{y_i} \partial_{y_j} \partial_{y_k} \Phi) T^k,$$

where T^0, \dots, T^m is the Poincaré dual basis to T_0, \dots, T_m . It is often useful to define g^{ij} to be the inverse matrix to

$$g_{ij} = \int_X T_i \cup T_j,$$

¹In the next section, we will redefine quantum cohomology as a family of products on the tangent bundle of a certain manifold.

so that $T^i = \sum g^{ij} T_j$, and

$$T_i * T_j = \sum_{a,k} (\partial_{y_i} \partial_{y_j} \partial_{y_a} \Phi) g^{ak} T_k.$$

It is not obvious that this product is associative; that it is associative is a fundamental observation and can be proved from additional properties of Gromov-Witten invariants. By writing out the equality

$$(T_i * T_j) * T_k = T_i * (T_j * T_k)$$

according to the definition above, one obtains the WDVV equation

$$\begin{aligned} \sum_{a,b} (\partial_{y_i} \partial_{y_j} \partial_{y_a} \Phi) g^{ab} (\partial_{y_b} \partial_{y_k} \partial_{y_\ell} \Phi) \\ = (-1)^{\deg T_i (\deg T_j + \deg T_k)} \sum_{a,b} (\partial_{y_j} \partial_{y_k} \partial_{y_a} \Phi) g^{ab} (\partial_{y_b} \partial_{y_i} \partial_{y_\ell} \Phi), \end{aligned}$$

a system of partial differential equations.

We shall not prove associativity here as this is covered in many basic references on quantum cohomology, see e.g., [28]. Instead, following [28], let's investigate a standard consequence of the WDVV equation, especially with regards to Gromov-Witten invariants for \mathbb{P}^2 .

First consider $\frac{1}{n!} \langle \gamma^n \rangle_{0,\beta}$ for $\beta = 0$. By the Point Mapping Axiom, this is zero except for $n = 3$, and then

$$\frac{1}{3!} \langle \gamma^3 \rangle_{0,0} = \frac{1}{3!} \int_X \gamma \cup \gamma \cup \gamma,$$

which is purely classical information. So split Φ as

$$\Phi = \Phi_{\text{classical}} + \Phi_{\text{quantum}}$$

with

$$\Phi_{\text{quantum}} = \sum_{n \geq 0} \sum_{\substack{\beta \in H_2(X, \mathbb{Z}) \\ \beta \neq 0}} \frac{1}{n!} \langle \gamma^n \rangle_{0,\beta}$$

and $\Phi_{\text{classical}}$ satisfies

$$\partial_{y_i} \partial_{y_j} \partial_{y_k} \Phi_{\text{classical}} = \int_X T_i \cup T_j \cup T_k,$$

and hence gives a contribution to $T_i * T_j$ of

$$\sum_k \left(\int_X T_i \cup T_j \cup T_k \right) T^k = T_i \cup T_j,$$

i.e., the classical cup product. Recalling that we are taking T_0 to be the fundamental class of X and T_1, \dots, T_p generators of $H^2(X, \mathbb{C})$, we can write, for $\beta \neq 0$,

$$\frac{1}{n!} \langle \gamma^n \rangle_{0,\beta} = \sum_{\substack{n_0, \dots, n_m \geq 0 \\ \sum n_j = n}} \epsilon(n_0, \dots, n_m) \langle T_0^{n_0}, \dots, T_m^{n_m} \rangle_{0,\beta} \frac{y_0^{n_0} \cdots y_m^{n_m}}{n_0! \cdots n_m!},$$

where the sign $\epsilon(n_0, \dots, n_m)$ is determined by the supercommutation relation

$$(y_0 T_0)^{n_0} \cdots (y_m T_m)^{n_m} = \epsilon(n_0, \dots, n_m) T_0^{n_0} \cdots T_m^{n_m} y_0^{n_0} \cdots y_m^{n_m}.$$

The Fundamental Class Axiom tells us that there is a non-zero contribution in the above sum only if $n_0 = 0$, and then the Divisor Axiom yields

$$\begin{aligned}
& \sum_{n \geq 0} \frac{1}{n!} \langle \gamma^n \rangle_{0, \beta} \\
= & \sum_{n_1, \dots, n_m \geq 0} \epsilon(0, n_1, \dots, n_m) \left(\prod_{i=1}^p \frac{1}{n_i!} \left(y_i \int_{\beta} T_i \right)^{n_i} \right) \langle T_{p+1}^{n_{p+1}}, \dots, T_m^{n_m} \rangle_{0, \beta} \\
(2.3) \quad & \cdot \frac{y_{p+1}^{n_{p+1}} \cdots y_m^{n_m}}{n_{p+1}! \cdots n_m!} \\
= & \sum_{n_{p+1}, \dots, n_m \geq 0} \epsilon(0, \dots, 0, n_{p+1}, \dots, n_m) \langle T_{p+1}^{n_{p+1}}, \dots, T_m^{n_m} \rangle_{0, \beta} \\
& \cdot \left(\prod_{i=1}^p e^{y_i \int_{\beta} T_i} \right) \frac{y_{p+1}^{n_{p+1}} \cdots y_m^{n_m}}{n_{p+1}! \cdots n_m!}.
\end{aligned}$$

Let us write this down for \mathbb{P}^2 , where T_0, T_1, T_2 are the cohomology classes of \mathbb{P}^2 , a line, and a point respectively. We have

$$\Phi_{\text{quantum}}(y_0, y_1, y_2) = \sum_{d > 0} \langle T_2^{3d-1} \rangle_{0, d[\ell]} e^{dy_1} \frac{y_2^{3d-1}}{(3d-1)!}.$$

Letting $\Gamma_{ijk} = \partial_{y_i} \partial_{y_j} \partial_{y_k} \Phi_{\text{quantum}}$, we obtain the following description of the quantum product:

$$\begin{aligned}
T_0 * T_i &= T_i \\
T_1 * T_1 &= T_2 + \Gamma_{111} T_1 + \Gamma_{112} T_0 \\
T_1 * T_2 &= \Gamma_{121} T_1 + \Gamma_{122} T_0 \\
T_2 * T_2 &= \Gamma_{221} T_1 + \Gamma_{222} T_0
\end{aligned}$$

with

$$\begin{aligned}
\Gamma_{111} &= \sum_{d \geq 1} d^3 \langle T_2^{3d-1} \rangle_{0, d[\ell]} e^{dy_1} \frac{y_2^{3d-1}}{(3d-1)!} \\
\Gamma_{112} &= \sum_{d \geq 1} d^2 \langle T_2^{3d-1} \rangle_{0, d[\ell]} e^{dy_1} \frac{y_2^{3d-2}}{(3d-2)!} \\
\Gamma_{122} &= \sum_{d \geq 1} d \langle T_2^{3d-1} \rangle_{0, d[\ell]} e^{dy_1} \frac{y_2^{3d-3}}{(3d-3)!} \\
\Gamma_{222} &= \sum_{d \geq 2} \langle T_2^{3d-1} \rangle_{0, d[\ell]} e^{dy_1} \frac{y_2^{3d-4}}{(3d-4)!}.
\end{aligned}$$

Keeping in mind that $\Phi_{\text{classical}} = \frac{y_0^2 y_2}{2} + \frac{y_0 y_2^2}{2}$, the WDVV equation with $i = j = 1$ and $k = l = 2$ gives us $\Gamma_{222} + \Gamma_{111} \Gamma_{122} = \Gamma_{112}^2$, or

$$\Gamma_{222} = \Gamma_{112}^2 - \Gamma_{111} \Gamma_{122}.$$

Comparing the coefficient of y_2^{3d-4} on the left- and right-hand sides of this last equation gives us the Vafa-Intriligator formula (first appearing in [60])

$$\langle T_2^{3d-1} \rangle_{0,d[\ell]} = \sum_{\substack{d_1, d_2 > 0 \\ d_1 + d_2 = d}} \langle T_2^{3d_1-1} \rangle_{0,d_1[\ell]} \langle T_2^{3d_2-1} \rangle_{0,d_2[\ell]} \cdot \left(d_1^2 d_2^2 \binom{3d-4}{3d_1-2} - d_1^3 d_2 \binom{3d-4}{3d_1-1} \right).$$

Since $\langle T_2^2 \rangle_{0,[\ell]} = 1$, all other genus zero Gromov-Witten invariants now follow from this formula recursively. So the WDVV equation is extremely powerful in this case.

2.1.3. Frobenius manifolds.

The quantum cohomology ring

$$H^*(X, \mathbb{C}[[y_0, \dots, y_m]])$$

can be viewed as a family of rings, parameterized by the formal completion of $0 \in H^*(X, \mathbb{C})$ (a formal supermanifold if X has odd cohomology). Noting that $H^*(X, \mathbb{C})$ is the tangent space to $H^*(X, \mathbb{C})$ at any point, we can view the formal completion as a formal manifold \mathcal{M} coming along with a ring structure on each fibre of $\mathcal{T}_{\mathcal{M}}$. These ring structures aren't arbitrary, but rather have properties turning \mathcal{M} into what is known as a *Frobenius manifold*.

To save rather tedious discussions concerning sign conventions, we will assume from now on that X has no odd cohomology, so that we only need to define the notion of a Frobenius manifold on an ordinary manifold, not a supermanifold.

In what follows, \mathcal{M} will be a complex manifold, a germ of a complex manifold, or a non-singular scheme. Later on, it will be a formal completion of a complex manifold along a complex submanifold.

DEFINITION 2.7. Let \mathcal{M} be as above. A *pre-Frobenius structure* on \mathcal{M} is a triple of data (∇, g, \mathcal{A}) where

- (1) $\nabla : \mathcal{T}_{\mathcal{M}} \rightarrow \mathcal{T}_{\mathcal{M}} \otimes \Omega_{\mathcal{M}}^1$ is a flat connection.
- (2) g is a metric on \mathcal{M} , i.e., a symmetric pairing $g : S^2(\mathcal{T}_{\mathcal{M}}) \rightarrow \mathcal{O}_{\mathcal{M}}$ which induces an isomorphism $\mathcal{T}_{\mathcal{M}} \cong \mathcal{T}_{\mathcal{M}}^*$. Furthermore, g must be compatible with ∇ , i.e.,

$$d(g(X, Y)) = g(\nabla X, Y) + g(X, \nabla Y).$$

- (3) $\mathcal{A} : S^3(\mathcal{T}_{\mathcal{M}}) \rightarrow \mathcal{O}_{\mathcal{M}}$ is a symmetric tensor.

Given this data, this defines a product on each tangent space of \mathcal{M} , by defining $X \circ Y$ via the formula

$$\mathcal{A}(X, Y, Z) = g(X \circ Y, Z).$$

By symmetry of \mathcal{A} , we also have

$$(2.4) \quad g(X \circ Y, Z) = g(X, Y \circ Z).$$

A pre-Frobenius manifold is a *Frobenius manifold* if this data satisfies two additional properties.

- (4) The product defined by \mathcal{A} is associative.
- (5) Locally on \mathcal{M} , there is a potential function Φ such that

$$\mathcal{A}(X, Y, Z) = XYZ\Phi.$$

EXAMPLE 2.8. For X a non-singular variety, with cohomology only in even degree and $H^*(X, \mathbb{C})$ generated by T_0, \dots, T_m , $\mathcal{M} = \text{Spec } \mathbb{C}[[y_0, \dots, y_m]]$ furnishes an example of a Frobenius manifold, assuming the Gromov-Witten potential Φ of X lies in $\mathbb{C}[[y_0, \dots, y_m]]$ (e.g., if X is Fano). Take the metric to be constant on \mathcal{M} , defined by

$$g(\partial_{y_i}, \partial_{y_j}) = \int_X T_i \cup T_j.$$

The connection ∇ is the trivial one, with $\partial_{y_0}, \dots, \partial_{y_m}$ flat sections. Finally,

$$\mathcal{A}(\partial_{y_i}, \partial_{y_j}, \partial_{y_k}) = \partial_{y_i} \partial_{y_j} \partial_{y_k} \Phi.$$

The product is then the quantum cohomology product, i.e.,

$$Y \circ Z = Y * Z.$$

We can in fact consider quantum cohomology as giving a Frobenius manifold structure on a much bigger manifold. Assume that we have chosen the generators T_1, \dots, T_p of $H^2(X, \mathbb{C})$ to in fact lie in $H^2(X, \mathbb{R})$ and to have the property that $\int_{\beta} T_i \geq 0$ for $1 \leq i \leq p$ and any class $\beta \in H_2(X, \mathbb{Z})$ represented by a stable curve. For example, if $H^{1,1}(X, \mathbb{C}) = H^2(X, \mathbb{C})$, we just need to choose T_1, \dots, T_p to lie in the closure of the Kähler cone of X . We then introduce new variables $\kappa_1, \dots, \kappa_p$, with the relation

$$e^{y_i} = \kappa_i, \quad 1 \leq i \leq p.$$

Now assume that for any n_{p+1}, \dots, n_m , there are a finite number of β with

$$\langle T_{p+1}^{n_{p+1}}, \dots, T_m^{n_m} \rangle_{0, \beta} \neq 0$$

(this holds e.g., if X is Fano). Noting that $\int_{\beta} T_1, \dots, \int_{\beta} T_p$ are all non-negative and determine β up to a finite number of choices, we see by (2.3) that

$$\Phi \in \mathbb{C}[\kappa_1, \dots, \kappa_p][[y_0, y_{p+1}, \dots, y_m]].$$

This ring is the ring of formal power series in the variables y_0, y_{p+1}, \dots, y_m with coefficients in the polynomial ring $\mathbb{C}[\kappa_1, \dots, \kappa_p]$.

Now set

$$(2.5) \quad \begin{aligned} \overline{\mathcal{M}} &= \text{Spec } \mathbb{C}[\kappa_1, \dots, \kappa_p][[y_0, y_{p+1}, \dots, y_m]] \\ \mathcal{M} &= \text{Spec } \mathbb{C}[\kappa_1^{\pm 1}, \dots, \kappa_p^{\pm 1}][[y_0, y_{p+1}, \dots, y_m]]. \end{aligned}$$

The \mathcal{M} used previously, $\text{Spec } \mathbb{C}[[y_0, \dots, y_m]]$, can be thought of as a formal neighbourhood of the point in the new \mathcal{M} with $\kappa_1, \dots, \kappa_p = 1$, $y_0 = y_{p+1} = \dots = y_m = 0$. The potential Φ is now a function on $\overline{\mathcal{M}}$. Note that $\partial_{y_i} = \kappa_i \partial_{\kappa_i}$, $1 \leq i \leq p$, so that the connection we previously defined extends to the flat connection on \mathcal{M} whose flat sections are $\overline{\partial}_{y_0}, \overline{\partial}_{y_{p+1}}, \dots, \overline{\partial}_{y_m}$ and $\kappa_1 \partial_{\kappa_1}, \dots, \kappa_p \partial_{\kappa_p}$. This connection does not extend across $\overline{\mathcal{M}} \setminus \mathcal{M}$. Similarly, the metric g and symmetric form \mathcal{A} extend to \mathcal{M} , always using $\partial_{y_i} = \kappa_i \partial_{\kappa_i}$, $1 \leq i \leq p$.

In general, when we talk about quantum cohomology for the remainder of this chapter, we are really referring to a Frobenius manifold structure either on the above \mathcal{M} or on closely related manifolds.

This example has some additional structure.

DEFINITION 2.9. If \mathcal{M} is a pre-Frobenius manifold, then

- (1) A vector field e on \mathcal{M} is called the *identity* if $e \circ Y = Y$ for all Y .

(2) A vector field E on \mathcal{M} is an *Euler vector field* if for all vector fields Y and Z ,

$$(2.6) \quad E(g(Y, Z)) - g([E, Y], Z) - g(Y, [E, Z]) = Dg(Y, Z)$$

for some constant D and

$$(2.7) \quad [E, Y \circ Z] - [E, Y] \circ Z - Y \circ [E, Z] = d_0 Y \circ Z$$

for some constant d_0 .

An Euler vector field can be used to produce a grading on vector fields: given a vector field Y , it is *homogeneous of degree d* if

$$[E, Y] = dY.$$

EXAMPLE 2.10. Continuing with Example 2.8, we make the same assumptions on the generators T_0, \dots, T_p , so that we have a Frobenius manifold structure on $\mathcal{M} = \text{Spec } \mathbb{C}[\kappa_1^{\pm 1}, \dots, \kappa_p^{\pm 1}][[y_0, y_{p+1}, \dots, y_m]]$. Then ∂_{y_0} is a flat identity: T_0 is the identity in quantum cohomology, as follows from the Fundamental Class Axiom for Gromov-Witten invariants.

There is also an Euler vector field, defined by

$$E = \sum_{i=0}^m \left(1 - \frac{\deg T_i}{2}\right) y_i \partial_{y_i} + \sum_{i=1}^p c_i \kappa_i \partial_{\kappa_i},$$

where $c_1(\mathcal{T}_X) = \sum_{i=1}^p c_i T_i$. Recall $\partial_{y_i} = \kappa_i \partial_{\kappa_i}$ for $1 \leq i \leq p$.

Note that with respect to this vector field E , a vector field Y of the form

$$Y = y_0^{a_0} \kappa_1^{a_1} \cdots \kappa_p^{a_p} y_{p+1}^{a_{p+1}} \cdots y_m^{a_m} \partial_{y_i}$$

is homogeneous, with the degree given by taking y_i to have degree

$$1 - (\deg T_i)/2$$

if $i \notin \{1, \dots, p\}$, κ_i having degree c_i if $i \in \{1, \dots, p\}$, and ∂_{y_i} having degree

$$(\deg T_i)/2 - 1.$$

This comes from the simple calculation that

$$[E, Y] = \left(- (1 - (\deg T_i)/2) + \sum_{j=1}^p c_j a_j + \sum_{\substack{j=0 \\ j \notin \{1, \dots, p\}}}^m (1 - (\deg T_j)/2) a_j \right) Y.$$

To check that E is an Euler vector field, one notes that both the left- and right-hand sides of (2.6) and (2.7) are tensors, i.e., are $\mathcal{O}_{\mathcal{M}}$ -linear, so we only need to check these for $Y = \partial_{y_i}$, $Z = \partial_{y_j}$. Checking (2.6) first, we see that

$$\begin{aligned} & E(g(\partial_{y_i}, \partial_{y_j})) - g([E, \partial_{y_i}], \partial_{y_j}) - g(\partial_{y_i}, [E, \partial_{y_j}]) \\ &= g\left(\left(1 - \frac{\deg T_i}{2}\right) \partial_{y_i}, \partial_{y_j}\right) + g\left(\partial_{y_i}, \left(1 - \frac{\deg T_j}{2}\right) \partial_{y_j}\right) \\ &= \left(2 - \frac{\deg T_i + \deg T_j}{2}\right) g(\partial_{y_i}, \partial_{y_j}) \\ &= (2 - \dim_{\mathbb{C}} X) g(\partial_{y_i}, \partial_{y_j}), \end{aligned}$$

as $g(\partial_{y_i}, \partial_{y_j}) \neq 0$ only if $\deg T_i + \deg T_j = 2 \dim_{\mathbb{C}} X$. Thus $D = 2 - \dim_{\mathbb{C}} X$.

Next we claim that (2.7) holds with $d_0 = 1$. Again taking $Y = \partial_{y_i}$, $Z = \partial_{y_j}$, we note that Y and Z are homogeneous of degrees $(\deg T_i)/2 - 1$ and $(\deg T_j)/2 - 1$ respectively. Thus for (2.7) to hold, $Y \circ Z$ must be homogeneous of degree $1 + ((\deg T_i)/2 - 1) + ((\deg T_j)/2 - 1)$. To see this, first note that Φ itself is homogeneous. Indeed, by (2.3) a term in Φ is of the form

$$\langle T_{p+1}^{n_{p+1}}, \dots, T_m^{n_m} \rangle_{0,\beta} \left(\prod_{i=1}^p \kappa_i^{\int_{\beta} T_i} \right) \frac{y_{p+1}^{n_{p+1}} \cdots y_m^{n_m}}{n_{p+1}! \cdots n_m!}$$

and is then of degree

$$\begin{aligned} & \sum c_i \int_{\beta} T_i + \sum_{j=p+1}^n \left(1 - \frac{\deg T_j}{2} \right) n_j \\ &= \int_{\beta} c_1(\mathcal{T}_X) + \sum_{j=p+1}^n n_j - \sum_{j=p+1}^n \left(\frac{\deg T_j}{2} \right) n_j. \end{aligned}$$

On the other hand, by (2.1), the Gromov-Witten invariant in this term is zero unless

$$\int_{\beta} c_1(\mathcal{T}_X) + \sum_{j=p+1}^n n_j - \sum_{j=p+1}^n \left(\frac{\deg T_j}{2} \right) n_j = 3 - \dim_{\mathbb{C}} X.$$

Hence Φ can be viewed as homogeneous of degree $3 - \dim_{\mathbb{C}} X$, and the degree of $\partial_{y_i} \partial_{y_j} \partial_{y_k} \Phi$ is

$$3 - \dim_{\mathbb{C}} X + ((\deg T_i)/2 - 1) + ((\deg T_j)/2 - 1) + ((\deg T_k)/2 - 1).$$

Putting this all together,

$$\partial_{y_i} \circ \partial_{y_j} = \sum_{k,l} (\partial_{y_i} \partial_{y_j} \partial_{y_k} \Phi) g^{kl} \partial_{y_l}$$

is of degree

$$\begin{aligned} & 3 - \dim_{\mathbb{C}} X + ((\deg T_i)/2 - 1) + ((\deg T_j)/2 - 1) \\ & + ((\deg T_k)/2 - 1) + ((\deg T_l)/2 - 1) \\ &= 3 - \dim_{\mathbb{C}} X - ((\deg T_i)/2 - 1) + ((\deg T_j)/2 - 1) \\ & \quad - 2 + \dim_{\mathbb{C}} X \\ &= 1 + ((\deg T_i)/2 - 1) + ((\deg T_j)/2 - 1) \end{aligned}$$

since $g^{kl} \neq 0$ only when $\deg T_k + \deg T_l = 2 \dim_{\mathbb{C}} X$. This shows (2.7). \square

If \mathcal{M} is a pre-Frobenius manifold with a vector field E and $d_0 \neq 0$, we can define a connection $\widehat{\nabla}$ on the vector bundle $p_1^* \mathcal{T}_{\mathcal{M}}$ on $\mathcal{M} \times \mathbb{C}^{\times}$, where $p_1 : \mathcal{M} \times \mathbb{C}^{\times} \rightarrow \mathcal{M}$ is the first projection, and $\mathbb{C}^{\times} = \mathbb{C} \setminus \{0\}$. Let \hbar be the coordinate on \mathbb{C}^{\times} . This connection is defined by

$$\begin{aligned} \widehat{\nabla}_X(Y) &= \nabla_X(Y) + \hbar^{-1} X \circ Y \\ d_0 \widehat{\nabla}_{\hbar \partial_{\hbar}}(Y) &= \hbar \partial_{\hbar} Y - \hbar^{-1} E \circ Y + \text{Gr}_E(Y), \end{aligned}$$

where Gr_E is the $\mathcal{O}_{\mathcal{M}}$ -linear map defined on flat vector fields Y by $Y \mapsto [E, Y]$. This connection is known as *the first structure connection*, or the *Dubrovin connection*. It has the following important property:

THEOREM 2.11. *The first structure connection is flat if and only if \mathcal{M} is Frobenius and E is an Euler vector field with*

$$[E, X \circ Y] - [E, X] \circ Y - X \circ [E, Y] = d_0 X \circ Y.$$

For a proof, see [76], Theorem I.2.5.2.

2.1.4. The quantum differential equation. Suppose we are given a non-singular variety X with cohomology only in even degrees. Let T_0, \dots, T_m be generators of $H^*(X, \mathbb{C})$ as usual, with T_0 the fundamental class of X and T_1, \dots, T_p generators of $H^2(X, \mathbb{R})$ non-negative on all $\beta \in H_2(X, \mathbb{Z})$ represented by stable maps. Let T^0, \dots, T^m be the Poincaré dual basis. This gives rise as in Example 2.8 to a Frobenius manifold \mathcal{M} with Euler vector field E and first structure connection $\widehat{\nabla}$. We will now describe solutions of the *quantum differential equation*

$$(2.8) \quad \widehat{\nabla}_{\partial_{y_i}} s = 0, \quad i = 0, \dots, m.$$

In particular, we will write down a fundamental set of solutions to this equation.

We first need

PROPOSITION 2.12. The Topological Recursion Relation (TRR).

$$\begin{aligned} & \langle \psi^{d_1} \alpha_1, \dots, \psi^{d_n} \alpha_n \rangle_{0, \beta} \\ &= \sum \langle \psi^{d_1-1} \alpha_1, \prod_{i \in S_1} \psi^{d_i} \alpha_i, T_e \rangle_{0, \beta_1} \langle T^e, \psi^{d_2} \alpha_2, \psi^{d_3} \alpha_3, \prod_{i \in S_2} \psi^{d_i} \alpha_i \rangle_{0, \beta_2} \end{aligned}$$

where the sum is over all $0 \leq e \leq m$, all splittings $\beta_1 + \beta_2 = \beta$, and S_1, S_2 disjoint index sets with $S_1 \cup S_2 = \{4, \dots, n\}$.

For a proof, see for example [90].

REMARK 2.13. In fact, the topological recursion relation actually shows that one can completely reconstruct the genus zero descendent invariants from the ordinary genus zero Gromov-Witten invariants. So genus zero descendent Gromov-Witten invariants do not actually carry new information, but it is useful to keep track of them here precisely because they give rise to solutions of the quantum differential equation.

Next, we introduce another notation: write, for $\alpha_1, \dots, \alpha_n \in H^*(X, \mathbb{C})$, and non-negative integers d_1, \dots, d_n ,

$$\langle \langle \psi^{d_1} \alpha_1, \dots, \psi^{d_n} \alpha_n \rangle \rangle := \sum_{k=0}^{\infty} \sum_{\beta \in H_2(X, \mathbb{Z})} \frac{1}{k!} \langle \psi^{d_1} \alpha_1, \dots, \psi^{d_n} \alpha_n, \gamma^k \rangle_{0, \beta}.$$

Here, as in the definition of the Gromov-Witten potential, $\gamma = \sum_{i=0}^m y_i T_i$, and γ^k means we take γ k times. Note this allows us to write the Gromov-Witten potential as

$$\Phi = \langle \langle 1 \rangle \rangle.$$

Furthermore,

$$\partial_{y_i} \langle \langle \psi^{d_1} \alpha_1, \dots, \psi^{d_n} \alpha_n \rangle \rangle = \langle \langle \psi^{d_1} \alpha_1, \dots, \psi^{d_n} \alpha_n, T_i \rangle \rangle.$$

We can then define, for $0 \leq i \leq m$,

$$s_i := T_i - \sum_{j=0}^m \langle \langle \frac{T_i}{\hbar + \psi}, T_j \rangle \rangle T^j.$$

Here $T_i/(\hbar + \psi)$ is viewed in terms of the power series expansion

$$\frac{T_i}{\hbar + \psi} = \sum_{n=0}^{\infty} (-1)^n \hbar^{-(n+1)} \psi^n T_i.$$

THEOREM 2.14. *The s_i 's satisfy the quantum differential equation (2.8).*

PROOF. The TRR implies

$$\begin{aligned}
 & \langle \langle \psi^{d_1+1} T_{j_1}, \psi^{d_2} T_{j_2}, \psi^{d_3} T_{j_3} \rangle \rangle \\
 &= \sum_{k=0}^{\infty} \sum_{\beta} \frac{1}{k!} \langle \psi^{d_1+1} T_{j_1}, \psi^{d_2} T_{j_2}, \psi^{d_3} T_{j_3}, \gamma^k \rangle_{0,\beta} \\
 &= \sum_{k=0}^{\infty} \sum_{i=0}^m \sum_{\substack{\beta_1, \beta_2 \\ S_1, S_2}} \frac{1}{k!} \langle \psi^{d_1} T_{j_1}, \gamma^{\#S_1}, T_i \rangle_{0,\beta_1} \langle T^i, \psi^{d_2} T_{j_2}, \psi^{d_3} T_{j_3}, \gamma^{\#S_2} \rangle_{0,\beta_2} \\
 (2.9) \quad &= \sum_{k=0}^{\infty} \sum_{i=0}^m \sum_{\substack{\beta_1, \beta_2 \\ a+b=k}} \frac{1}{a!b!} \langle \psi^{d_1} T_{j_1}, \gamma^a, T_i \rangle_{0,\beta_1} \langle T^i, \psi^{d_2} T_{j_2}, \psi^{d_3} T_{j_3}, \gamma^b \rangle_{0,\beta_2} \\
 &= \sum_{i=0}^m \langle \langle \psi^{d_1} T_{j_1}, T_i \rangle \rangle \langle \langle T^i, \psi^{d_2} T_{j_2}, \psi^{d_3} T_{j_3} \rangle \rangle.
 \end{aligned}$$

Note that the quantum differential equation (2.8) can be written, for $s = s_j$, as

$$(2.10) \quad \hbar \partial_{y_i}(s_j) = -T_i * s_j, \quad \text{for } 0 \leq i \leq m.$$

Now

$$\hbar \partial_{y_i}(s_j) = - \sum_{k=0}^m \langle \langle T_i, \frac{\hbar T_j}{\hbar + \psi}, T_k \rangle \rangle T^k.$$

On the other hand, $\partial_{y_i} \partial_{y_j} \partial_{y_k} \Phi = \langle \langle T_i, T_j, T_k \rangle \rangle$ and the quantum product is given by

$$T_i * T_j = \sum_k \langle \langle T_i, T_j, T_k \rangle \rangle T^k.$$

So

$$\begin{aligned}
 T_i * s_j &= T_i * T_j - \sum_{k=0}^m \langle \langle \frac{T_j}{\hbar + \psi}, T_k \rangle \rangle T_i * T^k \\
 &= \sum_{\ell=0}^m \langle \langle T_i, T_j, T_\ell \rangle \rangle T^\ell - \sum_{k,\ell=0}^m \langle \langle \frac{T_j}{\hbar + \psi}, T_k \rangle \rangle \langle \langle T_i, T^k, T_\ell \rangle \rangle T^\ell \\
 &= \sum_{\ell=0}^m \langle \langle T_i, T_j, T_\ell \rangle \rangle T^\ell - \sum_{\ell=0}^m \langle \langle \frac{\psi T_j}{\hbar + \psi}, T_i, T_\ell \rangle \rangle T^\ell \\
 &= \sum_{\ell=0}^m \langle \langle T_i, \frac{\hbar T_j}{\hbar + \psi}, T_\ell \rangle \rangle T^\ell.
 \end{aligned}$$

Here the next to last equality follows from (2.9). So the left- and right-hand sides of (2.10) agree. \square

It will be useful to use the axioms of descendent Gromov-Witten invariants to rewrite the s_i 's as follows.

PROPOSITION 2.15.

$$s_i = e^{-(\sum_{k=0}^p y_k T_k)/\hbar} \cup T_i - \sum_{\substack{\beta, j \\ n_k}} \left\langle \frac{e^{-(\sum_{k=0}^p y_k T_k)/\hbar} \cup T_i}{\hbar + \psi}, T_j, T_{p+1}^{n_{p+1}}, \dots, T_m^{n_m} \right\rangle_{0, \beta} T^j \\ \cdot \left(\prod_{\ell=1}^p e^{y_\ell \int_\beta T_\ell} \right) \frac{y_{p+1}^{n_{p+1}} \cdots y_m^{n_m}}{n_{p+1}! \cdots n_m!}$$

where the sum is over $0 \leq j \leq m$, $\beta \in H_2(X, \mathbb{Z})$, and $n_{p+1}, \dots, n_m \geq 0$. The expression $e^{-(\sum_{i=0}^p y_i T_i)/\hbar}$ is interpreted in the ring $H^*(X, \mathbb{C}) \otimes_{\mathbb{C}} \mathbb{C}[y_1, \dots, y_p][[y_0, \hbar^{-1}]]$, with the usual cup product.

PROOF. First note that by the Fundamental Class Axiom and then the Divisor Axiom for descendent invariants, if $\beta \neq 0$ or $\beta = 0$ and $\sum_{k=p+1}^m n_k \geq 1$,

$$\begin{aligned} & \left\langle \frac{T_i}{\hbar + \psi}, T_j, T_0^{n_0}, \dots, T_m^{n_m} \right\rangle_{0, \beta} \\ &= \left\langle \sum_{n=0}^{\infty} (-1)^n \hbar^{-(n+1)} \psi^n T_i, T_j, T_0^{n_0}, \dots, T_m^{n_m} \right\rangle_{0, \beta} \\ &= \left\langle \sum_{n=n_0}^{\infty} (-1)^n \hbar^{-(n+1)} \psi^{n-n_0} T_i, T_j, T_1^{n_1}, \dots, T_m^{n_m} \right\rangle_{0, \beta} \\ &= \sum_{n'_k} \left\langle \sum_n (-1)^n \hbar^{-(n+1)} \psi^{n-n_0-n'_1-\dots-n'_p} (T_i \cup T_1^{n'_1} \cup \dots \cup T_p^{n'_p}), T_j, T_{p+1}^{n_{p+1}}, \right. \\ & \quad \left. \dots, T_m^{n_m} \right\rangle_{0, \beta} \prod_{\ell=1}^p \binom{n_\ell}{n'_\ell} \left(\int_\beta T_\ell \right)^{n_\ell - n'_\ell} \end{aligned}$$

where the sum over n'_k is over all n'_1, \dots, n'_p with $0 \leq n'_k \leq n_k$ and the sum over n is over $n \geq n_0 + n'_1 + \dots + n'_p$. Note also that for $\beta = 0$, if $n_k = 0$ for $k > p$ but some $n_k \neq 0$, $0 \leq k \leq p$, then by the Fundamental Class Axiom, Divisor Axiom, and Point Mapping Axiom,

$$\begin{aligned} & \left\langle \frac{T_i}{\hbar + \psi}, T_j, T_0^{n_0}, \dots, T_p^{n_p} \right\rangle_{0, 0} \\ &= \langle (-1)^{n_0 + \dots + n_p - 1} \hbar^{-(n_0 + \dots + n_p)} (T_i \cup T_0^{n_0} \cup \dots \cup T_k^{n_k - 1} \cup \dots \cup T_p^{n_p}), T_j, T_k \rangle_{0, 0} \\ &= (-1)^{n_0 + \dots + n_p - 1} \hbar^{-(n_0 + \dots + n_p)} \int_X T_i \cup T_j \cup T_0^{n_0} \cup \dots \cup T_p^{n_p}. \end{aligned}$$

Thus if we expand $\langle \langle T_i / (\hbar + \psi), T_j \rangle \rangle$ by expanding γ^k , we split this up into two types of terms: terms involving $\beta = 0$ and no T_k 's for $k \geq p + 1$, and the other

terms. This gives

$$\begin{aligned}
& \langle \langle \frac{T_i}{\hbar + \psi}, T_j \rangle \rangle \\
&= \sum_{n_0 + \dots + n_p \geq 1} (-1)^{n_0 + \dots + n_p - 1} \hbar^{-(n_0 + \dots + n_p)} \int_X T_i \cup T_j \cup T_0^{n_0} \cup \dots \cup T_p^{n_p} \frac{y_0^{n_0} \dots y_p^{n_p}}{n_0! \dots n_p!} \\
&+ \sum_{\beta, n_k} \sum_{n'_k} \langle \sum_n (-1)^n \hbar^{-(n+1)} \psi^{n-n_0-n'_1-\dots-n'_p} (T_i \cup T_1^{n'_1} \cup \dots \cup T_p^{n'_p}), T_j, T_{p+1}^{n_{p+1}}, \\
&\quad \dots, T_m^{n_m} \rangle_{0, \beta} \left(\prod_{\ell=1}^p \binom{n_\ell}{n'_\ell} \left(\int_\beta T_\ell \right)^{n_\ell - n'_\ell} \right) \frac{y_0^{n_0} \dots y_m^{n_m}}{n_0! \dots n_m!},
\end{aligned}$$

where the sum in the second term is over all β, n_k such that we don't have both $\beta = 0$ and $n_{p+1} = \dots = n_m = 0$, and $0 \leq n'_k \leq n_k$. This is then the same thing as

$$\begin{aligned}
& \sum_{n_0 + \dots + n_p \geq 1} (-1)^{n_0 + \dots + n_p - 1} \hbar^{-(n_0 + \dots + n_p)} \int_X T_i \cup T_j \cup T_0^{n_0} \cup \dots \cup T_p^{n_p} \frac{y_0^{n_0} \dots y_p^{n_p}}{n_0! \dots n_p!} \\
&+ \sum_{\beta, n_k} \sum_{n'_k} \sum_{n=0}^{\infty} \\
&\langle \frac{(-1)^n \hbar^{-(n+1)} \psi^n (T_i \cup (y_0 T_0)^{n_0} \cup \bigcup_{\ell=1}^p (y_\ell T_\ell)^{n'_\ell}) (-\hbar)^{-(n_0 + n'_1 + \dots + n'_p)}}{n_0! (n'_1)! \dots (n'_p)!}, \\
&\quad T_j, T_{p+1}^{n_{p+1}}, \dots, T_m^{n_m} \rangle_{0, \beta} \prod_{\ell=1}^p \frac{1}{(n_\ell - n'_\ell)!} \left(y_\ell \int_\beta T_\ell \right)^{n_\ell - n'_\ell} \frac{y_{p+1}^{n_{p+1}} \dots y_m^{n_m}}{n_{p+1}! \dots n_m!} \\
&= - \int_X T_i \cup T_j \cup (e^{-(y_0 T_0 + \dots + y_p T_p)/\hbar} - 1) \\
&+ \sum_{\beta, n_k} \langle \frac{T_i \cup e^{-(y_0 T_0 + \dots + y_p T_p)/\hbar}}{\hbar + \psi}, T_j, T_{p+1}^{n_{p+1}}, \dots, T_m^{n_m} \rangle_{0, \beta} \\
&\quad \cdot \left(\prod_{\ell=1}^p e^{y_\ell \int_\beta T_\ell} \right) \frac{y_{p+1}^{n_{p+1}} \dots y_m^{n_m}}{n_{p+1}! \dots n_m!}.
\end{aligned}$$

This then gives the desired form for s_i . □

There are two issues we would still like to understand. First, as we described them, the solutions s_i to the quantum differential equation are purely formal solutions, and we would like to understand whether they converge. Second, we have shown that the s_i are flat sections on $\mathcal{M} \times \mathbb{C}^\times$ with respect to the connection $\widehat{\nabla}$ in the horizontal (\mathcal{M}) direction, but we would like to understand how they behave in the vertical (\mathbb{C}^\times) direction.

These two questions are closely connected, but first we shall introduce a replacement for \mathcal{M} which we shall work on from now on when discussing quantum cohomology.

Recall that $\mathcal{M} = \text{Spec } \mathbb{C}[\kappa_1^\pm, \dots, \kappa_p^\pm][[y_0, y_{p+1}, \dots, y_m]]$. We would like to think of this instead as a kind of formal completion of a complex analytic space. But since we in fact from now on will have to work with the universal cover $\widetilde{\mathcal{M}}$ of \mathcal{M} , we will describe $\widetilde{\mathcal{M}}$ directly.

DEFINITION 2.16. We define $\widetilde{\mathcal{M}}$ to be the ringed space $(\mathbb{C}^p, \mathcal{O}_{\widetilde{\mathcal{M}}})$ where the coordinates on \mathbb{C}^p are y_1, \dots, y_p and $\mathcal{O}_{\widetilde{\mathcal{M}}}$ denotes the structure sheaf on $\widetilde{\mathcal{M}}$: on an open set $U \subseteq \widetilde{\mathcal{M}}$, $\mathcal{O}_{\widetilde{\mathcal{M}}}(U)$ will consist of formal power series

$$(2.11) \quad \sum f_{i_0 i_{p+1} \dots i_m} y_0^{i_0} y_{p+1}^{i_{p+1}} \dots y_m^{i_m}$$

with $f_{i_0 i_{p+1} \dots i_m}$ a holomorphic function on the open set U .

There is a map $\widetilde{\mathcal{M}} \rightarrow \mathcal{M}$ given by $\kappa_i = e^{y_i}$ for $1 \leq i \leq p$ and $y_j = y_j$ for $j \notin \{1, \dots, p\}$. We will from now on, when we discuss structures arising from quantum cohomology, work in this formal setting.

We also define some additional notation.

DEFINITION 2.17. Let $\mathbb{C}\{\hbar, \hbar^{-1}\}$ denote the ring of Laurent series which converge on a punctured disk $\{\hbar \in \mathbb{C} \mid 0 < |\hbar| < \epsilon\}$ for some ϵ . Let $\mathbb{C}\{\hbar\}$ be the subring of functions holomorphic at $\hbar = 0$, and let $\mathcal{O}(\mathbb{P}^1 \setminus \{0\})$ be the ring of holomorphic functions on $\mathbb{P}^1 \setminus \{0\}$ (with coordinate \hbar). We can write

$$\mathbb{C}\{\hbar, \hbar^{-1}\} = \mathbb{C}\{\hbar\} \oplus \hbar^{-1} \mathcal{O}(\mathbb{P}^1 \setminus \{0\}).$$

We will now define

$$\mathcal{O}_{\widetilde{\mathcal{M}}}\{\hbar, \hbar^{-1}\}$$

to be the sheaf whose value on an open set U is the ring of formal series as in (2.11) such that each coefficient $f_{i_0 i_{p+1} \dots i_m}$ is a holomorphic function on

$$\{(y, \hbar) \in U \times \mathbb{C} \mid 0 < |\hbar| < \epsilon(y)\}$$

for some continuous map $\epsilon : U \rightarrow \mathbb{R}_{>0}$. We can then define

$$\mathcal{O}_{\widetilde{\mathcal{M}}}\{\hbar\} \subseteq \mathcal{O}_{\widetilde{\mathcal{M}}}\{\hbar, \hbar^{-1}\}$$

to be the subsheaf where the coefficients $f_{i_0 i_{p+1} \dots i_m}$ are functions holomorphic at $\hbar = 0$. We can similarly define $\mathcal{O}_{\mathcal{M}}\{\hbar\}$ and $\mathcal{O}_{\mathcal{M}}\{\hbar, \hbar^{-1}\}$ in the same way.

We will often want to consider the case when \mathcal{M} is just an ordinary complex manifold, as this is somewhat easier conceptually. In this case $\mathcal{O}_{\mathcal{M}}\{\hbar, \hbar^{-1}\}$ will be the sheaf which, on an open set $U \subseteq \mathcal{M}$, consists of functions f holomorphic on $\{(y, \hbar) \in U \times \mathbb{C} \mid 0 < \hbar < \epsilon(y)\}$ for some continuous $\epsilon : U \rightarrow \mathbb{R}_{>0}$ (depending on f). Furthermore, $\mathcal{O}_{\mathcal{M}}\{\hbar\}$ is the subsheaf of functions also holomorphic at $\hbar = 0$.

Returning to the question of the convergence of the s_i 's and the behaviour of the s_i 's in \hbar direction, let us introduce the map S which takes

$$\alpha \in H^*(X, \mathbb{C}) \otimes_{\mathbb{C}} \mathcal{O}_{\widetilde{\mathcal{M}}}\{\hbar, \hbar^{-1}\}$$

to

$$S(\alpha) = \alpha - \sum_{j=0}^m \langle \langle \frac{\alpha}{\hbar + \psi}, T_j \rangle \rangle T^j$$

We would like to consider S as a map

$$S : H^*(X, \mathbb{C}) \otimes_{\mathbb{C}} \mathcal{O}_{\widetilde{\mathcal{M}}}\{\hbar, \hbar^{-1}\} \rightarrow H^*(X, \mathbb{C}) \otimes_{\mathbb{C}} \mathcal{O}_{\widetilde{\mathcal{M}}}\{\hbar, \hbar^{-1}\},$$

but a priori, $S(\alpha)$ is a purely formal expression, so we will avoid describing the range of S at this point. However, we can use S to compute the effect of $\widehat{\nabla}_{\hbar \partial_{\hbar}}$ on s_i :

PROPOSITION 2.18.

$$\widehat{\nabla}_{\hbar\partial_{\hbar}} s_i = S(\mathrm{Gr}_E(T_i) - \hbar^{-1}c_1(\mathcal{T}_X) \cup T_i).$$

PROOF. We have

$$\widehat{\nabla}_{\hbar\partial_{\hbar}}(s_i) = \hbar\partial_{\hbar}s_i - \hbar^{-1}E * s_i + \mathrm{Gr}_E(s_i).$$

Since $0 = \nabla_E s_i + \hbar^{-1}E * s_i$, we can rewrite this as

$$\widehat{\nabla}_{\hbar\partial_{\hbar}}(s_i) = \hbar\partial_{\hbar}s_i + \nabla_E s_i + \mathrm{Gr}_E(s_i).$$

If we denote by $\beta \cup \bullet$ the map

$$\alpha \mapsto \beta \cup \alpha,$$

we can write

$$S = S' \circ (e^{-(\sum_{k=0}^p y_k T_k)/\hbar} \cup \bullet)$$

for a suitable $\mathcal{O}_{\widehat{\mathcal{M}}}\{\hbar, \hbar^{-1}\}$ -linear map S' , using the formula of Proposition 2.15. Define an element

$$s \in H^*(X, \mathbb{C}) \otimes_{\mathbb{C}} \mathcal{O}_{\widehat{\mathcal{M}}}\{\hbar, \hbar^{-1}\}$$

to be homogeneous of degree d if

$$(\hbar\partial_{\hbar} + \nabla_E + \mathrm{Gr}_E)(s) = ds.$$

Then note that an expression of the form

$$\hbar^{-(n+1)} T^j \kappa_1^{\int_{\beta} T_1} \cdots \kappa_p^{\int_{\beta} T_p} y_{p+1}^{n_{p+1}} \cdots y_m^{n_m}$$

is homogeneous of degree

$$-(n+1) + ((2 \dim_{\mathbb{C}} X - \deg T_j)/2 - 1) + \int_{\beta} c_1(\mathcal{T}_X) + \sum_{k=p+1}^m n_k (1 - (\deg T_k)/2).$$

On the other hand, note that

$$\langle \psi^n T_i, T_j, T_{p+1}^{n_{p+1}}, \dots, T_m^{n_m} \rangle_{0, \beta}$$

is non-zero only if, by (2.2),

$$2 + \sum_{k=p+1}^m n_k + \dim_{\mathbb{C}} X - 3 + \int_{\beta} c_1(\mathcal{T}_X) = n + \frac{\deg T_i + \deg T_j}{2} + \sum_{k=p+1}^m \frac{n_k \deg T_k}{2}.$$

This tells us that in fact $S'(T_i)$ is homogeneous of degree $(\deg T_i)/2 - 1$, and thus S' does not change degrees. So

$$(\hbar\partial_{\hbar} + \nabla_E + \mathrm{Gr}_E) \circ S' = S' \circ (\hbar\partial_{\hbar} + \nabla_E + \mathrm{Gr}_E),$$

and thus

$$(\hbar\partial_{\hbar} + \nabla_E + \mathrm{Gr}_E)s_i = S'((\hbar\partial_{\hbar} + \nabla_E + \mathrm{Gr}_E)(e^{-(\sum_{k=0}^p y_k T_k)/\hbar} \cup T_i)).$$

On the other hand,

$$\begin{aligned}
& (\hbar\partial_{\hbar} + E + \text{Gr}_E)(e^{-y_0/\hbar}(e^{-(\sum_{k=1}^p y_k T_k)/\hbar} \cup T_i)) \\
&= (\hbar\partial_{\hbar}(e^{-y_0/\hbar}) + E(e^{-y_0/\hbar}))(e^{-(\sum_{k=1}^p y_k T_k)/\hbar} \cup T_i) \\
&\quad + e^{-y_0/\hbar}(\hbar\partial_{\hbar} + E + \text{Gr}_E)(e^{-(\sum_{k=1}^p y_k T_k)/\hbar} \cup T_i) \\
&= \left(\frac{y_0}{\hbar}e^{-y_0/\hbar} - \frac{y_0}{\hbar}e^{-y_0/\hbar}\right)(e^{-(\sum_{k=1}^p y_k T_k)/\hbar} \cup T_i) \\
&\quad + e^{-y_0/\hbar}\left(\hbar\partial_{\hbar}(e^{-(\sum_{k=1}^p y_k T_k)/\hbar} \cup T_i)\right. \\
&\quad + \sum_{j=1}^p -c_j T_j \hbar^{-1} \cup e^{-(\sum_{k=1}^p y_k T_k)/\hbar} \cup T_i \\
&\quad + \sum_{n_1, \dots, n_p \geq 0} \left(\frac{y_1^{n_1} \cdots y_p^{n_p}}{n_1! \cdots n_p!} (-\hbar)^{-(n_1 + \cdots + n_p)} T_1^{n_1} \cup \cdots \cup T_p^{n_p} \cup T_i\right) \\
&\quad \left. \cdot \left((n_1 + \cdots + n_p) + \frac{\deg T_i}{2} - 1\right)\right) \\
&= e^{-y_0/\hbar}(\hbar\partial_{\hbar}(e^{-(\sum_{k=1}^p y_k T_k)/\hbar} \cup T_i) \\
&\quad - \hbar^{-1}c_1(\mathcal{T}_X) \cup e^{-(\sum_{k=1}^p y_k T_k)/\hbar} \cup T_i \\
&\quad + \left(\frac{\deg T_i}{2} - 1\right)e^{-(\sum_{k=1}^p y_k T_k)/\hbar} \cup T_i \\
&\quad - \hbar\partial_{\hbar}(e^{-(\sum_{k=1}^p y_k T_k)/\hbar} \cup T_i)) \\
&= e^{-(\sum_{k=0}^p y_k T_k)/\hbar} \cup \left(\left(\frac{\deg T_i}{2} - 1\right) - \hbar^{-1}c_1(\mathcal{T}_X) \cup T_i\right).
\end{aligned}$$

Thus

$$(2.12) \quad (\hbar\partial_{\hbar} + \nabla_E + \text{Gr}_E) \circ S = S \circ (\hbar\partial_{\hbar} + \nabla_E + \text{Gr}_E - \hbar^{-1}c_1(\mathcal{T}_X) \cup \bullet),$$

giving the desired result. \square

So the sections s_i are flat in the directions given by ∂_{y_i} , but not in the direction given by ∂_{\hbar} . This is not surprising, as the differential equation

$$\widehat{\nabla}_{\hbar\partial_{\hbar}} s = 0$$

in fact has singular points at $\hbar = 0$ and $\hbar = \infty$. As a result, one would expect that any solution to this equation would be multi-valued. Indeed, we can write down a multi-valued solution by allowing power series in $\log \hbar$ as follows.

PROPOSITION 2.19. *Let $\hbar^{-\text{Gr}_E} \hbar^{c_1(\mathcal{T}_X) \cup \bullet}$ be the endomorphism*

$$\begin{aligned}
& \exp(-\text{Gr}_E \log \hbar) \exp((c_1(\mathcal{T}_X) \cup \bullet) \log \hbar) : H^*(X, \mathbb{C}) \otimes_{\mathbb{C}} \mathcal{O}_{\widetilde{\mathcal{M}}}\{\hbar, \hbar^{-1}\}[[\log \hbar]] \\
& \rightarrow H^*(X, \mathbb{C}) \otimes_{\mathbb{C}} \mathcal{O}_{\widetilde{\mathcal{M}}}\{\hbar, \hbar^{-1}\}[[\log \hbar]].
\end{aligned}$$

Then

$$\widehat{\nabla}_{\partial_{y_i}} S(\hbar^{-\text{Gr}_E} \hbar^{c_1(\mathcal{T}_X) \cup \bullet} T_i) = 0$$

and

$$\widehat{\nabla}_{\hbar\partial_{\hbar}} S(\hbar^{-\text{Gr}_E} \hbar^{c_1(\mathcal{T}_X) \cup \bullet} T_i) = 0.$$

PROOF. The first vanishing holds because $\widehat{\nabla}_{\partial_{y_i}} S(T_j) = 0$ for all j . For the second, by (2.12), we have

$$\begin{aligned} & \widehat{\nabla}_{\hbar\partial_{\hbar}} (S(\hbar^{-\text{Gr}_E} \hbar^{c_1(\mathcal{T}_X) \cup \bullet} T_i)) \\ &= S((\hbar\partial_{\hbar} + \nabla_E + \text{Gr}_E - \hbar^{-1} c_1(\mathcal{T}_X) \cup \bullet)(\hbar^{-\text{Gr}_E} \hbar^{c_1(\mathcal{T}_X) \cup \bullet} T_i)) \\ &= S(-\text{Gr}_E \hbar^{-\text{Gr}_E} \hbar^{c_1(\mathcal{T}_X) \cup \bullet} T_i + \hbar^{-\text{Gr}_E} \hbar^{c_1(\mathcal{T}_X) \cup \bullet} (c_1(\mathcal{T}_X) \cup T_i) \\ &\quad + \text{Gr}_E \hbar^{-\text{Gr}_E} \hbar^{c_1(\mathcal{T}_X) \cup \bullet} T_i - \hbar^{-1} \hbar^{-\text{Gr}_E + 1} \hbar^{c_1(\mathcal{T}_X) \cup \bullet} (c_1(\mathcal{T}_X) \cup T_i)) \\ &= 0 \end{aligned}$$

by a straightforward computation, remembering that Gr_E and $c_1(\mathcal{T}_X) \cup \bullet$ don't commute! \square

It is worth exploring in greater detail what this means by considering the behaviour of the ordinary differential equation $\widehat{\nabla}_{\hbar\partial_{\hbar}} s = 0$ at $\hbar = \infty$. Rewriting this equation in terms of $q = \hbar^{-1}$, we get

$$q\partial_q s = -qE * s + \text{Gr}_E(s).$$

This equation has a *regular singular point* at $q = 0$. In general, an ordinary differential equation on a disk $\Delta \subseteq \mathbb{C}$ centered at the origin with coordinate q of the form

$$(2.13) \quad q\partial_q s = A(q) \cdot s,$$

where the unknown s is a holomorphic map with values in a vector space \mathbb{C}^n and A is a holomorphic map $A : \Delta \rightarrow \text{End}(\mathbb{C}^n)$, is said to have a *regular singular point* at $q = 0$. In our example, $A(q)$ is the endomorphism $-qE * \bullet + \text{Gr}_E$. Note that the equation $\widehat{\nabla}_{\hbar\partial_{\hbar}} s = 0$ does not have a regular singular point at $\hbar = 0$ as it takes the form $\hbar\partial_{\hbar} s = B(\hbar) \cdot s$ where $B(\hbar) = \hbar^{-1} E * \bullet - \text{Gr}$ has a pole at $\hbar = 0$.

The theory of ODEs with regular singular points is of course well-developed classically, and in particular, one can find n linearly independent multi-valued convergent solutions to (2.13) on $\Delta \setminus \{0\}$. In the case at hand, these must coincide with the formal solutions described in Proposition 2.19. This shows that these formal solutions are convergent, so in particular the solutions s_i to the quantum differential equation are convergent.

To read more about the theory of ODEs with singular points in the above context, [103] gives an excellent introduction.

As a consequence of the above discussion, we can view the map S as more than purely formal, viewing it as an $\mathcal{O}_{\widehat{\mathcal{M}}}\{\hbar, \hbar^{-1}\}$ -linear map

$$S : H^*(X, \mathbb{C}) \otimes_{\mathbb{C}} \mathcal{O}_{\widehat{\mathcal{M}}}\{\hbar, \hbar^{-1}\} \rightarrow H^*(X, \mathbb{C}) \otimes_{\mathbb{C}} \mathcal{O}_{\widehat{\mathcal{M}}}\{\hbar, \hbar^{-1}\}$$

with

$$S(\alpha) = \alpha - \sum_{j=0}^m \left\langle \left\langle \frac{\alpha}{\hbar + \psi}, T_j \right\rangle \right\rangle T^j.$$

2.1.5. Semi-infinite variations of Hodge structure. We will now introduce yet more structure into this picture. We will define the notion of a semi-infinite variation of Hodge structure, a notion introduced by Barannikov in [5]. These structures may seem to be, at this point, an unnecessarily complicated way of encoding the various data studied above, but as we shall see, it emerges naturally on the B-model side.

DEFINITION 2.20. A *semi-infinite variation of Hodge structure* parameterized by a space \mathcal{M} is a locally free $\mathcal{O}_{\mathcal{M}}\{\hbar\}$ -module \mathcal{E} of finite rank together with a flat connection

$$\nabla : \mathcal{E} \rightarrow \Omega_{\mathcal{M}}^1 \otimes \hbar^{-1}\mathcal{E}$$

and a pairing

$$(\cdot, \cdot)_{\mathcal{E}} : \mathcal{E} \times \mathcal{E} \rightarrow \mathcal{O}_{\mathcal{M}}\{\hbar\}$$

satisfying

- (1) $(s_1, s_2)_{\mathcal{E}}(\hbar) = (s_2, s_1)_{\mathcal{E}}(-\hbar)$. Here we view $(s_1, s_2)_{\mathcal{E}}$ as a function of \hbar ; on the right-hand side, we have replaced \hbar with $-\hbar$.
- (2) $(f(-\hbar)s_1, s_2)_{\mathcal{E}} = (s_1, f(\hbar)s_2)_{\mathcal{E}} = f(\hbar)(s_1, s_2)_{\mathcal{E}}$ for $f(\hbar) \in \mathcal{O}_{\mathcal{M}}\{\hbar\}$.
- (3) $Y(s_1, s_2)_{\mathcal{E}} = (\nabla_Y s_1, s_2)_{\mathcal{E}} + (s_1, \nabla_Y s_2)_{\mathcal{E}}$.
- (4) The pairing is *non-degenerate*, i.e., the induced pairing

$$(\mathcal{E}/\hbar\mathcal{E}) \otimes_{\mathcal{O}_{\mathcal{M}}} (\mathcal{E}/\hbar\mathcal{E}) \rightarrow \mathcal{O}_{\mathcal{M}}$$

is non-degenerate.

A *grading* on this semi-infinite variation of Hodge structure is a \mathbb{C} -linear endomorphism $\text{Gr} : \mathcal{E} \rightarrow \mathcal{E}$ such that there exists a vector field E on \mathcal{M} and a constant $D \in \mathbb{C}$ such that

- (5) $\text{Gr}(fs) = ((\hbar\partial_{\hbar} + E)f)s + f\text{Gr}(s)$ for $f \in \mathcal{O}_{\mathcal{M}}\{\hbar\}$, $s \in \mathcal{E}$.
- (6) $[\text{Gr}, \nabla_Y] = \nabla_{[E, Y]}$ for all vector fields Y on \mathcal{M} .
- (7) $(\hbar\partial_{\hbar} + E)(s_1, s_2)_{\mathcal{E}} = (\text{Gr}(s_1), s_2)_{\mathcal{E}} + (s_1, \text{Gr}(s_2))_{\mathcal{E}} + D(s_1, s_2)_{\mathcal{E}}$.

EXAMPLE 2.21. Quantum cohomology furnishes an example of a semi-infinite variation of Hodge structure. With T_0, \dots, T_m a basis for $H^*(X, \mathbb{C})$ as usual, with X having only even cohomology, we take

$$\mathcal{E} = H^*(X, \mathbb{C}) \otimes_{\mathbb{C}} \mathcal{O}_{\widetilde{\mathcal{M}}}\{\hbar\}.$$

We take the flat connection to be $\widehat{\nabla}$, i.e.,

$$\widehat{\nabla}_X(Y) = \nabla_X(Y) + \hbar^{-1}X * Y,$$

where ∇ denotes the standard flat connection on \mathcal{E} . The pairing is

$$(s_1, s_2)_{\mathcal{E}} = \int_X s_1(-\hbar) \cup s_2(\hbar).$$

Conditions (1) and (2) are obvious from this definition. For (3), note that for a vector field Y on $\widetilde{\mathcal{M}}$,

$$\begin{aligned} Y(s_1, s_2)_{\mathcal{E}} &= \int_X Y(s_1(-\hbar) \cup s_2(\hbar)) \\ &= \int_X (Ys_1(-\hbar)) \cup s_2(\hbar) + s_1(-\hbar) \cup (Ys_2(\hbar)) \end{aligned}$$

while

$$\begin{aligned} &(\widehat{\nabla}_Y s_1, s_2)_{\mathcal{E}} + (s_1, \widehat{\nabla}_Y s_2)_{\mathcal{E}} \\ &= \int_X (Ys_1(-\hbar) - \hbar^{-1}Y * s_1(-\hbar)) \cup s_2(\hbar) + s_1(-\hbar) \cup (Ys_2(\hbar) + \hbar^{-1}Y * s_2(\hbar)) \\ &= \int_X Ys_1(-\hbar) \cup s_2(\hbar) + s_1(-\hbar) \cup (Ys_2(\hbar)) \end{aligned}$$

using the fact that

$$\int_X (Y * Z) \cup W = \int_X Z \cup (Y * W)$$

by (2.4).

(4) is clear since $\mathcal{E}/\hbar\mathcal{E} \cong H^*(X, \mathbb{C}) \otimes_{\mathbb{C}} \mathcal{O}_{\widetilde{\mathcal{M}}}$, and the induced pairing is just the ordinary cup product.

The grading operator is

$$\text{Gr} = \hbar\partial_{\hbar} + E + \text{Gr}_E$$

where E is the Euler vector field on $\widetilde{\mathcal{M}}$ of Example 2.10 and as usual, Gr_E is $\mathcal{O}_{\widetilde{\mathcal{M}}}\{\hbar\}$ -linear with

$$\text{Gr}_E(T_i) = \left(\frac{\deg T_i}{2} - 1 \right) T_i.$$

Note that this operator has already appeared in the proof of Proposition 2.18. We check conditions (5)-(7). First, (5) is obvious from the Leibniz rule for $\hbar\partial_{\hbar}$ and ∇_E . For (6), first observe that if (6) holds for a vector field Y , it also holds for fY for $f \in \mathcal{O}_{\widetilde{\mathcal{M}}}$. Also, if the left-hand and right-hand sides agree when evaluated on a section s of \mathcal{E} , then they also agree when evaluated on a section fs of \mathcal{E} , for $f \in \mathcal{O}_{\widetilde{\mathcal{M}}}\{\hbar\}$. Thus it is enough to check (6) for $Y = \partial_{y_i}$ by evaluating on the section T_j . The left-hand-side is

$$\begin{aligned} [\text{Gr}, \widehat{\nabla}_{\partial_{y_i}}](T_j) &= \text{Gr}(\widehat{\nabla}_{\partial_{y_i}} T_j) - \widehat{\nabla}_{\partial_{y_i}}(\text{Gr}(T_j)) \\ &= \text{Gr}(\hbar^{-1} T_i * T_j) - \left(\frac{\deg T_j}{2} - 1 \right) \hbar^{-1} T_i * T_j \\ &= \hbar^{-1} \left(-T_i * T_j + \text{Gr}(T_i * T_j) - \left(\frac{\deg T_j}{2} - 1 \right) T_i * T_j \right) \\ &= \hbar^{-1} \left(\text{Gr}(T_i * T_j) - \frac{\deg T_j}{2} T_i * T_j \right). \end{aligned}$$

Now if we view $T_i * T_j$ as a vector field on $\widetilde{\mathcal{M}}$, after identifying $\widetilde{\mathcal{M}} \times H^*(X, \mathbb{C})$ with the tangent bundle to $\widetilde{\mathcal{M}}$, then we know that $T_i * T_j$ is homogeneous with respect to the Euler vector field E of degree $(\deg T_i + \deg T_j)/2 - 1$, as follows from the discussion of Example 2.10. So

$$\text{Gr}(T_i * T_j) = \left(\frac{\deg T_i + \deg T_j}{2} - 1 \right) T_i * T_j.$$

Thus the left-hand-side of (6) becomes

$$\hbar^{-1} \left(\frac{\deg T_i}{2} - 1 \right) (T_i * T_j).$$

The right-hand-side of (6) is

$$\widehat{\nabla}_{[E, \partial_{y_i}]} T_j = \hbar^{-1} \left(\frac{\deg T_i}{2} - 1 \right) T_i * T_j,$$

as desired.

Finally, for (7),

$$\begin{aligned}
& (\hbar\partial_{\hbar} + E)(s_1, s_2)_{\mathcal{E}} \\
&= \int_X (\hbar\partial_{\hbar} + E)(s_1(-\hbar)) \cup s_2(\hbar) + s_1(-\hbar) \cup (\hbar\partial_{\hbar} + E)(s_2(\hbar)) \\
&= (\text{Gr}(s_1), s_2)_{\mathcal{E}} + (s_1, \text{Gr}(s_2))_{\mathcal{E}} - \int_X \text{Gr}_E(s_1(-\hbar)) \cup s_2(\hbar) + s_1(-\hbar) \cup \text{Gr}_E(s_2(\hbar)) \\
&= (\text{Gr}(s_1), s_2)_{\mathcal{E}} + (s_1, \text{Gr}(s_2))_{\mathcal{E}} + (2 - \dim_{\mathbb{C}} X)(s_1, s_2)_{\mathcal{E}}.
\end{aligned}$$

□

2.1.6. The moving subspace realisation of a semi-infinite variation of Hodge structure. We can recast this definition in terms of a moving subspace inside a single space; this is more reminiscent of the usual notion of a variation of Hodge structure. Suppose we are given a semi-infinite variation of Hodge structure (\mathcal{E}, ∇) over a complex manifold \mathcal{M} , and assume \mathcal{M} is simply connected. Consider the space

$$(2.14) \quad \mathcal{H} = \{s \in \Gamma(\mathcal{M}, \mathcal{E} \otimes_{\mathcal{O}_{\mathcal{M}}\{\hbar\}} \mathcal{O}_{\mathcal{M}}\{\hbar, \hbar^{-1}\}) \mid \nabla s = 0\}.$$

Since ∇ is flat, one sees in fact that \mathcal{H} is a free $\mathbb{C}\{\hbar, \hbar^{-1}\}$ -module of the same rank as \mathcal{E} .

The pairing $(\cdot, \cdot)_{\mathcal{E}}$ on \mathcal{E} also defines a symplectic form on \mathcal{H} , by

$$(2.15) \quad \overline{\Omega}(s_1, s_2) := \text{Res}_{\hbar=0}(s_1, s_2)_{\mathcal{E}} d\hbar,$$

i.e., $\overline{\Omega}(s_1, s_2)$ is the coefficient of \hbar^{-1} in $(s_1, s_2)_{\mathcal{E}}$. Note that this is a constant in \mathbb{C} , since for any vector field Y on \mathcal{M} ,

$$Y(s_1, s_2)_{\mathcal{E}} = (\nabla_Y s_1, s_2)_{\mathcal{E}} + (s_1, \nabla_Y s_2)_{\mathcal{E}} = 0,$$

by Definition 2.20, (3), and flatness of s_1, s_2 .

For each point $x \in \mathcal{M}$, we now get an embedding of \mathcal{E}_x into \mathcal{H} , by sending $s \in \mathcal{E}_x$ to the section $s' \in \mathcal{H}$ satisfying $s'(x) = s$; in other words, we extend s to a flat section of $\mathcal{E} \otimes_{\mathcal{O}_{\mathcal{M}}\{\hbar\}} \mathcal{O}_{\mathcal{M}}\{\hbar, \hbar^{-1}\}$. Thus we get a varying family \mathcal{E}_x of subspaces of \mathcal{H} , parameterized by \mathcal{M} .

We can think about these varying subspaces via the Pressley-Segal Grassmannian of semi-infinite subspaces of \mathcal{H} . Suppose in the above semi-infinite variation of Hodge structures that \mathcal{E} is a free $\mathcal{O}_{\mathcal{M}}\{\hbar\}$ -module of rank N . Then the Pressley-Segal Grassmannian of \mathcal{H} is written as

$$L \text{GL}_N(\mathbb{C}) / L^+ \text{GL}_N(\mathbb{C}).$$

Here $L \text{GL}_N(\mathbb{C})$ is the loop group of $\text{GL}_N(\mathbb{C})$, consisting of smooth maps $\gamma : S^1 \rightarrow \text{GL}_N(\mathbb{C})$, where we think of $S^1 \subseteq \mathbb{C}$ as a circle with center the origin of some radius $\epsilon > 0$. We then have a subgroup $L^+ \text{GL}_N(\mathbb{C})$ consisting of those maps which are boundary values of holomorphic maps

$$\{\hbar \mid |\hbar| < \epsilon\} \rightarrow \text{GL}_N(\mathbb{C}).$$

The significance of this construction in our situation is as follows. Suppose that we have a rank N $\mathbb{C}\{\hbar\}$ -submodule of $\mathbb{C}^N \otimes_{\mathbb{C}} \mathbb{C}\{\hbar, \hbar^{-1}\}$. Then we can choose a $\mathbb{C}\{\hbar\}$ -basis $s_1(\hbar), \dots, s_N(\hbar)$ for this submodule, or equivalently, choose an $N \times N$ invertible matrix whose entries are elements of $\mathbb{C}\{\hbar, \hbar^{-1}\}$, and whose columns span the $\mathbb{C}\{\hbar\}$ -submodule. Now after choosing an $\epsilon > 0$ such that all the entries of this matrix converge on $\{\hbar \mid 0 < |\hbar| < 2\epsilon\}$, we can restrict each entry to the circle of

radius ϵ , getting an element of $LGL_N(\mathbb{C})$. An element of $L^+GL_N(\mathbb{C})$ is a matrix with entries in $\mathbb{C}\{\hbar\}$, hence acts by changing the basis of this $\mathbb{C}\{\hbar\}$ -submodule without changing the submodule. So in particular each $\mathbb{C}\{\hbar\}$ -submodule of \mathcal{H} of rank N is represented by a point in the Pressley-Segal Grassmannian. Thus the semi-infinite variation of Hodge structure \mathcal{E} induces a map

$$\mathcal{M} \rightarrow LGL_N(\mathbb{C})/L^+GL_N(\mathbb{C}).$$

This map can be described more explicitly as follows. Suppose that e_1, \dots, e_N are sections of \mathcal{E} such that for each point $x \in \mathcal{M}$, $e_1(x), \dots, e_N(x)$ descend to a \mathbb{C} -basis for $\mathcal{E}_x/\hbar\mathcal{E}_x$. On the other hand, let s_1, \dots, s_N be a basis for \mathcal{H} . Then there is a matrix $M = (M_{ij})_{1 \leq i, j \leq N}$ such that

$$s_i = \sum_{j=1}^N M_{ij} e_j$$

with $M_{ij} \in \mathcal{O}_{\mathcal{M}}\{\hbar, \hbar^{-1}\}$. At a point x_0 , we of course have

$$s_i(x_0) = \sum_{j=1}^N M_{ij}(x_0) e_j(x_0),$$

and if $M^{-1} = (M^{ij})_{1 \leq i, j \leq N}$, then

$$e_i(x_0) = \sum_{j=1}^N M^{ij}(x_0) s_j(x_0).$$

Thus \mathcal{E}_{x_0} is embedded in \mathcal{H} via the mapping

$$e_i(x_0) \mapsto \sum_{j=1}^N M^{ij}(x_0) s_j.$$

Thus, in the basis s_1, \dots, s_N of \mathcal{H} , the columns of M^{-1} generate the image of \mathcal{E}_{x_0} . Thus the map $\mathcal{M} \rightarrow LGL_N(\mathbb{C})/L^+GL_N(\mathbb{C})$ can be described explicitly by $x \mapsto M^{-1}(x)$.

The advantage of this description in terms of M is that this makes perfect sense in the situation arising from quantum cohomology, where we have the space $\widetilde{\mathcal{M}}$, which isn't quite a complex manifold because the variables y_0, y_{p+1}, \dots, y_m are formal. In this case, the varying family of subspaces \mathcal{E}_x in \mathcal{H} can be viewed as given by M^{-1} , an $N \times N$ matrix with entries in $\mathcal{O}_{\widetilde{\mathcal{M}}}\{\hbar, \hbar^{-1}\}$.

EXAMPLE 2.22. Returning to Example 2.21 given by quantum cohomology, we take the basis e_1, \dots, e_N to be T_0, \dots, T_m . We also have the flat sections s_0, \dots, s_m of §2.1.4. The matrix M is just given by the function S , and so the moving subspace realisation of the corresponding semi-infinite variation of Hodge structure is defined by S^{-1} . We define

$$\mathbb{J} = S^{-1} : H^*(X, \mathbb{C}) \otimes_{\mathbb{C}} \mathcal{O}_{\widetilde{\mathcal{M}}}\{\hbar, \hbar^{-1}\} \rightarrow H^*(X, \mathbb{C}) \otimes_{\mathbb{C}} \mathcal{O}_{\widetilde{\mathcal{M}}}\{\hbar, \hbar^{-1}\}.$$

Since the matrix for S defined using the basis $\{T_i\}$ of $H^*(X, \mathbb{C})$ has entries in $\mathcal{O}(\mathbb{P}^1 \setminus \{0\})$, and is invertible at every point of $\mathbb{P}^1 \setminus \{0\}$, the matrix for \mathbb{J} also has entries in $\mathcal{O}(\mathbb{P}^1 \setminus \{0\})$.

PROPOSITION 2.23. For $\alpha \in H^*(X, \mathbb{C})$,

$$\mathbb{J}(\alpha) = e^{(\sum_{k=0}^p y_k T_k)/\hbar} \cup \left(\alpha + \sum_{\beta, n_k} \sum_{i=0}^m \langle \alpha, T_{p+1}^{n_{p+1}}, \dots, T_m^{n_m}, \frac{T_i}{\hbar - \psi} \rangle_{0, \beta} T^i \right. \\ \left. \cdot \left(\prod_{k=1}^p e^{y_k \int_{\beta} T_k} \right) \frac{y_{p+1}^{n_{p+1}} \cdots y_m^{n_m}}{n_{p+1}! \cdots n_m!} \right).$$

PROOF. We will first prove, with $\mathcal{E} = H^*(X, \mathbb{C}) \otimes_{\mathbb{C}} \mathcal{O}_{\widetilde{\mathcal{M}}}\{\hbar\}$ as in Example 2.21, that for $\alpha, \beta \in H^*(X, \mathbb{C})$,

$$(S(\alpha), S(\beta))_{\mathcal{E}} = (\alpha, \beta)_{\mathcal{E}}.$$

Indeed, applying property (3) of Definition 2.20, we see that

$$\partial_{y_i}(S(\alpha), S(\beta))_{\mathcal{E}} = (\widehat{\nabla}_{\partial_{y_i}} S(\alpha), S(\beta))_{\mathcal{E}} + (S(\alpha), \widehat{\nabla}_{\partial_{y_i}} S(\beta))_{\mathcal{E}} \\ = 0,$$

since $S(\alpha), S(\beta)$ satisfy the quantum differential equation. Thus $(S(\alpha), S(\beta))_{\mathcal{E}}$ is a constant function on $\widetilde{\mathcal{M}}$. So we can try to compute it by setting $y_0 = y_{p+1} = \cdots = y_m = 0$ and letting $y_i \rightarrow -\infty$, i.e., letting $\kappa_i \rightarrow 0$ for $1 \leq i \leq p$. Taking this limit, we then get, using the formula of Proposition 2.15,

$$(S(\alpha), S(\beta))_{\mathcal{E}} = \int_X (e^{(\sum_{k=0}^p y_k T_k)/\hbar} \cup \alpha) \cup (e^{-(\sum_{k=0}^p y_k T_k)/\hbar} \cup \beta) \\ = \int_X \alpha \cup \beta.$$

This shows that S^{-1} is the adjoint to S with respect to the inner product $(\cdot, \cdot)_{\mathcal{E}}$. Thus we can compute \mathbb{J} using

$$(T_i, e^{-(\sum_{k=0}^p y_k T_k)/\hbar} \cup \mathbb{J}(T_j))_{\mathcal{E}} = (e^{(\sum_{k=0}^p y_k T_k)/\hbar} \cup T_i, \mathbb{J}(T_j))_{\mathcal{E}} \\ = (S(e^{(\sum_{k=0}^p y_k T_k)/\hbar} \cup T_i), T_j)_{\mathcal{E}},$$

thus reading off the coefficient of T^i in $e^{-(\sum_{k=0}^p y_k T_k)/\hbar} \cup \mathbb{J}(T_j)$ as the coefficient of T^j in $S(e^{(\sum_{k=0}^p y_k T_k)/\hbar} \cup T_i)(-\hbar)$, using the formula of Proposition 2.15. \square

One way to view \mathbb{J} is that it gives an embedding of bundles

$$\mathbb{J} : \mathcal{E} = H^*(X, \mathbb{C}) \otimes_{\mathbb{C}} \mathcal{O}_{\widetilde{\mathcal{M}}}\{\hbar\} \rightarrow \mathcal{H}_{\widetilde{\mathcal{M}}} = H^*(X, \mathbb{C}) \otimes_{\mathbb{C}} \mathcal{O}_{\widetilde{\mathcal{M}}}\{\hbar, \hbar^{-1}\}$$

taking flat sections of \mathcal{E} to constant sections of $\mathcal{H}_{\widetilde{\mathcal{M}}}$.

2.1.7. From variations of semi-infinite Hodge structure to Frobenius manifolds. We now demonstrate Barannikov's technique of going from a variation of semi-infinite Hodge structure to a Frobenius manifold. Let $(\mathcal{M}, \mathcal{E}, \nabla)$ be a variation of semi-infinite Hodge structure with \mathcal{M} a simply connected complex manifold, or a completion of a complex manifold along a simply connected submanifold, and fix a base-point $0 \in \mathcal{M}$. This determines a space \mathcal{H} as in (2.14) with an inclusion $\mathcal{E}_0 \hookrightarrow \mathcal{H}$ as a $\mathbb{C}\{\hbar\}$ -submodule.

An *opposite subspace* is an $\mathcal{O}(\mathbb{P}^1 \setminus \{0\})$ -submodule \mathcal{H}_- of \mathcal{H} such that the natural map

$$\mathcal{H}_- \oplus \mathcal{E}_0 \rightarrow \mathcal{H}$$

is an isomorphism.

First note that the projection $\mathcal{E}_0 \cap \hbar\mathcal{H}_- \rightarrow \mathcal{E}_0/\hbar\mathcal{E}_0$ is an isomorphism. Indeed, the map is injective since the kernel is $\hbar\mathcal{E}_0 \cap \hbar\mathcal{H}_- = \hbar(\mathcal{E}_0 \cap \mathcal{H}_-) = 0$. For surjectivity, if $s \in \mathcal{E}_0$, then $\hbar^{-1}s = s' + s''$ with $s' \in \mathcal{H}_-$ and $s'' \in \mathcal{E}_0$. Then $\hbar s' \in \mathcal{E}_0 \cap \hbar\mathcal{H}_-$ and $\hbar s' \equiv s \pmod{\hbar\mathcal{E}_0}$. Similarly, $\mathcal{E}_0 \cap \hbar\mathcal{H}_- \rightarrow \hbar\mathcal{H}_-/\mathcal{H}_-$ is an isomorphism.

These isomorphisms then give rise to isomorphisms

$$\mathcal{E}_0 \cong (\mathcal{E}_0 \cap \hbar\mathcal{H}_-) \otimes_{\mathbb{C}} \mathbb{C}\{\hbar\} \cong \left(\frac{\mathcal{E}_0}{\hbar\mathcal{E}_0} \right) \otimes_{\mathbb{C}} \mathbb{C}\{\hbar\} \cong \left(\frac{\hbar\mathcal{H}_-}{\mathcal{H}_-} \right) \otimes_{\mathbb{C}} \mathbb{C}\{\hbar\}.$$

Indeed, the map $(\mathcal{E}_0 \cap \hbar\mathcal{H}_-) \otimes_{\mathbb{C}} \mathbb{C}\{\hbar\} \rightarrow \mathcal{E}_0$ is the obvious one, taking $s \otimes f \mapsto s \cdot f$.

By §2.1.6, we can think of \mathcal{E} as a subbundle of the trivial bundle $\mathcal{H} \times \mathcal{M}$ on \mathcal{M} . In general, when we have a vector space V , we shall write $V_{\mathcal{M}}$ for the trivial vector bundle with fibre V , and an element $s \in V$ yields a constant section of $V_{\mathcal{M}}$ which we shall also denote by s . So we have the vector bundle $\mathcal{H}_{\mathcal{M}}$ containing \mathcal{E} as a subbundle, and also containing the trivial subbundle $\mathcal{H}_{-, \mathcal{M}}$. The projection $\mathcal{E} \cap \hbar\mathcal{H}_{-, \mathcal{M}} \rightarrow (\hbar\mathcal{H}_-/\mathcal{H}_-)_{\mathcal{M}}$ is then an isomorphism at 0, hence is an isomorphism in an open neighbourhood of 0. Replacing \mathcal{M} with this open neighbourhood, we obtain a trivialization

$$(2.16) \quad \tau : \left(\frac{\hbar\mathcal{H}_-}{\mathcal{H}_-} \right) \otimes_{\mathbb{C}} \mathcal{O}_{\mathcal{M}}\{\hbar\} \rightarrow \mathcal{E}$$

given, for $s \in \hbar\mathcal{H}_-$ with $s = s' + s''$ with s' a section of \mathcal{E} and s'' a section of $\mathcal{H}_{-, \mathcal{M}}$, by

$$\tau((s \bmod \mathcal{H}_-) \otimes f) = s' \cdot f.$$

PROPOSITION 2.24. *Given a choice of opposite subspace \mathcal{H}_- , we have:*

- (1) *Via the trivialization (2.16), the connection ∇ on \mathcal{E} yields a connection ∇ on $(\hbar\mathcal{H}_-/\mathcal{H}_-) \otimes_{\mathbb{C}} \mathcal{O}_{\mathcal{M}}\{\hbar\}$. We can write this connection as*

$$\nabla = d + \hbar^{-1}A,$$

where A is an $\text{End}(\hbar\mathcal{H}_-/\mathcal{H}_-)$ -valued 1-form on \mathcal{M} . For a vector field X on \mathcal{M} , we write A_X for the corresponding section of the trivial bundle $\text{End}(\hbar\mathcal{H}_-/\mathcal{H}_-)_{\mathcal{M}}$.

- (2) *If \mathcal{H}_- is isotropic with respect to $\overline{\Omega}$, defined in (2.15), then we obtain a symmetric pairing on the vector space $\hbar\mathcal{H}_-/\mathcal{H}_-$ given by*

$$(s_1, s_2)_{\hbar\mathcal{H}_-/\mathcal{H}_-} = (s_1, s_2)_{\mathcal{E}}|_{\hbar=\infty}.$$

(Here on the right we are thinking of $s_1, s_2 \in \hbar\mathcal{H}_-$ as flat sections of of $\mathcal{E} \otimes_{\mathcal{O}_{\mathcal{M}}\{\hbar\}} \mathcal{O}_{\mathcal{M}}\{\hbar, \hbar^{-1}\}$, on which the pairing on \mathcal{E} extends linearly.) Furthermore, this pairing is non-degenerate and

$$(2.17) \quad (A_X s_1, s_2)_{\hbar\mathcal{H}_-/\mathcal{H}_-} = (s_1, A_X s_2)_{\hbar\mathcal{H}_-/\mathcal{H}_-}.$$

- (3) *If the semi-infinite variation of Hodge structure has a grading operator Gr which preserves \mathcal{H}_- , then with respect to the trivialization (2.16), we can write*

$$\text{Gr} = \hbar\partial_{\hbar} + E + \text{Gr}_0$$

where $\text{Gr}_0 \in \text{End}(\hbar\mathcal{H}_-/\mathcal{H}_-)$ is defined, for $s \in \hbar\mathcal{H}_-$, by

$$\text{Gr}_0(s \bmod \mathcal{H}_-) = \text{Gr}(s) \bmod \mathcal{H}_-.$$

Furthermore, for $s_1, s_2 \in \hbar\mathcal{H}_-/\mathcal{H}_-$,

$$-D(s_1, s_2)_{\hbar\mathcal{H}_-/\mathcal{H}_-} = (\text{Gr}_0(s_1), s_2)_{\hbar\mathcal{H}_-/\mathcal{H}_-} + (s_1, \text{Gr}_0(s_2))_{\hbar\mathcal{H}_-/\mathcal{H}_-}.$$

PROOF. (1) Let us calculate $\nabla_X(\tau(s \otimes 1))$, where $s \in \hbar\mathcal{H}_-$ represents a constant section of $(\hbar\mathcal{H}_-/\mathcal{H}_-)_\mathcal{M}$. As sections, we can write $s = s' + s''$, where $s' \in \mathcal{E}$, $s'' \in \mathcal{H}_-, \mathcal{M}$, so that $\tau(s \otimes 1) = s' = s - s''$. Note that under the embedding $\mathcal{E} \subseteq \mathcal{H}_\mathcal{M}$, the trivial connection on $\mathcal{H}_\mathcal{M}$ induces the connection ∇ on \mathcal{E} defining the semi-infinite variation of Hodge structure since constant sections of $\mathcal{H}_\mathcal{M}$ are flat sections of $\mathcal{E} \otimes_{\mathcal{O}_\mathcal{M}\{\hbar\}} \mathcal{O}_\mathcal{M}\{\hbar, \hbar^{-1}\}$. So

$$\hbar\nabla_X(\tau(s \otimes 1)) = \hbar\nabla_X(s - s'') = -\hbar X(s''),$$

and as $X(s'') \in \mathcal{H}_-, \mathcal{M}$, we have $\hbar X(s'') \in \hbar\mathcal{H}_-, \mathcal{M}$. On the other hand,

$$\hbar\nabla_X(\tau(s \otimes 1)) \in \mathcal{E},$$

so

$$\hbar X(s'') \in \mathcal{E} \cap \hbar\mathcal{H}_-, \mathcal{M} \cong (\hbar\mathcal{H}_-/\mathcal{H}_-)_\mathcal{M}.$$

Thus $\hbar\nabla_X$ defines a section $s \mapsto -\hbar X(s'')$ of the trivial bundle $\text{End}(\hbar\mathcal{H}_-/\mathcal{H}_-)_\mathcal{M}$ and A is the 1-form with values in this bundle given by $A_X(s) = -\hbar X(s'')$.

(2) First note that the pairing given on $\hbar\mathcal{H}_-/\mathcal{H}_-$ makes sense. Indeed, given $s_1, s_2 \in \hbar\mathcal{H}_-/\mathcal{H}_-$, let $s'_i = \tau(s_i \otimes 1)$ be the corresponding sections of \mathcal{E} with $s_i = s'_i + s''_i$ as usual. Since s'_i is a section of \mathcal{E} , $(s'_1, s'_2)_\mathcal{E} \in \mathcal{O}_\mathcal{M}\{\hbar\}$. On the other hand, if t_1, t_2 are sections of $\hbar\mathcal{H}_-, \mathcal{M}$, we can write, for any $a_1, a_2 > 0$, $t_i = \hbar^{a_i} t'_i$ with t'_i a section of \mathcal{H}_- , so using the fact that \mathcal{H}_- is isotropic with respect to $\overline{\Omega}$, we have

$$0 = \overline{\Omega}(t'_1, t'_2) = \pm \text{Res}_{\hbar=0} \hbar^{-a_1 - a_2} (t_1, t_2)_\mathcal{E} d\hbar$$

from which we conclude that $(t_1, t_2)_\mathcal{E}$ only has terms of order ≤ 0 in \hbar . Similarly, if $t_1 \in \hbar\mathcal{H}_-, \mathcal{M}$, $t_2 \in \mathcal{H}_-, \mathcal{M}$, then $(t_1, t_2)_\mathcal{E}$ only has terms of order ≤ -1 in \hbar . Thus $(s'_1, s'_2)_\mathcal{E} = (s_1 - s''_1, s_2 - s''_2)_\mathcal{E}$ is in fact a section of $\mathcal{O}_\mathcal{M}$, and $(s'_1, s'_2)_\mathcal{E} = (s'_1, s'_2)_\mathcal{E}|_{\hbar=\infty}$ makes sense and is well-defined independently of the choice of representative for s_1, s_2 in $\hbar\mathcal{H}_-$. In addition, this is the same as $(s_1, s_2)_\mathcal{E}|_{\hbar=\infty}$, as the terms in this latter expression with non-negative powers of \hbar agree with the corresponding terms in $(s'_1, s'_2)_\mathcal{E}$. Finally, note that since s_1, s_2 are flat sections, it follow from Definition 2.20, (3) that $(s_1, s_2)_\mathcal{E}$ is in fact constant on \mathcal{M} , and hence $(s_1, s_2)_\mathcal{E}|_{\hbar=\infty}$ gives a well-defined element of \mathbb{C} .

In particular, again by Definition 2.20, (3),

$$\begin{aligned} 0 &= \hbar Y(s'_1, s'_2)_\mathcal{E} = -(\hbar\nabla_Y s'_1, s'_2)_\mathcal{E} + (s'_1, \hbar\nabla_Y s'_2)_\mathcal{E} \\ &= -(A_Y s_1, s_2)_{\hbar\mathcal{H}_-/\mathcal{H}_-} + (s_1, A_Y s_2)_{\hbar\mathcal{H}_-/\mathcal{H}_-}, \end{aligned}$$

by (1), giving the desired formula.

Non-degeneracy of the pairing on $\hbar\mathcal{H}_-/\mathcal{H}_-$ follows from the non-degeneracy of Definition 2.20, (4).

For (3), first note that by Definition 2.20, (6), Gr takes flat sections of \mathcal{E} to flat sections of \mathcal{E} . Thus Gr acts naturally on \mathcal{H} . Furthermore, by Definition 2.20, (5), if Gr preserves \mathcal{H}_- , it also preserves $\hbar\mathcal{H}_-$. Thus Gr defines a well-defined endomorphism Gr_0 of $\hbar\mathcal{H}_-/\mathcal{H}_-$. The description of Gr in terms of Gr_0 then follows from Definition 2.20, (5):

$$\begin{aligned} \text{Gr}(\tau(s \otimes f)) &= \text{Gr}(f \cdot s') = ((\hbar\partial_\hbar + E)f)s' + f \text{Gr}(s') \\ &= \tau(s \otimes (\hbar\partial_\hbar + E)f) + \text{Gr}_0(s \otimes f). \end{aligned}$$

The last statement follows immediately from Definition 2.20, (7). \square

DEFINITION 2.25. A semi-infinite variation of Hodge structure $(\mathcal{M}, \mathcal{E}, \nabla)$ is *miniversal* if there is a section s_0 of \mathcal{E} such that the $\mathcal{O}_{\mathcal{M}}$ -module homomorphism

$$\mathcal{T}_{\mathcal{M}} \rightarrow \mathcal{E}/\hbar\mathcal{E}$$

given by

$$X \mapsto \hbar\nabla_X s_0$$

is an isomorphism.

In the miniversal case, Proposition 2.24 along with one additional choice of data induces a Frobenius manifold structure on \mathcal{M} . More precisely, suppose we are given the following data:

- A semi-infinite variation of Hodge structure $\mathcal{E}, \nabla, (\cdot, \cdot)_{\mathcal{E}}$ with a grading Gr .
- An opposite subspace $\mathcal{H}_- \subseteq \mathcal{H}$ which is isotropic with respect to $\overline{\Omega}$ and preserved by Gr .
- An element $\Omega_0 \in \hbar\mathcal{H}_-$ which represents $[\Omega_0] \in \hbar\mathcal{H}_-/\mathcal{H}_-$, an eigenvector of Gr_0 . Furthermore, the corresponding section $s'_0 = \tau(\Omega_0 \otimes 1)$ of \mathcal{E} yields miniversality.

Then viewing \mathcal{E} as a subbundle of $\mathcal{H}_{\mathcal{M}}$ via the moving subspace construction, we identify s'_0 with a section of $\mathcal{H}_{\mathcal{M}}$ as usual, which can be viewed as a map

$$(2.18) \quad \Psi : \mathcal{M} \rightarrow \mathcal{H}.$$

This map can be described as follows. Because \mathcal{H} is a direct sum $\mathcal{H}_- \oplus \mathcal{E}_q$ for $q \in \mathcal{M}$, $\Omega_0 + \mathcal{H}_-$ intersects \mathcal{E}_q in a unique point, which is in fact $\Psi(q)$. This is Barannikov's *period map* of the semi-infinite variation of Hodge structure. We can then define a map

$$(2.19) \quad \psi : \mathcal{M} \rightarrow \hbar\mathcal{H}_-/\mathcal{H}_-$$

via

$$q \mapsto [\hbar(\Psi(q) - \Omega_0)].$$

Miniversality implies this is a local isomorphism. Indeed, at any point $q \in \mathcal{M}$, we need to check that the differential of ψ , $\psi_* : \mathcal{T}_{\mathcal{M},q} \rightarrow \hbar\mathcal{H}_-/\mathcal{H}_-$, is an isomorphism. Now the differential of Ψ is $\Psi_* : \mathcal{T}_{\mathcal{M},q} \rightarrow \mathcal{H}$ given by $X \mapsto \nabla_X s'_0 \in \mathcal{H}_-$, since, as was argued in the proof of Proposition 2.24, (1), the trivial connection on $\mathcal{H}_{\mathcal{M}}$ restricts to the connection ∇ on \mathcal{E} . Thus

$$\psi_*(X) = [\hbar\nabla_X s'_0] \in \hbar\mathcal{H}_-/\mathcal{H}_- \cong \mathcal{E}_q/\hbar\mathcal{E}_q.$$

So ψ_* is an isomorphism by miniversality. Note that

$$\psi_*(X) = A_X([\Omega_0]).$$

This now gives an identification of $\mathcal{T}_{\mathcal{M}}$ with the trivial bundle $(\hbar\mathcal{H}_-/\mathcal{H}_-)_{\mathcal{M}}$, given by

$$X \mapsto A_X([\Omega_0]).$$

This allows us to transport the product $(\cdot, \cdot)_{\hbar\mathcal{H}_-/\mathcal{H}_-}$ to $\mathcal{T}_{\mathcal{M}}$, which we call g . In addition, the trivial flat connection on $(\hbar\mathcal{H}_-/\mathcal{H}_-)_{\mathcal{M}}$ gives a flat connection $\nabla^{\mathcal{M}}$ on $\mathcal{T}_{\mathcal{M}}$.

We define a product on $\mathcal{T}_{\mathcal{M}}$ by the condition that

$$A_{X \circ Y}[\Omega_0] = A_X A_Y[\Omega_0]$$

and a tensor

$$\mathcal{A} : S^3 \mathcal{T}_{\mathcal{M}} \rightarrow \mathcal{O}_{\mathcal{M}}$$

given by

$$\mathcal{A}(X, Y, Z) = g(X \circ Y, Z).$$

That this is symmetric follows from flatness of ∇ , so that A_X and A_Y commute, and (2.17). Let e be the vector field on \mathcal{M} satisfying $A_e[\Omega_0] = [\Omega_0]$.

THEOREM 2.26. *\mathcal{M} with the data $(\nabla^{\mathcal{M}}, g, \mathcal{A})$ and vector fields e and E form a Frobenius manifold in which e is a flat identity and E is an Euler vector field satisfying (2.7) with constant $d_0 = 1$ and (2.6) with constant D being $2(\lambda + 1) + D$, where D is as in Definition 2.20, (7), and λ is the eigenvalue of $[\Omega_0]$ with respect to Gr_0 .*

PROOF. Clearly $(\nabla^{\mathcal{M}}, g, \mathcal{A})$ defines a pre-Frobenius manifold structure on \mathcal{M} . Also e is clearly an identity, flat since $[\Omega_0]$ is a constant, hence flat, section of $(\hbar\mathcal{H}_-/\mathcal{H}_-)_\mathcal{M}$.

By Theorem 2.11, it is enough to show that the first structure connection $\widehat{\nabla}^{\mathcal{M}}$, defined by

$$\begin{aligned}\widehat{\nabla}_X^{\mathcal{M}}(Y) &= \nabla_X^{\mathcal{M}}(Y) + \hbar^{-1}X \circ Y, \\ \widehat{\nabla}_{\hbar\partial_{\hbar}} &= \hbar\partial_{\hbar}Y - \hbar^{-1}E \circ Y + \text{Gr}_E(Y),\end{aligned}$$

is flat. Let us see what these operations correspond to on \mathcal{E} . First note that under the identification $\psi_* : \mathcal{T}_{\mathcal{M}} \rightarrow (\hbar\mathcal{H}_-/\mathcal{H}_-)_\mathcal{M}$, $\nabla^{\mathcal{M}}$ is the flat connection on $(\hbar\mathcal{H}_-/\mathcal{H}_-)_\mathcal{M}$, so

$$\begin{aligned}(\tau \circ \psi_*)(\widehat{\nabla}_X^{\mathcal{M}}(Y)) &= \tau(\psi_*(\nabla_X^{\mathcal{M}}(Y)) + \hbar^{-1}\psi_*(X \circ Y)) \\ &= \tau(d(\psi_*(Y))(X) + \hbar^{-1}A_{X \circ Y}([\Omega_0])) \\ &= \tau(d(A_Y([\Omega_0]))(X) + \hbar^{-1}A_X A_Y([\Omega_0])) \\ &= \nabla_X((\tau \circ \psi_*)(Y))\end{aligned}$$

by Proposition 2.24, (1).

Next, we compare Gr_E and Gr_0 . Let Y be a flat vector field, so that $\text{Gr}_E(Y) = [E, Y]$. Then

$$\begin{aligned}\text{Gr}_0(\psi_*(Y)) &= \text{Gr}_0(A_Y[\Omega_0]) \\ &= \text{Gr}(A_Y[\Omega_0]) \quad \text{mod } \mathcal{H}_- \\ &= \text{Gr}(\hbar\nabla_Y\Omega_0) \quad \text{mod } \mathcal{H}_- \\ &= \hbar\nabla_Y\Omega_0 + \hbar\text{Gr}(\nabla_Y\Omega_0) \quad \text{mod } \mathcal{H}_- \\ &= \hbar\nabla_Y\Omega_0 + \hbar\nabla_Y(\text{Gr}\Omega_0) + \hbar\nabla_{[E, Y]}\Omega_0 \quad \text{mod } \mathcal{H}_- \\ &= A_Y[\Omega_0] + \lambda A_Y[\Omega_0] + A_{\text{Gr}_E(Y)}[\Omega_0] \\ &= \psi_*((\lambda + 1 + \text{Gr}_E)(Y)).\end{aligned}$$

Here the third and sixth equalities hold from Proposition 2.24, (1), the fourth from Definition 2.20, (5), and the fifth equality follows from Definition 2.20, (6).

Consequently, for Y a flat vector field on \mathcal{M} and $f \in \mathcal{O}_{\mathcal{M}}\{\hbar\}$, we can compute

$$\begin{aligned}
\tau \circ \psi_*(\widehat{\nabla}_{\hbar\partial_{\hbar}} fY) &= \tau \circ \psi_*((\hbar\partial_{\hbar} f)Y - \hbar^{-1}fE \circ Y + f \operatorname{Gr}_E(Y)) \\
&= \tau((\hbar\partial_{\hbar} f)\psi_*(Y) - \hbar^{-1}fA_E\psi_*(Y) + f(\operatorname{Gr}_0 - \lambda - 1)\psi_*(Y)) \\
&= (\hbar\partial_{\hbar} f - f\nabla_E)(\tau \circ \psi_*)(Y) \\
&\quad + \tau(\operatorname{Gr}_0(\psi_*(fY))) - (\lambda + 1)\tau \circ \psi_*(fY) \\
&= (\hbar\partial_{\hbar} + E)(f)(\tau \circ \psi_*)(Y) + \tau(\operatorname{Gr}_0(\psi_*(fY))) \\
&\quad - \nabla_E(\tau \circ \psi_*)(fY) - (\lambda + 1)(\tau \circ \psi_*)(fY) \\
&= \operatorname{Gr}(\tau \circ \psi_*(fY)) - \nabla_E(\tau \circ \psi_*)(fY) - (\lambda + 1)(\tau \circ \psi_*)(fY),
\end{aligned}$$

the last equality by Proposition 2.24, (3). Thus $\widehat{\nabla}_{\hbar\partial_{\hbar}}$ coincides with $\operatorname{Gr} - \nabla_E - (\lambda + 1)$ on \mathcal{E} .

Now ∇ is a flat connection, and hence $\widehat{\nabla}^{\mathcal{M}}$ is flat in the directions tangent to \mathcal{M} . So we just need to show

$$(2.20) \quad [\widehat{\nabla}_{\hbar\partial_{\hbar}}^{\mathcal{M}}, \widehat{\nabla}_X^{\mathcal{M}}] = 0$$

for X a vector field on \mathcal{M} . Note that flatness of ∇ implies that for vector fields X and Y on \mathcal{M} ,

$$[\nabla_X, \nabla_Y] = \nabla_{[X, Y]}.$$

Transporting (2.20) to \mathcal{E} , we need to show that

$$[\operatorname{Gr} - \nabla_E - (\lambda + 1), \nabla_X] = 0.$$

Computing the left-hand side gives, by Definition 2.20, (6),

$$[\operatorname{Gr}, \nabla_X] - [\nabla_E, \nabla_X] = \nabla_{[E, X]} - \nabla_{[E, X]} = 0.$$

This shows that $\widehat{\nabla}^{\mathcal{M}}$ is flat, and hence \mathcal{M} is a Frobenius manifold.

Finally, we check the value of the constant D . We have, for X and Y flat vector fields on \mathcal{M} ,

$$\begin{aligned}
&E(g(X, Y)) - g([E, X], Y) - g(X, [E, Y]) \\
&= E((\psi_*(X), \psi_*(Y))_{\hbar\mathcal{H}_-/\mathcal{H}_-}) - (\psi_*(\operatorname{Gr}_E(X)), \psi_*(Y))_{\hbar\mathcal{H}_-/\mathcal{H}_-} \\
&\quad - (\psi_*(X), \psi_*(\operatorname{Gr}_E(Y)))_{\hbar\mathcal{H}_-/\mathcal{H}_-} \\
&= -((\operatorname{Gr}_0 - (\lambda + 1))\psi_*(X), \psi_*(Y))_{\hbar\mathcal{H}_-/\mathcal{H}_-} \\
&\quad - (\psi_*(X), (\operatorname{Gr}_0 - (\lambda + 1))\psi_*(Y))_{\hbar\mathcal{H}_-/\mathcal{H}_-} \\
&= (D + 2(\lambda + 1))(\psi_*(X), \psi_*(Y))_{\hbar\mathcal{H}_-/\mathcal{H}_-} \\
&= (D + 2(\lambda + 1))g(X, Y).
\end{aligned}$$

□

Noting that the map ψ is a local isomorphism between \mathcal{M} and the vector space $\hbar\mathcal{H}_-/\mathcal{H}_-$, as mentioned earlier, this gives a linear structure on \mathcal{M} . The flat vector fields with respect to $\nabla^{\mathcal{M}}$ are precisely the constant vector fields on $\hbar\mathcal{H}_-/\mathcal{H}_-$. Thus linear coordinates on $\hbar\mathcal{H}_-/\mathcal{H}_-$ induce coordinates on \mathcal{M} whose associated vector fields are flat. We call these induced coordinates *flat coordinates*, given to us canonically by the choice of data above.

EXAMPLE 2.27. We return to Example 2.21 and its continuation Example 2.22 coming from quantum cohomology. The space \mathcal{H} is identified with

$$H^*(X, \mathbb{C}) \otimes_{\mathbb{C}} \mathbb{C}\{\hbar, \hbar^{-1}\},$$

and we take

$$\mathcal{H}_- := H^*(X, \mathbb{C}) \otimes_{\mathbb{C}} \hbar^{-1} \mathcal{O}(\mathbb{P}^1 \setminus \{0\}) \subseteq \mathcal{H}.$$

Recall from Example 2.22 that the embedding of $\mathcal{E} = H^*(X, \mathbb{C}) \otimes_{\mathbb{C}} \mathbb{C}\{\hbar\}$ in $\mathcal{H}_{\widetilde{\mathcal{M}}}$ is given by the map \mathbb{J} , which takes flat sections of \mathcal{E} to constant sections of $\mathcal{H}_{\widetilde{\mathcal{M}}}$. At the point x of $\widetilde{\mathcal{M}}$ with coordinates $y_0 = \cdots = y_m = 0$, we have

$$\mathbb{J}(\alpha) = \alpha + O(\hbar^{-1}),$$

and hence the image of the fibre of \mathcal{E}_x under \mathbb{J} intersects \mathcal{H}_- only at 0. Thus $\mathcal{H} = \mathcal{H}_- \oplus \mathcal{E}_x$.

The induced isomorphism

$$\tau : (\hbar\mathcal{H}_-/\mathcal{H}_-) \otimes_{\mathbb{C}} \mathcal{O}_{\widetilde{\mathcal{M}}}\{\hbar\} \rightarrow \mathcal{E}$$

is a map

$$\tau : H^*(X, \mathbb{C}) \otimes_{\mathbb{C}} \mathcal{O}_{\widetilde{\mathcal{M}}}\{\hbar\} \rightarrow \mathcal{E},$$

which is defined by

$$\tau(s \otimes f) = f \cdot s',$$

where $s = s' + s''$ with s' a section of $\mathcal{E} \subseteq \mathcal{H}_{\widetilde{\mathcal{M}}}$ and s'' a section of $\mathcal{H}_{-, \widetilde{\mathcal{M}}}$. We do this as

$$s' = \mathbb{J}(s), \quad s'' = s - \mathbb{J}(s),$$

noting from the explicit formula for \mathbb{J} that $s - \mathbb{J}(s)$ is a section of $\mathcal{H}_{-, \widetilde{\mathcal{M}}}$ and $\mathbb{J}(s) \in \mathcal{E} \subseteq \mathcal{H}_{\widetilde{\mathcal{M}}}$ since \mathbb{J} gives the embedding of \mathcal{E} in $\mathcal{H}_{\widetilde{\mathcal{M}}}$.

To check that \mathcal{H}_- is isotropic with respect to $\overline{\Omega}$, we first need to extend the pairing on \mathcal{E} to elements of \mathcal{H}_- . As $\mathbb{J} : \mathcal{E} \otimes_{\mathbb{C}\{\hbar\}} \mathbb{C}\{\hbar, \hbar^{-1}\} \rightarrow \mathcal{H}_{\widetilde{\mathcal{M}}}$ gives the identification of $\mathcal{E} \otimes_{\mathbb{C}\{\hbar\}} \mathbb{C}\{\hbar, \hbar^{-1}\}$ with the bundle $\mathcal{H}_{\widetilde{\mathcal{M}}}$, we interpret $(s_1, s_2)_{\mathcal{E}}$ for $s_1, s_2 \in \mathcal{H}_{\widetilde{\mathcal{M}}}$ as $(\mathbb{J}^{-1}s_1, \mathbb{J}^{-1}s_2)_{\mathcal{E}}$. But $\mathbb{J}^{-1} = S$ by definition, and S^{-1} is the adjoint to S under this pairing by the argument of Proposition 2.23, so

$$(\mathbb{J}^{-1}s_1, \mathbb{J}^{-1}s_2)_{\mathcal{E}} = (s_1, s_2)_{\mathcal{E}}.$$

It is then clear that if $s_1, s_2 \in \mathcal{H}_-$, then

$$(s_1, s_2)_{\mathcal{E}} = (\text{constant})\hbar^{-2} + \text{higher order terms in } \hbar^{-1},$$

so $\overline{\Omega}(s_1, s_2) = 0$. Hence \mathcal{H}_- is isotropic with respect to $\overline{\Omega}$.

Next note that \mathcal{H}_- is preserved by Gr . Again, Gr is defined on \mathcal{E} , so given $s \in \mathcal{H}_-$, we should interpret $\text{Gr}(s)$ as $\mathbb{J} \circ \text{Gr} \circ \mathbb{J}^{-1}(s) = \mathbb{J} \circ \text{Gr} \circ S(s)$, which lies in \mathcal{H}_- by Proposition 2.18 and the definition of Gr . Furthermore, again by Proposition 2.18, note for $s \in \hbar\mathcal{H}_-$,

$$(2.21) \quad \text{Gr}(s) = \text{Gr}_E(s) \pmod{\mathcal{H}_-}.$$

In particular, the hypotheses of Proposition 2.24, (1)-(3) are satisfied.

We will now use this to produce a Frobenius manifold structure on $\widetilde{\mathcal{M}}$, by choosing $[\Omega_0] = T_0 \in \hbar\mathcal{H}_-/\mathcal{H}_-$ represented by $T_0 \in \hbar\mathcal{H}_-$. We need to verify:

- The corresponding section $\mathbb{J}(T_0)$ of \mathcal{E} yields miniversality. Note that $\mathbb{J}(T_0)$ is the description of this section of \mathcal{E} as a subbundle of $\mathcal{H}_{\widetilde{\mathcal{M}}}$. As an abstract bundle $\mathcal{E} \cong H^*(X, \mathbb{C}) \otimes_{\mathbb{C}} \mathcal{O}_{\widetilde{\mathcal{M}}}\{\hbar\}$, with the embedding $\mathbb{J} : \mathcal{E} \rightarrow \mathcal{H}_{\widetilde{\mathcal{M}}}$. Under this identification, the section $\mathbb{J}(T_0)$ of \mathcal{E} as a subbundle of $\mathcal{H}_{\widetilde{\mathcal{M}}}$ is identified with the section $\mathbb{J}^{-1}(\mathbb{J}(T_0)) = T_0$ of \mathcal{E} as an abstract bundle. Then

$$\begin{aligned} \hbar \nabla_{\partial_{y_i}}(T_0) &= \hbar \partial_{y_i}(T_0) + T_0 * T_i \\ &= T_i. \end{aligned}$$

Clearly T_0, \dots, T_m form a basis for $\mathcal{E}/\hbar\mathcal{E}$, hence we obtain miniversality.

- $[\Omega_0]$ is an eigenvector of Gr_0 . Note, using (2.21) for the second equality, that

$$\begin{aligned} \text{Gr}_0[\Omega_0] &= \text{Gr}(T_0) \pmod{\mathcal{H}_-} \\ &= \text{Gr}_E(T_0) \pmod{\mathcal{H}_-} \\ &= -T_0 \pmod{\mathcal{H}_-} \\ &= -[\Omega_0]. \end{aligned}$$

So $[\Omega_0]$ is an eigenvector with eigenvalue -1 .

So the chosen data gives a Frobenius manifold structure on $\widetilde{\mathcal{M}}$. Note also that

$$\psi : \widetilde{\mathcal{M}} \rightarrow \hbar\mathcal{H}_-/\mathcal{H}_-$$

is given by

$$(2.22) \quad \begin{aligned} (y_0, \dots, y_m) \in \widetilde{\mathcal{M}} &\mapsto \hbar(\mathbb{J}(T_0) - T_0) \pmod{\mathcal{H}_-} \\ &= \sum_{i=0}^m y_i T_i \end{aligned}$$

by the formula for \mathbb{J} of Proposition 2.23.

Let us check that this induced Frobenius manifold structure is in fact the Frobenius manifold of quantum cohomology. Indeed, from the definition of A_X , $A_{\partial_{y_i}} T_j = T_i * T_j$. Thus ψ_* is the identity, the multiplication \circ coincides with that given by $*$, and the connection $\nabla^{\widetilde{\mathcal{M}}}$ is the flat connection. The metric is clearly the correct one, T_0 is the identity, and the Euler vector field E is the standard one.

We see from the above discussion that the function $\mathbb{J}(T_0)$ plays a special role. This function has a name:

DEFINITION 2.28. The *Givental J-function* of X is

$$J_X := \mathbb{J}(T_0).$$

REMARK 2.29. Note that J_X satisfies the property

$$\hbar \partial_{y_i} J_X = \mathbb{J}(T_i).$$

This can be verified either by direct calculation, or by noting that, as \mathbb{J} gives the embedding of \mathcal{E} in $\mathcal{H}_{\widetilde{\mathcal{M}}}$,

$$\mathbb{J}(T_i) = \mathbb{J}(T_0 * T_i) = \mathbb{J}(\hbar \nabla_{\partial_{y_i}} T_0) = \hbar \partial_{y_i} \mathbb{J}(T_0) = \hbar \partial_{y_i} J_X.$$

Hence the function J_X , along with the pairing and the grading operator, completely determines the semi-infinite variation of Hodge structure.

EXAMPLE 2.30. Let us write down the Givental J -function for $X = \mathbb{P}^2$, using the formula of Proposition 2.23. For $\beta \neq 0$, the Fundamental Class Axiom says that

$$\langle T_0, T_2^m, \psi^\nu T_i \rangle_{0,\beta} = \langle T_2^m, \psi^{\nu-1} T_i \rangle_{0,\beta}.$$

Furthermore, by (2.2), this is non-zero unless $m + 3\beta = 2m + i + \nu - 1$, i.e., $m = 3\beta - i - \nu + 1$. On the other hand, for $\beta = 0$,

$$\langle T_0, T_2^m, \psi^\nu T_i \rangle_{0,0} = 0$$

unless $m + 1 = 2m + i + \nu$, i.e., $m = 1 - i - \nu$. But m must be at least 1 in order for this to be non-zero, since $\overline{\mathcal{M}}_{0,2}(X, 0)$ is empty, and hence we only get a contribution when $i = \nu = 0$. In particular, $\langle T_0, T_2, T_0 \rangle_{0,0} = \int_X T_0 \cup T_2 \cup T_0 = 1$. Putting this together into the formula of Proposition 2.23 and using $T^i = T_{2-i}$, we get

$$J_{\mathbb{P}^2} = e^{(y_0 T_0 + y_1 T_1)/\hbar} \cup \left(T_0 + \sum_{i=0}^2 \left(y_2 \hbar^{-1} \delta_{2,i} + \sum_{d \geq 1} \sum_{\nu \geq 0} \langle T_2^{3d+i-2-\nu}, \psi^\nu T_{2-i} \rangle_{0,d} \hbar^{-(\nu+2)} e^{dy_1} \frac{y_2^{3d+i-2-\nu}}{(3d+i-2-\nu)!} \right) T_i \right).$$

2.2. The B-model

We will now consider the mirror to \mathbb{P}^n , and explain the B-model version of the structures explored in the previous section.

The mirror to a Calabi-Yau manifold is a Calabi-Yau manifold. On the other hand, \mathbb{P}^n is a Fano manifold, and the mirror to a Fano manifold X is what is known as a *Landau-Ginzburg model*. In mathematical terms, this is a pair (\check{X}, W) where \check{X} is a variety and $W : \check{X} \rightarrow \mathbb{C}$ is a regular function. The regular function, known as the *Landau-Ginzburg potential*, plays a crucial role throughout.

Mirrors of toric Fano manifolds were first described by Givental in [33] and a physics derivation was given by Hori and Vafa in [58]. A crucial point will be the construction of a Frobenius manifold from the data (\check{X}, W) via the intermediary of a semi-infinite variation of Hodge structure. Mirror symmetry in this case is then the prediction that this B-model Frobenius manifold coincides with the Frobenius manifold arising from quantum cohomology of X . The construction of a Frobenius manifold from (\check{X}, W) , in some specific cases, was accomplished by Sabbah [103] and Barannikov [4]. We shall follow Barannikov's approach here. This should be viewed as a generalization of work of Kyoji Saito, who demonstrated in [104] how to associate a Frobenius manifold structure to a germ of a function W on \mathbb{C}^n .

Since we will deal primarily with the case of \mathbb{P}^n , and more specifically with \mathbb{P}^2 , let us give the mirror explicitly in this case so the reader can keep an example in mind. This construction will be explored in more detail in §2.2.3 and Chapter 5, but for the moment, the following will do.

We describe the mirror \check{X} to \mathbb{P}^n as

$$(2.23) \quad \check{X} = (\mathbb{C}^\times)^n \subseteq \mathbb{C}^{n+1}$$

given by the equation

$$x_0 \cdots x_n = 1,$$

where x_0, \dots, x_n are coordinates on \mathbb{C}^{n+1} . The potential W is then given by

$$W = x_0 + \dots + x_n.$$

2.2.1. The twisted de Rham complex. We now fix a non-singular variety X and a regular function $W : X \rightarrow \mathbb{C}$, a Landau-Ginzburg potential. The first question to consider is: what is the relevant cohomology group associated to (X, W) ? The answer takes several forms.

Consider the *twisted de Rham complex*

$$(\Omega_X^\bullet, d + dW \wedge).$$

This is the complex

$$\Omega_X^0 \rightarrow \Omega_X^1 \rightarrow \Omega_X^2 \rightarrow \dots$$

where the differential is $\omega \mapsto d\omega + dW \wedge \omega$. Here Ω_X^p is the sheaf of p -forms on X . There are two possible interpretations for Ω_X^p here: it could be the sheaf of algebraic p -forms in the Zariski topology, or it could be the sheaf of holomorphic p -forms in the analytic topology. We need the following theorem, originally due to Barannikov and Kontsevich, though unpublished, and then Sabbah [102] and Ogus-Vologodsky [89].

THEOREM 2.31. *If $W : X \rightarrow \mathbb{C}$ is projective, then*

$$\mathbb{H}_{\text{Zar}}^i(X, (\Omega_X^\bullet, d + dW \wedge)) \cong \mathbb{H}_{\text{Zar}}^i(X, (\Omega_X^\bullet, dW \wedge)) \cong \mathbb{H}_{\text{An}}^i(X, (\Omega_X^\bullet, d + dW \wedge)).$$

Here the subscripts *Zar* and *An* mean we are dealing with the algebraic version in the Zariski topology or the holomorphic version in the analytic topology. Furthermore, \mathbb{H}^i denotes hypercohomology.

The examples of Landau-Ginzburg potentials we are interested in, say the mirror to \mathbb{P}^n , are in fact not projective, but this turns out not to be too important: we can often find partial compactifications $X \subseteq \overline{X}$ such that W extends to $W : \overline{X} \rightarrow \mathbb{C}$ projective. While in general there may be some subtle difference between working with X versus \overline{X} , in fact in the cases we consider there will be no difference. We discuss this in detail in §3.5.

EXAMPLE 2.32. Assume we are in the situation of the above theorem and that W only has isolated critical points. Then it is not difficult to compute $\mathbb{H}^i(X, (\Omega_X^\bullet, dW \wedge))$ using the hypercohomology spectral sequence

$$E_2^{p,q} = H^p(X, \mathcal{H}^q(\Omega^\bullet, dW \wedge)) \Rightarrow \mathbb{H}^n(X, (\Omega_X^\bullet, dW \wedge)).$$

Here H^p is sheaf cohomology while \mathcal{H}^q denotes the cohomology of the complex $(\Omega_X^\bullet, dW \wedge)$.

First, let's compute $\mathcal{H}^q(\Omega_X^\bullet, dW \wedge)$. It is a basic fact of multilinear algebra that if V is a finite-dimensional vector space and $v \in V$ is non-zero, then the complex

$$\bigwedge^0 V \xrightarrow{\wedge v} \bigwedge^1 V \xrightarrow{\wedge v} \bigwedge^2 V \xrightarrow{\wedge v} \dots$$

is exact. Since $dW = 0$ exactly on the critical locus of W , which we call $\text{Crit}(W)$, the sheaf $\mathcal{H}^q(\Omega_X^\bullet, dW \wedge)$ is supported on $\text{Crit}(W)$. Furthermore, if $q < \dim X =: n$, then we will show inductively that $\mathcal{H}^q(\Omega_X^\bullet, dW \wedge) = 0$. Certainly $\mathcal{H}^0(\Omega_X^\bullet, dW \wedge) = 0$, since it is on the one hand supported on $\text{Crit}(W)$ and on the other hand contained

in the locally free sheaf Ω_X^0 . If $\mathcal{H}^{q'}(\Omega_X^\bullet, dW\wedge) = 0$ for $q' < q$, then we obtain a resolution

$$0 \rightarrow \Omega_X^0 \xrightarrow{dW\wedge} \cdots \xrightarrow{dW\wedge} \Omega_X^{q-1} \rightarrow \ker(\Omega_X^q \rightarrow \Omega_X^{q+1}) \rightarrow \mathcal{H}^q(\Omega_X^\bullet, dW\wedge) \rightarrow 0.$$

Now let $\mathcal{F} = \text{coker}(\Omega_X^{q-2} \xrightarrow{dW\wedge} \Omega_X^{q-1})$; the sheaf \mathcal{F} has projective dimension at most $q-1$, hence depth at least $n - (q-1) > 1$ (see [57], Proposition III 6.12A). On the other hand, $\ker(\Omega_X^q \rightarrow \Omega_X^{q+1})$ has depth at least 1, since it is torsion-free. But then, provided $\mathcal{H}^q(\Omega_X^\bullet, dW\wedge)$ is non-zero, the depth of $\mathcal{H}^q(\Omega_X^\bullet, dW\wedge)$ is at least 1, (see [57], III Exercise 3.4), contradicting it having support on the zero-dimensional set $\text{Crit}(W)$. Thus $\mathcal{H}^q(\Omega_X^\bullet, dW\wedge) = 0$.

Next, examining $\mathcal{H}^n(\Omega_X^\bullet, dW\wedge)$, if we have local coordinates x_1, \dots, x_n on X near a critical point, then we have $\Omega_X^{n-1} \rightarrow \Omega_X^n$ given by

$$dx_1 \wedge \cdots \wedge \widehat{dx_i} \wedge \cdots \wedge dx_n \mapsto \pm \frac{\partial W}{\partial x_i} dx_1 \wedge \cdots \wedge dx_n.$$

Hence the image of this map is locally just the Jacobian ideal of W .

Since $\mathcal{H}^q(\Omega_X^\bullet, dW\wedge)$ is now non-zero only for $q = n$ and for $q = n$ the support of this sheaf is zero-dimensional, the spectral sequence degenerates at the E_2 term and we get

$$\mathbb{H}^p(X, (\Omega_X^\bullet, dW\wedge)) = \begin{cases} 0 & p \neq n \\ \Gamma(X, \text{coker}(\Omega_X^{n-1} \xrightarrow{dW\wedge} \Omega_X^n)) & p = n. \end{cases}$$

□

EXAMPLE 2.33. Suppose (X, W) is the mirror to \mathbb{P}^n given in (2.23), with a slight modification. We in fact consider a possible variation of the Landau-Ginzburg potential by setting $x_0 \cdots x_n = \kappa$ where $\kappa \in \mathbb{C}^\times$, rather than $x_0 \cdots x_n = 1$. This gives a one-parameter family of Landau-Ginzburg models parameterized by κ . Of course, as mentioned before, we need to partially compactify X to apply the above theory, but in fact this won't add any additional critical points, so the computation will be the same. See §3.5 for more details on this partial compactification and computation. So without the partial compactification, we trivialize Ω_X^n via the n -form

$$\Omega = \frac{dx_1 \wedge \cdots \wedge dx_n}{x_1 \cdots x_n},$$

and then

$$\mathbb{H}^n(X, (\Omega_X^\bullet, dW\wedge)) = \frac{\mathbb{C}[x_1^\pm, \dots, x_n^\pm]}{(\partial W/\partial x_1, \dots, \partial W/\partial x_n)} \cdot \Omega.$$

Here we eliminate the variable $x_0 = \kappa/(x_1 \cdots x_n)$. The ring $\mathbb{C}[x_1^\pm, \dots, x_n^\pm]/\text{Jac}(W)$, where $\text{Jac}(W) = (\partial W/\partial x_1, \dots, \partial W/\partial x_n)$ is the Jacobian ideal, is known as the *Milnor ring* of W .

It is a finite dimensional \mathbb{C} -algebra; let us compute it. It is more convenient to use the generating set $x_i \partial W/\partial x_i$ of the Jacobian ideal, noting that

$$x_i \frac{\partial W}{\partial x_i} = x_i - \frac{\kappa}{x_1 \cdots x_n}.$$

Thus the critical points of W occur when

$$x_1 = \cdots = x_n = \frac{\kappa}{x_1 \cdots x_n},$$

i.e., $(x_1, \dots, x_n) = (\mu, \dots, \mu)$ with $\mu = \mu^{-n}\kappa$, i.e., μ is an $(n+1)$ -st root of κ . Furthermore, the Milnor ring has no nilpotents. So it is of dimension $n+1$ as a \mathbb{C} -vector space, the same as the dimension of $H^*(\mathbb{P}^n, \mathbb{C})$. \square

Returning to the twisted de Rham complex, let us note that we can modify this twisting to get a whole family of twisted de Rham complexes parameterized by $\hbar \in \mathbb{C}^\times$. We take the complex $(\Omega_X^\bullet, d + \hbar^{-1}dW \wedge)$. Of course, the dimension of $\mathbb{H}^p(X, (\Omega_X^\bullet, d + \hbar^{-1}dW \wedge))$ is independent of \hbar , as the complex $(\Omega_X^\bullet, \hbar^{-1}dW \wedge)$ is isomorphic to $(\Omega_X^\bullet, dW \wedge)$.

2.2.2. Homology. Given a Landau-Ginzburg model (X, W) , we obtain a natural homology theory dual to the cohomology of the twisted de Rham complex. To simplify the discussion, we will assume as in Example 2.32 that $\text{Crit}(W)$ is zero-dimensional, so that only $\mathbb{H}^n(X, (\Omega_X^\bullet, d + \hbar^{-1}dW \wedge))$ is non-zero, where $n = \dim X$. Furthermore, if X is affine, any element of $\mathbb{H}^n(X, (\Omega_X^\bullet, d + \hbar^{-1}dW \wedge))$ is represented by an n -form. Indeed, we can see this via a second hypercohomology spectral sequence,

$$E_1^{p,q} = H^p(H^q(X, \Omega_X^\bullet), d + \hbar^{-1}dW \wedge) \Rightarrow \mathbb{H}^n(X, (\Omega_X^\bullet, d + \hbar^{-1}dW \wedge)).$$

where we take the q -th sheaf cohomology of each entry in the complex Ω_X^\bullet , and then take the p -th cohomology of the resulting complex. Since X is affine, $H^q(X, \Omega_X^r) = 0$ for $q > 0$, and thus the spectral sequence degenerates, showing that

$$\mathbb{H}^n(X, (\Omega_X^\bullet, d + \hbar^{-1}dW \wedge)) \cong H^n(\Gamma(X, \Omega_X^\bullet), d + \hbar^{-1}dW \wedge).$$

In this situation, for a fixed $\hbar \in \mathbb{C}^\times$, we consider a homology group that we write as

$$H_n(X, \text{Re } W/\hbar \ll 0; \mathbb{C}).$$

By this, we mean a homology theory of cycles which are allowed to be unbounded, but only unbounded in the directions in which $\text{Re } W/\hbar \rightarrow -\infty$.

We will not be more precise here; it is possible to do so, but this for us is a technical point. The important point is that

$$H_n(X, \text{Re } W/\hbar \ll 0; \mathbb{C})$$

is naturally dual to

$$\mathbb{H}^n(X, (\Omega_X^\bullet, d + \hbar^{-1}dW \wedge)).$$

Indeed, given a cycle $\Xi \in H_n(X, \text{Re } W/\hbar \ll 0; \mathbb{C})$, and $\omega \in H^0(X, \Omega_X^n)$ algebraic representing an element of $\mathbb{H}^n(X, (\Omega_X^\bullet, d + \hbar^{-1}dW \wedge))$, we can compute the so-called oscillatory integral

$$\int_{\Xi} e^{W/\hbar} \omega.$$

With the right definition of $H_n(X, \text{Re } W/\hbar \ll 0; \mathbb{C})$, this integral converges precisely because $\text{Re } W/\hbar \rightarrow -\infty$ in the unbounded direction of Ξ . Furthermore, note that if η is an $(n-1)$ -form on X ,

$$d(e^{W/\hbar} \eta) = e^{W/\hbar} (d\eta + \hbar^{-1}dW \wedge \eta).$$

By Stokes' theorem,

$$\int_{\Xi} e^{W/\hbar} \omega = \int_{\Xi} e^{W/\hbar} \omega + d(e^{W/\hbar} \eta),$$

so ω and $\omega + (d + \hbar^{-1}dW \wedge)\eta$ give the same integrals. They also represent the same element of $\mathbb{H}^n(X, (\Omega_X^\bullet, d + \hbar^{-1}dW \wedge))$, so we obtain what turns out to be a natural perfect pairing

$$(2.24) \quad H_n(X, \operatorname{Re} W/\hbar \ll 0; \mathbb{C}) \times \mathbb{H}^n(X, (\Omega_X^\bullet, d + \hbar^{-1}dW \wedge)) \rightarrow \mathbb{C}.$$

There is also a natural intersection pairing

$$(2.25) \quad H_n(X, \operatorname{Re} W/\hbar \ll 0; \mathbb{C}) \times H_n(X, \operatorname{Re} W/(-\hbar) \ll 0; \mathbb{C}) \rightarrow \mathbb{C}.$$

Roughly, two cycles in the first and second homology groups on the left can only intersect in some bounded region of X , and the intersection number is well-defined as the cycles cannot be deformed so that intersection points run off to infinity.

EXAMPLE 2.34. We consider the mirror to \mathbb{P}^2 , with $(\check{X}, W) = ((\mathbb{C}^\times)^2, W)$. Here $(\mathbb{C}^\times)^2 = V(x_0x_1x_2 - 1) \subseteq \mathbb{C}^3$, and $W = x_0 + x_1 + x_2$. As we saw in Example 2.33, $\mathbb{H}^2(\check{X}, (\Omega_{\check{X}}^\bullet, d + \hbar dW \wedge))$ is three-dimensional. We will describe a basis Ξ_0, Ξ_1, Ξ_2 for $H_2(\check{X}, \operatorname{Re} W/\hbar \ll 0; \mathbb{C})$.

To describe this basis, let's identify $(\mathbb{C}^\times)^2$ with a trivial T^2 -bundle over \mathbb{R}^2 via the map

$$\begin{aligned} \operatorname{Log} : (\mathbb{C}^\times)^2 &\rightarrow \mathbb{R}^2 \\ \operatorname{Log}(x_1, x_2) &= (\log |x_1|, \log |x_2|). \end{aligned}$$

We take $\Xi_0 = \operatorname{Log}^{-1}(0, 0)$. This is a compact cycle, so certainly represents a class in $H_2(\check{X}, \operatorname{Re} W/\hbar \ll 0; \mathbb{C})$ for any \hbar . However, for this to actually be a cycle, we have to choose an orientation. Using the map Log , we can identify $(\mathbb{C}^\times)^2$ with $\mathbb{R}^2 \times T^2$, taking coordinates y_1, y_2 on \mathbb{R}^2 and coordinates θ_1, θ_2 on T^2 , so that $x_j = \exp(y_j + i\theta_j)$. Then we can orient Ξ_0 using $d\theta_1 \wedge d\theta_2$.

Next we define Ξ_1 . Let $\rho_0, \rho_1, \rho_2 \subseteq \mathbb{R}^2$ be the one-dimensional cones of the fan in Example 1.14. Let

$$\begin{aligned} S_0 &= \{(x_1, x_2) \in \operatorname{Log}^{-1}(\rho_0) \mid \arg(x_1^{-1}x_2^{-1}) = \arg(\hbar) + \pi\}, \\ S_1 &= \{(x_1, x_2) \in \operatorname{Log}^{-1}(\rho_1) \mid \arg x_1 = \arg(\hbar) + \pi\}, \\ S_2 &= \{(x_1, x_2) \in \operatorname{Log}^{-1}(\rho_2) \mid \arg x_2 = \arg(\hbar) + \pi\}. \end{aligned}$$

Each S_i is a cylinder with boundary on $\operatorname{Log}^{-1}(0, 0)$, which we depict as a square with opposite sides identified in Figure 2. We take S_3 to be a surface contained in $\operatorname{Log}^{-1}(0, 0)$ which bounds $\partial S_0 \cup \partial S_1 \cup \partial S_2$, and take Ξ_1 to be the piecewise smooth cycle

$$\Xi_1 = S_0 \cup S_1 \cup S_2 \cup S_3.$$

We orient, say, $S_1 \subseteq \Xi_1$ using $-dy_1 \wedge d\theta_2$. One can check that this extends to an orientation on all of Ξ_1 .

For example, Figure 2 depicts such a surface S_3 when $\hbar = 1$. To verify that Ξ_1 lives in $H_2(\check{X}, \operatorname{Re} W/\hbar \ll 0; \mathbb{C})$, note that on S_0 , $|x_1|$ and $|x_2|$ tend to zero, while $|x_0| = 1/|x_1x_2| \rightarrow \infty$ as we head in the non-compact direction of S_0 . As

$$\arg(x_1^{-1}x_2^{-1}) = \arg \hbar + \pi,$$

in fact $x_1^{-1}x_2^{-1}/\hbar$ is real and goes to $-\infty$. Thus this term dominates in W/\hbar , so $\operatorname{Re} W/\hbar \rightarrow -\infty$. The same argument works along S_1 and S_2 . Thus Ξ_1 lies in $H_2(\check{X}, \operatorname{Re} W/\hbar \ll 0; \mathbb{C})$.




FIGURE 2. The surface S_3 contained in a two-torus. The opposite sides of the square are identified.

Note that we had an arbitrary choice for S_3 , and this can always be modified by adding a multiple of $\text{Log}^{-1}(0, 0)$. In fact, if we make one choice, say for $\hbar = 1$, and vary \hbar continuously via $\hbar = e^{2\pi it}$, $0 \leq t \leq 1$, we can choose S_3 in a continuously varying way. As we complete the loop, we will find that Ξ_1 has been replaced by $\Xi_1 - 3\Xi_0$, if we orient Ξ_0 and Ξ_1 as indicated below. This demonstrates that there is no canonical choice for S_3 .

We will describe Ξ_2 only for $\hbar = -1$ to avoid complexity. Just take

$$\Xi_2 = \{(x_1, x_2) \in \check{X} \mid x_1, x_2 \text{ real}, x_1, x_2 > 0\}.$$

Then again $\text{Re } W/\hbar \rightarrow -\infty$ in the non-compact directions. We orient Ξ_2 using $dy_1 \wedge dy_2$,

We leave it to the reader to build a Ξ_2 which depends on \hbar as in the example of Ξ_1 . However, note that the Ξ_2 we have described actually works for any \hbar with $\text{Re } \hbar < 0$.

We can also calculate intersection numbers under the pairing (2.25). We just note the simplest one: the cycle Ξ_0 is independent of \hbar , and is a fibre of Log . Of course any two fibres are homologous. So $\Xi_0 \cdot \Xi_0 = 0$ and $\Xi_0 \cdot \Xi_1(\hbar) = 0$ for any \hbar , as Ξ_0 can be moved to be disjoint from $\Xi_1(\hbar)$. (Here we indicate the dependence of Ξ_1 on \hbar). On the other hand, Ξ_0 and Ξ_2 meet transversally at one point, and with the given choice of orientation, $\Xi_0 \cdot \Xi_2 = -1$, where we think of

$$\Xi_0 \in H_2(\check{X}, \text{Re } W \ll 0; \mathbb{C}) \text{ and } \Xi_2 \in H_2(\check{X}, \text{Re}(-W) \ll 0; \mathbb{C}).$$

□

The description of cycles given in the above example is not the conventional one; rather, we give it because these cycles are more closely related (but not identical) to the cycles mirror to T_0, T_1 and T_2 of $H^*(\mathbb{P}^2, \mathbb{C})$. The more conventional description of the generators of $H_n(X, \text{Re } W/\hbar \ll 0; \mathbb{C})$ is in terms of *Lefschetz thimbles*.

Suppose that the critical locus of W is reduced: the critical locus is given a scheme structure via the spectrum of the Milnor ring. Then each critical point is non-degenerate, in the sense that if X is n -dimensional, then in a neighbourhood of each critical point there are local holomorphic coordinates z_1, \dots, z_n such that $W = z_1^2 + \dots + z_n^2$. This is the case, for example, for the mirror (\check{X}, W) of \mathbb{P}^n .

Put a metric on X , and consider the function $\text{Re } W/\hbar$ as a Morse function on X . Locally near each critical point, there are real coordinates y_1, \dots, y_{2n} such that

$\operatorname{Re} W/\hbar = y_1^2 + \cdots + y_n^2 - y_{n+1}^2 - \cdots - y_{2n}^2$. So each critical point is a critical point of index n . The *stable manifold* Δ_p^+ of a critical point p is the union of gradient flow lines of $\operatorname{Re} W/\hbar$ which go to p as time $t \rightarrow +\infty$, and the *unstable manifold* Δ_p^- of a critical point p is the union of gradient flow lines of $\operatorname{Re} W/\hbar$ which go to p as time $t \rightarrow -\infty$. Of course, Δ_p^+ and Δ_p^- depend on \hbar , which we consider fixed in this discussion.

Both Δ_p^+ and Δ_p^- are n -dimensional submanifolds, and provided W is proper, clearly

$$\Delta_p^+ \in H_n(X, \operatorname{Re} W/\hbar \ll 0; \mathbb{C})$$

and

$$\Delta_p^- \in H_n(X, \operatorname{Re}(-W/\hbar) \ll 0; \mathbb{C}).$$

If W is not proper, we can choose a properification $X \subseteq \overline{X}$, $W : \overline{X} \rightarrow \mathbb{C}$ proper, and choose a metric so that $\Delta_p^\pm \subseteq X$ when $p \in \operatorname{Crit}(W)$ is contained in X . So we still get

$$\Delta_p^\pm \in H_n(X, \operatorname{Re}(\pm W/\hbar) \ll 0; \mathbb{C}).$$

These cycles are known as *Lefschetz thimbles*.

For a general choice of metric, $\Delta_p^\pm \cap \Delta_{p'}^\pm = \emptyset$ unless $p = p'$, and $\Delta_p^+ \cap \Delta_p^- = \{p\}$. Thus with the proper orientation, we can assume that under the intersection pairing (2.25),

$$(2.26) \quad \Delta_p^+ \cdot \Delta_{p'}^- = \delta_{pp'}.$$

We will need the following standard approximation for oscillatory integrals over these cycles describing the asymptotic behaviour of these integrals as $\hbar \rightarrow 0$. This approximation is known as the *stationary phase approximation*:

PROPOSITION 2.35. *Let X, W be as above, with W having only isolated non-degenerate critical points. Suppose $f(x, \hbar)$ is a holomorphic function on*

$$\check{X} \times \{|\hbar| < \epsilon\} \subseteq \check{X} \times \mathbb{C},$$

algebraic when restricted to $\check{X} \times \{\hbar\}$ for any \hbar with $|\hbar| < \epsilon$. Then for p a critical point of W ,

$$\int_{\Delta_p^+} f(x, \hbar) e^{W/\hbar} \omega = \frac{(-2\pi\hbar)^{n/2}}{\sqrt{\operatorname{Hess}(W)(p)}} e^{W(p)/\hbar} (f(p, 0) + O(\hbar))$$

where ω is a nowhere vanishing algebraic n -form on X , $n = \dim X$, f is a regular function, and the Hessian $\operatorname{Hess}(W)$ at p is $\det(\partial^2 W/\partial x_i \partial x_j)$, evaluated at p , with x_1, \dots, x_n local coordinates in which $\omega = dx_1 \wedge \cdots \wedge dx_n$.

PROOF. (Sketch) First note that we can write

$$\begin{aligned} \int_{\Delta_p^+} f(x, \hbar) e^{W/\hbar} \omega &= \int_{\Delta_p^+} f(x, \hbar) e^{(W-W(p))/\hbar} e^{W(p)/\hbar} \omega \\ &= e^{W(p)/\hbar} \int_{\Delta_p^+} f(x, \hbar) e^{(W-W(p))/\hbar} \omega, \end{aligned}$$

so we can assume $W(p) = 0$ at the expense of a factor of $e^{W(p)/\hbar}$. The main point now is that the asymptotic behaviour as $\hbar \rightarrow 0$ comes from the integral

$$(2.27) \quad \int_{\Delta_p^+ \cap B_\epsilon} f(x, \hbar) e^{W/\hbar} \omega,$$

where B_ϵ is a ball of radius ϵ centered at p . This is because

$$(2.28) \quad \int_{\Delta_p^+ \setminus B_\epsilon} f(x, \hbar) e^{W/\hbar} \omega \rightarrow 0$$

as $\hbar \rightarrow 0$, since $\operatorname{Re} W/\hbar$ is bounded away from zero and negative on $\Delta_p^+ \setminus B_\epsilon$. Thus $e^{W/\hbar} \rightarrow 0$ rapidly as $\hbar \rightarrow 0$, hence (2.28). The integral (2.27) can then be approximated by replacing $f(x, \hbar)$ by its value at $x = p, \hbar = 0$, and replacing W by the second order (leading) term in its Taylor expansion,

$$W = \sum_{i,j} \frac{1}{2} \frac{\partial^2 W}{\partial x_i \partial x_j} x_i x_j + \text{higher order terms.}$$

Letting $A = ((\partial^2 W / \partial x_i \partial x_j)(p))_{i,j}$, so that $\operatorname{Hess}(W)(p) = \det A$, we can make a change of coordinates $\mathbf{y} = \sqrt{A}\mathbf{x}$ for a choice of square root of A so that the second order part of W is now $(y_1^2 + \cdots + y_n^2)/2$ and locally near p ,

$$\omega = dy_1 \wedge \cdots \wedge dy_n / \sqrt{\operatorname{Hess}(W)(p)}.$$

Now take Δ_p^+ to be defined for $\hbar = 1$; we can then let $\hbar \rightarrow 0$ in an angular sector in the \hbar -plane containing 1. If we write $y_j = u_j + iv_j$, then

$$\operatorname{Re} W = \frac{1}{2}(u_1^2 + \cdots + u_n^2 - v_1^2 - \cdots - v_n^2),$$

so the stable manifold for $\operatorname{Re} W$ is given by $u_1 = \cdots = u_n = 0$. Thus, in this angular sector, we can approximate (2.27) by the standard Gaussian integral

$$\begin{aligned} & \frac{f(p, 0)}{\sqrt{\operatorname{Hess}(W)(p)}} \int_{\mathbb{R}^n} e^{-(v_1^2 + \cdots + v_n^2)/2\hbar} d(iv_1) \wedge \cdots \wedge d(iv_n) \\ &= \frac{f(p, 0)}{\sqrt{\operatorname{Hess}(W)(p)}} \sqrt{(2\pi\hbar)^n} i^n \\ &= f(p, 0) \frac{(-2\pi\hbar)^{n/2}}{\sqrt{\operatorname{Hess}(W)(p)}}. \end{aligned}$$

This gives the desired asymptotic behaviour. \square

2.2.3. The B-model semi-infinite variation of Hodge structure.

We first need to discuss the B-model moduli space on which the B-model semi-infinite variation of Hodge structure lives. In the case of the mirror of \mathbb{P}^n , we start with (\check{X}, W_0) , with $W_0 = x_0 + \cdots + x_n$, $\check{X} = V(x_0 \cdots x_n - 1) \subseteq \mathbb{C}^{n+1}$. The B-model moduli space is the *universal unfolding* of W_0 . In the context of complex manifolds, the universal unfolding is a germ of a complex manifold $0 \in \mathcal{M}$, along with a holomorphic function

$$W : \mathcal{M} \times \check{X} \rightarrow \mathbb{C},$$

such that:

- $W|_{\{0\} \times \check{X}} = W_0$.
- For any germ $0 \in Y$ of a complex space, and holomorphic function

$$W' : Y \times \check{X} \rightarrow \mathbb{C}$$

such that $W'|_{\{0\} \times \check{X}} = W_0$, there exists a commutative diagram

$$\begin{array}{ccc} Y \times \check{X} & \xrightarrow{\eta} & \mathcal{M} \times \check{X} \\ \downarrow & & \downarrow \\ Y & \xrightarrow{\varphi} & \mathcal{M} \end{array}$$

where $\varphi : (0 \in Y) \rightarrow (0 \in \mathcal{M})$ is a map of germs whose induced map on Zariski tangent spaces is unique, the vertical arrows are projections, $\eta|_{\{0\} \times \check{X}} : \{0\} \times \check{X} \rightarrow \{0\} \times \check{X}$ is the identity, and $W' = W \circ \eta$.

In fact, a standard fact (see e.g., [109]) says that the universal unfolding of W_0 , assuming W_0 has isolated critical points, is a germ of 0 in the Milnor ring,

$$\mathbb{C}[x_1^{\pm}, \dots, x_n^{\pm 1}] / (\partial W_0 / \partial x_1, \dots, \partial W_0 / \partial x_n),$$

thought of as an affine space. If f_0, \dots, f_n are elements of $\mathbb{C}[x_1^{\pm 1}, \dots, x_n^{\pm 1}]$ representing a basis for this vector space, giving coordinates t_0, \dots, t_n on a germ \mathcal{M} of the origin, then the equation for the universal unfolding W of W_0 on $\mathcal{M} \times \check{X}$ is

$$W := W_0 + \sum_{i=0}^n t_i f_i.$$

For example, we might take as a basis $f_i = W_0^i$, $0 \leq i \leq n$. We leave it to the reader to check that this is indeed a basis of the Milnor ring.

Ideally, we would like \mathcal{M} to be the space on which the B-model semi-infinite variation of Hodge structure lives. However, there is a problem which emerges when $n \geq 2$: when $t_i \neq 0$ for $i \geq 2$, $W_0 + \sum t_i W_0^i$ has more critical points than W_0 does.

For example, consider the case $n = 2$, and the perturbation

$$W = W_0 + t_2 W_0^2 = W_0(1 + t_2 W_0)$$

of W_0 for t_2 small. Then W factors as

$$\check{X} \xrightarrow{W_0} \mathbb{C} \xrightarrow{f} \mathbb{C},$$

with $f(x) = x(1 + t_2 x)$. Of course f is a $2 : 1$ cover, branched where $f'(x) = 0$, i.e., at $x = -1/2t_2$. So the fibres of \check{X} are disconnected in general, and $W_0 = -1/2t_2$ is a multiple fibre of W . This is drastically different behaviour than we wish, and is known as *wild* behaviour. Clearly as $t_2 \rightarrow 0$, the extra critical points move out to ∞ , which is why we didn't see this behaviour at $t_2 = 0$. But we need to avoid these sorts of potentials.

There are a number of different approaches to dealing with this problem. The approach we shall take here is to consider t_2, \dots, t_n only as formal variables. For example, we could take $\mathcal{M} = \text{Spec } \mathbb{C}[t_0, t_1][[t_2, \dots, t_n]]$. More in keeping with the discussion of the A-model, we will work directly with a moduli space $\widetilde{\mathcal{M}}$ similar to the one defined in §2.1.4.

DEFINITION 2.36. Define $\widetilde{\mathcal{M}}$ to be the ringed space $(\mathbb{C}, \mathcal{O}_{\widetilde{\mathcal{M}}})$ where \mathbb{C} is viewed as a complex manifold with coordinate t_1 , and sections of $\mathcal{O}_{\widetilde{\mathcal{M}}}$ over an open set $U \subseteq \mathbb{C}$ consist of formal power series

$$\sum f_{i_0 i_2 \dots i_n} t_0^{i_0} t_2^{i_2} \dots t_n^{i_n}$$

with $f_{i_0 i_2 \dots i_n}$ a holomorphic function on U . Let $\check{\mathcal{X}}$ be the subspace of $\widetilde{\mathcal{M}} \times \mathbb{C}^{n+1}$ defined by the equation

$$e^{t_1} = x_0 \cdots x_n,$$

with x_0, \dots, x_n coordinates on the \mathbb{C}^{n+1} factor. Setting

$$W_0 = x_0 + \cdots + x_n,$$

we can then define

$$(2.29) \quad W = t_0 + W_0 + \sum_{i=2}^n t_i W_0^i.$$

We denote by

$$\pi : \check{\mathcal{X}} \rightarrow \widetilde{\mathcal{M}}$$

the projection.

One can check that this is indeed universal in a neighbourhood of each point of $\widetilde{\mathcal{M}}$; the details of the particular choice of description of W won't be important for us. We will be in particular interested in the point $0 \in \widetilde{\mathcal{M}}$ where $t_0 = t_1 = \cdots = t_n = 0$.

We can then consider the subscheme $\text{Crit}(W) \subseteq \check{\mathcal{X}}$ of critical points of W . This is defined by the ideal generated by $\partial W / \partial x_0, \dots, \partial W / \partial x_n$. Computing, one finds

$$\frac{\partial W}{\partial x_i} = \frac{\partial W_0}{\partial x_i} \left(1 + \sum_{i=2}^n i t_i W_0^{i-1} \right).$$

Since t_2, \dots, t_n are formal parameters, $1 + \sum_{i=2}^n i t_i W_0^{i-1}$ is invertible, so the ideal generated by the $\partial W / \partial x_i$'s is the same as the ideal generated by the $\partial W_0 / \partial x_i$'s. Thus $\text{Crit}(W)$ is given by the locus $x_0 = \cdots = x_n = \mu$ for μ ranging over the $(n+1)$ -st roots of e^{t_1} , by Example 2.33. So $\text{Crit}(W)$ is étale over $\widetilde{\mathcal{M}}$, and $\pi_* \mathcal{O}_{\text{Crit}(W)}$ is a locally free sheaf of $\mathcal{O}_{\widetilde{\mathcal{M}}}$ -modules of rank $n+1$. On $\widetilde{\mathcal{M}}$, one can find sections $p_1, \dots, p_{n+1} : \widetilde{\mathcal{M}} \rightarrow \text{Crit}(W)$.

On $\widetilde{\mathcal{M}} \times \mathbb{C}^\times$, with coordinate \hbar on \mathbb{C}^\times , we define a local system R of \mathbb{C} -vector spaces as follows. We note that a local system is only concerned about the underlying topological space of $\widetilde{\mathcal{M}} \times \mathbb{C}^\times$, and has nothing to do with the structure sheaf $\mathcal{O}_{\widetilde{\mathcal{M}}}$. With this in mind, the fibre of R over $(t_1, \hbar) \in \mathbb{C} \times \mathbb{C}^\times$ will be

$$H_n(\pi^{-1}(t_1), \text{Re}(W|_{\pi^{-1}(t_1)})/\hbar) \ll 0; \mathbb{C}.$$

Here $W|_{\pi^{-1}(t_1)} = x_0 + \cdots + x_n$ with $x_0 \cdots x_n = e^{t_1}$, as we are ignoring the variables t_0, t_2, \dots, t_n . This function has $n+1$ critical points. Thus we do not need to worry about the dimension of this space jumping.

Let $\mathcal{R} = R \otimes_{\mathbb{C}} \mathcal{O}_{\widetilde{\mathcal{M}} \times \mathbb{C}^\times}$; this is a locally free sheaf of rank $n+1$ on $\widetilde{\mathcal{M}} \times \mathbb{C}^\times$, and carries a Gauss-Manin connection ∇^{GM} whose flat sections are the sections of R . We then have the dual locally free sheaf

$$\mathcal{R}^\vee = \text{Hom}_{\mathcal{O}_{\widetilde{\mathcal{M}} \times \mathbb{C}^\times}}(\mathcal{R}, \mathcal{O}_{\widetilde{\mathcal{M}} \times \mathbb{C}^\times})$$

with the dual local system $R^\vee \subseteq \mathcal{R}^\vee$, again given by flat sections of a connection ∇^{GM} .

Note that a section of \mathcal{R}^\vee on an open set $U \times V$, with $U \subseteq \widetilde{\mathcal{M}}$, $V \subseteq \mathbb{C}^\times$ open sets, can be given explicitly via differential forms

$$f\Omega,$$

where f is a holomorphic function on $\pi^{-1}(U) \times V$ with $f|_{\pi^{-1}(u) \times \{v\}}$ being algebraic for each $u \in U, v \in V$, and

$$\Omega = \frac{dx_1 \wedge \cdots \wedge dx_n}{x_1 \cdots x_n}.$$

The form $f\Omega$ defines a map $\mathcal{R} \rightarrow \mathcal{O}_{\widetilde{\mathcal{M}} \times \mathbb{C}^\times}$ over $U \times V$ by taking a section Ξ of \mathcal{R} over $U \times V$ to

$$\int_{\Xi} e^{W/\hbar} f\Omega \in \mathcal{O}_{\widetilde{\mathcal{M}} \times \mathbb{C}^\times}.$$

We write $[f\Omega]$ for the section of \mathcal{R}^\vee determined in this way.

Let us be more precise about what this means in the formal setting we are working in. Using the expression for W given in (2.29), we write

$$e^{W/\hbar} = e^{(x_0 + \cdots + x_n)/\hbar} g$$

where g is a formal power series in the variables t_0, t_2, \dots, t_n , with each coefficient appearing being an algebraic function in x_0, \dots, x_n and \hbar . Thus the integral of each term in this expansion converges, being

$$\int_{\Xi} e^{(x_0 + \cdots + x_n)/\hbar} \hbar \Omega$$

for h satisfying $h|_{\pi^{-1}(u) \times \{v\}}$ algebraic for each $u \in U, v \in V$. Indeed, the exponential decay of $e^{(x_0 + \cdots + x_n)/\hbar}$ dominates the polynomial growth of h . Thus we see that $f\Omega$ defines an element of \mathcal{R}^\vee .

There is a standard technique to describe the Gauss-Manin connection at the level of forms, which we sketch here. Given an n -form $f\Omega$ on $\pi^{-1}(U) \times V$, $f\Omega$ represents a section of \mathcal{R}^\vee . We now consider the total space $\check{X} \times \mathbb{C}^\times$ as a space with the Landau-Ginzburg potential $\check{h}^{-1}W$. If in fact $f\Omega$ were closed with respect to $d + d(\check{h}^{-1}W) \wedge$ on $\pi^{-1}(U) \times V$, then $f\Omega$ would in fact represent a *flat* section of \mathcal{R}^\vee . The Gauss-Manin connection measures the failure of $(d + d(\check{h}^{-1}W) \wedge)(f\Omega)$ to be zero on the total space by evaluating this $(n+1)$ -form on tangent vectors X on $\check{X} \times \mathbb{C}^\times$ lifted from $\widetilde{\mathcal{M}} \times \mathbb{C}^\times$. In particular, if X is in fact the lift of a tangent vector from $\widetilde{\mathcal{M}}$ to $\check{X} \times \mathbb{C}^\times$ which projects to zero in the \mathbb{C}^\times direction, then

$$\iota(X)(d + d(\check{h}^{-1}W) \wedge)(f\Omega) = (X(f) + \check{h}^{-1}X(W)f)\Omega,$$

while

$$\iota(\check{h}\partial_{\check{h}})(d + d(\check{h}^{-1}W) \wedge)(f\Omega) = (\check{h}\partial_{\check{h}}f - \check{h}^{-1}Wf)\Omega.$$

From this, we get

$$(2.30) \quad \nabla_X^{GM}[f\Omega] = [(X(f) + \check{h}^{-1}X(W)f)\Omega]$$

and

$$(2.31) \quad \nabla_{\check{h}\partial_{\check{h}}}^{GM}[f\Omega] = [(\check{h}\partial_{\check{h}}f - \check{h}^{-1}Wf)\Omega].$$

One checks easily that these classes only depend on the class of $f\Omega$ and not on $f\Omega$ itself or the lift X .

DEFINITION 2.37. We define the B-model semi-infinite variation of Hodge structure on $\widetilde{\mathcal{M}}$ as follows. The bundle \mathcal{E} is the $\mathcal{O}_{\widetilde{\mathcal{M}}}\{\hbar\}$ -module such that $\Gamma(U, \mathcal{E})$ consists of sections of \mathcal{R}^\vee over open sets of the form

$$U \times \{\hbar \mid |\hbar| < \epsilon\}$$

given by forms $f\Omega$ with f holomorphic for $|\hbar| < \epsilon$, including at $\hbar = 0$, and f algebraic when restricted to $\pi^{-1}(u) \times \{\hbar\}$. The connection

$$\nabla : \mathcal{E} \rightarrow \Omega_{\widetilde{\mathcal{M}}}^1 \otimes \hbar^{-1}\mathcal{E}$$

is given by the Gauss-Manin connection on \mathcal{R}^\vee , namely for a section s of \mathcal{E} ,

$$\nabla_X s = \nabla_X^{GM} s,$$

where X on the right-hand side is the lift of X to $\widetilde{\mathcal{M}} \times \mathbb{C}^\times$ which projects to zero in the \mathbb{C}^\times direction.

Alternatively, one can think of this as follows: we can extend the locally free sheaf \mathcal{R}^\vee from $\widetilde{\mathcal{M}} \times \mathbb{C}^\times$ to $\widetilde{\mathcal{M}} \times \mathbb{C}$, by dictating that a section $[f\Omega]$ of \mathcal{R}^\vee extends if f extends as a holomorphic function across $\hbar = 0$. Two forms $f_1\Omega$ and $f_2\Omega$ define the same section if for each \hbar in a neighbourhood of $0 \in \mathbb{C}$, $(f_1 - f_2)\Omega$ is in the image of $\hbar d + dW\wedge$; note for $\hbar \neq 0$ this is the same as being in the image of $d + \hbar^{-1}dW\wedge$. It then follows from [102] that this extension of \mathcal{R}^\vee is in fact a vector bundle. The fibre of \mathcal{R}^\vee over $(u, 0)$ as a vector bundle is then $\mathbb{H}^n(\pi^{-1}(u), (\Omega_{\pi^{-1}(u)}^n, dW\wedge))$. Sections of this vector bundle near $\hbar = 0$ give sections of \mathcal{E} . In particular, $\mathcal{E}/\hbar\mathcal{E}$ is a rank $n + 1$ vector bundle with fibres being $\mathbb{H}^n(\pi^{-1}(u), (\Omega_{\pi^{-1}(u)}^n, dW\wedge))$.

Let

$$(-) : \widetilde{\mathcal{M}} \times \mathbb{C}^\times \rightarrow \widetilde{\mathcal{M}} \times \mathbb{C}^\times$$

denote the map

$$(u, \hbar) \mapsto (u, -\hbar).$$

Then the pairing (2.25) induces a pairing

$$(\cdot, \cdot) : (-)^*\mathcal{R}^\vee \times \mathcal{R}^\vee \rightarrow \mathcal{O}_{\widetilde{\mathcal{M}} \times \mathbb{C}^\times}.$$

This gives a pairing $(\cdot, \cdot)_\mathcal{E}$ on \mathcal{E} defined by

$$(s_1, s_2)_\mathcal{E}(\hbar) = \frac{(-1)^{n(n+1)/2}}{(2\pi i \hbar)^n} ((-)^* s_1, s_2).$$

For the grading operator, we begin by taking a vector field E on $\widetilde{\mathcal{M}}$ given by

$$E = (n+1)\partial_{t_1} + \sum_{i=0}^n (1-i)t_i\partial_{t_i}.$$

We then define the grading operator by

$$(2.32) \quad \text{Gr}(s) = \nabla_{\hbar\partial_{\hbar} + E}^{GM}(s) - s.$$

□

PROPOSITION 2.38. *The data \mathcal{E} , ∇ and $(\cdot, \cdot)_\mathcal{E}$ yield a semi-infinite variation of Hodge structure. The operator Gr is a grading operator.*

PROOF. First note that ∇ indeed takes \mathcal{E} to $\Omega_{\widetilde{\mathcal{M}}}^1 \otimes \hbar^{-1}\mathcal{E}$, by (2.30). Furthermore, we can describe the pairing $(\cdot, \cdot)_\mathcal{E}$ using the basis $\{\Delta_p^\pm\}$ of Lefschetz thimbles. By (2.26), we can write

$$(2.33) \quad (s_1, s_2)_\mathcal{E} = \frac{(-1)^{n(n+1)/2}}{(2\pi i \hbar)^n} \sum_p \left(\int_{\Delta_p^-} s_1(-\hbar) \right) \left(\int_{\Delta_p^+} s_2(\hbar) \right).$$

Thus by the stationary phase approximation (Proposition 2.35), we have

$$\begin{aligned} ([f\Omega], [g\Omega])_{\mathcal{E}} &= \pm \sum_p \frac{f(p, 0)g(p, 0)}{\text{Hess}(W(p))} + O(\hbar) \\ &\in \mathcal{O}_{\widetilde{\mathcal{M}}}\{\hbar\}. \end{aligned}$$

So $(\cdot, \cdot)_{\mathcal{E}}$ takes values in $\mathcal{O}_{\widetilde{\mathcal{M}}}\{\hbar\}$, as needed.

As the Gauss-Manin connection is by definition flat, ∇ is flat. Axioms (1)-(3) of Definition 2.20 are obviously satisfied by the definitions. To see (4), we proceed as follows. Recall that we have $\text{Crit}(W) \subseteq \check{\mathcal{X}}$ and sections $p: \widetilde{\mathcal{M}} \rightarrow \text{Crit}(W)$ of the projection $\pi: \text{Crit}(W) \rightarrow \widetilde{\mathcal{M}}$, giving a family of critical points. There are $n+1$ such sections. For each p , choose a holomorphic function f_p on $\check{\mathcal{X}}$, regular on fibres of π , such that $f_p \circ p' = \delta_{pp'}$, so that as functions on $\text{Crit}(W)$, $\{f_p\}$ form a basis of sections for the vector bundle $\pi_*\mathcal{O}_{\text{Crit}(W)}$. Then

$$\begin{aligned} ([f_{p'}\Omega], [f_{p''}\Omega])_{\mathcal{E}} &= \pm \sum_p \frac{f_{p'}(p, 0)f_{p''}(p, 0)}{\text{Hess}(W(p))} + O(\hbar) \\ &= \pm \frac{\delta_{p'p''}}{\text{Hess } W(p')} + O(\hbar). \end{aligned}$$

Now by construction, the forms $f_p\Omega$ represent a basis for

$$\mathbb{H}^n(\pi^{-1}(u), (\Omega_{\pi^{-1}(u)}^\bullet, dW \wedge))$$

for each point $u \in \widetilde{\mathcal{M}}$, as the functions f_p clearly form a basis for the Milnor ring of $W|_{\pi^{-1}(u)}$. Thus the classes $[f_p\Omega]$ also give a local basis for the rank $n+1$ vector bundle $\mathcal{E}/\hbar\mathcal{E}$, as the fibres of this vector bundle are precisely

$$\mathbb{H}^n(\pi^{-1}(u), (\Omega_{\pi^{-1}(u)}^\bullet, dW \wedge)).$$

It is then clear that the non-degeneracy condition of Definition 2.20, (4), holds.

We next turn to the grading operator. First note that

$$(2.34) \quad \pi_*\left(\sum_{i=0}^n x_i \partial_{x_i}\right) = (n+1)\partial_{t_1},$$

using $e^{t_1} = x_0 \cdots x_n$. Thus we can lift E to the vector field

$$\sum_{i=0}^n x_i \partial_{x_i} + \sum_{i=0}^n (1-i)t_i \partial_{t_i}$$

on $\check{\mathcal{X}}$, and calculate

$$\begin{aligned} E(W) &= \left(\sum_{i=0}^n x_i \partial_{x_i} + \sum_{i=0}^n (1-i)t_i \partial_{t_i} \right) \left(W_0 + \sum_{\substack{i=0 \\ i \neq 1}}^n t_i W_0^i \right) \\ &= W_0 + \sum_{\substack{i=0 \\ i \neq 1}}^n (i + (1-i)t_i) W_0^i \\ &= W. \end{aligned}$$

In fact, it is easy to see that E is the only vector field on $\widetilde{\mathcal{M}}$ which has a lift to $\check{\mathcal{X}}$ which preserves W .

Applying Gr to a section $[f\Omega]$ of \mathcal{E} , using (2.30), (2.31), and $E(W) = W$, we see that

$$\text{Gr}([f\Omega]) = [(\hbar\partial_{\hbar}f + E(f) - f)\Omega].$$

So we see that indeed $\text{Gr}(\mathcal{E}) \subseteq \mathcal{E}$. Now axiom (5) of Definition 2.20 is obvious, (6) follows from the flatness of ∇^{GM} , and (7) is checked as follows. By (2.33), we have

$$\begin{aligned} (\hbar\partial_{\hbar} + E)(s_1, s_2)_{\mathcal{E}} &= (\hbar\partial_{\hbar} + E) \left(\frac{(-1)^{n(n+1)/2}}{(2\pi i \hbar)^n} \sum_p \left(\int_{\Delta_p^-} s_1(-\hbar) \right) \left(\int_{\Delta_p^+} s_2(\hbar) \right) \right) \\ &= -n(s_1, s_2)_{\mathcal{E}} \\ &\quad + \frac{(-1)^{n(n+1)/2}}{(2\pi i \hbar)^n} \left(\int_{\Delta_p^-} (\nabla_{\hbar\partial_{\hbar} + E}^{\text{GM}} s_1)(-\hbar) \right) \left(\int_{\Delta_p^+} s_2(\hbar) \right) \\ &\quad + \frac{(-1)^{n(n+1)/2}}{(2\pi i \hbar)^n} \left(\int_{\Delta_p^-} s_1(-\hbar) \right) \left(\int_{\Delta_p^+} \nabla_{\hbar\partial_{\hbar} + E}^{\text{GM}} s_2(\hbar) \right) \\ &= -n(s_1, s_2)_{\mathcal{E}} \\ &\quad + \frac{(-1)^{n(n+1)/2}}{(2\pi i \hbar)^n} \left(\int_{\Delta_p^-} (\text{Gr}(s_1) + s_1)(-\hbar) \right) \left(\int_{\Delta_p^+} s_2(\hbar) \right) \\ &\quad + \frac{(-1)^{n(n+1)/2}}{(2\pi i \hbar)^n} \left(\int_{\Delta_p^-} s_1(-\hbar) \right) \left(\int_{\Delta_p^+} (\text{Gr}(s_2) + s_2)(\hbar) \right) \\ &= (\text{Gr}(s_1), s_2)_{\mathcal{E}} + (s_1, \text{Gr}(s_2))_{\mathcal{E}} + (2-n)(s_1, s_2)_{\mathcal{E}}. \end{aligned}$$

Thus Gr is a grading operator with $D = 2 - \dim_{\mathbb{C}} X$. \square

2.2.4. The B-model Frobenius manifold. We now have the B-model semi-infinite variation of Hodge structure $\mathcal{E}, \nabla, (\cdot, \cdot)_{\mathcal{E}}$ with a grading operator Gr on $\widetilde{\mathcal{M}}$ from the previous section. According to §2.1.7, with an appropriate additional choice of certain data, we will obtain a Frobenius manifold structure on $\widetilde{\mathcal{M}}$. In order for this to coincide with the A-model Frobenius manifold arising from $X = \mathbb{P}^n$, this data cannot be chosen arbitrarily. However, in order to specify this data, we need to examine the nature of the local system R in more detail. In particular, we would like to understand the monodromy of this local system. As R is a local system on $\widetilde{\mathcal{M}} \times \mathbb{C}^{\times}$, and $\widetilde{\mathcal{M}}$ is simply connected, the only interesting monodromy comes from a loop in the \mathbb{C}^{\times} factor, so we can in fact restrict to the point $0 \in \widetilde{\mathcal{M}}$ and consider the local system on \mathbb{C}^{\times} whose fibres are $H_n(\check{X}, \text{Re } W_0/\hbar \ll 0; \mathbb{C})$, with $\check{X} = V(x_0 \cdots x_n - 1)$ and $W_0 = x_0 + \cdots + x_n$.

PROPOSITION 2.39. *Let $\check{X} = V(x_0 \cdots x_n - 1) \subseteq \text{Spec } \mathbb{C}[x_0, \dots, x_n]$, and $W_0 = x_0 + \cdots + x_n$. Then for any local section Ξ of the local system R on \mathbb{C}^{\times} with fibre over $\hbar \in \mathbb{C}^{\times}$ given by $H_n(\check{X}, \text{Re } W_0/\hbar \ll 0; \mathbb{C})$, the integral*

$$\psi = \int_{\Xi} e^{W_0/\hbar} \Omega$$

satisfies the ordinary differential equation

$$(2.35) \quad \left(-\frac{1}{n+1} \hbar \partial_{\hbar} \right)^{n+1} \psi = \hbar^{-(n+1)} \psi.$$

PROOF. Let $\omega_i = x_0 \cdots x_{i-1} \Omega$. Let $\Omega_j = \iota(x_j \partial_{x_j}) \Omega$. Then applying $d + \hbar^{-1} dW_0 \wedge$ to $x_0 \cdots x_{i-1} \Omega_j$, keeping in mind that $x_0 \cdots x_n = 1$, we obtain the n -form

$$((x_j \partial_{x_j} - x_0 \partial_{x_0})(x_0 \cdots x_{i-1}) + \hbar^{-1}(x_j - x_0)(x_0 \cdots x_{i-1})) \Omega,$$

so in particular,

$$((x_j \partial_{x_j} - x_i \partial_{x_i})(x_0 \cdots x_{i-1}) + \hbar^{-1}(x_j - x_i)(x_0 \cdots x_{i-1})) \Omega$$

is in the image of $d + \hbar^{-1} dW_0 \wedge$. So in $\mathbb{H}^n(\check{X}, (\Omega_{\check{X}}^\bullet, d + \hbar^{-1} dW_0 \wedge))$, we have

$$[(x_j \partial_{x_j})(x_0 \cdots x_{i-1}) \Omega] = -[\hbar^{-1}(x_j - x_i) \omega_i].$$

Summing over all j , $0 \leq j \leq n$, we obtain the equality

$$[i \omega_i] = -[\hbar^{-1}(W_0 - (n+1)x_i) \omega_i].$$

As $\nabla_{\hbar \partial_{\hbar}}^{GM}[\omega_i] = -[\hbar^{-1} W_0 \omega_i]$ by (2.31), we obtain

$$\frac{1}{n+1}(-\nabla_{\hbar \partial_{\hbar}}^{GM} + i)[\omega_i] = \hbar^{-1}[\omega_{i+1}].$$

From this it follows inductively that

$$\frac{1}{(n+1)^i}(-\nabla_{\hbar \partial_{\hbar}}^{GM})^i[\omega_0] = \hbar^{-i}[\omega_i].$$

Since $\omega_{n+1} = \omega_0$ as $x_0 \cdots x_n = 1$, we obtain

$$\frac{1}{(n+1)^{n+1}}(-\nabla_{\hbar \partial_{\hbar}}^{GM})^{n+1}[\omega_0] = \hbar^{-(n+1)}[\omega_0].$$

Thus the integral of ω_0 over Ξ satisfies the equation (2.35). \square

Note that (2.35) is an $(n+1)$ -st order ODE, so there will be $n+1$ independent solutions. We can write down a fundamental system of solutions quite easily.

PROPOSITION 2.40. *Working over the ring $\mathbb{C}[\alpha]/(\alpha^{n+1})$, write*

$$\hbar^{-(n+1)\alpha} = \exp(-(n+1)\alpha \log \hbar) = \sum_{i=0}^n \frac{(-(n+1))^i \log^i \hbar}{i!} \alpha^i.$$

Then the coefficients of $1, \alpha, \dots, \alpha^n$ in the expression

$$(2.36) \quad \xi(\hbar, \alpha) = \hbar^{-(n+1)\alpha} \sum_{d=0}^{\infty} \hbar^{-(n+1)d} \prod_{i=1}^d \frac{1}{(\alpha + i)^{n+1}}$$

form a fundamental system of solutions to (2.35).

PROOF. The coefficients of α^i , $i = 0, \dots, n$, are linearly independent, as can be seen by noting that the coefficient of α^i is of the shape

$$f \log^i \hbar + \text{lower order terms in } \log \hbar,$$

with $f \neq 0$. To see that the coefficients are in fact solutions, note that for a function φ on \mathbb{C}^\times with values in $\mathbb{C}[\alpha]/(\alpha^{n+1})$,

$$-\frac{1}{n+1} \hbar \partial_{\hbar} (\hbar^{-(n+1)\alpha} \varphi) = \alpha \hbar^{-(n+1)\alpha} \varphi - \frac{1}{n+1} \hbar^{-(n+1)\alpha} \hbar \partial_{\hbar} \varphi.$$

Thus (2.35) is satisfied by $\psi = \hbar^{-(n+1)\alpha} \varphi$ if φ satisfies the equation

$$\left(\alpha - \frac{1}{n+1} \hbar \partial_{\hbar} \right)^{n+1} \varphi = \hbar^{-(n+1)\alpha} \varphi.$$

Taking

$$\varphi = \sum_{d=0}^{\infty} \hbar^{-(n+1)d} \prod_{i=1}^d \frac{1}{(\alpha + i)^{n+1}},$$

one sees that

$$\left(\alpha - \frac{1}{n+1} \hbar \partial_{\hbar}\right) \varphi = \sum_{d=0}^{\infty} (\alpha + d) \hbar^{-(n+1)d} \prod_{i=1}^d \frac{1}{(\alpha + i)^{n+1}},$$

so

$$\left(\alpha - \frac{1}{n+1} \hbar \partial_{\hbar}\right)^{n+1} \varphi = \sum_{d=1}^{\infty} \hbar^{-(n+1)d} \prod_{i=1}^{d-1} \frac{1}{(\alpha + i)^{n+1}},$$

(noting that for $d = 0$, $(\alpha + d)^{n+1} = 0$) while

$$\hbar^{-(n+1)} \varphi = \sum_{d=0}^{\infty} \hbar^{-(n+1)(d+1)} \prod_{i=1}^d \frac{1}{(\alpha + i)^{n+1}}.$$

Comparing coefficients of $\hbar^{-(n+1)d}$, these agree. \square

A consequence of this proposition is the following:

LEMMA 2.41. *If Ξ_0, \dots, Ξ_n is a local basis of sections of R , the integrals $\int_{\Xi_i} e^{W_0/\hbar} \Omega$, $i = 0, \dots, n$, form a fundamental system of solutions of (2.35).*

PROOF. Recall the extension of \mathcal{R}^\vee across $\hbar = 0$ as described in Definition 2.37. As $1, W_0, \dots, W_0^n$ span the Jacobian ring of W_0 , the sections of the vector bundle $\mathcal{R}^\vee|_{\{0\} \times \mathbb{C}}$ (where $0 \in \widetilde{\mathcal{M}}$ is the origin) given by the classes $[W_0^i \Omega]$, $0 \leq i \leq n$, yield a basis for the fibre of $\mathcal{R}^\vee|_{\{0\} \times \mathbb{C}}$ at $\hbar = 0$, and hence yield a basis of sections of $\mathcal{R}^\vee|_{\{0\} \times \mathbb{C}}$ in some neighbourhood $V = \{\hbar \in \mathbb{C} \mid |\hbar| < \epsilon\}$. In particular, since (2.24) is a perfect pairing, the $(n+1) \times (n+1)$ matrix

$$\left(\int_{\Xi_i} W_0^j e^{W_0/\hbar} \right)_{0 \leq i, j \leq n}$$

has rank $n+1$ at every point of V . Now

$$(2.37) \quad \left(-\hbar^2 \frac{\partial}{\partial \hbar}\right)^j \int_{\Xi_i} e^{W_0/\hbar} \Omega = \int_{\Xi_i} W_0^j e^{W_0/\hbar} \Omega.$$

Thus the Wronskian of the system of solutions

$$\int_{\Xi_0} e^{W_0/\hbar} \Omega, \dots, \int_{\Xi_n} e^{W_0/\hbar} \Omega$$

of (2.35) is non-vanishing on V , and hence these integrals give a fundamental system of solutions. \square

In particular, one can choose a local basis Ξ_0, \dots, Ξ_n of R such that $\int_{\Xi_i} e^{W_0/\hbar} \Omega$ in fact coincides with the coefficient of α^i in ξ , i.e.,

$$(2.38) \quad \sum_{i=0}^n \alpha^i \int_{\Xi_i} e^{W_0/\hbar} \Omega = \hbar^{-(n+1)\alpha} \sum_{d=0}^{\infty} \hbar^{-(n+1)d} \prod_{i=1}^d \frac{1}{(\alpha + i)^{n+1}} \pmod{\alpha^{n+1}}.$$

From now on, we will denote by Ξ_0, \dots, Ξ_n this particular basis of R .

REMARK 2.42. We have seen several choices of bases of cycles. In general, we have the cycles Δ_p^+ , where p runs over critical points of W_0 . This basis has little to do with the basis Ξ_0, \dots, Ξ_n . On the other hand, in dimension two, we explicitly described a basis of $H_2(\check{X}, \text{Re } W_0/\hbar \ll 0; \mathbb{C})$ in Example 2.34. Denoting by Ξ'_0, Ξ'_1, Ξ'_2 the basis described there, in fact the relationship between the two bases can be understood via direct evaluation of the integrals. It is a somewhat enjoyable exercise to evaluate the oscillatory integrals explicitly for Ξ'_i . One finds the following relationship:

$$\begin{aligned}\Xi_0 &= \frac{1}{(2\pi i)^2} \Xi'_0 \\ \Xi_1 &= \frac{1}{2\pi i} \Xi'_1 + \frac{3\gamma}{(2\pi i)^2} \Xi'_0 \\ \Xi_2 &= \Xi'_2 + \frac{3\gamma}{2\pi i} \Xi'_1 - \frac{1}{(2\pi i)^2} \left(\frac{\pi^2}{4} + \frac{27}{2} \gamma^2 \right) \Xi'_0.\end{aligned}$$

Here

$$\gamma = \lim_{n \rightarrow \infty} \left(-\log(n) + \sum_{k=1}^n \frac{1}{k} \right)$$

is Euler's constant. The integral over Ξ'_0 is a triviality, coming from a residue calculation. The integral over Ξ'_1 is more complicated, involving exponential integrals, and the third integral is a real challenge. However, one may apply results of Iritani [61] to describe this third integral in terms of certain characteristic classes of the structure sheaf of \mathbb{P}^2 . \square

Now Ξ_0, \dots, Ξ_n are defined as local sections of the local system R on \mathbb{C}^\times . As we pass around a counterclockwise loop in the \hbar -plane, following the sections Ξ_0, \dots, Ξ_n , when we get back to the beginning of the loop, we will find that the Ξ_i 's will have been transformed into some linear combination of the Ξ_i 's. This linear transformation is the monodromy transformation for the local system R . To see exactly what this transformation is, we can study the behaviour of the integrals

$$\xi_i = \int_{\Xi_i} e^{W_0/\hbar} \Omega.$$

Since $e^{W_0/\hbar} \Omega$ of course gives a single-valued section of \mathcal{R}^\vee , the multi-valuedness of the integral must arise precisely from the monodromy in the Ξ_i .

This multi-valuedness is as follows. Note that $\xi = \sum_{i=0}^n \alpha^i \xi_i = \hbar^{-(n+1)\alpha} \varphi$ for some φ single-valued. As we follow a counterclockwise loop in the \mathbb{C}^\times plane around the origin, $\hbar^{-(n+1)\alpha} = \exp(-(n+1)\alpha \log \hbar)$ is replaced by

$$\exp(-(n+1)\alpha(\log \hbar + 2\pi i)) = \hbar^{-(n+1)\alpha} \exp(-(n+1)2\pi i \alpha).$$

From this, one sees that passing around this loop, ξ is replaced by $\xi e^{-(n+1)2\pi i \alpha}$ and ξ_k is transformed as

$$\xi_k \mapsto \sum_{j=0}^k \frac{(-(n+1)2\pi i)^{k-j}}{(k-j)!} \xi_j.$$

This means that the cycles Ξ_k are transformed in precisely the same way:

$$\Xi_k \mapsto \sum_{j=0}^k \frac{(-(n+1)2\pi i)^{k-j}}{(k-j)!} \Xi_j.$$

To describe the dual monodromy on R^\vee , it is convenient to write the dual basis to Ξ_0, \dots, Ξ_n as $1 = \alpha^0, \alpha^1, \dots, \alpha^n \in \mathbb{C}[\alpha]/(\alpha^{n+1})$. This allows us to write, for example, the section $[f\Omega]$ of \mathcal{R}^\vee in terms of this basis as

$$[f\Omega] = \sum_{i=0}^n \alpha^i \int_{\Xi_i} f e^{W/\hbar} \Omega.$$

The effect of monodromy on this dual basis is then given by the transpose inverse transformation, which is easily seen to be

$$\alpha^k \mapsto \exp((n+1)2\pi i \alpha) \alpha^k.$$

In particular, $\hbar^{-(n+1)\alpha} \alpha^k$, $k = 0, \dots, n$, is a single-valued, globally defined section of \mathcal{R}^\vee , as the multi-valuedness of α^k is cancelled by the multi-valuedness of $\hbar^{-(n+1)\alpha}$.

We now pass to the entire moduli space $\widetilde{\mathcal{M}}$. The local sections Ξ_0, \dots, Ξ_n now extend to multi-valued sections of the local system R on all of $\widetilde{\mathcal{M}} \times \mathbb{C}^\times$, and similarly the dual basis $1, \alpha, \dots, \alpha^n$, giving globally defined sections $\hbar^{-(n+1)\alpha} \alpha^k$ of \mathcal{R}^\vee on $\widetilde{\mathcal{M}} \times \mathbb{C}^\times$. So we can view $\hbar^{-(n+1)\alpha} \alpha^k$ as a section of $\mathcal{E} \otimes_{\mathcal{O}_{\widetilde{\mathcal{M}}}\{\hbar\}} \mathcal{O}_{\widetilde{\mathcal{M}}}\{\hbar, \hbar^{-1}\}$. Furthermore, since α^k is flat, $\nabla(\hbar^{-(n+1)\alpha} \alpha^k) = 0$, and hence

$$\{\hbar^{-(n+1)\alpha} \alpha^i \mid i = 0, \dots, n\}$$

form a basis for \mathcal{H} as a free $\mathbb{C}\{\hbar, \hbar^{-1}\}$ -module.

We can now be precise about the data we shall take to determine the Frobenius manifold structure on $\widetilde{\mathcal{M}}$:

- We take \mathcal{H}_- to be the $\mathcal{O}(\mathbb{P}^1 \setminus \{0\})$ -submodule of \mathcal{H} generated by

$$\{(\hbar\alpha)^k \hbar^{-1} \hbar^{-(n+1)\alpha} \mid 0 \leq k \leq n\}.$$

- We take Ω_0 to be the flat section of $\mathcal{E} \otimes_{\mathcal{O}_{\widetilde{\mathcal{M}}}\{\hbar\}} \mathcal{O}_{\widetilde{\mathcal{M}}}\{\hbar, \hbar^{-1}\}$ whose value at $0 \in \widetilde{\mathcal{M}}$ is $[\Omega]$.

PROPOSITION 2.43. (1) $\mathcal{H} = \mathcal{E}_0 \oplus \mathcal{H}_-$.

(2) \mathcal{H}_- is isotropic with respect to the symplectic form $\overline{\Omega}$ defined in (2.15).

(3) Gr preserves \mathcal{H}_- , the flat section $\Omega_0 \in \hbar\mathcal{H}_-$ represents an eigenvector of Gr_0 with eigenvalue -1 , and the corresponding section $s'_0 = \tau(\Omega_0 \otimes 1)$ of \mathcal{E} yields miniversality in a neighbourhood of $0 \in \widetilde{\mathcal{M}}$.

As a consequence, by Theorem 2.26, we obtain a Frobenius manifold structure in a neighbourhood of $0 \in \widetilde{\mathcal{M}}$ with a flat identity and Euler vector field E .

PROOF. \mathcal{E}_0 is given by sections of $\mathcal{R}^\vee|_{\{0\} \times \mathbb{C}}$ (using the extension of \mathcal{R}^\vee across $\hbar = 0$ given in Definition 2.37) in a neighbourhood of $\hbar = 0$. As observed in the proof of Lemma 2.41, there is a basis of such sections given by $[W_0^i \Omega]$, $0 \leq i \leq n$, and hence \mathcal{E}_0 is generated as a $\mathbb{C}\{\hbar\}$ -module by these classes. Now, by (2.37), we can write

$$\begin{aligned} \sum_{j=0}^n \alpha^j \int_{\Xi_j} W_0^i e^{W_0/\hbar} \Omega &= \left(-\hbar^2 \frac{\partial}{\partial \hbar}\right)^i \left(\sum_{j=0}^n \alpha^j \int_{\Xi_j} e^{W_0/\hbar} \Omega\right) \\ &= \left(-\hbar^2 \frac{\partial}{\partial \hbar}\right)^i (\xi). \end{aligned}$$

From the explicit formula (2.36) for ξ , the behaviour of $(\hbar^2 \partial_{\hbar})^i \xi$ is

$$(\alpha \hbar)^i \hbar^{-(n+1)\alpha} + \text{terms in } \mathcal{H}_-.$$

So clearly $\mathcal{H} = \mathcal{E}_0 \oplus \mathcal{H}_-$, and (1) holds.

Next, note that as α^i is a flat section of \mathcal{R}^\vee , $((-)^* \alpha^i, \alpha^j)$ is constant. Furthermore, as $\hbar^{-(n+1)\alpha} \alpha^i$ is a single-valued section of \mathcal{R}^\vee ,

$$((-)^*(\hbar^{-(n+1)\alpha} \alpha^i), \hbar^{-(n+1)\alpha} \alpha^j)$$

is single-valued. On the other hand, expanding $\hbar^{-(n+1)\alpha}$ in terms of powers of $\log \hbar$, one obtains a formula for $((-)^*(\hbar^{-(n+1)\alpha} \alpha^i), \hbar^{-(n+1)\alpha} \alpha^j)$ whose \hbar -dependence is as a polynomial in $\log \hbar$. Thus, in order for it to be single valued, it must in fact be independent of \hbar . Thus

$$\begin{aligned} 0 &= \hbar \partial_{\hbar} ((-)^*(\hbar^{-(n+1)\alpha} \alpha^i), \hbar^{-(n+1)\alpha} \alpha^j) \\ &= (-(-)^*((n+1)\hbar^{-(n+1)\alpha} \alpha^{i+1}), \hbar^{-(n+1)\alpha} \alpha^j) \\ &\quad + ((-)^*(\hbar^{-(n+1)\alpha} \alpha^i), -(n+1)\hbar^{-(n+1)\alpha} \alpha^{j+1}). \end{aligned}$$

So in fact

$$(2.39) \quad ((-)^*(\hbar^{-(n+1)\alpha} \alpha^{i+1}), \hbar^{-(n+1)\alpha} \alpha^j) = -((-)^*(\hbar^{-(n+1)\alpha} \alpha^i), \hbar^{-(n+1)\alpha} \alpha^{j+1}).$$

From this we conclude that

$$(2.40) \quad (\hbar^{-(n+1)\alpha} \alpha^i, \hbar^{-(n+1)\alpha} \alpha^j)_{\mathcal{E}} = \begin{cases} 0 & i+j > n, \\ (\text{Constant}) \cdot \hbar^{-n} & i+j \leq n. \end{cases}$$

So to prove \mathcal{H}_- is isotropic, we just note that for $k, \ell \geq 1$,

$$\begin{aligned} &\overline{\Omega}((\hbar \alpha)^i \hbar^{-(n+1)\alpha-k}, (\hbar \alpha)^j \hbar^{-(n+1)\alpha-\ell}) \\ &= \text{Res}_{\hbar=0}((\hbar \alpha)^i \hbar^{-(n+1)\alpha-k}, (\hbar \alpha)^j \hbar^{-(n+1)\alpha-\ell})_{\mathcal{E}} d\hbar \\ &= \begin{cases} 0 & i+j > n \\ \text{Res}_{\hbar=0} \text{Constant} \cdot \hbar^{i+j-k-\ell-n} d\hbar & i+j \leq n \end{cases} \\ &= 0. \end{aligned}$$

For (3), note that for $m \geq 1$,

$$\begin{aligned} \text{Gr}((\hbar \alpha)^k \hbar^{-(n+1)\alpha} \hbar^{-m}) &= \nabla_{\hbar \partial_{\hbar} + E}^{\text{GM}}(\hbar^{k-m} \alpha^k \hbar^{-(n+1)\alpha}) - (\hbar \alpha)^k \hbar^{-(n+1)\alpha} \hbar^{-m} \\ &= (k-m-(n+1)\alpha-1)(\hbar \alpha)^k \hbar^{-(n+1)\alpha} \hbar^{-m} \\ &\in \mathcal{H}_-, \end{aligned}$$

so Gr preserves \mathcal{H}_- . We see also that $\Omega_0 \in \hbar \mathcal{H}_-$, as ξ , which represents the flat section Ω_0 , takes the form $\xi = \hbar^{-(n+1)\alpha}(1 + O(\hbar^{-(n+1)d}))$. In addition,

$$(2.41) \quad \text{Gr}_0((\hbar \alpha)^k \hbar^{-(n+1)\alpha}) = (k-1)(\hbar \alpha)^k \hbar^{-(n+1)\alpha},$$

so $[\Omega_0]$ is an eigenvector of Gr_0 of eigenvalue -1 .

Finally, we need to show minversality. In particular, let $s'_0 = \tau([\Omega_0] \otimes 1)$. By definition, s'_0 is obtained by writing the constant section Ω_0 of $\mathcal{H}_{\widetilde{\mathcal{M}}}$ (whose value at $0 \in \widetilde{\mathcal{M}}$ is given by $[\Omega]$) as a sum $s'_0 + s''_0$, where s'_0 is a section of \mathcal{E} and s''_0 is a section of $\mathcal{H}_{-\widetilde{\mathcal{M}}}$. In other words, s'_0 will be represented by some $f\Omega$, with f a holomorphic function on $\check{\mathcal{X}} \times \{|\hbar| < \epsilon\}$ and regular on fibres of $\check{\mathcal{X}}$. Since

$\Omega_0 = \hbar^{-(n+1)\alpha} \bmod \mathcal{H}_-$, and we need $[f\Omega] \equiv \Omega_0 \bmod \mathcal{H}_-$, the requirement on $f\Omega$ is that

$$[f\Omega] = \hbar^{-(n+1)\alpha} \sum_{i=0}^n \varphi_i(\mathbf{t}, \hbar^{-1}) (\alpha \hbar)^i$$

where

$$\varphi_i(\mathbf{t}, \hbar^{-1}) = \delta_{0,i} + \sum_{j=1}^{\infty} \varphi_{i,j}(\mathbf{t}) \hbar^{-j}.$$

Here, we have written $\mathbf{t} = (t_0, \dots, t_n)$. Furthermore, as $[\Omega] \in \mathcal{E}_0$, in fact $s'_0 = [\Omega]$ at $0 \in \widetilde{\mathcal{M}}$, so we can assume that $f|_{\mathbf{t}=0} \equiv 1$. These two conditions uniquely determine the class represented by $f\Omega$.

Now once $s'_0 = [f\Omega]$ is given, we check miniversality at 0, by computing using (2.30) and (2.34) that

$$\hbar \nabla_{\partial/\partial t_i} [f\Omega] = \begin{cases} [(\hbar(\partial f/\partial t_i) + fW_0^i)\Omega] & i \neq 1 \\ [(\hbar(\partial f/\partial t_1) + \frac{f}{n+1}(W_0 + \sum_{i=2}^n it_i W_0^i))\Omega] & i = 1 \end{cases}$$

So in $\mathcal{E}/\hbar\mathcal{E}$ at $0 \in \widetilde{\mathcal{M}}$, this is just $[W_0^i\Omega]$ if $i \neq 1$ and $\frac{1}{n+1}[W_0\Omega]$ if $i = 1$. We know that $[W_0^i\Omega]$, $0 \leq i \leq n$, form a basis for $\mathcal{E}/\hbar\mathcal{E}$ at $0 \in \widetilde{\mathcal{M}}$ as W_0^i , $0 \leq i \leq n$, form a basis for the Milnor ring of W_0 . So we have miniversality at $0 \in \widetilde{\mathcal{M}}$, hence miniversality in a neighbourhood of $0 \in \widetilde{\mathcal{M}}$. \square

2.2.5. Mirror symmetry for \mathbb{P}^n . We can now, at last, state mirror symmetry for $X = \mathbb{P}^n$. To summarize, we now have two sets of data, the A-model data and the B-model data. The A-model data consists of a semi-infinite variation of Hodge structure $\mathcal{E}^A, \nabla^A, (\cdot, \cdot)_{\mathcal{E}^A}$ on $\widetilde{\mathcal{M}}^A$, with an opposite subspace $\mathcal{H}_-^A \subseteq \mathcal{H}^A$ at $0 \in \widetilde{\mathcal{M}}^A$ and an element $\Omega_0^A \in \hbar\mathcal{H}_-^A$. This data comes from quantum cohomology of \mathbb{P}^n . The B-model data consists of exactly the same sort of data, this time labelled with B's, arising from the construction of the previous two sections on $\widetilde{\mathcal{M}}^B$.

In particular, we have

$$\hbar\mathcal{H}_-^A/\mathcal{H}_-^A \cong H^*(X, \mathbb{C}) = \mathbb{C}[T_1]/(T_1^{n+1}),$$

the latter equality as a ring with the classical cup product; of course $T_i = T_1^i$, with T_1 the generator of $H^2(X, \mathbb{C})$. Also,

$$\hbar\mathcal{H}_-^B/\mathcal{H}_-^B \cong \mathbb{C}[\alpha]/(\alpha^{n+1}),$$

with α^i being identified with $\hbar^{-(n+1)\alpha}(\hbar\alpha)^i$. There is then an obvious choice of isomorphism

$$(2.42) \quad \begin{aligned} \hbar\mathcal{H}_-^A/\mathcal{H}_-^A &\cong \hbar\mathcal{H}_-^B/\mathcal{H}_-^B \\ T_1^i &\mapsto \hbar^{-(n+1)\alpha}(\hbar\alpha)^i. \end{aligned}$$

From this data we obtain Barannikov's period maps (2.18) Ψ^A and Ψ^B , hence maps of germs (2.19)

$$\psi^A : (\widetilde{\mathcal{M}}^A, 0) \rightarrow \hbar\mathcal{H}_-^A/\mathcal{H}_-^A$$

and

$$\psi^B : (\widetilde{\mathcal{M}}^B, 0) \rightarrow \hbar\mathcal{H}_-^B/\mathcal{H}_-^B$$

which are local isomorphisms by miniversality near $0 \in \widetilde{\mathcal{M}}$. Using the identification (2.42) of $\hbar\mathcal{H}_-^A/\mathcal{H}_-^A$ and $\hbar\mathcal{H}_-^B/\mathcal{H}_-^B$, we obtain the *mirror map*

$$m : (\widetilde{\mathcal{M}}^A, 0) \rightarrow (\widetilde{\mathcal{M}}^B, 0).$$

THEOREM 2.44 (Mirror symmetry for \mathbb{P}^n). *m identifies the A- and B-model semi-infinite variations of Hodge structure. More specifically, the identification (2.42) yields isomorphisms*

$$\mathcal{E}^A \xrightarrow{\tau_A^{-1}} (\hbar\mathcal{H}_-^A/\mathcal{H}_-^A) \otimes_{\mathbb{C}} \mathcal{O}_{\widetilde{\mathcal{M}}^A}\{\hbar\} \cong (\hbar\mathcal{H}_-^B/\mathcal{H}_-^B) \otimes_{\mathbb{C}} \mathcal{O}_{\widetilde{\mathcal{M}}^B}\{\hbar\} \xrightarrow{\tau_B} \mathcal{E}^B$$

which identifies the connections, the inner product, and the gradings. Furthermore, the opposite subspaces \mathcal{H}_-^A and \mathcal{H}_-^B and the elements Ω_0^A and Ω_0^B are identified under these isomorphisms.

We will not give the proof here. This is proved by Barannikov in [4], though not precisely in these words. He shows an isomorphism of the corresponding Frobenius manifolds, but the data of the semi-infinite variation of Hodge structure is recoverable from the Frobenius manifold structure. The basic idea in the proof is that the Frobenius manifold for the quantum cohomology of \mathbb{P}^n is *semi-simple* and as a consequence is completely determined by its behaviour at one point, say $0 \in \widetilde{\mathcal{M}}^A$. One finds the same structure on $\widetilde{\mathcal{M}}^B$ at 0, and hence these two Frobenius manifolds coincide in a neighbourhood of 0.

This statement is rather abstract, and for the purposes of Chapter 5, it is convenient to rewrite this statement in a much more down-to-earth form.

PROPOSITION 2.45. *Theorem 2.44 is equivalent to the following statement:*

There is a neighbourhood U of $0 \in \widetilde{\mathcal{M}}^B$ and a section s of \mathcal{E}^B over U defined by a form $f\Omega$ on $\pi^{-1}(U) \times \{|\hbar| < \epsilon\}$ with f a holomorphic function which is algebraic on fibres of π and $f|_{\pi^{-1}(0) \times \mathbb{C}} \equiv 1$, satisfying the following conditions:

(1) *If we write*

$$s(\mathbf{t}, \hbar) = \sum_{i=0}^n \alpha^i \int_{\Xi_i} e^{W/\hbar} f\Omega,$$

then

$$s(\mathbf{t}, \hbar) = \hbar^{-(n+1)\alpha} \sum_{i=0}^n \varphi_i(\mathbf{t}, \hbar^{-1})(\alpha\hbar)^i$$

for functions φ_i satisfying

$$\varphi_i(\mathbf{t}, \hbar^{-1}) = \delta_{0,i} + \sum_{j=1}^{\infty} \varphi_{i,j}(\mathbf{t}) \hbar^{-j}$$

for $0 \leq i \leq n$. If we set

$$y_i(\mathbf{t}) = \varphi_{i,1}(\mathbf{t}), \quad 0 \leq i \leq n,$$

y_0, \dots, y_n form a system of coordinates on $\widetilde{\mathcal{M}}^B$ in a neighbourhood of 0, which are called flat coordinates.

(2) *If we write the Givental J-function as*

$$J(y_0, \dots, y_n, \hbar^{-1}) = \sum_{i=0}^n J_i(y_0, \dots, y_n, \hbar^{-1}) T_i,$$

then in the \mathbb{C} -vector space $\mathbb{C}[[y_0, \dots, y_n, \hbar^{-1}]]$,

$$(2.43) \quad J_i = \varphi_i.$$

(3) Under the map $\widetilde{\mathcal{M}}^A \rightarrow \widetilde{\mathcal{M}}^B$ given by $y_i \mapsto y_i$, the Euler vector fields E_A and E_B are identified.

(4)

$$\int_{\mathbb{P}^n} T_i \cup T_j = (\hbar^{-(n+1)\alpha} (\hbar\alpha)^i, \hbar^{-(n+1)\alpha} (\hbar\alpha)^j)_{\mathcal{E}^B} \Big|_{\hbar=\infty}.$$

PROOF. First suppose that Theorem 2.44 holds. We already saw in the proof of Proposition 2.43, (3), that $s'_0 = \tau(\Omega_0 \otimes 1)$ is in fact given precisely by the description of the section s in (1). In particular, the map $\psi_B : (\widetilde{\mathcal{M}}^B, 0) \rightarrow \hbar\mathcal{H}_-^B/\mathcal{H}_-^B$ is given by $\psi_B(\mathbf{t}) = \hbar(s(\mathbf{t}, \hbar^{-1}) - \hbar^{-(n+1)\alpha}) \bmod \mathcal{H}_-^B$. Thus in terms of the basis $\{\hbar^{-(n+1)\alpha} (\hbar\alpha)^i\}$ of $\hbar\mathcal{H}_-^B/\mathcal{H}_-^B$, we see that

$$\psi_B(\mathbf{t}) = \sum_{i=0}^n \varphi_{i,1}(\mathbf{t}) \hbar^{-(n+1)\alpha} (\hbar\alpha)^i.$$

This shows (1). On the other hand, from (2.22),

$$\psi_A(y_0, \dots, y_n) = \sum_{i=0}^n y_i T_i.$$

Thus the mirror map $m : \widetilde{\mathcal{M}}^A \rightarrow \widetilde{\mathcal{M}}^B$ is just given by $y_i \mapsto y_i$. By Example 2.27, $\tau_A([T_0]) = J_{\mathbb{P}^n}$, and as $\tau_B([\Omega_0]) = s$, we obtain (2.43). In addition, (3) is clear since the Euler vector fields coincide under the mirror map. Furthermore, by Proposition 2.24 and Example 2.27, the left- and right-hand sides of the equality in (4) are the pairings on $\hbar\mathcal{H}_-^A/\mathcal{H}_-^A$ and $\hbar\mathcal{H}_-^B/\mathcal{H}_-^B$ induced by $(\cdot, \cdot)_{\mathcal{E}^A}$ and $(\cdot, \cdot)_{\mathcal{E}^B}$ respectively, so they must agree under the identification (2.42).

For the converse, suppose we are given $f\Omega$ giving a section of \mathcal{E}^B satisfying all the given properties. The map $y_i \mapsto y_i$ defines the mirror map

$$m : (\widetilde{\mathcal{M}}^A, 0) \rightarrow (\widetilde{\mathcal{M}}^B, 0).$$

As observed in Remark 2.29, $J_{\mathbb{P}^n}$ completely determines \mathbb{J} , hence the embedding

$$\mathbb{J} : \mathcal{E}^A \rightarrow \mathcal{H}_{\widetilde{\mathcal{M}}}^A,$$

and similarly, by miniversality of the section s of \mathcal{E}^B defined by $f\Omega$, s determines the embedding

$$\mathcal{E}^B \rightarrow \mathcal{H}_{\widetilde{\mathcal{M}}}^B.$$

The equality (2.43) then guarantees that after identifying

$$\begin{aligned} \mathcal{H}_{\widetilde{\mathcal{M}}}^A &= (\hbar\mathcal{H}_-^A/\mathcal{H}_-^A) \otimes_{\mathbb{C}} \mathcal{O}_{\widetilde{\mathcal{M}}}\{\hbar, \hbar^{-1}\} \\ \mathcal{H}_{\widetilde{\mathcal{M}}}^B &= (\hbar\mathcal{H}_-^B/\mathcal{H}_-^B) \otimes_{\mathbb{C}} \mathcal{O}_{\widetilde{\mathcal{M}}}\{\hbar, \hbar^{-1}\} \end{aligned}$$

using (2.42), the subbundles $\mathcal{E}^A \subseteq \mathcal{H}_{\widetilde{\mathcal{M}}}^A$ and $\mathcal{E}^B \subseteq \mathcal{H}_{\widetilde{\mathcal{M}}}^B$ are identified. By condition (4), $(\cdot, \cdot)_{\mathcal{E}^A}$ and $(\cdot, \cdot)_{\mathcal{E}^B}$ are identified. Next, $[\Omega_0^A] = T_0$ and $[\Omega_0^B] = \hbar^{-(n+1)\alpha}$ and $\mathcal{H}_-^A, \mathcal{H}_-^B$ are clearly identified under (2.42). Finally, by (3) and Proposition 2.24,

(3), in order to show that Gr^A and Gr^B coincide, it is enough to show that Gr_0^A and Gr_0^B coincide. But

$$\begin{aligned} \text{Gr}_0^A(T_i) &= \text{Gr}^A(T_i) \pmod{\mathcal{H}_-^A} \\ &= \text{Gr}_E(T_i) \pmod{\mathcal{H}_-^A} \\ &= \left(\frac{\deg T_i}{2} - 1 \right) T_i \\ &= (i-1)T_i \end{aligned}$$

while by (2.41),

$$\text{Gr}_0^B(\hbar^{-(n+1)\alpha}(\alpha\hbar)^i) = (i-1)\hbar^{-(n+1)\alpha}(\alpha\hbar)^i,$$

so Gr_0^A and Gr_0^B coincide. □

2.3. References and further reading

The discussion of Gromov-Witten invariants and quantum cohomology owes a great deal to the exposition by Fulton and Pandharipande [28]. Another useful source is the book of Cox and Katz [18]. For a great deal of information about Frobenius manifolds, see Manin's book, [76]. The expositions of Coates-Iritani-Tseng [17] and Iritani [61] were immensely useful for the discussion of the quantum differential equation and semi-infinite variations of Hodge structure.

For the B-model, I have mainly relied on Barannikov's paper [4], with help from Iritani [61] and Douai-Sabbah [20], [21]. For an important alternative point of view for the B-model, see Sabbah's book [103]. For a more general discussion of some of the structures involved in semi-infinite variations of Hodge structures, see the article of Katzarkov, Kontsevich and Pantev [67].

Log geometry

Log geometry was introduced by Illusie and Fontaine [59] and K. Kato [65]. The origins of log geometry, and in particular the term log, come from logarithmic differentials.

Suppose, for example, one wishes to study an open variety, say a non-singular quasi-projective variety X . If X is contained in \overline{X} , a projective variety with $D = \overline{X} \setminus X$ normal crossings¹, then $\Omega_{\overline{X}}^q(\log D)$ is defined to be the subsheaf of $i_*\Omega_X^q$ (where $i : X \hookrightarrow \overline{X}$ is the inclusion) locally generated by

$$\frac{dx_1}{x_1}, \dots, \frac{dx_p}{x_p}, dx_{p+1}, \dots, dx_n,$$

if locally D is given by $x_1 \cdots x_p = 0$. Of course $\frac{dx_i}{x_i} = d \log(x_i)$, hence the term logarithmic.

The initial importance of these sheaves was illustrated by Deligne's construction of a mixed Hodge structure on X . In particular, the exterior derivative takes logarithmic forms to logarithmic forms, and the complex $\Omega_{\overline{X}}^\bullet(\log D)$ with exterior derivative as differential in fact computes the cohomology of X : we have

$$\mathbb{H}^q(\overline{X}, \Omega_{\overline{X}}^\bullet(\log D)) \cong H^q(X, \mathbb{C}).$$

As in Chapter 2, \mathbb{H}^q denotes hypercohomology. This is an important step in putting a mixed Hodge structure on $H^q(X, \mathbb{C})$.

Another important application of the sheaf of log differentials came with Steenbrink's construction of the limiting mixed Hodge structure for a normal crossings degeneration. One considers a flat family $f : \mathcal{X} \rightarrow S$, where S is a non-singular one-dimensional scheme and for $s \in S$, \mathcal{X}_s is non-singular except for a closed point $0 \in S$. Furthermore, f is always given in suitable local coordinate charts as $(x_1, \dots, x_n) \mapsto x_1 \cdots x_p$ for some $p \leq n$. One then has the sheaf of relative log q -forms,

$$\Omega_{\mathcal{X}/S}^q(\log \mathcal{X}_0) := \Omega_{\mathcal{X}}^q(\log \mathcal{X}_0) / (f^* \Omega_S^1(\log 0) \wedge \Omega_{\mathcal{X}}^{q-1}(\log \mathcal{X}_0)).$$

In particular, $\Omega_{\mathcal{X}/S}^1(\log \mathcal{X}_0)$ is locally generated by

$$d \log x_1, \dots, d \log x_p, dx_{p+1}, \dots, dx_n$$

modulo the relation $d \log x_1 + \cdots + d \log x_p = 0$. Restricting $\Omega_{\mathcal{X}/S}^q(\log \mathcal{X}_0)$ to \mathcal{X}_0 gives a sheaf we will write, for the moment, as $\Omega_{\mathcal{X}_0^\dagger}^q$ to distinguish it from the ordinary sheaf of differentials $\Omega_{\mathcal{X}_0}^q$. Now one can check easily that $\Omega_{\mathcal{X}/S}^q(\log \mathcal{X}_0)$ is locally free, and hence so is $\Omega_{\mathcal{X}_0^\dagger}^q$, whereas $\Omega_{\mathcal{X}_0}^q$ is not locally free. The exterior derivative

¹A normal crossings divisor D is a divisor such that there are coordinates x_1, \dots, x_n on \overline{X} in a neighbourhood of any point of D such that D is given by $x_1 \cdots x_p = 0$ for some $p \leq n$.

still makes sense, either on $\Omega_{\mathcal{X}/S}^q(\log \mathcal{X}_0)$ or on $\Omega_{\mathcal{X}_0^\dagger}^q$. If f is proper, one can show (see [107]) that $R^p f_* \Omega_{\mathcal{X}/S}^q(\log \mathcal{X}_0)$ is locally free. Away from \mathcal{X}_0 , $\Omega_{\mathcal{X}/S}^q(\log \mathcal{X}_0)$ is just the ordinary sheaf of q -forms, so away from 0, $R^p f_* \Omega_{\mathcal{X}/S}^q(\log \mathcal{X}_0)$ is the vector bundle whose fibre at s is just the Dolbeault cohomology group $H^p(\mathcal{X}_s, \Omega_{\mathcal{X}_s}^q)$. The fibre at 0 coincides with $H^p(\mathcal{X}_0, \Omega_{\mathcal{X}_0^\dagger}^q)$. These groups are part of the data determining the limiting mixed Hodge structure associated to this degeneration.

The point here is that the sheaf of logarithmic differentials is much better behaved than the ordinary sheaf of differentials, and somehow gives \mathcal{X}_0 some of the properties of a smooth variety.

The concept of a log structure is an abstraction of this idea. Essentially, it is a way of making certain singular schemes behave as if they were smooth; as Kazuya Kato described it, a log structure is a magic powder (*poudre magique*) which makes a singular variety smooth.

We will give the precise definition of a log structure in §3.2, and explore its meaning. First, we give a brief review of toric geometry. For a more detailed but pleasant introduction to the subject, see Fulton's book [27].

3.1. A brief review of toric geometry

3.1.1. Monoids. We will work frequently with monoids and monoid rings. For us, a *monoid* P is a set with an associative and commutative operation with an identity element. The operation, $P \times P \rightarrow P$, is usually, but not always, written additively, in which case the identity is written as $0 \in P$.

Given a field \mathbb{k} , we define the monoid ring

$$\mathbb{k}[P] := \bigoplus_{p \in P} \mathbb{k}z^p,$$

where z^p is a symbol, and multiplication is \mathbb{k} -bilinear and is determined by

$$z^p \cdot z^{p'} = z^{p+p'}.$$

A standard example is the additive monoid of natural numbers

$$\mathbb{N} := \{0, 1, 2, \dots\}.$$

Then $\mathbb{k}[\mathbb{N}] \cong \mathbb{k}[x]$ and $\mathbb{k}[\mathbb{N}^r] \cong \mathbb{k}[x_1, \dots, x_r]$.

A homomorphism of monoids is a map $f : P \rightarrow Q$ between monoids such that $f(0) = 0$ and $f(p + p') = f(p) + f(p')$.

We say a monoid P is *finitely generated* if there is a surjective homomorphism $\mathbb{N}^r \rightarrow P$ for some r .

The (Grothendieck) group of a monoid P is the group

$$P^{\text{gp}} := \{p - p' \mid p, p' \in P\} / \sim,$$

where $p - p'$ is a formal symbol and the equivalence relation is given by $p - p' \sim q - q'$ if $p + q' = q + p'$. The group operations are given by $(p - p') + (q - q') = (p + q) - (p' + q')$ and $-(p - p') = p' - p$.

A monoid P is *integral* if the canonical map $P \rightarrow P^{\text{gp}}$ is injective.

One can divide monoids by *congruence relations*: given a monoid P , an equivalence relation $E \subseteq P \times P$ is a congruence relation if whenever $(q_1, q_2) \in E$, $p \in P$, we have $(q_1 + p, q_2 + p) \in E$. The set of equivalence classes of a congruence relation is then easily seen to inherit a monoid structure from P . If $Q \subseteq P$ is a submonoid,

then Q induces a congruence relation on P by $a \sim b$ if $a + q = b + r$ for some $q, r \in Q$. Thus in this case the quotient P/Q makes sense as a monoid.

Given homomorphisms $f_i : Q \rightarrow P_i$, $i = 1, 2$ of monoids, one can construct the fibred coproduct $P_1 \oplus_Q P_2$ as the quotient of $P_1 \oplus P_2$ by the congruence relation $(p_1, p_2) \sim (p'_1, p'_2)$ if there exists $q, r \in Q$ with

$$p_1 + f_1(r) = p'_1 + f_1(q) \text{ and } p_2 + f_2(q) = p'_2 + f_2(r).$$

For many other facts about monoids, see [88].

3.1.2. Toric varieties from fans. We fix in this section data

$$\begin{aligned} M &:= \mathbb{Z}^n, & N &:= \text{Hom}_{\mathbb{Z}}(M, \mathbb{Z}) \\ M_{\mathbb{R}} &:= M \otimes_{\mathbb{Z}} \mathbb{R}, & N_{\mathbb{R}} &:= N \otimes_{\mathbb{Z}} \mathbb{R}. \end{aligned}$$

Let $\sigma \subseteq M_{\mathbb{R}}$ be a strictly convex rational polyhedral cone, as in Definition 1.7, with dual cone $\sigma^{\vee} \subseteq N_{\mathbb{R}}$ given by

$$\sigma^{\vee} = \{n \in N_{\mathbb{R}} \mid \langle n, m \rangle \geq 0 \quad \forall m \in \sigma\}$$

as in Example 1.31, (1). Then $\sigma^{\vee} \cap N$ is a monoid. It is a standard fact, known as Gordan's Lemma, which follows from Carathéodory's theorem, that this monoid is finitely generated. Thus $\mathbb{k}[\sigma^{\vee} \cap N]$ is a finitely generated algebra and

$$X_{\sigma} := \text{Spec } \mathbb{k}[\sigma^{\vee} \cap N]$$

is the *affine toric variety* defined by the cone σ .

EXAMPLES 3.1. (1) Let σ be the first quadrant. Then σ^{\vee} is also the first quadrant, and $\sigma^{\vee} \cap N$ is the monoid \mathbb{N}^2 . The monoid ring $\mathbb{k}[\sigma^{\vee} \cap N]$ is then a polynomial ring $\mathbb{k}[x, y]$, so $X_{\sigma} \cong \mathbb{A}_{\mathbb{k}}^2$.

(2) Let σ be the cone in \mathbb{R}^2 generated by $(1, 0)$ and $(1, e)$ for e a positive integer. Then σ^{\vee} is generated by $(0, 1)$ and $(e, -1)$. Note that the monoid $\sigma^{\vee} \cap \mathbb{Z}^2$ isn't generated by $(0, 1)$ and $(e, -1)$, however, since $(1, 0)$ cannot be expressed as an integral linear combination of these two vectors. In fact, as a monoid, it is generated by $(0, 1)$, $(e, -1)$ and $(1, 0)$. These generators satisfy the obvious relation $(0, 1) + (e, -1) = e(1, 0)$. As a consequence, the monoid ring is isomorphic to $\mathbb{k}[x, y, t]/(xy - t^e)$, where $x = z^{(0,1)}$, $y = z^{(e,-1)}$, and $t = z^{(1,0)}$. So V_{σ} is a surface with an A_{e-1} singularity.

Now let Σ be a fan in $M_{\mathbb{R}}$, as defined in Definition 1.7. Then, for each cone $\sigma \in \Sigma$, we obtain an affine toric variety. Furthermore, if $\tau, \sigma \in \Sigma$ with τ a face of σ , then $\sigma^{\vee} \subseteq \tau^{\vee}$, giving an inclusion of monoid rings. This in fact describes X_{τ} as an open subset of X_{σ} . For $\sigma_1, \sigma_2 \in \Sigma$ with $\tau = \sigma_1 \cap \sigma_2$, we can then glue together X_{σ_1} and X_{σ_2} along the common open subset $X_{\tau} \subseteq X_{\sigma_1}$ and $X_{\tau} \subseteq X_{\sigma_2}$. After performing these gluings, we obtain the toric variety X_{Σ} , a separated algebraic variety over \mathbb{k} .

Properties of the fan Σ are reflected in properties of the variety X_{Σ} . For example,

- (1) X_{Σ} is proper over $\text{Spec } \mathbb{k}$ if and only if Σ is a complete fan. (See Definition 1.7).
- (2) X_{Σ} is non-singular if and only if each cone $\sigma \in \Sigma$ is *standard*, i.e., there exists a basis e_1, \dots, e_n of M such that σ is generated by e_1, \dots, e_p for some $p \leq n$.

EXAMPLE 3.2. Consider the fan Σ for \mathbb{P}^2 described in Example 1.14. There are seven cones in the fan, $0, \rho_0, \rho_1, \rho_2$ and $\sigma_{0,1}, \sigma_{1,2}, \sigma_{2,0}$ as depicted in Figure 14 of Chapter 1. The dual cones are

$$\begin{aligned} 0^\vee &= N_{\mathbb{R}} \\ \rho_0^\vee &= \{(a, b) | a + b \leq 0\} \\ \rho_1^\vee &= \{(a, b) | a \geq 0\} \\ \rho_2^\vee &= \{(a, b) | b \geq 0\} \\ \sigma_{0,1}^\vee &= \{(a, b) | a + b \leq 0, a \geq 0\} \\ \sigma_{1,2}^\vee &= \{(a, b) | a \geq 0, b \geq 0\} \\ \sigma_{2,0}^\vee &= \{(a, b) | a + b \leq 0, b \geq 0\}. \end{aligned}$$

Note that $\sigma_{i,i+1}^\vee \cap N \cong \mathbb{N}^2$ for each i (indices taken modulo 3), from which we conclude that X_Σ is covered by three copies of $\mathbb{A}_{\mathbb{k}}^2$. To see X_Σ is \mathbb{P}^2 , we define maps $X_{\sigma_{i,i+1}} \rightarrow \mathbb{P}^2$ which are compatible with the gluing maps. In particular, if \mathbb{P}^2 has homogeneous coordinates x_0, x_1, x_2 and U_i is the standard affine subset with $x_i \neq 0$, then we identify $X_{\sigma_{i,i+1}}$ with U_{i+2} via the identifications

$$\begin{aligned} \mathbb{k}[\sigma_{0,1}^\vee \cap N] &= \mathbb{k}[x_0/x_2, x_1/x_2], & \text{with } z^{(0,-1)} &= x_0/x_2, z^{(1,-1)} = x_1/x_2, \\ \mathbb{k}[\sigma_{1,2}^\vee \cap N] &= \mathbb{k}[x_1/x_0, x_2/x_0], & \text{with } z^{(1,0)} &= x_1/x_0, z^{(0,1)} = x_2/x_0, \\ \mathbb{k}[\sigma_{2,0}^\vee \cap N] &= \mathbb{k}[x_2/x_1, x_0/x_1], & \text{with } z^{(-1,1)} &= x_2/x_1, z^{(-1,0)} = x_0/x_1. \end{aligned}$$

One checks easily that these identifications are compatible with the inclusions $X_{\rho_i} \subseteq X_{\sigma_{i-1,i}}, X_{\sigma_{i,i+1}}$. For example, we can identify $\mathbb{k}[\rho_1^\vee \cap N]$ with $\mathbb{k}[x_1/x_0, x_2/x_0, x_0/x_2]$ with $z^{(1,0)} = x_1/x_0, z^{(0,1)} = x_2/x_0$ and $z^{(0,-1)} = x_0/x_2$. Then the obvious inclusions

$$\mathbb{k}[x_0/x_2, x_1/x_2] \subseteq \mathbb{k}[x_1/x_0, x_2/x_0, x_0/x_2] \supseteq \mathbb{k}[x_1/x_0, x_2/x_0]$$

coincide, under the above identifications, with

$$\mathbb{k}[\sigma_{0,1}^\vee \cap N] \subseteq \mathbb{k}[\rho_1^\vee \cap N] \supseteq \mathbb{k}[\sigma_{1,2}^\vee \cap N].$$

This shows that X_Σ is isomorphic to \mathbb{P}^2 . □

The construction of toric varieties from fans is functorial in the following sense. Suppose that M_1 and M_2 are lattices, and we have a group homomorphism $\varphi : M_1 \rightarrow M_2$. If Σ_1 is a fan in $M_1 \otimes \mathbb{R}$ and Σ_2 is a fan in $M_2 \otimes \mathbb{R}$, we say φ is a *map of fans* if for every $\sigma_1 \in \Sigma_1$, there exists a $\sigma_2 \in \Sigma_2$ such that $\varphi(\sigma_1) \subseteq \sigma_2$. This implies that the transpose map ${}^t\varphi : N_2 \rightarrow N_1$ satisfies ${}^t\varphi(\sigma_2^\vee) \subseteq \sigma_1^\vee$. Hence ${}^t\varphi$ induces a map $\mathbb{k}[\sigma_2^\vee \cap N_2] \rightarrow \mathbb{k}[\sigma_1^\vee \cap N_1]$, hence a map $X_{\sigma_1} \rightarrow X_{\sigma_2}$. These maps are compatible for various choices of σ_1 , so they patch to give a morphism

$$\varphi : X_{\Sigma_1} \rightarrow X_{\Sigma_2}.$$

This morphism is proper if the induced map $\varphi_{\mathbb{R}} : M_1 \otimes_{\mathbb{Z}} \mathbb{R} \rightarrow M_2 \otimes_{\mathbb{Z}} \mathbb{R}$ satisfies the condition that $\varphi_{\mathbb{R}}^{-1}(|\Sigma_2|) = |\Sigma_1|$.

This morphism is birational provided it induces an isomorphism on the affine subsets corresponding to the cones $0 \in \Sigma_1$ and $0 \in \Sigma_2$; this only happens if $\varphi : M_1 \rightarrow M_2$ is an isomorphism. Thus φ is a proper birational morphism if and only if φ is an isomorphism and Σ_1 is a refinement of the fan Σ_2 .

EXAMPLE 3.3. Consider the cone $\sigma \subseteq M_{\mathbb{R}} \cong \mathbb{R}^n$ generated by a basis e_1, \dots, e_n . This gives an affine toric variety $X_{\sigma} \cong \mathbb{A}_{\mathbb{k}}^n$. Now subdivide the cone σ by considering the fan Σ given by

$$\Sigma = \{\tau \mid \tau \text{ a proper face of } \sigma\} \cup \{\tau + \mathbb{R}_{\geq 0}(e_1 + \dots + e_n) \mid \tau \text{ a proper face of } \sigma\}.$$

We have introduced one new ray in the fan, generated by $e_1 + \dots + e_n$. We then have a birational morphism $X_{\Sigma} \rightarrow X_{\sigma}$. As an exercise, show that this morphism is the blow-up of X_{σ} at the origin.

The lattices M and N have the following interpretation in the context of toric varieties. First, since a fan always contains the cone 0 , and $X_0 = \text{Spec } \mathbb{k}[N]$ is an algebraic torus, X_{Σ} always contains an open subset isomorphic to an algebraic torus. This subset is often referred to as the *big torus* in X_{Σ} , or the *big torus orbit*.

In general, if L is a lattice, we use the notation

$$\mathbb{G}(L) := \text{Spec } \mathbb{k}[\text{Hom}_{\mathbb{Z}}(L, \mathbb{Z})]$$

to denote the algebraic torus determined by L . The set of \mathbb{k} -valued points of $\mathbb{G}(L)$ is $L \otimes_{\mathbb{Z}} \mathbb{k}^{\times}$.

Now the torus $\mathbb{G}(M)$ acts on X_{Σ} . This action $\mathbb{G}(M) \times_{\mathbb{k}} X_{\Sigma} \rightarrow X_{\Sigma}$ is given on an affine set X_{σ} by the ring map

$$\mathbb{k}[\sigma^{\vee} \cap N] \rightarrow \mathbb{k}[N] \otimes_{\mathbb{k}} \mathbb{k}[\sigma^{\vee} \cap N]$$

given by

$$z^n \mapsto z^n \otimes z^n.$$

This action is clearly compatible on different open sets X_{σ} , hence we obtain a torus action on X_{Σ} .

The cones of the fan Σ are in fact in one-to-one correspondence with orbits of $\mathbb{G}(M)$ acting on X_{Σ} . Indeed, the torus orbit corresponding to $\tau \in \Sigma$ is

$$X_{\tau} \setminus \bigcup_{\omega \subsetneq \tau} X_{\omega},$$

the union being over all proper faces of τ . For a point x in this orbit, the subtorus $\mathbb{G}(\mathbb{R}\tau \cap M) \subseteq \mathbb{G}(M)$ is the stabilizer of x . We denote the closure of this orbit as D_{τ} . The correspondence $\tau \mapsto D_{\tau}$ is inclusion reversing. We call the subvarieties D_{τ} , for $\tau \in \Sigma$, the *toric strata* of X_{Σ} .

D_{τ} itself is a toric variety, whose fan is given by the quotient fan $\Sigma(\tau)$ (see Definition 1.10). Note that if ρ is a ray, then D_{ρ} is a (Weil) divisor on X_{Σ} .

Next consider N . This may be interpreted as the character lattice of the algebraic torus $\mathbb{G}(M)$, i.e., $N = \text{Hom}(\mathbb{G}(M), \mathbb{G}_m)$. Here $\mathbb{G}_m = \mathbb{G}(\mathbb{Z})$ is the multiplicative group. An element $n \in N$ gives a map $\mathbb{Z} \rightarrow N$, hence a map $\mathbb{k}[\mathbb{Z}] \rightarrow \mathbb{k}[N]$, yielding the character $\mathbb{G}(M) \rightarrow \mathbb{G}_m$. This character is usually written as z^n , which can be thought of as the regular function which on the set of \mathbb{k} -valued points $M \otimes_{\mathbb{Z}} \mathbb{k}^{\times}$ of $\mathbb{G}_m(M)$ is given by $m \otimes z \mapsto z^{\langle n, m \rangle}$.

Given a fan Σ , z^n is a regular function on $\mathbb{G}(M) \subseteq X_{\Sigma}$, hence defines a rational function on X_{Σ} . Since toric varieties are always normal, it makes sense to talk of the order of vanishing of z^n on a Weil divisor D_{ρ} , for $\rho \in \Sigma^{[1]}$, the set of one-dimensional cones in Σ . This order of zero is in fact $\langle n, m_{\rho} \rangle$, where m_{ρ} is a primitive integral generator of the ray ρ . Thus

$$\sum_{\rho \in \Sigma^{[1]}} \langle n, m_{\rho} \rangle D_{\rho}$$

is a principal divisor.

One can show that the Weil divisor class group of X_Σ is given as follows. Let T_Σ denote the free abelian group generated by $\Sigma^{[1]}$ as in §1.3, with generators t_ρ for each $\rho \in \Sigma^{[1]}$. Recall the map $r : T_\Sigma \rightarrow M$ defined by $r(t_\rho) = m_\rho$, where m_ρ is a primitive generator of ρ . The transpose map ${}^t r : N \rightarrow T_\Sigma^\vee$ is given by

$$n \mapsto \sum_{\rho \in \Sigma^{[1]}} \langle n, m_\rho \rangle t_\rho^*.$$

Then $\text{Cl}(X_\Sigma) \cong \text{coker } {}^t r$. Here, an element $\psi : T_\Sigma \rightarrow \mathbb{Z}$ of T_Σ^\vee corresponds to the Weil divisor $\sum_{\rho \in \Sigma^{[1]}} \psi(t_\rho) D_\rho$, and ${}^t r(n)$ is the divisor of zeroes and poles of z^n . If r is in fact surjective, which is the case, for example, if X_Σ is non-singular and proper over $\text{Spec } \mathbb{k}$, we get an exact sequence

$$(3.1) \quad 0 \rightarrow K_\Sigma \rightarrow T_\Sigma \rightarrow M \rightarrow 0.$$

Then the dual exact sequence is

$$(3.2) \quad 0 \rightarrow N \rightarrow T_\Sigma^\vee \rightarrow \text{Cl}(X_\Sigma) \rightarrow 0,$$

A divisor induced by $\psi \in T_\Sigma^\vee$ is Cartier if ψ is induced by a PL function $\varphi : |\Sigma| \rightarrow \mathbb{R}$ with $\varphi(m_\rho) = \psi(t_\rho)$.

Given a Cartier divisor D defined by a PL function φ , the divisor D is very ample if and only if φ is strictly convex.

Of course, a Cartier divisor defines a line bundle. Given a Cartier divisor specified by a PL function φ , it is easy to describe this line bundle \mathcal{L}_φ in terms of a trivialisation. For each maximal cone $\sigma \in \Sigma$, $\varphi|_\sigma$ is given by some $n_\sigma \in N$. Then \mathcal{L}_φ is naturally identified with the trivial line bundle $\mathcal{O}_{X_\sigma} \cdot z^{-n_\sigma}$. Using this trivialization, the transition map from X_σ to $X_{\sigma'}$ is given by multiplication by $z^{n_{\sigma'} - n_\sigma}$.

If $\tau \in \Sigma$, then the restriction of the line bundle \mathcal{L}_φ to the toric stratum D_τ is easily described: it is given by the PL function $\varphi(\tau)$ on the fan $\Sigma(\tau)$ (see Definition 1.10). This is well-defined up to a choice of a linear function, which is sufficient for specifying the line bundle $\mathcal{L}_\varphi|_{D_\tau}$.

3.1.3. Toric varieties from polyhedra. In the case of toric varieties projective over affine toric varieties, there is in fact a dual point of view which is very important.

Let $\Delta \subseteq N_\mathbb{R}$ be a lattice polyhedron with at least one vertex. We can define a variety as follows. First, let $C(\Delta) \subseteq N_\mathbb{R} \oplus \mathbb{R}$ be defined by

$$C(\Delta) = \overline{\{(rn, r) | n \in \Delta, r \geq 0\}}.$$

This is the *cone over* Δ . The overline denotes closure. This cone is rational polyhedral, with

$$C(\Delta) \cap (N_\mathbb{R} \oplus \{0\}) = \text{Asym}(\Delta),$$

the asymptotic cone of Δ , defined to be the Hausdorff limit of $r\Delta$ as $r \rightarrow 0$. If Δ is compact, then of course $\text{Asym}(\Delta) = \{0\}$.

Now $\mathbb{k}[C(\Delta) \cap (N \oplus \mathbb{Z})]$ is a finitely generated \mathbb{k} -algebra which has a natural grading given by the projection $N \oplus \mathbb{Z} \rightarrow \mathbb{Z}$, i.e. $\deg z^{(n,d)} = d$. The degree 0 piece of this ring is $\mathbb{k}[\text{Asym}(\Delta) \cap N]$, so we get a variety

$$\mathbb{P}_\Delta := \text{Proj } \mathbb{k}[C(\Delta) \cap (N \oplus \mathbb{Z})]$$

which is projective over $\text{Spec } \mathbb{k}[\text{Asym}(\Delta) \cap N]$.

\mathbb{P}_Δ is in fact a toric variety: it is isomorphic to $X_{\check{\Sigma}_\Delta}$, where $\check{\Sigma}_\Delta$ is the normal fan to Δ as defined in Definition 1.9.

EXAMPLE 3.4. Let $\Delta = \text{Conv}\{(0, 0), (1, 0), (0, 1)\}$. Then $\mathbb{k}[C(\Delta) \cap (N \oplus \mathbb{Z})] \cong \mathbb{k}[x_0, x_1, x_2]$, via $z^{(0,0,1)} \mapsto x_0$, $z^{(1,0,1)} \mapsto x_1$, and $z^{(0,1,1)} \mapsto x_2$. The grading is the standard one. So $\mathbb{P}_\Delta \cong \mathbb{P}^2$. Note that the normal fan to Δ is the fan for \mathbb{P}^2 given in Example 3.2. More generally, the standard n -simplex gives rise to \mathbb{P}^n .

Using polyhedra in this way to define toric varieties gives a more geometric picture. There is a one-to-one inclusion reversing correspondence between faces of Δ and cones in the normal fan of Δ , hence a one-to-one correspondence between faces of Δ and toric strata of \mathbb{P}_Δ which is now inclusion preserving. As a result, one can think of the faces of Δ as indicating in a geometric way the combinatorics of how various strata intersect. Also, if σ is a face of Δ , then $D_{\check{\sigma}} \cong \mathbb{P}_\sigma$, where $\check{\sigma} = N_\Delta(\sigma)$, the normal cone to Δ along σ . Thus it is very easy to read off the geometry of the toric variety and its toric strata from this picture.

REMARK 3.5. We note several basic facts here. First, since \mathbb{P}_Δ is defined as a Proj, it comes along with a natural line bundle $\mathcal{O}_{\mathbb{P}_\Delta}(1)$. This line bundle is determined by the PL function $\varphi_\Delta : |\check{\Sigma}_\Delta| \rightarrow \mathbb{R}$ as defined in Definition 1.9. There is a basis for $\Gamma(\mathbb{P}_\Delta, \mathcal{O}_{\mathbb{P}_\Delta}(1))$ indexed by the integral points of Δ , i.e., by the set $\Delta \cap N$. This is clear since these integral points index a basis of the degree one part of $\mathbb{k}[C(\Delta) \cap (N \oplus \mathbb{Z})]$. Second, a face $\sigma \subseteq \Delta$ determines a cone $N_\Delta(\sigma)$ in the normal fan to Δ , and by (1.3) the dual of this cone is precisely the tangent wedge $T_\sigma\Delta$. Hence a face $\sigma \subseteq \Delta$, as well as determining a closed toric stratum of \mathbb{P}_Δ , also determines an open affine subset of \mathbb{P}_Δ given by $\text{Spec } \mathbb{k}[(T_\sigma\Delta) \cap N]$.

This is a convenient description for trivializing the line bundle $\mathcal{O}_{\mathbb{P}_\Delta}(1)$. For each vertex v of Δ , we can trivialize $\mathcal{O}_{\mathbb{P}_\Delta}(1)$ on $\text{Spec } \mathbb{k}[(T_v\Delta) \cap N]$ in such a way that the section corresponding to some $n \in \Delta \cap N$ becomes the monomial z^{n-v} . The transition map between trivializations on open sets corresponding to vertices v and v' is then given by multiplication by $z^{v-v'}$.

The following is a key example for this book.

EXAMPLE 3.6. *The Mumford degeneration.* Consider the following situation, which already appeared in §1.1. Let $\Delta \subseteq N_\mathbb{R}$ be a compact lattice polyhedron, \mathcal{P} a polyhedral decomposition of Δ into lattice polyhedra, and $\varphi : \Delta \rightarrow \mathbb{R}$ a PL convex function with integral slopes. We then have similarly as in §1.1 the polyhedron

$$\tilde{\Delta} = \{(n, r) \in N_\mathbb{R} \oplus \mathbb{R} \mid n \in \Delta, r \geq \varphi(n)\}.$$

The asymptotic cone to $\tilde{\Delta}$ is $0 \times \mathbb{R}_{\geq 0} \subseteq N_\mathbb{R} \times \mathbb{R}$, so $\mathbb{k}[C(\tilde{\Delta}) \cap (N \oplus \mathbb{Z} \oplus \mathbb{Z})]$ is a $\mathbb{k}[\mathbb{N}]$ -algebra. The toric variety $\mathbb{P}_{\tilde{\Delta}}$ is then equipped with a projective morphism $\pi : \mathbb{P}_{\tilde{\Delta}} \rightarrow \mathbb{A}_\mathbb{k}^1$.

To get a clearer idea of what π does for us, let us first describe the normal fan $\check{\Sigma}_{\tilde{\Delta}}$ to $\tilde{\Delta}$. Note that $\tilde{\Delta}$ has two sorts of faces. If $p : N_\mathbb{R} \oplus \mathbb{R} \rightarrow N_\mathbb{R}$ denotes the projection, then there are *vertical faces* of $\tilde{\Delta}$ which project non-homeomorphically via p to faces of Δ , and *horizontal faces* of $\tilde{\Delta}$, which project homeomorphically to elements of \mathcal{P} . If σ is a vertical face, then the normal cone $\check{\sigma} := N_{\tilde{\Delta}}(\sigma)$ lies in $M_\mathbb{R} \times \{0\}$, and in fact is a cone in the normal fan to Δ . On the other hand, the maximal horizontal faces σ , for which $\varphi|_{p(\sigma)}$ has slope $m_\sigma \in M$, have normal cone a ray generated by $(-m_\sigma, 1)$.




FIGURE 1

Now the morphism $\pi : \mathbb{P}_{\tilde{\Delta}} \rightarrow \mathbb{A}_{\mathbb{k}}^1$ is of course a regular function and is defined by the monomial z^ρ where $\rho = (0, 1) \in \text{Asym}(\tilde{\Delta}) \subseteq N_{\mathbb{R}} \oplus \mathbb{R}$. Since the primitive generators of the rays of $\tilde{\Sigma}_{\tilde{\Delta}}$ are either of the form $(m, 0)$ or $(m, 1)$ for various $m \in M$, z^ρ does not vanish on divisors corresponding to rays of the first type, but vanishes precisely to order one on divisors corresponding to rays of the second type. Hence $\pi^{-1}(0)$ is isomorphic to a union of toric divisors of $\mathbb{P}_{\tilde{\Delta}}$ corresponding to codimension one horizontal faces. These divisors are in one-to-one correspondence with elements of \mathcal{P}_{\max} , and we can write

$$\pi^{-1}(0) = \bigcup_{\sigma \in \mathcal{P}_{\max}} \mathbb{P}_{\sigma}.$$

Of course these toric varieties intersect precisely as dictated by \mathcal{P} . For example, suppose Δ and \mathcal{P} are as depicted in Figure 1. Then $\pi^{-1}(0)$ is a union of three \mathbb{P}^2 's and one $\mathbb{P}^1 \times \mathbb{P}^1$.

Next, observe that $\mathbb{P}_{\tilde{\Delta}} \setminus \pi^{-1}(0)$ is isomorphic to $\mathbb{P}_{\Delta} \times \mathbb{G}_m$. Indeed, we just need to localize the ring $\mathbb{k}[C(\tilde{\Delta}) \cap (N \oplus \mathbb{Z} \oplus \mathbb{Z})]$ at the element $z^{(0,1,0)}$. This is the same thing as replacing $\tilde{\Delta}$ with $\Delta \times \mathbb{R}$, and $\mathbb{P}_{\Delta \times \mathbb{R}} = \mathbb{P}_{\Delta} \times_{\mathbb{k}} \text{Spec } \mathbb{k}[\mathbb{Z}] = \mathbb{P}_{\Delta} \times_{\mathbb{k}} \mathbb{G}_m$. This shows that $\mathbb{P}_{\tilde{\Delta}} \setminus \pi^{-1}(0) \cong \mathbb{P}_{\Delta} \times_{\mathbb{k}} \mathbb{G}_m$.

Putting this all together, we see that π is a degeneration of toric varieties, with $\pi^{-1}(t) \cong \mathbb{P}_{\Delta}$ for $t \neq 0$, and $\pi^{-1}(0)$ is a union of toric varieties whose intersections are described by \mathcal{P} .

3.2. Log schemes

We begin immediately with the definition.

DEFINITION 3.7. A *pre-log structure* on a scheme X is a sheaf of monoids \mathcal{M}_X on X along with a homomorphism of sheaves of monoids

$$\alpha_X : \mathcal{M}_X \rightarrow \mathcal{O}_X$$

where the monoid structure on \mathcal{O}_X is given by multiplication.

A pre-log structure is a *log structure* if

$$\alpha_X : \alpha_X^{-1}(\mathcal{O}_X^\times) \rightarrow \mathcal{O}_X^\times$$

is an isomorphism. Here \mathcal{O}_X^\times denotes the sheaf of invertible elements of \mathcal{O}_X .

A *log scheme* is a scheme X equipped with a log structure $\alpha_X : \mathcal{M}_X \rightarrow \mathcal{O}_X$. We will usually write a log scheme as X^\dagger , with \mathcal{M}_X, α_X implicit in the notation.

A morphism of log schemes $f : X^\dagger \rightarrow Y^\dagger$ is a morphism of schemes $f : X \rightarrow Y$ along with a homomorphism of sheaves of monoids $f^\# : f^{-1}\mathcal{M}_Y \rightarrow \mathcal{M}_X$ such that the diagram

$$\begin{array}{ccc} f^{-1}\mathcal{M}_Y & \xrightarrow{f^\#} & \mathcal{M}_X \\ \downarrow \alpha_Y & & \downarrow \alpha_X \\ f^{-1}\mathcal{O}_Y & \xrightarrow{f^*} & \mathcal{O}_X \end{array}$$

is commutative. Here, f^* is the usual pull-back of regular functions defined by the morphism f .

These definitions, while straightforward, are ones I have always found hard to really absorb and internalize, and learning to think about log structures can take some time. I will try to aid this process by giving a large number of examples.

Before going on to these examples, however, there are several comments to make. First, we should mention briefly the connection between this definition and the discussion of log differentials. Essentially, the idea is that the sheaf of monoids \mathcal{M}_X specifies “things we are allowed to take $d \log$ of”. We will make this precise when we talk about log differentials in §3.3.

Second, we have not specified in which topology the sheaf \mathcal{M}_X lives. In general, when working with schemes, one should work with sheaves in the étale topology rather than the Zariski topology, as there are certainly important examples of log structures in which significant information would be lost if one only considers Zariski open subsets; see Example 3.14. If one instead works with complex analytic spaces, then one can always work comfortably with the analytic topology, and the reader uncomfortable with the notion of étale topology is best off thinking about things in the analytic topology. Nevertheless, for most of the log structures which will appear in this book, the Zariski topology will be sufficient.

Since we will, however, officially be working in the étale topology, it is worth making a few remarks for those unfamiliar with it. Many more details can be found in Milne’s book [83]. In the étale topology, one replaces the notion of an open set $U \subseteq X$ with an étale morphism $U \rightarrow X$. An étale neighbourhood of a point $x \in X$ is an étale map $U \rightarrow X$ whose image contains x . A sheaf \mathcal{F} in the étale topology then associates an abelian group, (or monoid, etc.), $\mathcal{F}(U)$ to each étale morphism $U \rightarrow X$. These come with restriction maps as usual: given $U \rightarrow V \rightarrow X$, we have a map $\mathcal{F}(V) \rightarrow \mathcal{F}(U)$. We then require \mathcal{F} to satisfy the usual sheaf axioms, appropriately stated (see, e.g., [83], Chapter II, §1).

When we talk about stalks of sheaves, we need to choose a *geometric point* of the scheme X . Let $x \in X$ have residue field $k(x)$, and fix a separable closure $k(\bar{x})$ of $k(x)$. Let $\bar{x} = \text{Spec } k(\bar{x})$, and let $\bar{x} \rightarrow X$ be the induced map. This is called a geometric point of X . Then we define the stalk $\mathcal{F}_{\bar{x}}$ of \mathcal{F} at \bar{x} as the direct limit of groups (monoids, etc.,) $\mathcal{F}(U)$ running over diagrams

$$\begin{array}{ccc} \bar{x} & \longrightarrow & U \\ & \searrow & \downarrow \\ & & X \end{array}$$

For example, the sheaf \mathcal{O}_X , in the étale topology, is defined by $\mathcal{O}_X(U) = \Gamma(U, \mathcal{O}_U)$, where \mathcal{O}_U is the usual structure sheaf on U . The stalks $\mathcal{O}_{X, \bar{x}}$ are considerably bigger than the usual local rings $\mathcal{O}_{X, x}$.

These details will not play a particularly important role for us in this text, but it seems necessary to introduce this notation so that the statements given here are technically correct.

EXAMPLES 3.8. (0) If X is a scheme, and $\mathcal{M}_X = \mathcal{O}_X^\times$, α_X the inclusion, we obtain a log structure on X known as the *trivial log structure*. A log morphism between schemes carrying the trivial log structure is the same thing as a morphism of schemes.

(1) *The fundamental example.* Let X be a scheme, $D \subseteq X$ a closed subset of pure codimension one. Let $j : X \setminus D \hookrightarrow X$ be the inclusion, and consider

$$\mathcal{M}_{(X, D)} := (j_* \mathcal{O}_{X \setminus D}^\times) \cap \mathcal{O}_X.$$

This is the sheaf of regular functions on X which are invertible on $X \setminus D$. We take $\alpha_X : \mathcal{M}_{(X, D)} \hookrightarrow \mathcal{O}_X$ to be the inclusion. This is obviously a log structure, and is known as the *divisorial log structure induced by D* .

(2) Let \mathbb{k} be a field, $X = \text{Spec } \mathbb{k}$. Let $\mathcal{M}_X = \mathbb{k}^\times \oplus Q$, where Q is a monoid whose only invertible element is $0 \in Q$. Let $\alpha_X : \mathbb{k}^\times \oplus Q \rightarrow \mathbb{k}$ be given by

$$\alpha_X(x, q) := \begin{cases} x & q = 0 \\ 0 & q \neq 0 \end{cases}$$

This is a log structure. There are two special cases: if $Q = \{0\}$, then we get the trivial log structure on X . If $Q = \mathbb{N}$, we get what is known as the *standard log point*. We usually write the standard log point as $\text{Spec } \mathbb{k}^\dagger$.

There is a relationship between divisorial log structures and the standard log point, which arises from the following general construction.

DEFINITION 3.9. Let $\alpha : P_X \rightarrow \mathcal{O}_X$ be a pre-log structure on X . Then the *log structure associated to this pre-log structure* is given by the sheaf of monoids

$$\mathcal{M}_X := \frac{P_X \oplus \mathcal{O}_X^\times}{\{(p, \alpha(p)^{-1}) \mid p \in \alpha^{-1}(\mathcal{O}_X^\times)\}}$$

and $\alpha_X : \mathcal{M}_X \rightarrow \mathcal{O}_X$ is given by

$$\alpha_X(p, f) := \alpha(p) \cdot f.$$

This is clearly a log structure.

DEFINITION 3.10. If $f : X \rightarrow Y$ is a morphism of schemes and Y has a log structure $\alpha_Y : \mathcal{M}_Y \rightarrow \mathcal{O}_Y$, then the *pull-back log structure* on X is the log structure associated to the pre-log structure given by the composition

$$\alpha : f^{-1}(\mathcal{M}_Y) \xrightarrow{\alpha_Y} f^{-1}(\mathcal{O}_Y) \xrightarrow{f^*} \mathcal{O}_X.$$

We will write the sheaf of monoids of this pull-back log structure as $f^* \mathcal{M}_Y$.

EXAMPLE 3.11. Let $X = \text{Spec } \mathbb{k}$, $Y = \text{Spec } \mathbb{k}[x] = \mathbb{A}_{\mathbb{k}}^1$, and let $f : X \rightarrow Y$ map X to the origin in Y . Consider the divisorial log structure on $\mathbb{A}_{\mathbb{k}}^1$ given by $\{0\} \subseteq Y$, and pull it back to X via f . Now, $f^{-1} \mathcal{M}_Y$ is just the stalk of the sheaf \mathcal{M}_Y at 0;

this consists of germs of functions of the form $\varphi \cdot x^n$, $n \geq 0$, where φ is an invertible function defined in an (étale) neighbourhood of 0. The map

$$\alpha : f^{-1}\mathcal{M}_Y \rightarrow f^{-1}\mathcal{O}_Y \rightarrow \mathcal{O}_X = \mathbb{k}$$

is then evaluation at $x = 0$ of $\varphi \cdot x^n$, and hence

$$\alpha(\varphi \cdot x^n) = \begin{cases} \varphi(0) & n = 0, \\ 0 & n > 0. \end{cases}$$

This defines a pre-log structure on X . Taking the associated log structure, we consider

$$\frac{f^{-1}\mathcal{M}_Y \oplus \mathbb{k}^\times}{\{(\varphi, \varphi(0)^{-1}) \mid \varphi \in \mathcal{O}_{Y,0}^\times\}}.$$

This is in fact isomorphic to $\mathbb{k}^\times \oplus \mathbb{N}$ via the map

$$f^{-1}\mathcal{M}_Y \oplus \mathbb{k}^\times \ni (\varphi \cdot x^n, s) \mapsto (\varphi(0) \cdot s, n).$$

Furthermore,

$$\alpha_X(\varphi \cdot x^n, s) = \begin{cases} \varphi(0) \cdot s & n = 0 \\ 0 & n > 0 \end{cases}$$

and hence the pull-back log structure on X yields the standard log point. \square

It is often useful to think of log structures by considering the exact sequence

$$1 \longrightarrow \mathcal{O}_X^\times \xrightarrow{\alpha_X^{-1}} \mathcal{M}_X \longrightarrow \overline{\mathcal{M}}_X \longrightarrow 0.$$

The sheaf of monoids $\overline{\mathcal{M}}_X$, written additively, is called the *ghost sheaf* of X^\dagger , and should be viewed as containing combinatorial information about the log structure.

EXAMPLE 3.12. Suppose X is locally Noetherian and normal, $D \subseteq X$ a closed subset of pure codimension one. Then for $x \in X$, let r be the number of irreducible components of D containing x . There is a map

$$q : \mathcal{M}_{(X,D),\bar{x}} \rightarrow \mathbb{N}^r$$

given by associating to a regular function $f \in \mathcal{M}_{(X,D),\bar{x}}$ the vanishing orders of f along the r components of D containing x . This map clearly factors through $\overline{\mathcal{M}}_{(X,D),\bar{x}}$. On the other hand, if $f_1, f_2 \in \mathcal{M}_{(X,D),\bar{x}} \subseteq \mathcal{O}_{X,\bar{x}}$ have the same vanishing orders along D , then $f_1 f_2^{-1}$ vanishes to order zero on every prime divisor in a neighbourhood of \bar{x} . Since X is normal, it then follows (see e.g., the argument given in [57], page 132) that $f_1 = f_2 \cdot h$ for $h \in \mathcal{O}_{X,\bar{x}}^\times$. Thus $\overline{\mathcal{M}}_{(X,D),\bar{x}}$ is a submonoid of \mathbb{N}^r . \square

Note that if $f : X \rightarrow Y$ is a morphism of schemes and Y carries a log structure yielding a pull-back log structure on X , then $\overline{\mathcal{M}}_X = f^{-1}\overline{\mathcal{M}}_Y$. Indeed, we have a

diagram

$$\begin{array}{ccccccc}
 & & 1 & & 1 & & \\
 & & \downarrow & & \downarrow & & \\
 & & K & \xrightarrow{=} & K & & \\
 & & \downarrow & & \downarrow & & \\
 1 & \longrightarrow & f^{-1}\mathcal{O}_Y^\times \oplus \mathcal{O}_X^\times & \longrightarrow & f^{-1}\mathcal{M}_Y \oplus \mathcal{O}_X^\times & \longrightarrow & f^{-1}\overline{\mathcal{M}}_Y \longrightarrow 0 \\
 & & \downarrow & & \downarrow & & \downarrow \\
 & & \mathcal{O}_X^\times & \longrightarrow & \mathcal{M}_X & \longrightarrow & \overline{\mathcal{M}}_X \longrightarrow 0 \\
 & & \downarrow & & \downarrow & & \\
 & & 1 & & 1 & &
 \end{array}$$

where K is both the kernel of the map $f^{-1}\mathcal{O}_Y^\times \oplus \mathcal{O}_X^\times \rightarrow \mathcal{O}_X^\times$, given by $(\varphi_1, \varphi_2) \mapsto f^*(\varphi_1) \cdot \varphi_2$, and the kernel of the map $f^{-1}\mathcal{M}_Y \oplus \mathcal{O}_X^\times \rightarrow \mathcal{M}_X$. (Kernels do not in general exist for morphisms of monoids, but do if the morphism is given as the quotient by a subgroup.) This latter kernel, by definition, is given by

$$\{(\varphi_1, \varphi_2) \in f^{-1}\mathcal{M}_Y \oplus \mathcal{O}_X^\times \mid f^*\alpha_Y(\varphi_1) = \varphi_2^{-1}\}.$$

One checks easily that the snake lemma works in this context.

EXAMPLE 3.13. Let $X = \mathbb{A}_{\mathbb{k}}^2 = \text{Spec } \mathbb{k}[x, y]$, and let $D \subseteq X$ be given by $xy = 0$, i.e., $D = V(xy)$. Then we obtain the divisorial log structure $\mathcal{M}_{(X,D)}$ on X , and $\overline{\mathcal{M}}_{(X,D)}$ is the sheaf $(i_{1*}\mathbb{N}) \oplus (i_{2*}\mathbb{N})$, where $i_1 : V(x) \hookrightarrow X$ and $i_2 : V(y) \hookrightarrow X$ are the inclusions of the two coordinate axes. The map $\mathcal{M}_{(X,D)} \rightarrow (i_{1*}\mathbb{N}) \oplus (i_{2*}\mathbb{N})$ takes a regular function f on an open set U invertible on $U \setminus D$ to the order of vanishings of f on $U \cap V(x)$ and $U \cap V(y)$. As in Example 3.12, this map is surjective and has kernel \mathcal{O}_X^\times .

If we pull-back $\mathcal{M}_{(X,D)}$ via the inclusion $D \hookrightarrow X$, we obtain a log structure on D , with $\overline{\mathcal{M}}_D = i_{1*}\mathbb{N} \oplus i_{2*}\mathbb{N}$ (now thinking of i_1, i_2 as the inclusions of the coordinate axes into D). In particular, there are sections of \mathcal{M}_D which cannot be thought of as functions on D . One should think of the log structure on D as remembering some information about how D sits inside X .

EXAMPLE 3.14. Suppose we have a divisor $D \subseteq X$ of the following form. D is an irreducible, non-normal surface whose normalization \tilde{D} is isomorphic to $E \times \mathbb{P}^1$, where E is an elliptic curve. Let $\tau \in E$ be a two-torsion point and $P \in \mathbb{P}^1$ a point. We assume that the map $\tilde{D} \rightarrow D$ identifies the point $e \times \{P\}$ with the point $(e + \tau) \times \{P\}$. Let E' be the quotient of E by translation by τ . Then there is a map $D \rightarrow E'$, a fibre of which is a union of two \mathbb{P}^1 's meeting at a point. The singular locus of D is isomorphic to E' . Now we have the divisorial log structure $\mathcal{M}_{(X,D)}$ on X , which we can restrict to a log structure $\mathcal{M}_{E'}$ on E' using the identification of E' with the singular locus of D . Note that $\overline{\mathcal{M}}_{E'}$ is now a locally constant sheaf with fibre \mathbb{N}^2 , but only in the étale topology or the analytic topology, not in the Zariski topology. Locally at a point in E' in either of these two topologies, D has two irreducible components, but globally, these components get interchanged.

So the locally constant sheaf $\overline{\mathcal{M}}_{E'}$ has non-trivial monodromy, interchanging the generators of \mathbb{N}^2 . We can't see this monodromy in the Zariski topology.

This explains why, in general, the Zariski topology is not sufficient. However, this type of example will never occur in the cases needed in this book.

EXAMPLE 3.15. Our next, rather extended, example explores the data carried by log morphisms. We will find that a log morphism can carry a lot of extra data. Let X^\dagger be the affine plane of Example 3.13. We can think of this example as follows. Fix $M = \mathbb{Z}^2$, $M_{\mathbb{R}}, N, N_{\mathbb{R}}$ as usual, and let $\sigma \subseteq M_{\mathbb{R}}$ be the first quadrant. Then $X = X_\sigma$, since $\sigma^\vee \cap N = \mathbb{N}^2$.

Now consider a log morphism $f : \text{Spec } \mathbb{k}^\dagger \rightarrow X^\dagger$ from the standard log point, with image the origin. The additional information associated to the log morphism is a commutative diagram

$$\begin{array}{ccc} \mathcal{M}_{(X,D),\bar{0}} & \xrightarrow{f^\#} & \mathbb{k}^\times \oplus \mathbb{N} \\ \alpha_X \downarrow & & \downarrow \alpha_{\text{Spec } \mathbb{k}^\dagger} \\ \mathcal{O}_{X,\bar{0}} & \xrightarrow{f^*} & \mathbb{k} \end{array}$$

Here f^* is just evaluation of a germ of a function at 0, and $\mathcal{M}_{(X,D),\bar{0}}$ consists of germs of functions invertible on $X \setminus D$, i.e., locally of the form $\varphi \cdot z^n$ for $n \in \sigma^\vee \cap N$ and $\varphi \in \mathcal{O}_{X,\bar{0}}^\times$. Note that $f^\#$ induces a map $\mathbb{N}^2 = \overline{\mathcal{M}}_{(X,D),\bar{0}} \rightarrow \overline{\mathcal{M}}_{\text{Spec } \mathbb{k}^\dagger} = \mathbb{N}$, a monoid homomorphism, so this is discrete data determined by the morphism. This map is given by $n \mapsto \langle n, m \rangle$ for some $m \in \sigma \cap M$. Furthermore, $m \in \text{Int}(\sigma)$. Indeed, suppose $m = (a, 0)$ for some $a \in \mathbb{N}$. Then $f^\#(z^{(0,1)}) = (x, 0) \in \mathbb{k}^\times \oplus \mathbb{N}$ for some $x \in \mathbb{k}^\times$, and $x = \alpha_{\text{Spec } \mathbb{k}^\dagger}(f^\#(z^{(0,1)})) = f^* \alpha_X(z^{(0,1)}) = 0$, a contradiction. Similarly, $m = (0, a)$ leads to a contradiction.

Note also that if φ is invertible, then $f^* \alpha_X(\varphi) = \varphi(0)$, so $f^\#(\varphi) = (\varphi(0), 0)$. From this we conclude that $f^\#$ is completely determined by m and a map of monoids $g : \sigma^\vee \cap N \rightarrow \mathbb{k}^\times$, so that

$$(3.3) \quad f^\#(\varphi z^n) = (\varphi(0)g(n), \langle n, m \rangle).$$

So we see the map is completely specified by $m \in \text{Int}(\sigma) \cap M$ and $g \in M \otimes_{\mathbb{Z}} \mathbb{k}^\times$.

Let us elaborate on this to attempt to explain to a certain extent what this additional data means. Consider a subdivision of the cone σ obtained by introducing another ray $\mathbb{R}_{\geq 0} m'$, where $m' \in \text{Int}(\sigma) \cap M$ is primitive. This gives a fan Σ with three rays, $\rho_1 = \mathbb{R}_{\geq 0}(1, 0)$, $\rho_2 = \mathbb{R}_{\geq 0}(0, 1)$, $\rho = \mathbb{R}_{\geq 0} m'$, and two two-dimensional cones. The identity map on M induces a map of fans from Σ to σ , since each cone of Σ is contained in σ , and hence we obtain a morphism

$$\pi : \tilde{X} := X_\Sigma \rightarrow X_\sigma = X.$$

As an exercise in toric geometry, one can check that this map is proper, and induces an isomorphism $\tilde{X} \setminus D_\rho \cong X \setminus \{(0, 0)\}$. This map is called a toric, or weighted, blow-up, and if $m' = (1, 1)$, this is just the ordinary blow-up of $\mathbb{A}_{\mathbb{k}}^2$ at the origin. (See Example 3.3.)

Let $\tilde{D} \subseteq \tilde{X}$ be the union of the toric divisors on \tilde{X} ; note that $\tilde{D} = \pi^{-1}(D)$. Thus we obtain a log structure $\mathcal{M}_{(\tilde{X}, \tilde{D})}$ on \tilde{X} , and a log morphism $\pi : \tilde{X}^\dagger \rightarrow X^\dagger$, since functions on X which only vanish on D pull-back to functions on \tilde{X} which

only vanish on \tilde{D} . Now given $f : \text{Spec } \mathbb{k}^\dagger \rightarrow X^\dagger$ as described above, let us ask: when does this map lift to give a commutative diagram

$$\begin{array}{ccc} & & \tilde{X}^\dagger \\ & \nearrow \tilde{f} & \downarrow \pi \\ \text{Spec } \mathbb{k}^\dagger & \xrightarrow{f} & X^\dagger \end{array}$$

such that the image of $\text{Spec } \mathbb{k}^\dagger$ lies in $D_\rho \setminus \text{Sing}(\tilde{D})$? (This is the open torus orbit of D_ρ .)

To answer this question, first consider what it means to give the map \tilde{f} . This works in the same way as giving the map f . Choose a basis n_1, n_2 of N such that $\langle n_1, m' \rangle = 1$ and $\langle n_2, m' \rangle = 0$, so that the cone ρ^\vee is generated by n_1 and $\pm n_2$ and z^{n_1} vanishes to order one along D_ρ . We specify the image of \tilde{f} to be a point $x \in D_\rho \setminus \text{Sing}(\tilde{D})$, and now $\mathcal{M}_{(\tilde{X}, \tilde{D}), \tilde{x}}$ locally consists of functions of the form $\varphi \cdot z^{an_1}$ for $a \in \mathbb{N}$ and φ invertible in a neighbourhood of \tilde{x} . Then the map $\tilde{f}^\#$ is given by

$$\tilde{f}^\#(\varphi \cdot z^{an_1}) = (\varphi(x)\tilde{g}(a), ab) \in \mathbb{k}^\times \oplus \mathbb{N},$$

for some $\tilde{g} : \mathbb{N} \rightarrow \mathbb{k}^\times$ and $b \in \mathbb{N}$, $b \neq 0$.

For $\pi \circ \tilde{f}$ to agree with f , we need $\tilde{f}^\# \circ \pi^\# = f^\#$. Now for

$$\varphi \cdot z^{\beta_1 n_1 + \beta_2 n_2} \in \mathcal{M}_{(X, D), 0},$$

we have

$$\begin{aligned} \tilde{f}^\# \circ \pi^\#(\varphi \cdot z^{\beta_1 n_1 + \beta_2 n_2}) &= \tilde{f}^\#((\varphi \circ \pi) \cdot z^{\beta_1 n_1 + \beta_2 n_2}) \\ &= (\varphi(\pi(x)) \cdot z^{\beta_2 n_2}(x) \cdot \tilde{g}(\beta_1), \beta_1 b). \end{aligned}$$

Here, $z^{\beta_2 n_2}(x)$ means the value of the rational function $z^{\beta_2 n_2}$ at x . Note that z^{n_2} in fact is a coordinate on D_ρ . From the previous formula for $f^\#$, we see first that we need $\beta_1 b = \beta_1 \langle n_1, m \rangle + \beta_2 \langle n_2, m \rangle$ for all β_1, β_2 such that $\beta_1 n_1 + \beta_2 n_2 \in \sigma^\vee$. So in particular, we must have $\langle n_2, m \rangle = 0$, i.e., m is proportional to m' , and $b = \langle n_1, m \rangle$. So we must have $m = bm'$. Second, we need $g(\beta_1 n_1 + \beta_2 n_2) = (z^{n_2}(x))^{\beta_2} \cdot \tilde{g}(\beta_1)$. So in particular, g determines $z^{n_2}(x)$, hence the point x itself.

So morally, the choice of log point is specifying a toric blow-up along with a point on the exceptional divisor of that toric blow-up. There is still some residual information, namely the choice of b and $\tilde{g}(1)$. However, in Example 3.17, we will consider a situation where this extra degree of freedom disappears.

EXAMPLE 3.16. Consider next a morphism $X^\dagger \rightarrow \text{Spec } \mathbb{k}^\dagger$ from a log scheme X^\dagger to the standard log point. First of all X is a scheme over \mathbb{k} , and we have a map $f^\# : \mathbb{k}^\times \times \mathbb{N} \rightarrow \mathcal{M}_X$. The map $\mathbb{k}^\times \rightarrow \mathcal{M}_X$ is completely determined, given by $c \mapsto \alpha_X^{-1}(c)$. So $f^\#$ is completely determined by a map $\mathbb{N} \rightarrow \mathcal{M}_X$. This in turn is determined by giving a section $\rho \in \Gamma(X, \mathcal{M}_X)$ which is the image of $1 \in \mathbb{N}$. Note that we require $\alpha_X(\rho) = f^*(\alpha_{\text{Spec } \mathbb{k}^\dagger}(0, 1)) = 0$.

So a log scheme over the standard log point is just a log scheme over \mathbb{k} along with a section ρ of \mathcal{M}_X with $\alpha_X(\rho) = 0$.

EXAMPLE 3.17. Returning to the situation of Example 3.15, consider the map $\pi : X \rightarrow \mathbb{A}_{\mathbb{k}}^1$ given by the monomial $z^{(1,1)}$. The fibre over $0 \in \mathbb{A}_{\mathbb{k}}^1$ is just D . We have the divisorial log structure $\mathcal{M}_{(\mathbb{A}_{\mathbb{k}}^1, 0)}$, and a log morphism $\pi : X^\dagger \rightarrow (\mathbb{A}_{\mathbb{k}}^1)^\dagger$,

again because functions on $\mathbb{A}_{\mathbb{k}}^1$ invertible outside of 0 pull-back to functions on X invertible outside of D . Restricting these log structures to D and 0 respectively gives a log morphism

$$\pi_0 : D^\dagger \rightarrow \mathrm{Spec} \mathbb{k}^\dagger.$$

In terms of Example 3.16, this log morphism is determined by the section of \mathcal{M}_D given by restricting the section $z^{(1,1)}$ of $\mathcal{M}_{(X,D)}$ to D .

For $d > 0$ an integer, let $\eta_d : \mathrm{Spec} \mathbb{k}^\dagger \rightarrow \mathrm{Spec} \mathbb{k}^\dagger$ be the log morphism such that

$$\eta_d^\#(c, \beta) = (c, d\beta).$$

This can be thought of as a degree d base-change, as it is induced by the morphism of log schemes

$$\xi_d : (\mathbb{A}_{\mathbb{k}}^1)^\dagger \rightarrow (\mathbb{A}_{\mathbb{k}}^1)^\dagger$$

given, on the level of schemes, by the ring homomorphism $x \mapsto x^d$.

Now consider log points $f : \mathrm{Spec} \mathbb{k}^\dagger \rightarrow D^\dagger$ giving a commutative diagram

$$\begin{array}{ccc} & & D^\dagger \\ & \nearrow f & \downarrow \pi_0 \\ \mathrm{Spec} \mathbb{k}^\dagger & \xrightarrow{\eta_d} & \mathrm{Spec} \mathbb{k}^\dagger \end{array}$$

Assuming that the image of f is the origin in D , one can check easily that the data determining $f^\#$, before imposing the condition that this diagram commutes, is exactly the data appearing in (3.3). The condition that the above diagram commutes means that $f^\# \circ \pi_0^\# = \eta_d^\#$. But for $(c, \beta) \in \mathbb{k}^\times \oplus \mathbb{N}$,

$$\begin{aligned} f^\# \circ \pi_0^\#(c, \beta) &= f^\#(cz^{(\beta, \beta)}) \\ &= (cg(\beta, \beta), \langle m, (\beta, \beta) \rangle). \end{aligned}$$

Thus for this to be $\eta_d^\#(c, \beta)$, we require $g(1, 1) = 1$ and $\langle m, (1, 1) \rangle = d$. This tells us that $g \in ((1, 1)^\perp \cap M) \otimes_{\mathbb{Z}} \mathbb{k}^\times$, a one-dimensional torus, and m is subject to the constraint $\langle m, (1, 1) \rangle = d$. In particular, for $d = 1$, there are no such maps, since, as in Example 3.15, we need $m \in \mathrm{Int}(\sigma)$.

Roughly put, this can be viewed as detecting certain residual data about multi-sections of $\pi : X \rightarrow \mathbb{A}^1$. Suppose we are given a commutative diagram

$$\begin{array}{ccc} & & X \\ & \nearrow \bar{f} & \downarrow \pi \\ \mathbb{A}_{\mathbb{k}}^1 & \xrightarrow{\xi_d} & \mathbb{A}_{\mathbb{k}}^1 \end{array}$$

with $\bar{f}(0) = 0$. If $d = 1$, this is in fact impossible, since a section of the map π can never intersect a singular point of a fibre. On the other hand, for $d > 1$, this is possible, and there is always a unique toric blow-up $\tilde{X} \rightarrow X$ such that the map \bar{f} lifts to a map $\tilde{f} : \mathbb{A}_{\mathbb{k}}^1 \rightarrow \tilde{X}$ such that $\tilde{f}(0)$ is in the big torus orbit of the exceptional divisor. Restricting all these maps to D^\dagger and $\mathrm{Spec} \mathbb{k}^\dagger$ then gives a map $f : \mathrm{Spec} \mathbb{k}^\dagger \rightarrow D^\dagger$ as above, and the log morphism f is remembering which toric blow-up is necessary and what $\tilde{f}(0)$ is. \square

The category of log schemes is in general far too broad: there are many very perverse examples. A crucial notion is that of a fine log structure.

DEFINITION 3.18. Let P be a monoid, X a scheme. Denote by \underline{P} the constant sheaf on X with stalk P , and suppose we have a pre-log structure $\underline{P} \rightarrow \mathcal{O}_X$. Then we say $\underline{P} \rightarrow \mathcal{O}_X$ is a *chart* for the associated log structure. Two charts $\alpha_1 : \underline{P} \rightarrow \mathcal{O}_X$, $\alpha_2 : \underline{P} \rightarrow \mathcal{O}_X$ are said to be *equivalent* if there exists a map $\varphi : \underline{P} \rightarrow \mathcal{O}_X^\times$ such that

$$\alpha_2(p) = \alpha_1(p) \cdot \varphi(p) \quad \forall p \in P.$$

Equivalent charts are easily seen to induce isomorphic log structures.

A log structure on X is *fine* if there is an étale open cover $\{U_i\}$ such that on U_i , there is a finitely generated integral monoid P_i and a pre-log structure $\underline{P}_i \rightarrow \mathcal{O}_{U_i}$ whose associated log structure is isomorphic to the log structure pulled back from X .

A monoid P is *saturated* if it is integral and whenever $p \in P^{\text{gp}}$ with $mp \in P$ for some positive integer m , one has $p \in P$.

A log structure on X is *fine saturated* if it is fine and $\overline{\mathcal{M}}_{X,\bar{x}}$ is saturated for all $\bar{x} \in X$.

EXAMPLE 3.19. Let P be a toric monoid, i.e., a monoid of the form $P = \sigma^\vee \cap N$ for some $\sigma \subseteq M_{\mathbb{R}}$ a strictly convex rational polyhedral cone, and let $X = X_\sigma = \text{Spec } \mathbb{k}[P]$. The map

$$P \rightarrow \Gamma(X, \mathcal{O}_X) = \mathbb{k}[P]$$

given by

$$p \mapsto z^p$$

induces a morphism of sheaves of monoids

$$\underline{P} \rightarrow \mathcal{O}_X,$$

hence an associated log structure on X .

If ∂X denotes the *toric boundary* of X , i.e., the complement of the big torus orbit in X , then in fact this log structure on X is the divisorial log structure coming from $\partial X \subseteq X$. Indeed, the map $\underline{P} \rightarrow \mathcal{O}_X$ factors as

$$\underline{P} \rightarrow \mathcal{M}_{(X, \partial X)} \xrightarrow{\alpha_X} \mathcal{O}_X$$

via $P \ni p \mapsto z^p \in \mathcal{M}_{(X, \partial X)}$, since z^p is a regular function on X invertible on the big torus orbit $X \setminus \partial X$. This gives a map

$$\underline{P} \oplus \mathcal{O}_X^\times \rightarrow \mathcal{M}_{(X, \partial X)}$$

whose kernel on an open set $U \subseteq X$ consists precisely of pairs (p, z^{-p}) such that z^{-p} is invertible on U . Furthermore, locally at $x \in X$, a function defined in a neighbourhood U of x which is invertible on $U \setminus \partial X$ is always of the form $\varphi \cdot z^p$ for φ on U invertible and $p \in P$. So $\underline{P} \oplus \mathcal{O}_X^\times \rightarrow \mathcal{M}_{(X, \partial X)}$ is surjective, showing that the log structure associated to this chart is the divisorial log structure. \square

This example demonstrates that a chart on an (étale) open subset U of X can be thought of as a map $U \rightarrow \text{Spec } \mathbb{Z}[P]$ such that the log structure on U is the pull-back of the divisorial log structure on $\text{Spec } \mathbb{Z}[P]$ induced by the toric boundary of $\text{Spec } \mathbb{Z}[P]$.

EXAMPLE 3.20. Here is an example of a log structure which occurs naturally, but is not fine. Let $X = \text{Spec } \mathbb{k}[x, y, w, t]/(xy - wt)$, and let $D \subseteq X$ be the divisor given by $t = 0$, inducing a divisorial log structure $\mathcal{M}_{(X, D)}$. I claim that this log structure is not fine at $0 = (0, 0, 0, 0) \in X$. Indeed, suppose there is an étale

neighbourhood $U \rightarrow X$ of 0 and a chart $\underline{P} \rightarrow \mathcal{O}_U$ for $\mathcal{M}_{(X,D)}$. Denoting the pull-back of $\mathcal{M}_{(X,D)}$ to U by \mathcal{M}_U , this means we have an isomorphism $(\underline{P} \oplus \mathcal{O}_U^\times)/K \rightarrow \mathcal{M}_U$, where dividing out by K defines the log structure associated to the pre-log structure $\underline{P} \rightarrow \mathcal{O}_U$. In particular, there is a surjective map $\underline{P} \rightarrow \overline{\mathcal{M}}_U$. This is surjective on stalks and for any geometric point $\bar{x} \in U$ we have a commutative diagram

$$\begin{array}{ccc} P = \Gamma(U, \underline{P}) & \longrightarrow & \Gamma(U, \overline{\mathcal{M}}_U) \\ \downarrow \cong & & \downarrow \\ P = \underline{P}_{\bar{x}} & \longrightarrow & \overline{\mathcal{M}}_{U, \bar{x}} \end{array}$$

so $\Gamma(U, \overline{\mathcal{M}}_U) \rightarrow \overline{\mathcal{M}}_{U, \bar{0}}$ must be surjective. However, if $\varphi \in \mathcal{O}_{X, \bar{0}}$ is a function whose zero locus is contained in D , then $\varphi = \psi \cdot t^n$ with $\psi \in \mathcal{O}_{X, \bar{0}}^\times$. Indeed, in a neighbourhood of 0, the only Cartier divisor with support in D is a multiple of D . Thus there is some non-negative integer n such that $t^{-n} \cdot \varphi$ does not vanish along D and hence, by normality, $t^{-n} \cdot \varphi$ is an invertible regular function. Thus $\overline{\mathcal{M}}_{X, \bar{0}} = \mathbb{N} \hookrightarrow \mathbb{N}^2$ via the diagonal embedding as in Example 3.12. On the other hand, for p a point in X with coordinates $x = y = 0$ and $w \neq 0$, both of the irreducible components of D are Cartier, and $\overline{\mathcal{M}}_{(X,D), \bar{p}} = \mathbb{N}^2$. Thus there is some étale neighbourhood of 0 on which $\Gamma(U, \overline{\mathcal{M}}_{(X,D)}) = \mathbb{N}$, and this cannot surject to $\overline{\mathcal{M}}_{(X,D), \bar{p}} = \mathbb{N}^2$. Thus the log structure can't be fine. \square

The following notion is often useful:

DEFINITION 3.21. A morphism $f : X^\dagger \rightarrow Y^\dagger$ is *strict* if $f^\# : f^{-1}\mathcal{M}_Y \rightarrow \mathcal{M}_X$ induces an isomorphism between the pull-back of the log structure on Y to X and the log structure on X .

EXAMPLES 3.22. (1) Let $f : X \rightarrow \mathbb{A}_{\mathbb{k}}^1$ be a *smooth* morphism, and consider the divisorial log structures on X and $\mathbb{A}_{\mathbb{k}}^1$ given by $f^{-1}(0) \subseteq X$ and $0 \in \mathbb{A}_{\mathbb{k}}^1$. Then f is a strict morphism of log schemes.

(2) Let X be obtained by blowing up $\mathbb{A}_{\mathbb{k}}^1 \times \mathbb{P}_{\mathbb{k}}^1$ at the point $(0, (1 : 0))$, and let $f : X \rightarrow \mathbb{A}_{\mathbb{k}}^1$ be the composition of the blow-down followed by projection to $\mathbb{A}_{\mathbb{k}}^1$. Give X the divisorial log structure given by the total transform of

$$(\mathbb{A}_{\mathbb{k}}^1 \times \{(1 : 0)\}) \cup (\{0\} \times \mathbb{P}_{\mathbb{k}}^1) \cup (\mathbb{A}_{\mathbb{k}}^1 \times \{(0 : 1)\}),$$

and give $\mathbb{A}_{\mathbb{k}}^1$ the divisorial log structure coming from $0 \in \mathbb{A}_{\mathbb{k}}^1$. Then f , as a morphism of log schemes, is strict everywhere *except* at the proper transforms of $\mathbb{A}_{\mathbb{k}}^1 \times \{(1 : 0)\}$, $\mathbb{A}_{\mathbb{k}}^1 \times \{(0 : 1)\}$, and the double point of $f^{-1}(0)$. \square

We next define the notion of log smoothness. Like ordinary morphisms of schemes, where smoothness can be defined using an infinitesimal lifting criterion, we can give the same condition in the log case. However, for practical purposes, it is better to give an equivalent definition.

DEFINITION 3.23. A morphism $f : X^\dagger \rightarrow Y^\dagger$ of fine log schemes is *log smooth* if étale locally on X and Y it fits into a commutative diagram

$$\begin{array}{ccc} X & \longrightarrow & \text{Spec } \mathbb{Z}[P] \\ \downarrow & & \downarrow \\ Y & \longrightarrow & \text{Spec } \mathbb{Z}[Q] \end{array}$$

with the following properties:

- (1) The horizontal maps induce charts $\underline{P} \rightarrow \mathcal{O}_X$ and $\underline{Q} \rightarrow \mathcal{O}_Y$ for X^\dagger and Y^\dagger .
- (2) The induced morphism

$$X \rightarrow Y \times_{\mathrm{Spec} \mathbb{Z}[Q]} \mathrm{Spec} \mathbb{Z}[P]$$

is a smooth morphism of schemes.

- (3) The right-hand vertical arrow is induced by a monoid homomorphism $Q \rightarrow P$ with $\ker(Q^{\mathrm{gp}} \rightarrow P^{\mathrm{gp}})$ and the torsion part of $\mathrm{coker}(Q^{\mathrm{gp}} \rightarrow P^{\mathrm{gp}})$ finite groups of orders invertible on X . Here P^{gp} denotes the Grothendieck group of P , see §3.1.1.

In this book, we will be exclusively concerned about schemes over fields of characteristic zero, so the last item is not so significant for us. In addition, most of the time P and Q will be toric monoids, hence P^{gp} and Q^{gp} will be torsion-free.

EXAMPLES 3.24. (1) Take $Y = \mathrm{Spec} \mathbb{k}$ with the trivial log structure, $Q = 0$, $X = \mathrm{Spec} \mathbb{k}[P]$, P a toric monoid, so that $X \cong Y \times_{\mathrm{Spec} \mathbb{Z}[Q]} \mathrm{Spec} \mathbb{Z}[P]$. Give X the log structure described in Example 3.19, given by the chart $P \rightarrow \mathbb{k}[P]$. Then X^\dagger is log smooth over $\mathrm{Spec} \mathbb{k}$. Of course, the toric variety X need not be smooth over $\mathrm{Spec} \mathbb{k}$ in the usual sense, so this is a first example where the log structure “makes” X smooth.

(2) Let $X = \mathrm{Spec} \mathbb{k}[P]$, $Y = \mathrm{Spec} \mathbb{k}[N]$, and suppose we are given $\rho \in P$, $\rho \neq 0$, defining a map $N \rightarrow P$ via $1 \mapsto \rho$. This induces a morphism $X \rightarrow Y$, which is obtained via base-change from $\mathrm{Spec} \mathbb{Z}[P] \rightarrow \mathrm{Spec} \mathbb{Z}[N]$, hence is log smooth with the log structures on X and Y as in Example 3.19. We can make a further base-change

$$\begin{array}{ccc} X_0 & \longrightarrow & X \\ \downarrow & & \downarrow \\ \mathrm{Spec} \mathbb{k} & \longrightarrow & Y \end{array}$$

where $\mathrm{Spec} \mathbb{k}$ maps to $0 \in Y = \mathbb{A}_{\mathbb{k}}^1$. Pulling back the log structure on X and Y gives us a morphism of log schemes $X_0^\dagger \rightarrow \mathrm{Spec} \mathbb{k}^\dagger$, where $\mathrm{Spec} \mathbb{k}^\dagger$ is the standard log point. Again,

$$X_0 \cong \mathrm{Spec} \mathbb{k} \times_{\mathrm{Spec} \mathbb{Z}[N]} \mathrm{Spec} \mathbb{Z}[P],$$

so this map is log smooth.

This is the crucial example for us: it explains how the log structure on the singular scheme X_0^\dagger actually makes X_0^\dagger smooth over $\mathrm{Spec} \mathbb{k}^\dagger$, even though X_0 may be reducible or even non-reduced!

(3) Globalizing (2), consider the degeneration of toric varieties $\pi : \mathbb{P}_{\Delta} \rightarrow \mathbb{A}_{\mathbb{k}}^1$ of Example 3.6. Then we obtain divisorial log structures induced by $\partial \mathbb{P}_{\Delta} \subseteq \mathbb{P}_{\Delta}$ and $0 \in \mathbb{A}_{\mathbb{k}}^1$ on \mathbb{P}_{Δ} and $\mathbb{A}_{\mathbb{k}}^1$ respectively, giving a log morphism

$$\pi : \mathbb{P}_{\Delta}^\dagger \rightarrow (\mathbb{A}_{\mathbb{k}}^1)^\dagger.$$

The map $\pi^\# : \pi^{-1} \mathcal{M}_{(\mathbb{A}_{\mathbb{k}}^1, 0)} \rightarrow \mathcal{M}_{(\mathbb{P}_{\Delta}, \partial \mathbb{P}_{\Delta})}$ is just pull-back of functions. Then one checks easily that f is log smooth on standard affine subsets of \mathbb{P}_{Δ} , the map π coinciding with (2) above on these affine open subsets. Furthermore, restricting the log structures to $\pi^{-1}(0)$ and 0 gives a log smooth morphism

$$\pi^{-1}(0)^\dagger \rightarrow \mathrm{Spec} \mathbb{k}^\dagger.$$

REMARK 3.25. Unlike in the case of ordinary schemes, log smooth morphisms need not be flat. For example, consider the map

$$X = \operatorname{Spec} \mathbb{Z}[\mathbb{N}^2] \rightarrow Y = \operatorname{Spec} \mathbb{Z}[\mathbb{N}^2]$$

defined by the homomorphism of monoids $\mathbb{N}^2 \rightarrow \mathbb{N}^2$ given by $(a, b) \mapsto (a + b, b)$. This describes X as an open patch of the blow-up of Y at the origin, hence is not flat. In general, a morphism $\operatorname{Spec} \mathbb{Z}[P] \rightarrow \operatorname{Spec} \mathbb{Z}[Q]$ induced by a homomorphism of integral monoids $h : Q \rightarrow P$ will be flat if f is *integral*. A homomorphism of integral monoids is integral if whenever $q_1, q_2 \in Q$, $p_1, p_2 \in P$, and $h(q_1) + p_1 = h(q_2) + p_2$, there exists $q_3, q_4 \in Q$ and $p \in P$ such that

$$p_1 = h(q_3) + p, \quad p_2 = h(q_4) + p, \quad q_1 + q_3 = q_2 + q_4.$$

We say a morphism of fine log schemes $f : X^\dagger \rightarrow Y^\dagger$ is *integral* if the induced morphisms $f^\# : \overline{\mathcal{M}}_{Y, f(\bar{x})} \rightarrow \overline{\mathcal{M}}_{X, \bar{x}}$ are integral for all $\bar{x} \in X$.

EXAMPLE 3.26. In Chapter 4, we will need to make use of log smooth curves. A local description of such curves was given by F. Kato in [64]. The precise description is as follows. Let $f : C^\dagger \rightarrow W^\dagger$ be log smooth and integral of relative dimension one, $W = \operatorname{Spec}(A)$, where (A, \mathfrak{m}) a complete local ring over an algebraically closed field \mathbb{k} . Suppose also C^\dagger and W^\dagger are fine saturated log schemes. Let $0 \in W$ be the closed point, $Q = \overline{\mathcal{M}}_{W, 0}$. There is necessarily a chart $\sigma : Q \rightarrow A$ defining the log structure on W . Let C_0 be the fibre of f over $0 \in W$ and \bar{x} a geometric point of C_0 . Then, étale locally at \bar{x} , C^\dagger is isomorphic to one of the following log schemes:

- (1) $V = \operatorname{Spec} A[u]$, where the log structure is induced by the chart

$$Q \rightarrow \mathcal{O}_V, \quad q \mapsto f^* \sigma(q).$$

- (2) $V = \operatorname{Spec} A[u, v]/(uv - t)$ for some $t \in \mathfrak{m}$, with the log structure induced by the chart

$$\mathbb{N}^2 \oplus_{\mathbb{N}} Q \rightarrow \mathcal{O}_V, \quad ((a, b), q) \mapsto u^a v^b f^* \sigma(q).$$

Here the fibred sum is defined using the diagonal map $\mathbb{N} \rightarrow \mathbb{N}^2$ and $\mathbb{N} \rightarrow Q$ a homomorphism determined by f given by $1 \mapsto \alpha \in Q$, with $\sigma(\alpha) = t$.

- (3) $V = \operatorname{Spec} A[u]$ with the log structure induced by the chart

$$\mathbb{N} \oplus Q \rightarrow \mathcal{O}_V, \quad (a, q) \mapsto u^a f^* \sigma(q).$$

The three cases should be viewed as follows. In a neighbourhood of type (1), the morphism is actually smooth, not just log smooth, and there is no interesting information, locally, given by the log structure. The log structure is just the pull-back of the log structure on the base. For neighbourhoods of type (2), C_0 is nodal. For neighbourhoods of type (3), $u = 0$ is the image of a section $W \rightarrow C$, which we can think of as a marked point. The log structure is the sum of the pull-back log structure on the base and the divisorial log structure associated to the divisor $u = 0$.

We shall call points of types (1), (2) and (3) on C *smooth points*, *double points*, and *log marked points* respectively. \square

EXAMPLE 3.27. Let us consider the double point case in the above example when $W = \operatorname{Spec} \mathbb{k}$, equipped with the standard log structure coming from the chart $\mathbb{N} \rightarrow \mathbb{k}$ given by

$$n \mapsto \begin{cases} 1 & n = 0, \\ 0 & n > 0. \end{cases}$$

According to Example 3.26, (2), the description of the log structure on $V = \text{Spec } \mathbb{k}[u, v]/(uv)$ depends on the choice of the element $\alpha \in Q = \mathbb{N}$. We take this to be an integer $e > 0$ (the case $e = 0$ is ruled out by the requirement that $\sigma(\alpha) = t$, and t is zero in our case).

Let

$$S_e := \mathbb{N}^2 \oplus_{\mathbb{N}} \mathbb{N},$$

where $\mathbb{N} \rightarrow \mathbb{N}^2$ is the diagonal map and $\mathbb{N} \rightarrow \mathbb{N}$ is multiplication by e . This monoid will prove to be especially important in the next chapter. It can be described in terms of generators $\alpha_1 = ((1, 0), 0)$, $\alpha_2 = ((0, 1), 0)$, and $\rho = ((0, 0), 1)$, with $\alpha_1 + \alpha_2 = e\rho$. It is also the monoid given by $\sigma^\vee \cap N$, where $\sigma \subseteq M_{\mathbb{R}}$ is the cone generated by $(1, 0)$ and $(1, e)$. See Example 3.1, (2). Yet another description is as the submonoid of \mathbb{N}^2 generated by $\alpha_1 = (e, 0)$, $\alpha_2 = (0, e)$ and $\rho = (1, 1)$.

Then the chart given by Example 3.26, (2), in this case is

$$S_e \rightarrow \mathbb{k}[u, v]/(uv), \quad ((a, b), c) \mapsto \begin{cases} u^a v^b & c = 0, \\ 0 & c \neq 0. \end{cases}$$

Note that the log structure induced by this chart can be thought of as follows. Let

$$V_e := \text{Spec } \mathbb{k}[S_e].$$

This is a toric surface with an A_{e-1} singularity, with the toric boundary given by the equation $z^\rho = 0$. Then the chart above induces an embedding

$$V \hookrightarrow V_e,$$

identifying V with ∂V_e . The log structure on V is then the restriction of the divisorial log structure given by $V \subseteq V_e$ to V .

EXAMPLE 3.28. Fix positive integers e_1, e_2 and w , and suppose $e_2 = we_1$. Let ζ be a w -th root of unity. Consider the two curves C_1, C_2 given as follows:

$$\begin{aligned} C_1 &= \text{Spec } \mathbb{k}[x, y]/(xy), \\ C_2 &= \text{Spec } \mathbb{k}[u, v]/(uv), \end{aligned}$$

with the log structure on C_1 and C_2 given by charts as follows:

$$\begin{aligned} S_{e_1} \rightarrow \mathbb{k}[x, y]/(xy), \quad ((a, b), c) &\mapsto \begin{cases} (\zeta x)^a y^b & c = 0, \\ 0 & c \neq 0; \end{cases} \\ S_{e_2} \rightarrow \mathbb{k}[u, v]/(uv), \quad ((a, b), c) &\mapsto \begin{cases} u^a v^b & c = 0, \\ 0 & c \neq 0. \end{cases} \end{aligned}$$

Note that \mathcal{M}_{C_i} is a quotient of $S_{e_i} \oplus \mathcal{O}_{C_i}^\times$, so given $\beta \in S_{e_i}$, $(\beta, 1)$ represents an element s_β of \mathcal{M}_{C_i} . Thus we can write any section of \mathcal{M}_{C_i} as $\varphi \cdot s_\beta$ for some invertible φ and $\beta \in S_{e_i}$. Using this notation, let $C_i^\dagger \rightarrow \text{Spec } \mathbb{k}^\dagger$ be given by the section s_{ρ_i} of \mathcal{M}_{C_i} corresponding to $\rho_i = ((0, 0), 1) \in S_{e_i}$. Despite the fact that the chart for C_1 depends on the choice of ζ , in fact the log scheme C_1^\dagger does not depend on this choice, as any two choices of ζ yield equivalent charts. However, the log morphism $C_1^\dagger \rightarrow \text{Spec } \mathbb{k}^\dagger$ *does* depend on this chart.

So we now have specified log curves C_i^\dagger , $i = 1, 2$, and log smooth morphisms $C_i^\dagger \rightarrow \text{Spec } \mathbb{k}^\dagger$. Consider the map $f : C_1 \rightarrow C_2$ of schemes given by $u \mapsto x^w$,

$v \mapsto y^w$. We wish to lift f to a morphism of log schemes $f^\dagger : C_1^\dagger \rightarrow C_2^\dagger$, giving a commutative diagram

$$(3.4) \quad \begin{array}{ccc} C_1^\dagger & \xrightarrow{f^\dagger} & C_2^\dagger \\ \downarrow & & \downarrow \\ \mathrm{Spec} \mathbb{k}^\dagger & \xrightarrow{=} & \mathrm{Spec} \mathbb{k}^\dagger \end{array}$$

To do this, we need to identify maps $f^\#$ making the diagram

$$\begin{array}{ccc} f^{-1}\mathcal{M}_{C_2} & \xrightarrow{f^\#} & \mathcal{M}_{C_1} \\ \alpha_{C_2} \downarrow & & \downarrow \alpha_{C_1} \\ f^{-1}\mathcal{O}_{C_2} & \xrightarrow{f^*} & \mathcal{O}_{C_1} \end{array}$$

commute.

Since we want (3.4) to be commutative, we need $f^\#(s_{\rho_2}) = s_{\rho_1}$. Also, note that

$$\begin{aligned} f^* \alpha_{C_2}(s_{((1,0),0)}) &= f^* u = x^w, \\ f^* \alpha_{C_2}(s_{((0,1),0)}) &= f^* v = y^w. \end{aligned}$$

So we must have

$$\begin{aligned} f^\#(s_{((1,0),0)}) &= \varphi_x \cdot s_{((w,0),0)}, \\ f^\#(s_{((0,1),0)}) &= \varphi_y \cdot s_{((0,w),0)}, \end{aligned}$$

where φ_x, φ_y are invertible functions on C_1 with $\varphi_x = 1$ when $x \neq 0$ and $\varphi_y = 1$ when $y \neq 0$. But since

$$\begin{aligned} s_{\rho_1}^{e_2} &= f^\#(s_{\rho_2}^{e_2}) = f^\#(s_{((1,0),0)} \cdot s_{((0,1),0)}) \\ &= \varphi_x \varphi_y s_{((w,w),0)} \\ &= \varphi_x \varphi_y s_{((0,0),we_1)} \\ &= \varphi_x \varphi_y s_{\rho_1}^{e_2}, \end{aligned}$$

we must have $\varphi_x \cdot \varphi_y = 1$, so $\varphi_x = \varphi_y = 1$.

Thus the log morphism $f^\dagger : C_1^\dagger \rightarrow C_2^\dagger$ is determined by the data $C_1^\dagger \rightarrow \mathrm{Spec} \mathbb{k}^\dagger$ and $C_2^\dagger \rightarrow \mathrm{Spec} \mathbb{k}^\dagger$.

How do we interpret the existence of the w different choices for this data? These choices reflect different ways in which we can deform C_1 covering a standard deformation of C_2 . We can deform C_2 via

$$X_2 = \mathrm{Spec} \mathbb{k}[u, v, t]/(uv - t^{e_2}) \rightarrow \mathbb{A}_{\mathbb{k}}^1 = \mathrm{Spec} \mathbb{k}[t].$$

If we substitute $u \mapsto x^w, v \mapsto y^w$, we can factor

$$x^w y^w - t^{e_2} = \prod_{i=1}^w (xy - \zeta^i t^{e_1})$$

where ζ is a primitive w -th root of unity. Thus we have commutative diagrams

$$\begin{array}{ccc} X_1^i = \text{Spec } \mathbb{k}[x, y, t]/(xy - \zeta^i t^{e_1}) & \longrightarrow & X_2 \\ \downarrow & & \downarrow \\ \mathbb{A}_{\mathbb{k}}^1 & \xrightarrow{=} & \mathbb{A}_{\mathbb{k}}^1 \end{array}$$

for each $1 \leq i \leq w$. Here the vertical maps to $\mathbb{A}_{\mathbb{k}}^1 = \text{Spec } \mathbb{k}[t]$ are given by $t \mapsto t$. This gives different deformations of C_1 over the given deformation of C_2 . We leave it as an exercise to the reader that, restricting this diagram to obtain a diagram as in (3.4), we obtain the w possibilities for (3.4) obtained from the w choices of ζ .

This will prove to be important in Chapter 4, where these different choices of log structure yield genuinely different stable curves.

3.3. Log derivations and differentials

We finally return to the original motivation for introducing log structures. We will define log derivations, leading to the logarithmic tangent bundle in the log smooth case, and to the sheaf of log differentials.

Let $\pi : X^\dagger = (X, \mathcal{M}_X) \rightarrow S^\dagger = (S, \mathcal{M}_S)$ be a morphism of log schemes.

DEFINITION 3.29. A *log derivation* on X^\dagger over S^\dagger with values in a sheaf of \mathcal{O}_X -modules \mathcal{E} is a pair (D, Dlog) , where $D : \mathcal{O}_X \rightarrow \mathcal{E}$ is an ordinary derivation of X over S and $\text{Dlog} : \mathcal{M}_X^{\text{gp}} \rightarrow \mathcal{E}$ is a homomorphism of sheaves of abelian groups with $\text{Dlog} \circ \pi^\# = 0$; these fulfill the compatibility condition

$$(3.5) \quad D(\alpha_X(m)) = \alpha_X(m) \cdot \text{Dlog}(m)$$

for all $m \in \mathcal{M}_X$, where $\alpha_X : \mathcal{M}_X \rightarrow \mathcal{O}_X$ is the log structure.

We denote by $\Theta_{X^\dagger/S^\dagger}$ the sheaf of log derivations of X^\dagger over S^\dagger with values in \mathcal{O}_X . \square

In many cases a log derivation (D, Dlog) is already determined by D .

PROPOSITION 3.30. *Assume that there is an open, dense subset $U \subseteq X$ such that $\pi|_U : U^\dagger \rightarrow S^\dagger$ is strict, and that \mathcal{E} has no sections with support in $X \setminus U$. Then the forgetful map*

$$(D, \text{Dlog}) \longmapsto D$$

from the sheaf of log derivations on X^\dagger/S^\dagger with values in \mathcal{E} to the sheaf of usual derivations on X/S with values in \mathcal{E} is injective.

PROOF. Let $V \subseteq X$ be an open subset. Then each $m \in \mathcal{M}_X(U \cap V)$ may be written as $h \cdot \pi^\#(n)$ for $h \in \mathcal{O}_X^\times$ and $n \in \mathcal{M}_S$. Hence $\text{Dlog}(m)$ is determined by D via (3.5). Thus if $D = 0$ then $\text{Dlog}|_U = 0$, which under the assumption on \mathcal{E} implies $\text{Dlog} = 0$. \square

We may thus often think of log derivations as usual derivations with certain vanishing behaviour determined by the log structure:

EXAMPLE 3.31. Let X be a normal integral scheme over a field \mathbb{k} , and $Y \subseteq X$ a pure codimension one subscheme, yielding the divisorial log structure $\mathcal{M}_{(X,Y)}$ on X . Then $\Theta_{X^\dagger/\mathbb{k}}$ consists of the usual derivations of X which preserve the ideal of Y . To see this, first note that if $U \subseteq X \setminus Y$, $\mathcal{M}_X|_U = \mathcal{O}_X^\times$, so the hypothesis of Proposition 3.30 is satisfied and $\Theta_{X^\dagger/\mathbb{k}}$ is a subsheaf of $\Theta_{X/\mathbb{k}}$. Now if (D, Dlog)

is a log derivation and $f \in \mathcal{I}_{Y/X}$, then at the generic point η of an irreducible component of Y , we can write $f = f' \cdot t^p$ for t a generator of $\mathcal{I}_{Y/X}$ at η , $p > 0$, and f' a regular function. Then t defines an element of $\mathcal{M}_{(X,Y)}$ in a neighbourhood of η , so $Df = t^p Df' + p f' t^{p-1} Dt = t^p (Df' + p f' D \log t)$ is in $\mathcal{I}_{Y/X}$ in a neighbourhood of η . Thus Df vanishes along every component of Y , so is in $\mathcal{I}_{Y/X}$.

Conversely, if D is an ordinary derivation preserving $\mathcal{I}_{Y/X}$, then for $f \in \mathcal{M}_{(X,Y)}$ we define $D \log f$ as $\frac{Df}{f}$; that this is a regular function is immediately checked again as above at the generic points of Y . \square

EXAMPLES 3.32. (1) Let P be a toric monoid, $X = \text{Spec } \mathbb{k}[P]$, with the log structure on X given by the chart $P \rightarrow \mathbb{k}[P]$, or alternatively, the divisorial log structure $\partial X \subseteq X$. We have a map $X^\dagger \rightarrow \text{Spec } \mathbb{k}$, where $\text{Spec } \mathbb{k}$ has the trivial log structure. Let us compute $\Theta_{X^\dagger/\mathbb{k}}$. By Example 3.31, this is the subsheaf of the sheaf of ordinary derivations $\Theta_{X/\mathbb{k}}$ preserving the ideal of ∂X . We will first see that the ideal of ∂X is generated by monomials z^p with $p \in \text{Int}(P)$. Here, by $\text{Int}(P)$ we mean the following. Since P is a toric monoid, it is of the form $\sigma \cap P^{\text{gp}}$ for some cone $\sigma \subseteq P^{\text{gp}} \otimes_{\mathbb{Z}} \mathbb{R}$. Then $\text{Int}(P) = \text{Int}(\sigma) \cap P^{\text{gp}}$. To see why this is the correct ideal, consider the dual cone σ^\vee . The toric divisors of X are in one-to-one correspondence with the one-dimensional faces of σ^\vee , and for $p \in P$, z^p vanishes on a toric divisor if and only if p is non-zero on the corresponding edge of σ^\vee . Thus z^p vanishes on all toric divisors if and only if p is nowhere zero on $\sigma^\vee \setminus \{0\}$, i.e., $p \in \text{Int}(P)$.

Now consider ordinary derivations on the big torus orbit of X , $X^o = \text{Spec } \mathbb{k}[M]$, where $M = P^{\text{gp}}$, $N = \text{Hom}_{\mathbb{Z}}(M, \mathbb{Z})$. The module of derivations of $\mathbb{k}[M]$ can be identified with

$$\bigoplus_{m \in M} z^m (N \otimes_{\mathbb{Z}} \mathbb{k}),$$

where $z^m (n \otimes 1)$ is the derivation, written as $z^m \partial_n$, which acts by

$$z^m \partial_n (z^{m'}) = \langle n, m' \rangle z^{m+m'}.$$

Now such a derivation extends to a derivation of $\mathbb{k}[P]$ if it preserves $\mathbb{k}[P]$, and it yields a log derivation if it preserves the ideal of ∂X , which is generated by $\{z^p \mid p \in \text{Int}(P)\}$. So let us test a derivation $z^m \partial_n$.

It is clear that if $m \in P$, then $z^m \partial_n$ both preserves $\mathbb{k}[P]$ and preserves the ideal of ∂X , as P and $\text{Int}(P)$ are closed under addition by elements of P .

Next, suppose $m \notin P$. As above, assume P comes from a cone σ . Then there is some edge of σ^\vee with primitive generator n such that $\langle n, m \rangle < 0$. Suppose first that $n' \in N$ is not proportional to n , and consider $z^m \partial_{n'}$. Then we can find $m' \in P \cap n^\perp$ such that $\langle n', m' \rangle \neq 0$, and then $z^m \partial_{n'} z^{m'} = \langle n', m' \rangle z^{m'+m} \neq 0$, and $\langle n, m' + m \rangle < 0$, so $z^m \partial_{n'}$ does not preserve $\mathbb{k}[P]$.

If n' is proportional to n , we can take $n' = n$. There exists an $m' \in \text{Int}(P)$ such that $\langle n, m' \rangle = 1$. But then $z^m \partial_n z^{m'} = z^{m+m'}$, and $\langle n, m + m' \rangle \leq 0$. Thus $z^m \partial_n$ does not preserve the ideal of ∂X .

So we see that the module of log derivations on X is

$$\bigoplus_{p \in P} z^p (N \otimes_{\mathbb{Z}} \mathbb{k}) = \mathbb{k}[P] \otimes_{\mathbb{Z}} N.$$

Thus, as a sheaf,

$$\Theta_{X^\dagger/\mathbb{k}} = \mathcal{O}_X \otimes_{\mathbb{Z}} N.$$

(2) Next suppose we have a monomial $z^\rho \in \mathbb{k}[P]$ inducing a map

$$X \rightarrow \operatorname{Spec} \mathbb{k}[\mathbb{N}] = \mathbb{A}_{\mathbb{k}}^1$$

with the log structure on $\mathbb{A}_{\mathbb{k}}^1$ induced by the chart $\mathbb{N} \rightarrow \mathbb{k}[\mathbb{N}]$ (or equivalently, by the divisor $0 \in \mathbb{A}_{\mathbb{k}}^1$). We obtain a log smooth morphism $X^\dagger \rightarrow (\mathbb{A}_{\mathbb{k}}^1)^\dagger$. Then

$$\Theta_{X^\dagger/(\mathbb{A}_{\mathbb{k}}^1)^\dagger} = \mathcal{O}_X \otimes_{\mathbb{Z}} \rho^\perp.$$

Indeed, $\Theta_{X^\dagger/(\mathbb{A}_{\mathbb{k}}^1)^\dagger}$ is precisely the submodule of derivations annihilating pull-backs of functions on $\mathbb{A}_{\mathbb{k}}^1$, i.e., annihilating z^ρ . \square

There also exists a universal log derivation, which brings in the sheaf of log differentials:

LEMMA 3.33. *Given a morphism $\pi : X^\dagger \rightarrow S^\dagger$ of log schemes, let*

$$\Omega_{X^\dagger/S^\dagger}^1 = (\Omega_{X/S}^1 \oplus (\mathcal{O}_X \otimes_{\mathbb{Z}} \mathcal{M}_X^{\text{gp}})) / \mathcal{K},$$

with \mathcal{K} the \mathcal{O}_X -module generated by

$$(d\alpha_X(m), -\alpha_X(m) \otimes m) \quad \text{and} \quad (0, 1 \otimes \pi^*(n)),$$

for $m \in \mathcal{M}_X$, $n \in \mathcal{M}_S$. Then the pair (d, dlog) of natural maps

$$d : \mathcal{O}_X \xrightarrow{d} \Omega_{X/S}^1 \longrightarrow \Omega_{X^\dagger/S^\dagger}^1, \quad \text{dlog} : \mathcal{M}_X^{\text{gp}} \xrightarrow{1 \otimes} \mathcal{O}_X \otimes \mathcal{M}_X^{\text{gp}} \longrightarrow \Omega_{X^\dagger/S^\dagger}^1,$$

is a universal log derivation. In other words, for any \mathcal{O}_X -module \mathcal{E} and a log derivation (D, Dlog) on X^\dagger over S^\dagger with values in \mathcal{E} , there exists a map $\Phi : \Omega_{X^\dagger/S^\dagger}^1 \rightarrow \mathcal{E}$ of \mathcal{O}_X -modules such that $D = \Phi \circ d$ and $\text{Dlog} = \Phi \circ \text{dlog}$.

PROOF. We first verify that (d, dlog) is a log derivation. Of course d is an ordinary derivation. Also,

$$(d(\alpha_X(m)), -\alpha_X(m) \otimes m), (0, 1 \otimes (\pi^\# n)) \in \mathcal{K},$$

and hence $d(\alpha_X(m)) = \alpha_X(m) \text{dlog}(m)$ and $\text{dlog} \circ \pi^\# = 0$, so (d, dlog) is a log derivation.

We next verify the universal property. Let (D, Dlog) be a log derivation with values in the sheaf of \mathcal{O}_X -modules \mathcal{E} :

$$D : \mathcal{O}_X \longrightarrow \mathcal{E}, \quad \text{Dlog} : \mathcal{M}_X^{\text{gp}} \longrightarrow \mathcal{E}.$$

By the universal property of $\Omega_{X/S}^1$ there is a unique morphism $\varphi : \Omega_{X/S}^1 \rightarrow \mathcal{E}$ fulfilling

$$D = \varphi \circ d.$$

Define

$$\Phi : \Omega_{X/S}^1 \oplus (\mathcal{O}_X \otimes_{\mathbb{Z}} \mathcal{M}_X^{\text{gp}}) \longrightarrow \mathcal{E}, \quad \Phi(\gamma, h \otimes m) = \varphi(\gamma) + h \cdot \text{Dlog}(m).$$

This descends to the quotient by \mathcal{K} because

$$\varphi(d\alpha_X(m)) - \alpha_X(m) \cdot \text{Dlog}(m) = 0, \quad \text{Dlog}(\pi^\#(n)) = 0,$$

for every $m \in \mathcal{M}_X$, $n \in \mathcal{M}_S$. Thus we obtain $\Phi : \Omega_{X^\dagger/S^\dagger}^1 \rightarrow \mathcal{E}$. Clearly $\Phi \circ d = D$ and $\Phi \circ \text{dlog} = \text{Dlog}$. Uniqueness follows since $\Omega_{X/S}^1 \oplus (\mathcal{O}_X \otimes_{\mathbb{Z}} \mathcal{M}_X^{\text{gp}})$ is generated as \mathcal{O}_X -modules by $\Omega_{X/S}^1$ and by $1 \otimes \mathcal{M}_X^{\text{gp}}$. On these subsets Φ is determined by φ and by Dlog respectively. \square

The \mathcal{O}_X -module $\Omega_{X^\dagger/S^\dagger}^1$ is the sheaf of *log differentials*.

REMARK 3.34. First, we remark that by the universal property, we have the relation

$$\Theta_{X^\dagger/S^\dagger} = \mathcal{H}om_{\mathcal{O}_X}(\Omega_{X^\dagger/S^\dagger}^1, \mathcal{O}_X).$$

Second, if $\alpha : \underline{P} \rightarrow \mathcal{O}_U$ is a chart for the log structure on X , then in the formula for $\Omega_{X^\dagger/S^\dagger}^1$ one may replace $\mathcal{M}_X^{\text{gp}}$ by P^{gp} and α_X by α . In fact, any $h \in \mathcal{O}_X^\times$ gives a relation

$$(dh, -h \otimes \alpha_X^{-1}(h)) \in \mathcal{K}.$$

Therefore, for any $m \in P^{\text{gp}}$ the log differential $(0, 1 \otimes (\alpha_X^{-1}(h) \cdot m))$ may be written as $h^{-1}(dh, 0 \otimes 1) + (0, 1 \otimes m)$, which is the sum of an ordinary differential and a log differential involving only m . \square

DEFINITION 3.35. Given a morphism $f : X^\dagger \rightarrow Y^\dagger$ of log schemes over S^\dagger , there is a natural functorial morphism

$$f^* : f^* \Omega_{Y^\dagger/S^\dagger}^1 \rightarrow \Omega_{X^\dagger/S^\dagger}^1$$

induced by $f^* : f^* \Omega_{Y/S}^1 \rightarrow \Omega_{X/S}^1$ and $f^* \otimes f^\# : f^{-1}(\mathcal{O}_Y \otimes \mathcal{M}_Y^{\text{gp}}) \rightarrow \mathcal{O}_X \otimes \mathcal{M}_X^{\text{gp}}$.

Dually, this gives a morphism

$$f_* : \Theta_{X^\dagger/S^\dagger} \rightarrow f^* \Theta_{Y^\dagger/S^\dagger}.$$

EXAMPLES 3.36. (1) As in Example 3.32, take $X = \text{Spec } \mathbb{k}[P]$ with P a toric monoid and with X carrying the usual log structure given by the standard chart $P \rightarrow \mathbb{k}[P]$. Consider

$$\Omega_{X^\dagger/\mathbb{k}}^1 = (\Omega_{X/\mathbb{k}}^1 \oplus (\mathcal{O}_X \otimes_{\mathbb{Z}} P^{\text{gp}}))/\mathcal{K}.$$

Note that $\Omega_{X/\mathbb{k}}^1$ is generated by differentials of the form $d(z^p)$, and

$$(d(z^p), -z^p \otimes p) \in \mathcal{K},$$

so under the equivalence relation induced by \mathcal{K} , every element of $\Omega_{X/\mathbb{k}}^1$ is equivalent to an element of $\mathcal{O}_X \otimes P^{\text{gp}}$. Furthermore, there is no non-zero element of \mathcal{K} in $0 \oplus (\mathcal{O}_X \otimes_{\mathbb{Z}} P^{\text{gp}})$, so $\mathcal{O}_X \otimes_{\mathbb{Z}} P^{\text{gp}}$ injects into $\Omega_{X^\dagger/\mathbb{k}}^1$. Thus

$$\Omega_{X^\dagger/\mathbb{k}}^1 = \mathcal{O}_X \otimes_{\mathbb{Z}} P^{\text{gp}} = \mathcal{O}_X \otimes_{\mathbb{Z}} M,$$

where as before $M = P^{\text{gp}}$, $N = \text{Hom}_{\mathbb{Z}}(M, \mathbb{Z})$. Of course, this gives another, much easier, computation of

$$\Theta_{X^\dagger/\mathbb{k}}^1 = \mathcal{H}om_{\mathcal{O}_X}(\Omega_{X^\dagger/\mathbb{k}}^1, \mathcal{O}_X) = \mathcal{O}_X \otimes_{\mathbb{Z}} N.$$

Note that, under this identification, ∂_n is the element of the dual of $\Omega_{X^\dagger/\mathbb{k}}^1$ which takes the value $\langle m, n \rangle$ on $d \log m$. Since $d(z^m) = z^m d \log m$, we see $\partial_n(d(z^m)) = \langle m, n \rangle z^m$. This fits with the description of ∂_n given in Example 3.32, (1).

(2) Similarly, given $\rho \in P$, we get a regular function $z^\rho : X \rightarrow \mathbb{A}_{\mathbb{k}}^1$, which gives a log morphism $X^\dagger \rightarrow (\mathbb{A}_{\mathbb{k}}^1)^\dagger$. We obtain $\Omega_{X^\dagger/(\mathbb{A}_{\mathbb{k}}^1)^\dagger}^1$ by dividing out by the image of $\mathcal{O}_X \otimes_{\mathbb{Z}} \mathbb{N}^{\text{gp}}$ in $\mathcal{O}_X \otimes_{\mathbb{Z}} M$, i.e.,

$$\Omega_{X^\dagger/(\mathbb{A}_{\mathbb{k}}^1)^\dagger}^1 = \mathcal{O}_X \otimes_{\mathbb{Z}} M/\mathbb{Z}\rho.$$

(3) We can further restrict the morphism of (2) to the fibre over zero, to get a log morphism $X_0^\dagger \rightarrow \text{Spec } \mathbb{k}^\dagger$. We then note that $\Omega_{X_0^\dagger/\mathbb{k}}^1 = \mathcal{O}_{X_0} \otimes_{\mathbb{Z}} M$ via exactly the same argument as in (1). Then

$$\Omega_{X_0^\dagger/\mathbb{k}^\dagger}^1 = \mathcal{O}_{X_0} \otimes_{\mathbb{Z}} M/\mathbb{Z}\rho,$$

exactly as in (2).

(4) If $f : X^\dagger \rightarrow Y^\dagger$ is log smooth, then $\Omega_{X^\dagger/Y^\dagger}^1$ is locally free: we leave the details to the reader, based on the above examples.

(5) If X is a non-singular variety over a field \mathbb{k} , $D \subseteq X$ a normal-crossings divisor, then we obtain the divisorial log structure $\mathcal{M}_{(X,D)}$ on X , and $\Omega_{X^\dagger/\mathbb{k}}^1 \cong \Omega_X^1(\log D)$. This is the subsheaf of $j_*\Omega_{(X \setminus D)/\mathbb{k}}^1$ (with $j : X \setminus D \hookrightarrow X$ the inclusion) locally generated by $\frac{dx_1}{x_1}, \dots, \frac{dx_p}{x_p}, dx_{p+1}, \dots, dx_n$, where x_1, \dots, x_n are local coordinates such that D is given by $x_1 \cdots x_p = 0$.

(6) Let $C^\dagger \rightarrow \text{Spec } \mathbb{k}^\dagger$ be a log smooth curve, as in Example 3.26. Then $\Theta_{C^\dagger/\mathbb{k}^\dagger}$ is a line bundle. Let $C_0 \subseteq C$ be an irreducible component. We will describe $\Theta_{C^\dagger/\mathbb{k}^\dagger}$ by describing $\Theta_{C^\dagger/\mathbb{k}^\dagger}|_{C_0}$. (This does not completely determine $\Theta_{C^\dagger/\mathbb{k}^\dagger}$ if C is not of genus zero, as there is a \mathbb{k}^\times choice for gluing these line bundles at the double points of C .) Let $x_1, \dots, x_n \in C_0$ be the points of C_0 which are either double points in C or log marked points. Then we claim that

$$\Theta_{C^\dagger/\mathbb{k}^\dagger}|_{C_0} \cong \Theta_{C_0/\mathbb{k}}(-\sum_{i=1}^n x_i).$$

Indeed, by Proposition 3.30, $\Theta_{C^\dagger/\mathbb{k}^\dagger}$ is a subsheaf of $\Theta_{C/\mathbb{k}}$. It is clear that one obtains equality at smooth points of $C^\dagger \rightarrow \text{Spec } \mathbb{k}^\dagger$.

On the other hand, let's understand $\Theta_{C^\dagger/\mathbb{k}^\dagger}$ at special points. At a double point $x \in C$, we have $\overline{\mathcal{M}}_{C,x} = S_e$ for some e , in the notation of Example 3.28. Then C and its log structure can be described as follows. Let $N = \mathbb{Z}^2$, and $\sigma \subseteq N_{\mathbb{R}}$ be generated by $(0, 1)$ and $(e, 1)$. Then the corresponding toric variety X_σ has a regular function z^ρ with $\rho = (0, 1) \in M$, giving a morphism $X_\sigma \rightarrow \mathbb{A}_{\mathbb{k}}^1$. Locally near x , C is the fibre over $0 \in \mathbb{A}_{\mathbb{k}}^1$, with the log structure induced by the divisorial log structure on X_σ given by $\partial X_\sigma \subseteq X_\sigma$. So we are precisely in the situation of Example 3.36, (3), above, with, locally,

$$\Theta_{C^\dagger/\mathbb{k}^\dagger} = \text{Hom}(\Omega_{C^\dagger/\mathbb{k}^\dagger}^1, \mathcal{O}_C) = \mathcal{O}_C \otimes_{\mathbb{Z}} \rho^\perp.$$

Of course, $\rho^\perp = \mathbb{Z}(1, 0) \subseteq N$.

If $n \in N$, we write as ∂_n the corresponding log derivation in $\Theta_{C^\dagger/\mathbb{k}^\dagger}$, with $\partial_n(z^m) = \langle m, n \rangle z^m$. Thus in particular, if we write, in our local description, $C = \text{Spec } \mathbb{k}[z, w]/(zw) \subseteq \text{Spec } \mathbb{k}[z, w, t]/(zw - t^e)$, with $z = z^{(1,0)}$, $w = z^{(-1,e)}$, $t = z^\rho$, then $\partial_{(1,0)}z = z$, $\partial_{(1,0)}w = -w$. Thus if we restrict this derivation to an irreducible component C_0 of C , which, say, we can take to be $V(w)$, we get the derivation $z\partial_z$ on $C_0 = \text{Spec } \mathbb{k}[z]$. Thus the restriction of $\Theta_{C^\dagger/\mathbb{k}^\dagger}$ to C_0 is locally given by the sheaf of vector fields with a zero at the origin.

We perform a similar analysis at the log marked points. Here C is locally given by specifying the cone $\sigma \subseteq N_{\mathbb{R}}$ generated by $(1, 0)$ and $(0, 1)$, and taking $\rho = (0, 1) \in M$ so that again $\Theta_{C^\dagger/\mathbb{k}^\dagger}$ is locally generated by $\partial_{(1,0)}$. Taking $z = z^{(1,0)}$, $t = z^\rho$, we see that $\partial_{(1,0)} = z\partial_z$ as ordinary derivations, and $C = C_0 = \text{Spec } \mathbb{k}[z, t]/(t)$, so we see locally $\Theta_{C^\dagger/\mathbb{k}^\dagger}|_{C_0}$ again consists of the sheaf of vector fields with a zero at the log marked point. This proves the claim. \square

3.4. Log deformation theory

In this section we will briefly survey log deformation theory, referring the reader to [65] and [63] for full details and a more general setup. Here, we try to keep the discussion relatively simple, just doing what we need for Chapter 4.

First, let us review the basic idea behind deformation theory for smooth varieties. Suppose we are given a smooth, separated variety X over $\text{Spec } \mathbb{k}$, and we would like to construct a k -th order deformation of X . By this we mean the following. Let $R_k = \mathbb{k}[t]/(t^{k+1})$, and let $O_k = \text{Spec } R_k$, so $O_0 = \text{Spec } \mathbb{k}$. The natural surjections $R_k \rightarrow R_\ell$ for $\ell < k$ give closed embeddings $O_\ell \rightarrow O_k$. Given a morphism $X_k \rightarrow O_k$, the *restriction of X_k to O_ℓ* is the base-change $X_k \times_{O_k} O_\ell \rightarrow O_\ell$. So a k -th order deformation of $X \rightarrow \text{Spec } \mathbb{k}$ is a flat morphism $X_k \rightarrow O_k$ such that the restriction of X_k to O_0 is isomorphic to $X \rightarrow \text{Spec } \mathbb{k}$.

If we wish, then, to construct a deformation of X , we proceed step-by-step. Suppose we have constructed a $(k-1)$ -st order deformation $X_{k-1} \rightarrow O_{k-1}$, and we wish to lift this to a deformation $X_k \rightarrow O_k$ which restricts to X_{k-1} over O_{k-1} . This can be done as follows.

Begin by choosing an affine open cover $\{U_i\}$ of X . Let $U_{ij} = U_i \cap U_j$. Now the first fact we need is that the $(k-1)$ -st order deformation of any non-singular affine variety U is trivial, i.e., isomorphic to $U \times_{\mathbb{k}} O_{k-1}$. But given a $(k-1)$ -st order deformation X_{k-1} of X , the restriction of $\mathcal{O}_{X_{k-1}}$ to U_i induces a $(k-1)$ -st order deformation U_i^{k-1} of U_i . This must be trivial, i.e., there is an isomorphism $\theta_i^{k-1} : U_i^{k-1} \rightarrow U_i \times_{\mathbb{k}} O_{k-1}$. This gives rise to gluing maps

$$\theta_{ij}^{k-1} : U_{ij} \times_{\mathbb{k}} O_{k-1} \rightarrow U_{ij} \times_{\mathbb{k}} O_{k-1}$$

which are defined by the composition

$$U_{ij} \times_{\mathbb{k}} O_{k-1} \xrightarrow{(\theta_i^{k-1})^{-1}} U_{ij}^{k-1} \xrightarrow{\theta_j^{k-1}} U_{ij} \times_{\mathbb{k}} O_{k-1}.$$

Note that

$$\theta_{j\ell}^{k-1} \circ \theta_{ij}^{k-1} = \theta_{i\ell}^{k-1}$$

on $U_{ij\ell} = U_i \cap U_j \cap U_\ell$.

To construct X_k , we just try to lift these identifications, by choosing identifications

$$\theta_{ij}^k : U_{ij} \times_{\mathbb{k}} O_k \rightarrow U_{ij} \times O_k$$

which restrict to the given θ_{ij}^{k-1} . There are many such liftings, and the choice of liftings form a torsor over $\Gamma(U_{ij}, \Theta_{X/\mathbb{k}})$, where $\Theta_{X/\mathbb{k}}$ is the tangent bundle of X , or module of derivations with values in \mathcal{O}_X . To see this explicitly, consider a local situation where $X = \text{Spec } A$, and we are given a homomorphism $f^{k-1} : A[t]/(t^k) \rightarrow A[t]/(t^k)$ of R_{k-1} -algebras which is the identity modulo t . Let f^k, \bar{f}^k be two lifts of f^{k-1} to homomorphisms of R_k -algebras $A[t]/(t^{k+1}) \rightarrow A[t]/(t^{k+1})$. Then for $a \in A$, we have

$$\bar{f}^k(a) - f^k(a) = t^k \psi(a),$$

for some $\psi(a) \in A$. It is not difficult to see that ψ must be a derivation of A (see in fact the proof of Proposition 3.39). Furthermore, given f^k and a derivation ψ , the above equation defines \bar{f}^k , and the fact that this is a homomorphism of R_k -algebras follows from the fact that ψ is a derivation for A over \mathbb{k} . Note also that if f^{k-1} is the identity, then there is a canonical lifting f^k of f^{k-1} , namely the identity, so

the set of all possible liftings of the identity homomorphism $A[t]/(t^k) \rightarrow A[t]/(t^k)$ to $A[t]/(t^{k+1}) \rightarrow A[t]/(t^{k+1})$ is identified canonically with $\Gamma(X, \Theta_{X/\mathbb{k}})$.

Having chosen one set of lifts θ_{ij}^k , we now want to impose the condition

$$\theta_{j\ell}^k \circ \theta_{ij}^k = \theta_{i\ell}^k$$

on $U_{ij\ell}$, in order for these morphisms to induce compatible gluings. Of course,

$$\theta_{j\ell}^k \circ \theta_{ij}^k \circ (\theta_{i\ell}^k)^{-1} \equiv \text{id} \pmod{t^k}.$$

Thus the automorphism $\theta_{j\ell}^k \circ \theta_{ij}^k \circ (\theta_{i\ell}^k)^{-1}$ of $U_{ij\ell} \times O_k$ is a lifting of the identity, and hence is determined by an element $\psi_{ij\ell} \in \Gamma(U_{ij\ell}, \Theta_{X/\mathbb{k}})$. Of course $(U_{ij\ell}, \psi_{ij\ell})$ is a Čech 2-cocycle defining an element $\text{ob}(X_{k-1}/O_{k-1}) \in H^2(X, \Theta_{X/\mathbb{k}})$.

This element is precisely the obstruction to lifting X_{k-1} to X_k . Indeed, this cohomology class is zero if and only if there exists a Čech cochain (U_{ij}, ψ_{ij}) for $\Theta_{X/\mathbb{k}}$ such that $\psi_{ij} + \psi_{j\ell} - \psi_{i\ell} = \psi_{ij\ell}$. Then if we modify our initial choice of lifting θ_{ij}^k by subtracting $t^k \psi_{ij}$ from $(\theta_{ij}^k)^*$, we can replace $\psi_{ij\ell}$ with $\psi_{ij\ell} - (\psi_{ij} + \psi_{j\ell} - \psi_{i\ell}) = 0$. Thus with this new choice, $\theta_{j\ell}^k \circ \theta_{ij}^k = \theta_{i\ell}^k$ and the open sets $U_i \times O_k$ glue. Conversely, if we can find liftings θ_{ij}^k which glue, then the 2-cocycle $\text{ob}(X_{k-1}/O_{k-1})$ clearly vanishes in $H^2(X, \Theta_{X/\mathbb{k}})$.

Furthermore, we have a choice of liftings. For any Čech one-cocycle (U_{ij}, ψ_{ij}) for $\Theta_{X/\mathbb{k}}$, we can replace a given choice of lifting θ_{ij}^k by adding $t^k \psi_{ij}$ to $(\theta_{ij}^k)^*$ without affecting the compatibility of gluings.

Any such choice of lifting can be modified by the Čech differential of a one-cocycle (U_i, ψ_i) and this will give an isomorphic lifting X_k . This isomorphism is induced via the identifications $U_i^k \rightarrow U_i^k$ which are the identity modulo t^k and are hence determined by the derivation $t^k \psi_i$.

We conclude that the set of liftings of $X_{k-1} \rightarrow O_{k-1}$ to $X_k \rightarrow O_k$, up to isomorphism, forms a torsor over $H^1(X, \Theta_{X/\mathbb{k}})$.

In conclusion, this argument sketches a proof of

PROPOSITION 3.37. *Given a non-singular separated variety X/\mathbb{k} and a lifting of X/\mathbb{k} to a flat deformation X_{k-1}/O_{k-1} , there exists an element*

$$\text{ob}(X_{k-1}/O_{k-1}) \in H^2(X, \Theta_{X/\mathbb{k}})$$

such that X_{k-1}/O_{k-1} lifts to X_k/O_k if and only if $\text{ob}(X_{k-1}/O_{k-1}) = 0$. Furthermore, if there exists a lifting X_k/O_k of X_{k-1}/O_{k-1} , the set of all such liftings is a torsor over $H^1(X, \Theta_{X/\mathbb{k}})$.

This is really just a special case of deformation theory, and a more in-depth study would explain a lot more. However, this example forms a good model for what we will need.

Keeping the above framework in mind, let's consider the log version. Suppose we are given a log smooth morphism $X^\dagger \rightarrow \text{Spec } \mathbb{k}^\dagger$. We would like to understand log smooth deformations of the form $X_k^\dagger \rightarrow O_k^\dagger$.

Now the very first issue that distinguishes this situation from the ordinary situation described above is that there is not a unique choice of log structure on O_k : we will restrict to fine log structures defined by charts $\mathbb{N} \rightarrow R_k$ whose composition $\mathbb{N} \rightarrow R_k \rightarrow \mathbb{k}$ sends 0 to 1 and all other elements of \mathbb{N} to 0. Thus the restriction of the induced log structure on O_k to $\text{Spec } \mathbb{k}$ is the standard log structure on $\text{Spec } \mathbb{k}$.

The second issue is that in log smooth deformation theory, deformations are usually not locally trivial. For example, the map $\mathrm{Spec} \mathbb{k}[x, y] \rightarrow \mathrm{Spec} \mathbb{k}[t]$ given by $t \mapsto xy$ is a log smooth deformation of the singular (but log smooth) variety $xy = 0$; of course this is not locally trivial.

The main point, however, is that if we want a log smooth deformation $X_k^\dagger \rightarrow O_k^\dagger$ of $X^\dagger \rightarrow \mathrm{Spec} \mathbb{k}^\dagger$ over O_k^\dagger for some choice of log structure O_k^\dagger on O_k , then the *local* structure of the deformation $X_k^\dagger \rightarrow O_k^\dagger$ is determined completely by the log structure on O_k . This follows from the following proposition, which we state without proof. See [65], Proposition 3.14 or [63], Proposition 8.3.

PROPOSITION 3.38. *Let $X^\dagger \rightarrow \mathrm{Spec} \mathbb{k}^\dagger$ be log smooth, with X affine. Given a fine log structure O_k^\dagger on O_k whose restriction to $O_0 = \mathrm{Spec} \mathbb{k}$ is the standard log point, there is, up to isomorphism, a unique log smooth lifting $X_k^\dagger \rightarrow O_k^\dagger$ of $X^\dagger \rightarrow \mathrm{Spec} \mathbb{k}^\dagger$.*

For example, if $X = \mathrm{Spec} \mathbb{k} \times_{\mathrm{Spec} \mathbb{Z}[\mathbb{N}]} \mathbb{Z}[P]$ for some monoid homomorphism $\mathbb{N} \rightarrow P$, we take $X_k = O_k \times_{\mathrm{Spec} \mathbb{Z}[\mathbb{N}]} \mathrm{Spec} \mathbb{Z}[P]$, where the map $O_k \rightarrow \mathrm{Spec} \mathbb{Z}[\mathbb{N}]$ is determined by the chart $\mathbb{N} \rightarrow R_k$ defining the log structure on O_k . If, say, this chart is trivial, sending 1 to 0, then X_k is a trivial deformation of X , but if the chart is given, say, by $1 \mapsto t$, then X_k is a k -th order smoothing of X .

We can now proceed as before. Suppose we are given $X^\dagger \rightarrow \mathrm{Spec} \mathbb{k}^\dagger$ log smooth, and log structures $O_{k-1}^\dagger, O_k^\dagger$ on O_{k-1}, O_k respectively, with the log structure O_k^\dagger restricting to O_{k-1}^\dagger and the standard log structure on $\mathrm{Spec} \mathbb{k}$. Suppose also we have a log smooth lifting $X_{k-1}^\dagger \rightarrow O_{k-1}^\dagger$ of $X^\dagger \rightarrow \mathrm{Spec} \mathbb{k}^\dagger$, and we want to lift this to $X_k^\dagger \rightarrow O_k^\dagger$.

As before, we choose an open affine covering $\{U_i^\dagger\}$ of X^\dagger . By the above proposition, each $U_i^\dagger \rightarrow \mathrm{Spec} \mathbb{k}^\dagger$ lifts uniquely to $(U_i^{k-1})^\dagger \rightarrow O_{k-1}^\dagger$, and hence X_{k-1}^\dagger is obtained by gluing the log schemes $(U_i^{k-1})^\dagger$ via log automorphisms

$$\theta_{ij}^{k-1} : (U_{ij}^{k-1})^\dagger \rightarrow (U_{ij}^{k-1})^\dagger$$

over O_{k-1}^\dagger . We then choose lifts

$$\theta_{ij}^k : (U_{ij}^k)^\dagger \rightarrow (U_{ij}^k)^\dagger.$$

Here, $(U_i^k)^\dagger$ is again the unique lift of $(U_i^{k-1})^\dagger$ to a log smooth scheme over O_k^\dagger . We then apply the following proposition:

PROPOSITION 3.39. *Let $X^\dagger \rightarrow \mathrm{Spec} \mathbb{k}^\dagger$ be log smooth with X affine, and suppose we are given log smooth $X_{k-1}^\dagger \rightarrow O_{k-1}^\dagger$ which restricts to $X^\dagger \rightarrow \mathrm{Spec} \mathbb{k}^\dagger$ and with a given log smooth lifting $X_k^\dagger \rightarrow O_k^\dagger$. Then:*

- (1) *Given a log automorphism $\theta_{k-1} : X_{k-1}^\dagger \rightarrow X_{k-1}^\dagger$ over O_{k-1}^\dagger which restricts to the identity on X^\dagger , the set of lifts of θ_{k-1} to a log automorphism $\theta_k : X_k^\dagger \rightarrow X_k^\dagger$ over O_k^\dagger is a torsor over $\Gamma(X, \Theta_{X^\dagger/\mathbb{k}^\dagger})$.*
- (2) *The set of log automorphisms $\theta_k : X_k^\dagger \rightarrow X_k^\dagger$ over O_k^\dagger which restrict to the identity on X_{k-1}^\dagger is canonically isomorphic to $\Gamma(X, \Theta_{X^\dagger/\mathbb{k}^\dagger})$.*

PROOF. We will first observe without proof that θ_{k-1} always lifts to an automorphism θ_k . This follows from the infinitesimal lifting criterion for log smoothness: see [65], Proposition 3.4 and Corollary 3.11, or [63], Theorem 4.1.

A log automorphism of X_k^\dagger over O_k^\dagger is induced by an automorphism of sheaves of R_k -algebras,

$$\theta_k^* : \mathcal{O}_{X_k} \rightarrow \mathcal{O}_{X_k}$$

and an automorphism

$$\theta_k^\# : \mathcal{M}_{X_k} \rightarrow \mathcal{M}_{X_k}$$

compatible with θ_k^* and the log structure on O_k^\dagger . Suppose we are given two such automorphisms, θ_k and $\bar{\theta}_k$, which restrict to the same automorphism θ_{k-1} of X_{k-1}^\dagger . Then define a log derivation (D, Dlog) on X_k with values in $t^k \mathcal{O}_{X_k} = t^k \mathcal{O}_X$ as follows. For f a section of \mathcal{O}_{X_k} , define

$$Df = \bar{\theta}_k^*(f) - \theta_k^*(f) \in t^k \mathcal{O}_{X_k} = t^k \mathcal{O}_X.$$

Then D is an ordinary derivation:

$$\begin{aligned} D(f \cdot g) &= \bar{\theta}_k^*(f) \cdot \bar{\theta}_k^*(g) - \theta_k^*(f) \cdot \theta_k^*(g) \\ &= \bar{\theta}_k^*(f) \cdot (\bar{\theta}_k^*(g) - \theta_k^*(g)) \\ &\quad + (\bar{\theta}_k^*(f) - \theta_k^*(f)) \cdot \theta_k^*(g) \\ &= fDg + (Df)g. \end{aligned}$$

Here the third equality follows since Df and Dg lie in $t^k \mathcal{O}_{X_k}$ and θ_k^* and $\bar{\theta}_k^*$ are the identity modulo t .

Next define Dlog as follows. For $m \in \mathcal{M}_{X_k}$, we have $\bar{\theta}_k^\#(m) \equiv \theta_k^\#(m) \pmod{t^k}$, and since $\theta_k^\#$ and $\bar{\theta}_k^\#$ induce the identity on $\overline{\mathcal{M}}_{X_k} = \overline{\mathcal{M}}_X$, we have

$$\bar{\theta}_k^\#(m) = h_m \cdot \theta_k^\#(m),$$

where $h_m \in 1 + t^k \mathcal{O}_{X_k}$. Define $\text{Dlog}(m) = h_m - 1$. Then (D, Dlog) forms a log derivation over O_k^\dagger : we have

$$\begin{aligned} D(\alpha_{X_k}(m)) &= \bar{\theta}_k^*(\alpha_{X_k}(m)) - \theta_k^*(\alpha_{X_k}(m)) \\ &= \alpha_{X_k}(\bar{\theta}_k^\#(m)) - \alpha_{X_k}(\theta_k^\#(m)) \\ &= \alpha_{X_k}((1 + \text{Dlog}(m))\theta_k^\#(m)) - \alpha_{X_k}(\theta_k^\#(m)) \\ &= \text{Dlog}(m) \cdot \alpha_{X_k}(\theta_k^\#(m)) \\ &= \text{Dlog}(m) \cdot \alpha_{X_k}(m). \end{aligned}$$

The last line again follows since $\text{Dlog}(m) \in t^k \mathcal{O}_{X_k}$. Furthermore, if $\pi : X_k^\dagger \rightarrow O_k^\dagger$ is the given morphism, then $\text{Dlog}(\pi^\#(m)) = 0$ since $\bar{\theta}_k^\#(\pi^\#(m)) = \pi^\#(m) = \theta_k^\#(\pi^\#(m))$, as θ_k and $\bar{\theta}_k$ are automorphisms over O_k^\dagger .

Note that $D(f)$ and $\text{Dlog}(m)$ only depend on f and m modulo t . Thus (D, Dlog) can be viewed as a log derivation in $\Gamma(X, \Theta_{X^\dagger/\mathbb{k}^\dagger})$ (or more precisely, in $t^k \Gamma(X, \Theta_{X^\dagger/\mathbb{k}^\dagger})$). Furthermore, going backwards, giving θ_k and (D, Dlog) determines $\bar{\theta}_k$.

This gives (1). For (2), we take θ_k to be the identity, so any $\bar{\theta}_k$ is canonically determined by a derivation in $\Gamma(X, \Theta_{X^\dagger/\mathbb{k}^\dagger})$. \square

We can then follow the same argument we gave in the usual deformation theory setting to obtain:

PROPOSITION 3.40. *Suppose we are given a log smooth morphism $X^\dagger \rightarrow \mathrm{Spec} \mathbb{k}^\dagger$ with X separated over \mathbb{k} and a lifting of this morphism to a log smooth morphism $X_{k-1}^\dagger \rightarrow O_{k-1}^\dagger$, for some choice of log structure on O_{k-1} . Given a lift of this log structure to O_k , there exists an element $\mathrm{ob}(X_{k-1}^\dagger/O_{k-1}^\dagger) \in H^2(X, \Theta_{X^\dagger/\mathbb{k}^\dagger})$ such that $X_{k-1}^\dagger \rightarrow O_{k-1}^\dagger$ lifts to a log smooth morphism $X_k^\dagger \rightarrow O_k^\dagger$ if and only if $\mathrm{ob}(X_{k-1}^\dagger/O_{k-1}^\dagger) = 0$. Furthermore, the set of such liftings, up to isomorphism, is a torsor over $H^1(X, \Theta_{X^\dagger/\mathbb{k}^\dagger})$.*

The nicest situation in which to apply this proposition is when $H^2(X, \Theta_{X^\dagger/\mathbb{k}^\dagger}) = 0$, in which case there are no obstructions. This is the case, for example, if X is a curve.

For the purposes of Chapter 4, we will need to consider a more complicated deformation theory problem. Suppose we have a toric variety X with the standard log structure, induced by $\partial X \subseteq X$, along with a morphism $\pi : X^\dagger \rightarrow (\mathbb{A}_{\mathbb{k}}^1)^\dagger$ induced by a monomial function on X . We will keep this fixed in the following discussion. Consider also a log smooth curve $C_0^\dagger \rightarrow \mathrm{Spec} \mathbb{k}^\dagger = O_0^\dagger$, and suppose we have a commutative diagram

$$(3.6) \quad \begin{array}{ccc} C_0^\dagger & \xrightarrow{f_0} & X^\dagger \\ \downarrow & & \downarrow \pi \\ O_0^\dagger & \xrightarrow{\alpha_0} & (\mathbb{A}_{\mathbb{k}}^1)^\dagger \end{array}$$

Here $\alpha_0 : O_0 \rightarrow \mathbb{A}_{\mathbb{k}}^1$ maps to the origin. We would like to understand liftings of the diagram (3.6) to diagrams

$$(3.7) \quad \begin{array}{ccc} C_k^\dagger & \xrightarrow{f_k} & X^\dagger \\ \downarrow & & \downarrow \pi \\ O_k^\dagger & \xrightarrow{\alpha_k} & (\mathbb{A}_{\mathbb{k}}^1)^\dagger \end{array}$$

where the bottom horizontal arrow determines the log structure on O_k : indeed, giving a morphism $\alpha_k : O_k \rightarrow \mathbb{A}_{\mathbb{k}}^1$ is the same thing as giving a map $\mathbb{N} \rightarrow R_k$, which yields the chart for the log structure on O_k .

We shall use the notation $[f_k : C_k/O_k \rightarrow X]$ for the data of a diagram (3.7).

THEOREM 3.41. *Suppose C_0 is rational. Consider the map*

$$f_{0*} : \Theta_{C_0^\dagger/\mathbb{k}^\dagger} \rightarrow f_0^* \Theta_{X^\dagger/(\mathbb{A}^1)^\dagger}$$

dual to the functorial map of sheaves of differentials

$$f_0^* : f_0^* \Omega_{X^\dagger/(\mathbb{A}^1)^\dagger}^1 \rightarrow \Omega_{C_0^\dagger/\mathbb{k}^\dagger}^1.$$

Suppose that f_{0} is injective, with cokernel \mathcal{N}_{f_0} , the logarithmic normal sheaf of f_0 . Let*

$$[f_{k-1} : C_{k-1}/O_{k-1} \rightarrow X]$$

be a lift of

$$[f_0 : C_0/O_0 \rightarrow X].$$

Then the set of isomorphism classes of lifts

$$[f_k : C_k/O_k \rightarrow X]$$

restricting to

$$[f_{k-1} : C_{k-1}/O_{k-1} \rightarrow X]$$

is a torsor over $H^0(C_0, \mathcal{N}_{f_0})$.

PROOF. First note that since C_0 is a curve, $H^2(C_0, \Theta_{C_0^\dagger/\mathbb{k}^\dagger}) = 0$, so there are no obstructions to lifting $C_{k-1}^\dagger \rightarrow O_{k-1}^\dagger$ to $C_k^\dagger \rightarrow O_k^\dagger$. The set of such liftings forms a torsor over $H^1(C_0, \Theta_{C_0^\dagger/\mathbb{k}^\dagger})$. Choose one such lifting. As above, cover C_0 with affine open sets U_i , with unique thickenings $(U_i^{k-1})^\dagger$, $(U_i^k)^\dagger$ determined by O_{k-1}^\dagger and O_k^\dagger respectively. We have gluing maps

$$(3.8) \quad \begin{aligned} \theta_{ij}^{k-1} &: (U_{ij}^{k-1})^\dagger \rightarrow (U_{ij}^{k-1})^\dagger \\ \theta_{ij}^k &: (U_{ij}^k)^\dagger \rightarrow (U_{ij}^k)^\dagger \end{aligned}$$

with $\theta_{ij}^{k-1} = \theta_{ij}^k \pmod{t^k}$. We also have maps $f_i^{k-1} : (U_i^{k-1})^\dagger \rightarrow X^\dagger$ determined by f_{k-1} ; these must satisfy compatibility on U_{ij}^{k-1} :

$$f_i^{k-1} = f_j^{k-1} \circ \theta_{ij}^{k-1}.$$

As usual in deformation theory, we first choose lifts $f_i^k : (U_i^k)^\dagger \rightarrow X^\dagger$ of f_i^{k-1} . The fact that such a lift exists is actually the infinitesimal lifting criterion for log smoothness: see [65], Proposition 3.4, or [63], Theorem 4.1. Again, we state this without proof.

The set of choices of lifts f_i^k is in fact a torsor over $H^0(U_i, f_0^* \Theta_{X^\dagger/(\mathbb{A}^1)^\dagger})$ (see [65], Proposition 3.9). Indeed, giving a section of $f_0^* \Theta_{X^\dagger/(\mathbb{A}^1)^\dagger}$ is the same as giving a log derivation on X with values in $f_0^* \mathcal{O}_{C_0}$. Now given two lifts f_i^k and \bar{f}_i^k of f_i^{k-1} , we obtain such a derivation (D, Dlog) by

$$Dg = (\bar{f}_i^k)^*(g) - (f_i^k)^*(g)$$

for a function g on X , and $\text{Dlog}(m) = h_m - 1$, for m a section of \mathcal{M}_X , where $h_m \in \mathcal{O}_{U_i}^\times$ is defined by

$$(\bar{f}_i^k)^\#(m) = h_m (f_i^k)^\#(m).$$

This is a log derivation just as in the proof of Proposition 3.39, and vanishes on the pull-backs of functions on \mathbb{A}^1 , hence gives a log derivation of $X^\dagger/(\mathbb{A}^1)^\dagger$ with values in $t^k \mathcal{O}_{U_i} \cong \mathcal{O}_{U_i}$. This is then an element of $\Gamma(U_i, f_0^* \Theta_{X^\dagger/(\mathbb{A}^1)^\dagger})$. Conversely, giving f_i^k and such a log derivation (D, Dlog) clearly determines \bar{f}_i^k . If $(D, \text{Dlog}) =: \psi_i$, we write

$$\bar{f}_i^k = f_i^k + t^k \psi_i.$$

Once we choose liftings f_i^k of f_i^{k-1} , we can then compare, for each i and j , f_i^k with $f_j^k \circ \theta_{ij}^k$ on U_{ij}^k . These are two different liftings of the same map $f_i^{k-1} = f_j^{k-1} \circ \theta_{ij}^{k-1}$, and hence differ by an element $\psi_{ij} \in \Gamma(U_{ij}, f_0^* \Theta_{X^\dagger/(\mathbb{A}^1)^\dagger})$. These yield a Čech 1-cocycle for $f_0^* \Theta_{X^\dagger/(\mathbb{A}^1)^\dagger}$ on C_0 . However, $\Theta_{X^\dagger/(\mathbb{A}^1)^\dagger}$ is a trivial vector bundle since X is a toric variety, by Example 3.32, (2). Since C_0 is assumed to be rational, $H^1(C_0, f_0^* \Theta_{X^\dagger/(\mathbb{A}^1)^\dagger}) = 0$. Thus there is no obstruction to lifting, and we

obtain a lifting $[f_k : C_k/O_k \rightarrow X]$. It only remains to classify all choices of such liftings, up to isomorphism.

Consider then two sets of liftings of the data $\theta_{ij}^{k-1}, f_i^{k-1}$, given by θ_{ij}^k, f_{ij}^k and $\bar{\theta}_{ij}^k, \bar{f}_{ij}^k$. Then $\bar{\theta}_{ij}^k$ and θ_{ij}^k differ by some $t^k \psi_{ij}$ for $\psi_{ij} \in H^0(U_{ij}, \Theta_{C_0^\dagger/\mathbb{k}^\dagger})$, and \bar{f}_i^k differs from f_i^k by $t^k \eta_i$ for some $\eta_i \in H^0(U_i, f_0^* \Theta_{X^\dagger/(\mathbb{A}^1)^\dagger})$. The compatibility condition

$$\bar{f}_i^k = \bar{f}_j^k \circ \bar{\theta}_{ij}^k$$

becomes

$$f_i^k + t^k \eta_i = f_j^k \circ \theta_{ij}^k + (f_{0*}(\psi_{ij}) + \eta_j) t^k.$$

Hence $\eta_i = f_{0*}(\psi_{ij}) + \eta_j$, so the data $(\eta_i)_i$ determine a section of \mathcal{N}_{f_0} .

Conversely, given a section of \mathcal{N}_{f_0} determined by $(\bar{\eta}_i)_i$, $\bar{\eta}_i \in \Gamma(U_i, \mathcal{N}_{f_0})$, we can lift these to $\eta_i \in \Gamma(U_i, f_0^* \Theta_{X^\dagger/(\mathbb{A}^1)^\dagger})$, with $\eta_i - \eta_j = f_{0*}(\psi_{ij})$ for some ψ_{ij} . Since f_{0*} is assumed to be injective, this determines the ψ_{ij} 's, and hence determines $\bar{\theta}_{ij}^k, \bar{f}_i^k$ from θ_{ij}^k, f_i^k . A different set of lifts (η_i) will change the Čech one-cocycle $\{(U_{ij}, \psi_{ij})\}$ by a coboundary, and hence determine an isomorphic lift. Hence the set of all possible lifts is a torsor over $H^0(C_0, \mathcal{N}_{f_0})$. \square

We have more decorations needed for the next chapter, as follows. We wish to mark C_0 , considering points $x_1^0, \dots, x_s^0 \in C_0$, which we think of as sections $x_i^0 : O_0 \rightarrow C_0$ of $C_0 \rightarrow O_0$. We consider liftings $[f_k : C_k/O_k \rightarrow X]$ along with sections $x_i^k : O_k \rightarrow C_k$ of $C_k \rightarrow O_k$ lifting x_i^0 . We write the data of such a lift as $[f_k : C_k/O_k \rightarrow X, \mathbf{x}^k]$.

THEOREM 3.42. *Given $[f_0 : C_0/O_0 \rightarrow X, \mathbf{x}^0]$ such that the points $x_1^0, \dots, x_s^0 \in C_0^\dagger$ are all log smooth points, with C_0 rational and f_{0*} injective as in Theorem 3.41, let $\mathcal{N}_{f_0, \mathbf{x}^0}$ be defined to be the cokernel of the composition*

$$\Theta_{C_0^\dagger/\mathbb{k}^\dagger} \left(- \sum_{i=1}^s x_i^0 \right) \hookrightarrow \Theta_{C_0^\dagger/\mathbb{k}^\dagger} \xrightarrow{f_{0*}} f_0^* \Theta_{X^\dagger/(\mathbb{A}^1)^\dagger}.$$

Given a lifting

$$[f_{k-1} : C_{k-1}/O_{k-1} \rightarrow X, \mathbf{x}^{k-1}],$$

the set of liftings

$$[f_k : C_k/O_k \rightarrow X, \mathbf{x}^k]$$

of

$$[f_{k-1} : C_{k-1}/O_{k-1} \rightarrow X, \mathbf{x}^{k-1}]$$

is a torsor over $H^0(C_0, \mathcal{N}_{f_0, \mathbf{x}^0})$.

PROOF. We just need to modify the argument of Theorem 3.41, taking into account the additional data. First, we understand liftings $C_k/O_k, \mathbf{x}^k$, without taking into account the map f_0 . Fix some $C_k^\dagger \rightarrow O_k^\dagger$ lifting $C_{k-1}^\dagger \rightarrow O_{k-1}^\dagger$. Now the set of lifts $x_i^k : O_k \rightarrow C_k$ of $x_i^{k-1} : O_{k-1} \rightarrow C_{k-1}$ is in fact a torsor over the Zariski tangent space to C_0 at x_i^0 . Indeed, first of all, a lift exists, as $C_k \rightarrow O_k$ is smooth at x_i^0 . (This is a standard fact about smooth morphisms: it is just the formal lifting criterion for smoothness; see for example [53], §17). Next, given two lifts x_i^k, \bar{x}_i^k of x_i^{k-1} and $g \in \mathfrak{m}_{C_k, x_i^0}$, the maximal ideal of \mathcal{O}_{C_k, x_i^0} , we have

$$(\bar{x}_i^k)^*(g) - (x_i^k)^*(g) \in (t^k) \subseteq R_k.$$

This induces a linear functional $\mathfrak{m}_{C_k, x_i^0} \rightarrow t^k \mathbb{k}$, which clearly descends to give a linear functional $\mathfrak{m}_{C_0, x_i^0} / \mathfrak{m}_{C_0, x_i^0}^2 \rightarrow t^k \mathbb{k}$. This is an element of the Zariski tangent space to C_0 at x_i^0 . Conversely, given such a linear function, we can modify $(x_j^k)^*$ to obtain $(\bar{x}_i^k)^*$.

Note that as x_i^0 is always a log smooth point, in fact the fibres of the vector bundles $\Theta_{C_0^\dagger/\mathbb{k}^\dagger}$ and $\Theta_{C_0/\mathbb{k}}$ agree at x_i^0 , and both fibres coincide with the Zariski tangent space to C_0 at x_i^0 .

Suppose we are given a log automorphism $\psi : C_k^\dagger \rightarrow C_k^\dagger$ over O_k^\dagger which is the identity modulo t^k , hence induced by $\psi \in \Gamma(C_0, \Theta_{C_0^\dagger/\mathbb{k}^\dagger})$. Furthermore, suppose we are given lifts $x_i^k : O_k \rightarrow C_k$ of x_i^{k-1} , and set

$$\bar{x}_i^k = \psi \circ x_i^k.$$

Then the difference between \bar{x}_i^k and x_i^k is the image of ψ in the Zariski tangent space of C_0 at x_i^0 .

This can be summarized by the exact sequence

$$(3.9) \quad 0 \rightarrow \Theta_{C_0^\dagger/\mathbb{k}^\dagger}(-\mathbf{x}^0) \rightarrow \Theta_{C_0^\dagger/\mathbb{k}^\dagger} \rightarrow \bigoplus_{\ell=1}^s \Theta_{C_0^\dagger/\mathbb{k}^\dagger} \otimes k(x_\ell^0) \rightarrow 0.$$

Here $\Theta_{C_0^\dagger/\mathbb{k}^\dagger}(-\mathbf{x}^0)$ denotes the sheaf $\Theta_{C_0^\dagger/\mathbb{k}^\dagger}$ twisted by the line bundle

$$\mathcal{O}_{C_0}(-\sum_{\ell=1}^s x_\ell^0),$$

and $k(x_\ell^0)$ denotes the residue field of C_0 at x_ℓ^0 . Taking sections over an open affine set U of C_0 , with induced thickenings $(U^{k-1})^\dagger, (U^k)^\dagger$, the space of liftings of the maps $\{x_\ell^{k-1} \mid x_\ell^0 \in U\}$ to maps x_ℓ^k for $x_\ell^0 \in U$ is a torsor over

$$\Gamma(U, \bigoplus_{\ell=1}^s \Theta_{C_0^\dagger/\mathbb{k}^\dagger} \otimes k(x_\ell^0)).$$

The set of liftings of an automorphism of $(U^{k-1})^\dagger$ to $(U^k)^\dagger$ is a torsor over

$$\Gamma(U, \Theta_{C_0^\dagger/\mathbb{k}^\dagger}),$$

and the set of such liftings which leave given liftings x_ℓ^k of x_ℓ^{k-1} fixed for all ℓ with $x_\ell^0 \in U$ is a torsor over $\Gamma(U, \Theta_{C_0^\dagger/\mathbb{k}^\dagger}(-\mathbf{x}^0))$.

Now as usual, cover C_{k-1}^\dagger with affine open subsets $(U_i^{k-1})^\dagger$ and fix lifts $(U_i^k)^\dagger$. For each i , let j_1, \dots, j_{s_i} be the indices such that $x_{j_\ell}^0 \in U_i^{k-1}$ for $1 \leq \ell \leq s_i$. Assuming we have lifts $x_{j_\ell}^{k-1} : O_{k-1} \rightarrow U_i^{k-1}$, we can extend these to lifts $x_{j_\ell}^k : O_k \rightarrow U_i^k$. On the other hand, consider the map coming from (3.9)

$$\Psi_i : \Gamma(U_i, \Theta_{C_0^\dagger/\mathbb{k}^\dagger}) \rightarrow \bigoplus_{\ell=1}^{s_i} \Theta_{C_0^\dagger/\mathbb{k}^\dagger} \otimes k(x_{j_\ell}^0).$$

This map is surjective since U_i is affine. Thus, up to automorphisms of $(U_i^k)^\dagger$ which are the identity modulo t^k , there is a unique choice of liftings $x_{j_\ell}^k$.

So we can now assume we have fixed open sets $(U_i^k)^\dagger$ and liftings $x_{j_\ell}^k : O_k \rightarrow U_i^k$. We just need to glue these together. We follow the usual procedure. However, we need to match up those liftings x_ℓ^k which land in $U_{i_j}^k$. One follows the usual

deformation theory argument as given in Proposition 3.40 and sees that the set of all liftings $C_k/O_k, \mathbf{x}^k$ of $C_{k-1}/O_{k-1}, \mathbf{x}^{k-1}$ form a torsor over

$$H^1(C_0, \Theta_{C_0^\dagger/\mathbb{A}^1}(-\mathbf{x}^0)).$$

Repeating the argument of Theorem 3.41 then gives the desired result. \square

The final decoration is as follows. We now assume that we are also given, on top of the above data, sections $\sigma_1, \dots, \sigma_s : \mathbb{A}^1 \rightarrow X$ of π . Assume we start with $[f_0 : C_0/O_0 \rightarrow X, \mathbf{x}^0]$ chosen so that $f_0(x_i^0) = \sigma_i(0)$ for $1 \leq i \leq s$. In other words, we are imposing conditions on where the marked points of C_0 map. Note we can write this condition also as $f_0 \circ x_i^0 = \sigma_i \circ \alpha_0$. We then wish to find liftings $[f_k : C_k/O_k \rightarrow X, \mathbf{x}^k]$ which satisfy

$$f_k \circ x_i^k = \sigma_i \circ \alpha_k.$$

(Recall that $\alpha_k : O_k \rightarrow \mathbb{A}_k^1$ is the given map.)

THEOREM 3.43. *Suppose we are given the same hypotheses as Theorem 3.42, with the additional hypothesis that $f_0 \circ x_i^0 = \sigma_i \circ \alpha_0$ for $1 \leq i \leq s$. Denote by $T_{X/\mathbb{A}^1, \sigma_i}$ the fibre of the vector bundle $\Theta_{X^\dagger/(\mathbb{A}^1)^\dagger}$ at $\sigma_i(0)$, or equivalently, the fibre of the vector bundle $f_0^* \Theta_{X^\dagger/(\mathbb{A}^1)^\dagger}$ at x_i^0 . These can be written either as $\Theta_{X^\dagger/(\mathbb{A}^1)^\dagger} \otimes k(\sigma_i(0))$ or as $f_0^*(\Theta_{X^\dagger/(\mathbb{A}^1)^\dagger}) \otimes k(x_i^0)$. Then there is a natural map*

$$\Xi : H^0(C_0, \mathcal{N}_{f_0, \mathbf{x}^0}) \rightarrow \prod_{i=1}^s T_{X/\mathbb{A}^1, \sigma_i(0)}.$$

The map $H^0(C_0, \mathcal{N}_{f_0, \mathbf{x}^0}) \rightarrow T_{X/\mathbb{A}^1, \sigma_i(0)}$ is given by lifting a section of $\mathcal{N}_{f_0, \mathbf{x}^0}$ locally in a neighbourhood of x_i^0 to a section of $f_0^* \Theta_{X^\dagger/(\mathbb{A}^1)^\dagger}$ and then evaluating the section at x_i^0 to get an element of $T_{X/\mathbb{A}^1, \sigma_i(0)}$. Then given a lift

$$[f_{k-1} : C_{k-1}/O_{k-1} \rightarrow X, \mathbf{x}^{k-1}]$$

of

$$[f_0 : C_0/O_0 \rightarrow X, \mathbf{x}^0]$$

with $f_{k-1} \circ x_i^{k-1} = \sigma_i \circ \alpha_{k-1}$ for all i , there exists a lift

$$[f_k : C_k/O_k \rightarrow X, \mathbf{x}^k]$$

of

$$[f_{k-1} : C_{k-1}/O_{k-1} \rightarrow X, \mathbf{x}^{k-1}]$$

with $f_k \circ x_i^k = \sigma_i \circ \alpha_k$ for all i if Ξ is surjective. Furthermore, if a lift exists, the set of such lifts is a torsor over $\ker \Xi$.

PROOF. We first note that Ξ is well-defined. Indeed, two local lifts of a section s of $\mathcal{N}_{f_0, \mathbf{x}^0}$ to $f_0^* \Theta_{X^\dagger/(\mathbb{A}^1)^\dagger}$ differ by a vector field on C_0 vanishing at x_i^0 . Applying f_{0*} to this vector field gives a section of $f_0^* \Theta_{X^\dagger/(\mathbb{A}^1)^\dagger}$ vanishing at x_i^0 , and hence Ξ is well-defined.

To obtain the result, one now just considers thickenings $(U_i^k)^\dagger$ as usual, along with fixed maps $x_\ell^k : O_k \rightarrow U_i^k$ for various ℓ . All possible lifts f_k, \mathbf{x}^k are then obtained by modifying the gluings of the sets $(U_i^k)^\dagger$, as well as modifying the map $f_k : (U_i^k)^\dagger \rightarrow X$. Note that modifications of the gluings don't affect the maps $f_k \circ x_\ell^k$. So if $f_k : (U_i^k)^\dagger \rightarrow X^\dagger$ is replaced by \bar{f}_k which differs from f_k by a section

ψ of $f_0^* \Theta_{X^\dagger/(\mathbb{A}^1)^\dagger}$, we see that for a function g on X defined in a neighbourhood of $\sigma_\ell(0) = f_0(x_\ell^0)$,

$$\begin{aligned} (x_\ell^k)^* \circ \bar{f}_k^*(g) &= (x_\ell^k)^*(f_k^*(g) + t^k \psi(g)) \\ &= (x_\ell^k)^* \circ f_k^*(g) + t^k \psi(g)(x_\ell^0). \end{aligned}$$

So this changes $(x_\ell^k)^* \circ f_k^*$ by the element of $T_{X/\mathbb{A}^1, x_\ell^0}$ given by evaluating ψ at x_ℓ^0 . Since we wish to ensure that

$$(x_\ell^k)^* \circ (f_k)^* = \alpha_k^* \circ \sigma_\ell^*,$$

and the right-hand side is completely determined, the result follows. \square

3.5. The twisted de Rham complex revisited

This section doesn't properly belong here. It explains a technical point we evaded in §2.2.1. Namely, Theorem 2.31 only applies in the case that $W : X \rightarrow \mathbb{C}$ is projective.

Since this is not the case for the chief example of interest, namely the mirror Landau-Ginzburg model for \mathbb{P}^r , we proceed as follows. If $W : X \rightarrow \mathbb{C}$ is not projective, but only quasi-projective, then we find a non-singular variety \bar{X} containing X such that $D := \bar{X} \setminus X$ is normal crossings and W extends to give $W : \bar{X} \rightarrow \mathbb{C}$ projective. Then we consider the twisted log de Rham complex

$$(\Omega_{\bar{X}}^\bullet(\log D), d + dW \wedge)$$

with $\Omega_{\bar{X}}^p(\log D) = \bigwedge^p \Omega_{\bar{X}}^1(\log D)$, and the differential $d + dW \wedge$, as before, given by $\omega \mapsto d\omega + dW \wedge \omega$. Then we have, in analogy with Theorem 2.31, proved by Sabbah in [102],

THEOREM 3.44. *In the above situation,*

$$\begin{aligned} \mathbb{H}_{\text{Zar}}^i(\bar{X}, (\Omega_{\bar{X}}^\bullet(\log D), d + dW \wedge)) &\cong \mathbb{H}_{\text{Zar}}^i(\bar{X}, (\Omega_{\bar{X}}^\bullet(\log D), dW \wedge)) \\ &\cong \mathbb{H}_{\text{An}}^i(\bar{X}, (\Omega_{\bar{X}}^\bullet(\log D), d + dW \wedge)). \end{aligned}$$

In fact, when $W : X \rightarrow \mathbb{C}$ is not proper, the correct cohomology groups to use are the ones appearing in the above theorem.

We will now carry out this procedure for the mirror of \mathbb{P}^r . In the end, we will see we get the same answer we did in §2.2.1, so this calculation may be safely skipped, but it does give a good example of toric techniques in action.

Let $M = \mathbb{Z}^r$, $N = \text{Hom}_{\mathbb{Z}}(M, \mathbb{Z})$ etc., as usual, and let $\Delta \subseteq N_{\mathbb{R}}$ be the lattice polytope given by

$$\Delta := \text{Conv}(\{e_0, e_1, \dots, e_r\}),$$

where e_1, \dots, e_r is the standard basis of N and

$$e_0 = - \sum_{i=1}^r e_i.$$

Note that Δ is in fact a reflexive polytope (see Example 1.28) with dual

$$\Delta^* := \text{Conv}\{e_0^*, e_0^* + (r+1)e_1^*, \dots, e_0^* + (r+1)e_r^*\},$$

where e_1^*, \dots, e_r^* is the dual basis for M and

$$e_0^* = - \sum_{i=1}^r e_i^*.$$

Now Δ defines a projective toric variety \mathbb{P}_Δ . We define

$$X \cong (\mathbb{C}^\times)^r$$

to be the open torus orbit of \mathbb{P}_Δ . We identify this X with the X used in Example 2.33. In that example, we took $W = z^{e_1} + \dots + z^{e_r} + \kappa z^{e_0}$ in the current notation. Now any $n \in \Delta \cap N$, via Remark 3.5, can be thought of as a section of $\mathcal{O}_{\mathbb{P}_\Delta}(1)$. This line bundle can be trivialized on the open torus orbit of \mathbb{P}_Δ in such a way that a section determined by $n \in \Delta \cap N$ corresponds to the regular function z^n on X . Since $0 \in \Delta$, this determines a section which corresponds to the regular function $z^0 = 1$ on X . Thus the expression for W on X can also be written as

$$W = \frac{z^{e_1} + \dots + z^{e_r} + \kappa z^{e_0}}{z^0},$$

and by writing it as this ratio, it in fact extends to \mathbb{P}_Δ as a ratio of two sections of $\mathcal{O}_{\mathbb{P}_\Delta}(1)$, and hence defines a rational function

$$W : \mathbb{P}_\Delta \dashrightarrow \mathbb{P}_k^1.$$

There are now two problems that have be dealt with. First, \mathbb{P}_Δ is quite singular, and these singularities need to be resolved. Second, W is only rational, so some further blow-up will be necessary to resolve the singularities of W .

To address these problems, we first resolve the singularities of \mathbb{P}_Δ torically. The normal fan of Δ is in fact the set of cones generated by proper faces of Δ^* . To resolve \mathbb{P}_Δ , we can refine the normal fan by choosing a polyhedral decomposition \mathcal{P} of $\partial\Delta^*$ into standard simplices. Then the fan Σ consisting of cones generated by elements of \mathcal{P} is a refinement of the normal fan of Δ , and defines a non-singular toric variety X_Σ along with a proper birational map $\pi : X_\Sigma \rightarrow \mathbb{P}_\Delta$. Then $W \circ \pi$ can be expressed as the ratio of two sections s_1/s_0 of $\pi^*\mathcal{O}_{\mathbb{P}_\Delta}(1)$. We write this function again as

$$W : X_\Sigma \dashrightarrow \mathbb{P}_k^1.$$

We note that W fails to be defined when $s_1 = s_0 = 0$.

To understand how to resolve the singularities of W , we look at an affine open subset of X_Σ , determined by a cone $\sigma \in \Sigma$ which is generated by $\bar{\sigma} = \sigma \cap \partial\Delta^*$. Without loss of generality, we can assume that $\bar{\sigma}$ is a maximal cell in \mathcal{P} . Note that for each maximal proper face $\bar{\omega}$ of Δ^* , there is a unique vertex v of Δ such that $\langle v, m \rangle = -1$ for all $m \in \bar{\omega}$. Take $\bar{\omega}$ to be the unique face of Δ^* containing $\bar{\sigma}$ and let ω be the cone generated by $\bar{\omega}$. Then in particular there is a unique i such that $\langle e_i, m \rangle = -1$ for all $m \in \bar{\sigma}$. Note that X_ω is the open affine subset of \mathbb{P}_Δ corresponding to the vertex e_i of Δ , and π maps X_σ into X_ω . Thus we can study W by using the trivialization of $\mathcal{O}_{\mathbb{P}_\Delta}(1)$ on the open set X_ω given in Remark 3.5, i.e.,

$$W = \frac{z^{e_1 - e_i} + \dots + z^{e_r - e_i} - \kappa z^{e_0 - e_i}}{z^{-e_i}}.$$

Now the cone σ is generated by some $m_1, \dots, m_r \in M$ with $m_1, \dots, m_r \in \bar{\sigma}$ forming a basis for M , and $X_\sigma \cong \mathbb{A}_k^r$. Since $\langle -e_i, m_j \rangle = 1$ for each j , the denominator of W vanishes precisely once on each toric divisor of \mathbb{A}_k^r , i.e., each coordinate hyperplane. On the other hand, we can write the numerator of W as a polynomial

$$c_0 + \sum_{i=1}^r c_i z^{n_i}$$

where n_1, \dots, n_r form a basis for $N_{\mathbb{R}}$ (but not N) and each c_i is either 1 or κ , hence non-zero.

Claim. The zero locus of $c_0 + \sum_{i=1}^r c_i z^{n_i}$ is non-singular and intersects each toric stratum of X_{σ} transversally.

PROOF. It is enough to show that if we restrict $c_0 + \sum_{i=1}^r c_i z^{n_i}$ to a torus orbit of X_{σ} , we get a non-zero equation which defines a non-singular hypersurface in this torus orbit. Note that the restriction is non-zero as the constant term c_0 is non-zero. Suppose, without loss of generality, we restrict to the toric stratum D_{τ} of X_{σ} corresponding to the cone τ generated by m_1, \dots, m_p . Let $N_p := \{m_1, \dots, m_p\}^{\perp} \subseteq N$. Then the big torus orbit in this stratum is isomorphic to $\text{Spec } \mathbb{k}[N_p]$, and a monomial z^{n_i} on X_{σ} restricts to a non-zero monomial on D_{τ} if and only if $n_i \in N_p$. Assume, again without loss of generality, that $n_i \in N_p$ only for $1 \leq i \leq q$. Then the restriction of the numerator of W to this torus orbit is

$$c_0 + \sum_{i=1}^q c_i z^{n_i}.$$

Now let N'_p be the sublattice of N_p generated by n_1, \dots, n_q . Then the canonical inclusion $N'_p \hookrightarrow N_p$ defines a smooth map of algebraic tori $\text{Spec } \mathbb{k}[N_p] \rightarrow \text{Spec } \mathbb{k}[N'_p]$. Writing $\mathbb{k}[N'_p] = \mathbb{k}[x_1^{\pm 1}, \dots, x_q^{\pm 1}]$, with $x_i = z^{n_i}$, we see that the numerator of W is the pull-back under this map of tori of the function $c_0 + \sum_{i=1}^q c_i x_i$. The zero-set of the latter function is a hyperplane, hence is non-singular, and the pull-back then also defines a non-singular hypersurface. \square

This now shows that $s_1 = 0$ defines a non-singular hypersurface in X_{Σ} . On the other hand, s_0 , as mentioned above, vanishes precisely on the toric boundary of X_{Σ} . By the above claim, $s_1 = 0$ intersects the toric boundary transversally, so we now have a local description of how $s_1 = 0$ and $s_0 = 0$ intersect. This allows us to describe the blow-up of $s_1 = s_0 = 0$ inside X_{Σ} . In particular, (étale or analytically) locally, we can describe $s_0 = 0$ as the locus $y_1 \cdots y_p = 0$ and $s_1 = 0$ as the locus $y_{p+1} = 0$, for some local coordinates y_1, \dots, y_r . So we can describe the local structure of the blow-up using these equations, blowing up $y_1 \cdots y_p = y_{p+1} = 0$. In $\mathbb{A}_{\mathbb{k}}^r \times \mathbb{P}_{\mathbb{k}}^1$, with u, v homogeneous coordinates on $\mathbb{P}_{\mathbb{k}}^1$, this blow-up is given by the equation

$$uy_1 \cdots y_p = vy_{p+1}.$$

In the coordinate chart where $v = 1$, we have $y_{p+1} = uy_1 \cdots y_p$, and eliminating the variable y_{p+1} , W becomes the regular function u on a non-singular variety. On the other hand, if $u = 1$, we get $y_1 \cdots y_p = vy_{p+1}$, which defines a singular hypersurface, but $y_1 \cdots y_p = 0$ at any singular point. In particular, $W = 1/v$ has a pole at such points, but W is well-defined everywhere.

Thus, if \tilde{X}_{Σ} is the blow-up of X_{Σ} along the locus $s_1 = s_2 = 0$, $W : \tilde{X}_{\Sigma} \rightarrow \mathbb{P}_{\mathbb{k}}^1$ is now a morphism. Let $\overline{X} = W^{-1}(\mathbb{A}_{\mathbb{k}}^1)$, with $\mathbb{A}_{\mathbb{k}}^1 = \mathbb{P}_{\mathbb{k}}^1 \setminus \{\infty\}$. What we have seen above is that \overline{X} is non-singular and $W : \overline{X} \rightarrow \mathbb{C}$ is proper. Furthermore, we had identified X with the big torus orbit of X_{Σ} , and \tilde{X}_{Σ} was obtained via a blow-up of a locus contained in the toric boundary of X_{Σ} . Hence we still have $X \subseteq \overline{X} \subseteq \tilde{X}_{\Sigma}$ naturally. This is the desired compactification.

To finish this story, we need to compute $D := \overline{X} \setminus X$ and understand the complex $(\Omega_{\overline{X}}^{\bullet}(\log D), dW \wedge)$. First, note that $s_0 = 0$ defines the toric boundary of

X_Σ , so that D is the total transform of $s_0 = 0$ restricted to \overline{X} . So in the two charts given above for the local description of the blow-up, D is given by $y_1 \cdots y_p = 0$, i.e., is normal crossings, and W is given by u or v^{-1} . Now $dW = du$ or $d(v^{-1})$, and since we are away from the pole of W so that $v \neq 0$, we see in fact dW is locally non-vanishing in $\Omega_{\overline{X}}^1(\log D)$. Thus, as in §2.2.1, the complex $(\Omega_{\overline{X}}^\bullet(\log D), dW \wedge)$ is exact in a neighbourhood of D . Since this complex, when restricted to X , gives the usual complex $(\Omega_X^\bullet, dW \wedge)$, we see that the cohomology of the complex satisfies

$$\mathcal{H}^i(\Omega_{\overline{X}}^\bullet(\log D), dW \wedge) \cong \mathcal{H}^i(\Omega_X^\bullet, dW \wedge),$$

so from the hypercohomology spectral sequence,

$$\mathbb{H}^i(\overline{X}, (\Omega_{\overline{X}}^\bullet(\log D), dW \wedge)) \cong \mathbb{H}^i(X, (\Omega_X^\bullet, dW \wedge)).$$

Hence the calculation of §2.2.1 in fact remains valid.

3.6. References and further reading

Fulton's book [27] provides an excellent introduction to toric varieties; Oda's book [87] also covers more advanced topics. The foundational material on log schemes can all be found in the original paper of K. Kato [65]; a much more in depth exposition of log geometry is given in [88]. In addition, material on log smooth deformation theory and log smooth curves can be found in [63] and [64].

Part 2

Example: \mathbb{P}^2 .

Mikhalkin's curve counting formula

4.1. The statement and outline of the proof

We will first give the statement of Mikhalkin's Theorem, whose proof will be given in the remainder of the chapter.

We fix, as usual, $M = \mathbb{Z}^n$, $N = \text{Hom}_{\mathbb{Z}}(M, \mathbb{Z})$, etc., and also fix a complete fan Σ in $M_{\mathbb{R}}$. Using the notation of §1.3, this gives the lattice T_{Σ} and a map $r : T_{\Sigma} \rightarrow M$; pick $\Delta = \sum_{\rho \in \Sigma[1]} d_{\rho} t_{\rho}$ in the kernel of r . If $n = 2$, then in §1.3 we defined the number $N_{\Delta, \Sigma}^{0, \text{trop}}$ counting (with multiplicity) tropical curves in the moduli space $\mathcal{M}_{0, |\Delta| - 1}(\Sigma, \Delta)$ passing through $|\Delta| - 1$ general points.

We first need to define the analogous holomorphic or algebro-geometric count. We will count stable curves as in Chapter 2, but we place certain restrictions on the curves we consider. In addition, we need to specify the homology class of the curve. For the first point, we use

DEFINITION 4.1. A curve $C \subseteq X_{\Sigma}$ is *torically transverse* if it is disjoint from all toric strata of the toric variety X_{Σ} of codimension > 1 . (In particular, C does not have an irreducible component contained in a codimension one stratum, as then this irreducible component would also intersect a higher codimension stratum.)

If $f : C \rightarrow X_{\Sigma}$ is a stable map, we say it is *torically transverse* if $f(C) \subseteq X_{\Sigma}$ is torically transverse and no irreducible component of C maps into ∂X_{Σ} . \square

For the degree, we use the following:

PROPOSITION 4.2. *If X_{Σ} is non-singular, then*

$$H_2(X_{\Sigma}, \mathbb{Z}) \cong \ker r.$$

The isomorphism takes $\beta \in H_2(X_{\Sigma}, \mathbb{Z})$ to

$$\sum_{\rho \in \Sigma[1]} (\beta \cdot D_{\rho}) t_{\rho} \in T_{\Sigma}.$$

PROOF. From (3.2), we see that $\text{Cl}(X_{\Sigma}) = \text{Hom}_{\mathbb{Z}}(\ker r, \mathbb{Z})$, but since X_{Σ} is non-singular (and all toric varieties are rational), $\text{Cl}(X_{\Sigma}) \cong \text{Pic}(X_{\Sigma}) \cong H^2(X_{\Sigma}, \mathbb{Z})$. Thus $\ker r \cong H_2(X_{\Sigma}, \mathbb{Z})$. More explicitly, an element of T_{Σ}^{\vee} of the form $\sum a_{\rho} t_{\rho}^*$ represents the divisor in $\text{Cl}(X_{\Sigma})$ given by $\sum a_{\rho} D_{\rho}$. Furthermore, the pairing $H_2(X_{\Sigma}, \mathbb{Z}) \times H^2(X_{\Sigma}, \mathbb{Z}) \rightarrow \mathbb{Z}$ which makes $H_2(X_{\Sigma}, \mathbb{Z})$ the dual of $H^2(X_{\Sigma}, \mathbb{Z})$ is given by the intersection pairing $(\beta, D_{\rho}) \mapsto \beta \cdot D_{\rho}$. This gives the explicit description of the isomorphism. \square

So $\Delta \in \ker r$ determines $\Delta \in H_2(X_{\Sigma}, \mathbb{Z})$. Thus we can write, as in Chapter 2, $\overline{\mathcal{M}}_{0, |\Delta| - 1}(X_{\Sigma}, \Delta)$ for the moduli space of stable maps of genus zero into X_{Σ} with $|\Delta| - 1$ marked points and representing the homology class Δ .

DEFINITION 4.3. Suppose $Q_1, \dots, Q_{|\Delta|-1} \in X_\Sigma$ are general points. Define

$$N_{\Delta, \Sigma}^{0, \text{hol}} = \# \left\{ f \in \overline{\mathcal{M}}_{0, |\Delta|-1}(X_\Sigma, \Delta) \left| \begin{array}{l} f : (C, x_1, \dots, x_{|\Delta|-1}) \rightarrow X_\Sigma \text{ is torically} \\ \text{transverse and } f(x_i) = Q_i \text{ for} \\ 1 \leq i \leq |\Delta| - 1 \end{array} \right. \right\}.$$

It is not yet clear that this is a finite number, but this shall be shown in the course of proving the main theorem. However, once we know this, then we know it does not depend on the choice of $Q_1, \dots, Q_{|\Delta|-1}$ (provided they are chosen generally) as the set being counted fits into one algebraic family as $Q_1, \dots, Q_{|\Delta|-1}$ vary, so the cardinality is the same for any two general choices of points. We can then state the main theorem:

THEOREM 4.4. *Suppose $n = 2$. Then $N_{\Delta, \Sigma}^{0, \text{hol}}$ is finite, and*

$$N_{\Delta, \Sigma}^{0, \text{trop}} = N_{\Delta, \Sigma}^{0, \text{hol}}.$$

Note that we do not know yet that $N_{\Delta, \Sigma}^{0, \text{trop}}$ is independent of the choice of points in $M_{\mathbb{R}}$; this follows from the above theorem. However, we shall see a different proof of this independence in the next chapter which is purely combinatorial.

The proof we give is not the original proof of Mikhalkin, but rather the proof of Siebert and Nishinou [86]. The latter proof works in all dimensions, whereas Mikhalkin's original proof only works in two dimensions. We will, however, restrict to two dimensions. The definition of $N_{\Delta, \Sigma}^{0, \text{trop}}$ in higher dimensions is somewhat more complicated, as is some of the combinatorics involved in the proof. We don't actually save very much by restricting to the two-dimensional case though, so after reading this chapter, it should be easy to consult [86] for the general case.

The basic idea is to construct a one-parameter degeneration of the toric variety X_Σ very similar in flavour to Example 3.6. This degeneration will be adapted for the particular choice of points $P_1, \dots, P_{|\Delta|-1} \in M_{\mathbb{R}}$ used to define $N_{\Delta, \Sigma}^{0, \text{trop}}$. We will then construct a correspondence between tropical curves in $M_{\mathbb{R}}$ passing through $P_1, \dots, P_{|\Delta|-1}$, "log stable curves" in the central fibre of the degeneration, and ordinary stable curves in the general fibre of this degeneration.

More specifically, choose points $P_1, \dots, P_s \in M_{\mathbb{Q}} = M \otimes_{\mathbb{Z}} \mathbb{Q}$, with $s = |\Delta| - 1$. Note here that we are only considering points with rational, rather than real, coordinates. This turns out not to change the earlier discussion: the general position arguments of Lemma 1.20 work just as well over \mathbb{Q} , so we can assume these points have been chosen so that there are only a finite number of marked genus zero tropical curves $h : (\Gamma, x_1, \dots, x_{|\Delta|-1}) \rightarrow M_{\mathbb{R}}$ in X_Σ with $h(x_i) = P_i$, and these curves are all simple. Denote the set of these curves by

$$\mathcal{M}_{0, s}(\Sigma, \Delta, P_1, \dots, P_s).$$

DEFINITION 4.5. Given the data of Σ and $P_1, \dots, P_{|\Delta|-1} \in M_{\mathbb{Q}}$ general with $\dim M_{\mathbb{Q}} = 2$, a finite polyhedral decomposition \mathcal{P} of $M_{\mathbb{R}}$ is said to be *good* if it satisfies the following properties:

- (1) For $\sigma \in \mathcal{P}$, σ has faces of rational slope and vertices in $M_{\mathbb{Q}}$. Furthermore, each $\sigma \in \mathcal{P}$ has at least one vertex.
- (2) If $\sigma \in \mathcal{P}$, then $\text{Asym}(\sigma)$, the asymptotic cone to σ , is an element of the fan Σ . Furthermore, every cone of Σ appears as the asymptotic cone of some $\sigma \in \mathcal{P}$.




FIGURE 1

- (3) The image of any $h \in \mathcal{M}_{0,s}(\Sigma, \Delta, P_1, \dots, P_s)$ lies in the one-dimensional skeleton of \mathcal{P} .
- (4) Each P_i is a vertex of \mathcal{P} .

PROPOSITION 4.6. *There exists a good polyhedral decomposition.*

PROOF. If \mathcal{P} and \mathcal{P}' are two polyhedral decompositions satisfying conditions (1) and (2), then

$$\mathcal{P}'' = \{\sigma \cap \sigma' \mid \sigma \in \mathcal{P}, \sigma' \in \mathcal{P}'\}$$

also satisfies these two conditions. We apply this as follows. Take for \mathcal{P} the polyhedral decomposition obtained by taking the maximal cells to be the closures of the connected components of $M_{\mathbb{R}} \setminus S$, where S is the union of images of parameterized tropical curves in $\mathcal{M}_{0,s}(\Sigma, \Delta, P_1, \dots, P_s)$ and elements of $\Sigma^{[1]}$. (The latter is not necessary if $d_{\rho} > 0$ for all ρ , with $\Delta = \sum_{\rho} d_{\rho} t_{\rho}$.) This \mathcal{P} satisfies (1), (2), and (3). Take \mathcal{P}' to be the polyhedral decomposition obtained by taking the maximal cells to be the closures of the connected components of $M_{\mathbb{R}} \setminus S'$, where

$$S' = \bigcup_{\rho \in \Sigma^{[1]}} \bigcup_{i=1}^s (P_i + \rho).$$

Then \mathcal{P}' satisfies conditions (1), (2), and (4). The refinement \mathcal{P}'' of \mathcal{P} , \mathcal{P}' given above satisfies conditions (1)-(4). \square

For example, if we take five points in $M_{\mathbb{R}}$ contained in one tropical curve of degree 2 in \mathbb{P}^2 depicted on the left in Figure 1, we can take \mathcal{P} to be as given on the right.

Without loss of generality, we can in fact assume that \mathcal{P} is a lattice polyhedral decomposition by replacing M with the superlattice $\frac{1}{p}M$, where p is the common denominator of coordinates of all vertices of \mathcal{P} . This of course makes the points P_1, \dots, P_s lie in M too, but does not change the number $N_{\Delta, \Sigma}^{0, \text{trop}}$. We will fix this good \mathcal{P} for the remainder of our discussion.

We now introduce the notation

$$\widetilde{M} = M \oplus \mathbb{Z}, \quad \widetilde{N} = \text{Hom}_{\mathbb{Z}}(\widetilde{M}, \mathbb{Z}),$$

and define the fan $\Sigma_{\mathcal{P}}$ in $\widetilde{M}_{\mathbb{R}}$ as follows. For $\sigma \in \mathcal{P}$, let

$$C(\sigma) = \overline{\{(rm, r) \mid r \geq 0, m \in \sigma\}} \subseteq \widetilde{M}_{\mathbb{R}}$$

be the cone over σ , as in §3.1.3. Then we set $\Sigma_{\mathcal{P}}$ to be the fan consisting of all faces of cones $C(\sigma)$ for $\sigma \in \mathcal{P}$. It is easy to see $\Sigma_{\mathcal{P}}$ is a fan. In addition, note that $C(\sigma) \cap (M_{\mathbb{R}} \oplus \{0\})$ is $\text{Asym}(\sigma)$, as in §3.1.3. Also, by property (2) of Definition 4.5, $\text{Asym}(\sigma) \in \Sigma$. Furthermore, every cone of Σ arises in this way. Thus

$$\Sigma = \{\tau \in \Sigma_{\mathcal{P}} \mid \tau \subseteq M_{\mathbb{R}} \oplus \{0\}\}.$$

Putting this together, we now have a toric variety

$$X := X_{\Sigma_{\mathcal{P}}},$$

which comes along with a regular function $\pi : X \rightarrow \mathbb{A}_{\mathbb{k}}^1$ induced by the monomial $z^{(0,1)}$, $(0,1) \in \tilde{N} = N \oplus \mathbb{Z}$. Note that $\pi^{-1}(0)$ is the union of toric divisors on which $z^{(0,1)}$ vanishes. These are precisely the divisors corresponding to cones of $\Sigma_{\mathcal{P}}$ of the form $C(v)$ for $v \in \mathcal{P}$ a vertex. In particular $\pi^{-1}(0)$ is a union of toric varieties. On the other hand, if one removes these divisors from $X_{\Sigma_{\mathcal{P}}}$, one obtains the toric variety corresponding to the fan Σ as a fan, not in $M_{\mathbb{R}}$, but in $\widetilde{M}_{\mathbb{R}}$. As a result, $X \setminus \pi^{-1}(0) \cong X_{\Sigma} \times \mathbb{G}_m$.

The degeneration $\pi : X \rightarrow \mathbb{A}_{\mathbb{k}}^1$ is the degeneration of X_{Σ} we shall work with.

We will use the following notation for the toric strata of $X_0 := \pi^{-1}(0)$. Each cell $\tau \in \mathcal{P}$ yields a cone $C(\tau) \in \Sigma_{\mathcal{P}}$, hence a toric stratum $D_{C(\tau)}$. This stratum is contained in X_0 . To simplify notation, we will write this as D_{τ} . The fan defining D_{τ} is written as Σ_{τ} , and lives in $M_{\mathbb{R}}/T_{\tau}\tau$. This can be defined in terms of the polyhedra in \mathcal{P} containing τ as

$$\Sigma_{\tau} := \{T_{\tau}\sigma/T_{\tau}\tau \mid \sigma \in \mathcal{P} \text{ such that } \tau \subseteq \sigma\}.$$

We will denote by ∂X_0 the union of one-dimensional toric strata of X_0 not contained in $\text{Sing}(X_0)$. This is precisely the union

$$(4.1) \quad \partial X_0 := \bigcup_{\substack{\omega \in \mathcal{P}^{[1]} \\ \omega \text{ non-compact}}} D_{\omega},$$

as it is precisely the D_{ω} 's for ω non-compact which are not contained in two irreducible components of X_0 . Alternatively, one can write

$$\partial X_0 = \overline{\partial X \setminus X_0} \cap X_0.$$

Next, we explain the role of the points P_1, \dots, P_s . Let

$$L_i = \mathbb{Z}(P_i, 1) \subseteq \widetilde{M}$$

be the rank one sublattice generated by $(P_i, 1)$. Then we have an inclusion of tori $\mathbb{G}(L_i) \subseteq \mathbb{G}(\widetilde{M})$, which acts on X . Now choose points Q_1, \dots, Q_s in the open torus orbit of X , and consider the closed subset of X given by $\overline{\mathbb{G}(L_i) \cdot Q_i}$. Here we are using the canonical action of $\mathbb{G}(\widetilde{M})$, hence $\mathbb{G}(L_i)$, on X .

PROPOSITION 4.7. $\pi|_{\overline{\mathbb{G}(L_i) \cdot Q_i}}$ is an isomorphism of $\overline{\mathbb{G}(L_i) \cdot Q_i}$ with $\mathbb{A}_{\mathbb{k}}^1$.

PROOF. Note that the composition $\mathbb{G}(L_i) \rightarrow \mathbb{G}(\widetilde{M}) \xrightarrow{\pi} \mathbb{G}(\mathbb{Z})$ is an isomorphism, since it is induced by the isomorphism of lattices given by the composition $L_i \rightarrow \widetilde{M} \rightarrow \mathbb{Z}$. Thus $\pi|_{\overline{\mathbb{G}(L_i) \cdot Q_i}}$ is an isomorphism onto its image. Note that $\overline{\mathbb{G}(L_i) \cdot Q_i}$ is closed in $\pi^{-1}(\mathbb{A}_{\mathbb{k}}^1 \setminus \{0\})$. Since π is proper, when we take the closure of $\overline{\mathbb{G}(L_i) \cdot Q_i}$ in X , we get precisely one additional point sitting in $\pi^{-1}(0)$, and $\pi|_{\overline{\mathbb{G}(L_i) \cdot Q_i}}$ is an isomorphism. \square

Thus, in our setup, the choice of $P_1, \dots, P_s \in M$ yields sections

$$\sigma_1, \dots, \sigma_s : \mathbb{A}_{\mathbb{k}}^1 \rightarrow X$$

of π . So we get a one-parameter family of choices of s points in X_Σ , degenerating to choices of points in X_0 .

We can now describe precisely the three worlds which appear in this picture.

The tropical world. One considers marked parameterized tropical curves of genus 0 through the points P_1, \dots, P_s in $M_{\mathbb{R}}$.

The log world. One considers log morphisms $f : (C^\dagger, x_1, \dots, x_s) \rightarrow X_0^\dagger$ from genus zero log curves C^\dagger with marked points x_1, \dots, x_s such that

$$f(x_i) = \sigma_i(0) \in X_0$$

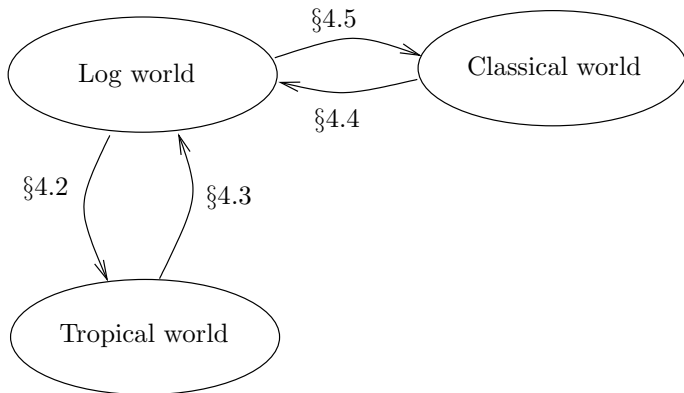
and X_0 is given the log structure induced by the divisorial log structure coming from the inclusion $\partial X \subseteq X$.

The classical world. The numbers $N_{\Delta, \Sigma}^{0, \text{hol}}$ in fact make sense over any algebraically closed field of characteristic zero, and do not depend on the choice of field. So we can consider the field $K = \overline{\mathbb{k}((t))}$, the algebraic closure of $\mathbb{k}((t))$. The inclusion $\mathbb{k}[t] \subseteq K$ gives a map $\text{Spec } K \rightarrow \mathbb{A}_{\mathbb{k}}^1$, and

$$X \times_{\mathbb{A}_{\mathbb{k}}^1} K \cong X_\Sigma \times_{\mathbb{k}} K$$

is the toric variety defined by Σ over the field K . Furthermore, the sections $\sigma_1, \dots, \sigma_s : \mathbb{A}_{\mathbb{k}}^1 \rightarrow X$ determine points $\sigma_1, \dots, \sigma_s : \text{Spec } K \rightarrow X_\Sigma \times_{\mathbb{k}} K$. Using these s points, one considers $N_{\Delta, \Sigma}^{0, \text{hol}}$. A curve contributing to this count is a torically transverse curve over $\text{Spec } K$. There will then turn out to be a natural correspondence between these curves and curves in the log world.

A diagram of the proof of Theorem 4.4 is then



In other words, we will first show, in §4.2, how certain log curves give rise to tropical curves via the “dual intersection complex” construction. In §4.3, we show how to build log curves from tropical curves, in fact showing that for a given tropical curve h , there are $\text{Mult}(h)$ log curves associated to it. In particular, this demonstrates the origin of the multiplicity.

In §4.4, we show how each torically transverse curve over $\text{Spec } K$ gives rise to a log curve, and in §4.5, via logarithmic deformation theory, we will show each log curve deforms to give rise to a torically transverse curve over $\text{Spec } K$. Finally, in §4.6, we put this all together to give a proof of Theorem 4.4.

So on with the details.

4.2. Log world \rightarrow tropical world

We fix in this section the data we have described in the previous section: a fan Σ in $M_{\mathbb{R}}$ two-dimensional, a degree $\Delta \in \ker r$, $s = |\Delta| - 1$, points $P_1, \dots, P_s \in M$, a good lattice polyhedral decomposition \mathcal{P} , and general points $Q_1, \dots, Q_s \in \mathbb{G}(\overline{M})$.

We then obtain from this data a degeneration of toric varieties $\pi : X \rightarrow \mathbb{A}_{\mathbb{k}}^1$ with sections $\sigma_1, \dots, \sigma_s : \mathbb{A}_{\mathbb{k}}^1 \rightarrow X$ with the image of σ_i being $\overline{\mathbb{G}(L_i) \cdot Q_i}$.

As before, let $X_0 = \pi^{-1}(0)$, equipped with the log structure induced by the divisorial log structure on X coming from the inclusion $\partial X \subseteq X$. The morphism $\pi_0 : X_0^\dagger \rightarrow \text{Spec } \mathbb{k}^\dagger$ obtained by restricting $\pi : X^\dagger \rightarrow (\mathbb{A}_{\mathbb{k}}^1)^\dagger$ as in Example 3.24, (3), is log smooth.

DEFINITION 4.8. A *torically transverse log curve* in X_0^\dagger is a diagram

$$\begin{array}{ccc} C^\dagger & \xrightarrow{f} & X_0^\dagger \\ g \downarrow & & \downarrow \pi_0 \\ \text{Spec } \mathbb{k}^\dagger & \xrightarrow{=} & \text{Spec } \mathbb{k}^\dagger \end{array}$$

of log morphisms where g is log smooth and integral, C is a curve with C^\dagger fine saturated, the scheme morphism underlying f is a stable map, and for every vertex $v \in \mathcal{P}$, $f^{-1}(D_v) \rightarrow D_v$ is a torically transverse stable map.

The following gives crucial restrictions on log curves which allow us to make the connection with tropical geometry.

PROPOSITION 4.9. *Let $f : C^\dagger \rightarrow X_0^\dagger$ be a torically transverse log curve. If $x \in C$ is a point such that $f(x) \in \text{Sing}(X_0)$, then*

- (1) x is a double point of C , contained in two distinct irreducible components C_1, C_2 of C , and $f(C_i) \subseteq D_{v_i}$, $i = 1, 2$, for $v_1, v_2 \in \mathcal{P}$ distinct vertices joined by an edge ω of \mathcal{P} .
- (2) Let w_i be the multiplicity of the intersection of C_i at x with $D_\omega \subseteq D_{v_i}$. By this, we mean the order of vanishing of the regular function $f^*(\varphi)$ on C_i at x , where φ is a regular function defined in a neighbourhood of $f(x)$ in D_{v_i} such that $\varphi = 0$ defines D_ω . Then $w_1 = w_2$.
- (3) There is an $e \geq 1$ such that

$$\overline{\mathcal{M}}_{C,x} = S_e := \mathbb{N}^2 \oplus_{\mathbb{N}} \mathbb{N}$$

where the fibred sum is defined by the diagonal map $\mathbb{N} \rightarrow \mathbb{N}^2$ and $\mathbb{N} \rightarrow \mathbb{N}$ multiplication by e . Alternatively, S_e can be described as the monoid generated by α_1, α_2 and ρ subject to the relation $\alpha_1 + \alpha_2 = e\rho$. (See Example 3.27.)

- (4) If ℓ is the affine length of ω , then $e w_i = \ell$.

PROOF. Let us first describe the local situation in X near $f(x)$. By the toric transversality condition, $f(x)$ cannot lie in a zero-dimensional stratum. So $f(x)$ must lie in a one-dimensional stratum, say D_ω for $\omega \in \mathcal{P}$ an edge with endpoints v_1 and v_2 . Without loss of generality, let us write $v_1 = (0, 0) \in M$ and $v_2 = (\ell, 0) \in M$,

where ℓ is the affine length of ω . Then the cone $C(\omega)$ defines an affine open subset $X_{C(\omega)} \subseteq X$. For convenience, we write this affine open subset as X_ω .

Note that the cone $C(\omega)^\vee$ is generated in $\tilde{N}_\mathbb{R} = \mathbb{R}^3$ by $(0, \pm 1, 0)$, $(1, 0, 0)$, $(-1, 0, \ell)$, and $C(\omega)^\vee \cap \tilde{N} = \mathbb{Z} \oplus S_\ell$, with $(0, 1, 0)$ generating \mathbb{Z} and $\alpha_1 = (1, 0, 0)$, $\alpha_2 = (-1, 0, \ell)$, $\rho = (0, 0, 1)$ generating S_ℓ . Note that $X_0 \cap X_\omega = D_1 \cup D_2$, the two divisors corresponding to v_1 and v_2 .

Replace C with an open neighbourhood of x whose image under f is contained in X_ω . We can assume that this open neighbourhood is small enough so that no point other than x maps into $\text{Sing}(X_0)$ and C has at most two irreducible components. Here we are using toric transversality of f to ensure that no component of C maps into $\text{Sing}(X_0)$. So we now have a commutative diagram

$$\begin{array}{ccccc} C^\dagger & \xrightarrow{f} & (X_\omega \cap X_0)^\dagger & \longrightarrow & X_\omega^\dagger \\ g \downarrow & & \pi_0 \downarrow & & \downarrow \pi \\ \text{Spec } \mathbb{k}^\dagger & \xrightarrow{=} & \text{Spec } \mathbb{k}^\dagger & \longrightarrow & (\mathbb{A}_\mathbb{k}^1)^\dagger \end{array}$$

Replace f with the composition of f with the inclusion $(X_\omega \cap X_0)^\dagger \rightarrow X_\omega^\dagger$, so we now have a commutative diagram

$$(4.2) \quad \begin{array}{ccc} C^\dagger & \xrightarrow{f} & X_\omega^\dagger \\ g \downarrow & & \downarrow \pi \\ \text{Spec } \mathbb{k}^\dagger & \longrightarrow & (\mathbb{A}_\mathbb{k}^1)^\dagger \end{array}$$

First suppose that $f(C)$ is contained in, say, D_1 , with $f(x) \in D_1 \cap D_2$. Note that $z^{\alpha_1} \in \Gamma(X_\omega, \mathcal{M}_{X_\omega})$, so $f^\#(z^{\alpha_1}) \in \Gamma(C, \mathcal{M}_C)$. On the other hand, z^{α_1} vanishes on D_2 but not on D_1 , so $z^{\alpha_1} \in \Gamma(X_\omega \setminus D_2, \mathcal{O}_{X_\omega}^\times)$, so in fact $\alpha_C(f^\#z^{\alpha_1}) = f^*\alpha_X(z^{\alpha_1}) = f^*(z^{\alpha_1}) \in \Gamma(C \setminus \{x\}, \mathcal{O}_C^\times)$. Noting that $f^*(z^{\alpha_1})$ is not invertible at x , we conclude that the image of $f^\#(z^{\alpha_1})$ in $\overline{\mathcal{M}}_C$ has support exactly at the point x .

Now there are three possibilities for the behaviour of C^\dagger at x , by Example 3.26. In Case (1) of that example, where x is a smooth point of C , $\overline{\mathcal{M}}_C$ is locally the constant sheaf $\underline{\mathbb{N}}$, hence there are no such sections with support only at x .

In Case (2), where x is a double point, $\overline{\mathcal{M}}_C$ again has no sections with support only at x . To see this, note that from the description in Example 3.26, (2), C locally looks like $\text{Spec } \mathbb{k}[u, v]/(uv)$ near x . Furthermore, the description in Example 3.27 tells us the log structure. It is induced by the inclusion $C \subseteq V_e$, where

$$V_e = \text{Spec } \mathbb{k}[u, v, t]/(uv - t^e)$$

for some $e > 0$ and $C = V(t) \subseteq V_e$. This inclusion induces a divisorial log structure on V_e , which pulls back to the log structure on C . Since $\overline{\mathcal{M}}_{V_e}$ is supported on C , $\overline{\mathcal{M}}_C$ has a section with support only at x only if $\overline{\mathcal{M}}_{V_e}$ does. But as sections of $\overline{\mathcal{M}}_{V_e}$ have support along Cartier divisors of V_e supported on C , there can be no section with support only at x . Hence the second case is ruled out.

In Case (3) of Example 3.26, where x is a log marked point, locally near x we have $\overline{\mathcal{M}}_C = \underline{\mathbb{N}} \oplus \mathbb{N}_x$, which does have a section supported only at x . However, we rule this case out as follows. First note that by Example 3.26, (3), the map $g : C^\dagger \rightarrow \text{Spec } \mathbb{k}^\dagger$ is determined by some $\rho_C \in \Gamma(C, \mathcal{M}_C)$ which maps to $\bar{\rho}_C \in \Gamma(C, \overline{\mathcal{M}}_C)$,

locally of the form $(1, 0)$ in $\mathbb{N} \oplus \mathbb{N}_x$. Second, by the commutativity of (4.2), since π is given by $z^\rho \in \Gamma(X_\omega, \mathcal{M}_{X_\omega})$, we have $f^\# z^\rho = \rho_C$. But $f^\#(z^{\ell\rho}) = f^\#(z^{\alpha_1} z^{\alpha_2})$. Consider the image $\overline{f^\#(z^{\alpha_1})}$ of $f^\#(z^{\alpha_1})$ in $\overline{\mathcal{M}_{C, \bar{x}}}$. Since $\overline{f^\#(z^{\alpha_1})}$ has support at x and is non-zero, we must have

$$\overline{f^\#(\alpha_1)} = (0, a)$$

for some $a > 0$. Then there is no choice for the image $\overline{f^\#(z^{\alpha_2})}$ of $f^\#(z^{\alpha_2})$ in $\overline{\mathcal{M}_{C, \bar{x}}}$ to get

$$(\ell, 0) = \ell \bar{\rho}_C = (0, a) + \overline{f^\#(z^{\alpha_2})}$$

in $\overline{\mathcal{M}_{C, \bar{x}}}$. This rules out the third case of Example 3.26.

In conclusion, $f(C)$ is not contained in either D_1 or D_2 . The only way this can happen is if x is a double point of C , the intersection of two irreducible components C_1 and C_2 with $f(C_1) \subseteq D_1$, $f(C_2) \subseteq D_2$. This demonstrates (1).

For (2)-(4), note that $\overline{\mathcal{M}_{C, \bar{x}}} = S_e$ for some $e \geq 0$. The diagram (4.2), at the level of stalks of the ghost sheaves, is

$$(4.3) \quad \begin{array}{ccc} S_e = \overline{\mathcal{M}_{C, \bar{x}}} & \xleftarrow{\bar{f}^\#} & \overline{\mathcal{M}_{X_\omega, \overline{f(x)}}} = S_\ell \\ \uparrow 1 \mapsto \bar{\rho}_C & & \uparrow 1 \mapsto \rho \\ \mathbb{N} & \xleftarrow{=} & \overline{\mathcal{M}_{\mathbb{A}_k^1, \bar{0}}} = \mathbb{N} \end{array}$$

Let w_1, w_2 be the integers defined in the statement of (2). We note that, in the local description of C near x given by Example 3.26, (2), with $C_1 = V(v)$ and $C_2 = V(u)$, we have

$$(4.4) \quad \begin{aligned} f^*(z^{\alpha_1}) &= u^{w_1} \cdot \varphi_1 \\ f^*(z^{\alpha_2}) &= v^{w_2} \cdot \varphi_2 \end{aligned}$$

with φ_1, φ_2 invertible functions on C near x . Indeed, this is precisely the definition of w_1 and w_2 .

Let us use this information to compute $f^\#(z^{\alpha_1})$ and $f^\#(z^{\alpha_2})$. Given the description of the chart defining the log structure on C given in Example 3.26, (2), $f^\#(z^{\alpha_i})$ can be written, locally near x , as a section of $\underline{S_e} \oplus \mathcal{O}_C^\times$, say $((a_i, b_i), q_i), \psi_i$, with

$$f^\#(z^{\alpha_i}) = \alpha_C f^\#(z^{\alpha_i}) = \begin{cases} u^{a_i} v^{b_i} \psi_i & q_i = 0 \\ 0 & q_i \neq 0 \end{cases}$$

Thus in fact from (4.4),

$$\begin{aligned} ((a_1, b_1), q_1) &= ((w_1, 0), 0), \\ ((a_2, b_2), q_2) &= ((0, w_2), 0). \end{aligned}$$

This tells us that the map $\bar{f}^\#$ in (4.3) must satisfy

$$(4.5) \quad \begin{aligned} \bar{f}^\#((1, 0), 0) &= ((w_1, 0), 0) \\ \bar{f}^\#((0, 1), 0) &= ((0, w_2), 0) \\ \bar{f}^\#((0, 0), 1) &= ((0, 0), 1), \end{aligned}$$

the latter by commutativity of (4.3). But since in S_e , $((0, 0), e) = ((1, 1), 0)$ and in S_ℓ , $((0, 0), \ell) = ((1, 1), 0)$, we have in S_e ,

$$((\ell, \ell), 0) = ((0, 0), \ell e) = \bar{f}^\#((0, 0), \ell e) = \bar{f}^\#((e, e), 0) = ((ew_1, ew_2), 0).$$

From this we conclude $w_1 = w_2 = \ell/e$.

This gives the remainder of the proposition. \square

Having understood these restrictions on a torically transverse log curve, we can now define its associated tropical curve.

DEFINITION 4.10. Let $f : C^\dagger \rightarrow X_0^\dagger$ be a torically transverse log curve over $\text{Spec } \mathbb{k}^\dagger$. Two irreducible components of C are said to be *indistinguishable* if they intersect in a node not mapping into $\text{Sing}(X_0)$.

We define an equivalence relation on the set of irreducible components of C , with two irreducible components C_1 and C_n said to be *equivalent* if there is a chain C_1, C_2, \dots, C_n of irreducible components with C_i and C_{i+1} indistinguishable for $1 \leq i \leq n-1$.

Let $\tilde{\Gamma}_f$ be the weighted graph (with some non-compact edges and possibly with some bivalent vertices) such that:

- (1) The set of vertices of $\tilde{\Gamma}_f$ are the irreducible components of C modulo equivalence. If C' is an irreducible component of C , we write $V_{C'}$ for the vertex corresponding to the equivalence class of C' .
- (2) The set of bounded edges of $\tilde{\Gamma}_f$ are in one-to-one correspondence with nodes of C mapping into $\text{Sing}(X_0)$. If y is such a node, denote by E_y the corresponding edge. If $y \in C_1 \cap C_2$ with C_1, C_2 distinct components of C (these components exist by Proposition 4.9) then E_y has endpoints V_{C_1} and V_{C_2} . If $f(x) \in D_\omega$ with $\omega \in \mathcal{P}$ an edge of affine length ℓ , and $\overline{\mathcal{M}}_{C,x} = S_e$, then the weight $w(E_x)$ of the edge E_x is ℓ/e . By Proposition 4.9, (4), this is a positive integer and coincides with either of the two integers w_1, w_2 defined in (2) of that proposition.
- (3) The set of unbounded edges of $\tilde{\Gamma}_f$ is in one-to-one correspondence with the set $f^{-1}(\partial X_0)$, where ∂X_0 is defined in (4.1). Let $p \in f^{-1}(\partial X_0)$ correspond to an unbounded edge E_p . If $p \in C'$, an irreducible component of C , then the endpoint of E_p is $V_{C'}$, and the weight $w(E_p)$ is the intersection multiplicity of C' with ∂X_0 at $f(p)$.

Given this weighted graph, we then define a tropical curve

$$h : \tilde{\Gamma}_f \rightarrow M_{\mathbb{R}}$$

by

- (1) $h(V_{C'}) = v \in \mathcal{P}$ if $f(C') \subseteq D_v$.
- (2) $h(E_y)$ is the line segment joining $h(V_{C_1}), h(V_{C_2})$, for $y \in C_1 \cap C_2$.
- (3) $h(E_p)$ is the edge $\omega \in \mathcal{P}$ corresponding to the one-dimensional toric stratum of ∂X_0 containing $f(p)$.

Note that $\tilde{\Gamma}_f$ can contain bivalent vertices, which we did not allow in our definition of tropical curve (Definition 1.11). These bivalent vertices arise from irreducible components of C (or equivalence classes) which only contain two nodes mapping under f to $\text{Sing}(X_0)$. To get an actual tropical curve, one can remove the bivalent vertices, replacing a chain of edges connected via bivalent vertices with a

single edge. This gives a new graph, Γ_f , which has the same homeomorphism type as $\tilde{\Gamma}_f$. So we can view h as giving a map

$$h : \Gamma_f \rightarrow M_{\mathbb{R}}.$$

There are two issues which one might have to worry about. First, given a chain of edges connected by bivalent vertices, one needs to check that these edges all have the same weight, so that there is a well-defined choice of weight for the edge of Γ_f corresponding to this chain. Second, one needs to check that this entire chain is mapped by h to an affine line segment, so that h maps the new edge in Γ_f to an affine line segment. Both of these facts will follow from balancing for $\tilde{\Gamma}_f$, in the following proposition.

PROPOSITION 4.11. *If $f : C^\dagger \rightarrow X_0^\dagger$ is a torically transverse log curve, then $h : \Gamma_f \rightarrow M_{\mathbb{R}}$ is a parameterized tropical curve.*

PROOF. We will first show that the balancing condition (2) of Definition 1.11 holds for $h : \tilde{\Gamma}_f \rightarrow M_{\mathbb{R}}$. Once we know this, then any two edges adjacent to a bivalent vertex of $\tilde{\Gamma}_f$ have the same weight and map into the same affine line of $M_{\mathbb{R}}$. This then shows that the corresponding $h : \Gamma_f \rightarrow M_{\mathbb{R}}$ is now a parameterized tropical curve. Indeed, each edge of Γ_f now has a well-defined weight, as discussed above, and maps to an affine line segment. Any vertex in Γ_f is already a vertex in $\tilde{\Gamma}_f$, hence the balancing condition still holds for $h : \Gamma_f \rightarrow M_{\mathbb{R}}$.

To show balancing for $\tilde{\Gamma}_f$, let V be a vertex, and $C' \subseteq C$ the union of irreducible components in the equivalence class corresponding to the vertex V . Necessarily $h(C') \subseteq D_v$ for some vertex $v \in \mathcal{P}$. Let Σ_v be the fan in $M_{\mathbb{R}}$ defining D_v , with $m_1, \dots, m_r \in M$ primitive generators of the rays of Σ_v , corresponding to toric divisors D_1, \dots, D_r of D_v . These divisors are Cartier away from the zero-dimensional toric strata of D_v , hence $(f|_{C'})^*(D_i)$ is a Cartier divisor on C' .

Let

$$w_i := \deg(f|_{C'})^*(D_i).$$

Clearly w_i is the sum of weights of edges of $\tilde{\Gamma}_f$ adjacent to V mapping into $V + \mathbb{R}_{\geq 0}m_i$, by the definition of the weights. So we need to show that $\sum_{i=1}^r w_i m_i = 0$. To show this we show for any $n \in N$ that

$$\sum_{i=1}^r w_i \langle n, m_i \rangle = 0.$$

Note that $\langle n, m_i \rangle$ is the order of vanishing of z^n on D_i , so $\sum w_i \langle n, m_i \rangle$ is the number of zeroes and poles of $(f|_{C'})^*(z^n)$ on C' . Of course, this is zero. \square

Finally, we consider a decoration of the above situation. Recall we have fixed points $P_1, \dots, P_s \in M$, $Q_1, \dots, Q_s \in \mathbb{G}(\tilde{M})$, giving sections σ_i of π . Let $q_i = \sigma_i(0)$.

DEFINITION 4.12. Let $(C^\dagger, x_1, \dots, x_s)$ denote a log curve over $\text{Spec } \mathbb{k}^\dagger$ along with a choice of points $x_1, \dots, x_s \in C$ which are all smooth points (not log marked points). Then a *torically transverse marked log curve passing through q_1, \dots, q_s* is a log morphism $f : (C^\dagger, x_1, \dots, x_s) \rightarrow X_0^\dagger$ over $\text{Spec } \mathbb{k}^\dagger$ such that f is a torically transverse log curve in X_0^\dagger with $f(x_i) = q_i$.

Given such a map, we associate to f a graph $\tilde{\Gamma}_f$ as in Definition 4.10, but with one additional unbounded edge E_{x_i} for each x_i , with E_{x_i} having endpoint $V_{C'}$ if

$x_i \in C' \subseteq C$. We set $w(E_{x_i}) = 0$, and $h : \tilde{\Gamma}_f \rightarrow M_{\mathbb{R}}$ is defined exactly as before, with $h|_{E_{x_i}}$ constant.

As before, we can remove any remaining bivalent vertices (there may be fewer in this marked case, since what would have been a bivalent vertex in the unmarked case may now have acquired a third, unbounded, marked edge adjacent to it). This gives

$$h : \Gamma_f \rightarrow M_{\mathbb{R}}.$$

We also denote by

$$h : \hat{\Gamma}_f \rightarrow M_{\mathbb{R}}$$

the tropical curve given by Definition 4.10 associated to the unmarked map $f : C^\dagger \rightarrow X_0^\dagger$. This is obtained from Γ_f by removing the marked edges E_{x_i} and removing any additional bivalent vertices produced.

Finally, we observe that $h : \Gamma_f \rightarrow M_{\mathbb{R}}$ is now a marked tropical curve with $h(x_i) = P_i$:

PROPOSITION 4.13. (1) *Keeping in mind that P_i is a vertex of \mathcal{P} , we have $q_i \in \text{Int}(D_{P_i})$.*

(2) *$h : \Gamma_f \rightarrow M_{\mathbb{R}}$ is a marked parameterized tropical curve with $h(E_{x_i}) = P_i$.*

PROOF. (1) Consider the affine open subset $X_{P_i} \subseteq X$ determined by the cone $C(P_i) \in \Sigma_{\mathcal{P}}$. Let $n_1, n_2 \in \tilde{N}$ be a basis for $(P_i, 1)^\perp$, and $n_3 = (0, 0, 1) \in \tilde{N}$. Then $\pm n_1, \pm n_2$ and n_3 generate $C(P_i)^\vee \cap \tilde{N}$, so

$$X_{P_i} \cong \text{Spec } \mathbb{k}[x_1^{\pm 1}, x_2^{\pm 1}, x_3],$$

with $x_i = z^{n_i}$. Furthermore $X_{P_i} \cap D_{P_i}$ is given by $x_3 = 0$.

Now let Q_i be given in these coordinates by $Q_i = (a_1, a_2, a_3)$, $a_i \in \mathbb{k}^\times$, $i = 1, 2, 3$. Then for $z \in \mathbb{G}(L_i) = \mathbb{k}^\times$, $z \cdot Q_i = (a_1, a_2, za_3)$. Clearly $(a_1, a_2, 0) \in \overline{\mathbb{G}(L_i) \cdot Q_i} \cap X_{P_i}$, so $q_i = (a_1, a_2, 0) \in D_{P_i}$.

(2) If $x_i \in C'$, then $h(E_{x_i}) = h(V_{C'}) = v$ for some $v \in \mathcal{P}$, and of course then $f(C') \subseteq D_v$. Since $f(x_i) = q_i \in \text{Int}(D_{P_i})$, we must have $v = P_i$. \square

4.3. Tropical world \rightarrow log world

We continue with the data $P_1, \dots, P_s \in M$, a good lattice polyhedral decomposition \mathcal{P} , and general points $Q_1, \dots, Q_s \in \mathbb{G}(\bar{M})$. This gives points $q_1, \dots, q_s \in X_0$ with $q_i = \overline{\mathbb{G}(L_i) \cdot Q_i} \cap X_0$.

Now suppose we are given a marked tropical curve

$$h : (\Gamma, x_1, \dots, x_s) \rightarrow M_{\mathbb{R}}$$

with $h(x_i) = P_i$.

We would like to find all log curves $f : (C^\dagger, x_1, \dots, x_s) \rightarrow X_0^\dagger$ passing through q_1, \dots, q_s whose associated tropical curve is h . Now in general, such a curve h could be quite complicated. Each vertex of the curve can be represented by a union of curves of various genera, and these can move in families. The simplest case, however, is that h is a *simple* curve, in the sense of Definition 1.19. In this case, we can be much more precise.

So now assume h is a simple curve of genus zero. To ensure that we only have a finite number of such curves, we are assuming as usual that $s = |\Delta| - 1$, where

Δ is the degree of h . Since \mathcal{P} contains all rational tropical curves passing through P_1, \dots, P_s , the image of h is contained in the one-skeleton of \mathcal{P} .

As we noted in the previous section, the initial construction of the tropical curve associated to a log curve yields a graph with possibly bivalent vertices. By subdividing edges of Γ to obtain a graph $\tilde{\Gamma}$, along with $h : \tilde{\Gamma} \rightarrow M_{\mathbb{R}}$, we can assume that for a point $y \in \tilde{\Gamma}$, $h(y)$ is a vertex of \mathcal{P} if and only if y is a vertex of $\tilde{\Gamma}$ or y is contained in a marked unbounded edge. Of course, $\tilde{\Gamma}$ may now have bivalent vertices.

The main theorem of the section is:

THEOREM 4.14. *Given the above situation, suppose that for each edge E of $\tilde{\Gamma}$, the affine length of $h(E)$ is divisible by the weight $w(E)$. Then the number of torically transverse marked log curves*

$$f : (C^\dagger, x_1, \dots, x_s) \rightarrow X_0^\dagger$$

with $f(x_i) = q_i$ whose associated tropical curve is h is $\text{Mult}(h)$.

REMARK 4.15. The hypothesis about the divisibility of the affine length of each edge is necessary by the construction of the tropical curve associated to a log curve. On the other hand, after a further rescaling of the lattice M , one can always achieve this divisibility. We will make this rescaling step more precise later.

Before embarking on the proof, we first examine the nature of irreducible components of C , noting that simplicity of Γ implies that all vertices of $\tilde{\Gamma}$ are either bivalent or trivalent.

DEFINITION 4.16. Let Y be a complete toric surface. A *line* on Y is a non-constant, torically transverse map $\varphi : \mathbb{P}^1 \rightarrow Y$ such that $\#\varphi^{-1}(\partial Y) \leq 3$ and $\#\varphi^{-1}(D) \leq 1$ for any toric divisor D of Y .

Given a line, fix the following notation. Let $u_1, \dots, u_p \in M$ be the primitive generators of the rays in the fan defining Y corresponding to those toric divisors $D_1, \dots, D_p \subseteq Y$ such that $\varphi^{-1}(D_i) \neq \emptyset$ for $1 \leq i \leq p$. Furthermore, let w_i be the order of vanishing of $\varphi^*(z_i)$, where z_i is a local equation for D_i . This gives the order of tangency of φ with D_i . Then as in the proof of Proposition 4.11, the balancing condition holds:

$$\sum_i w_i u_i = 0.$$

Since $p \leq 3$, this implies in particular that $p = 2$ or 3 . We refer to these two cases as the *bivalent* and *trivalent* cases respectively.

We say that a line with associated data $\mathbf{u} = (u_1, u_2, \dots)$, $\mathbf{w} = (w_1, w_2, \dots)$, is of type (\mathbf{u}, \mathbf{w}) , and we write $\mathcal{L}_{(\mathbf{u}, \mathbf{w})}$ for the set of all lines in the toric surface Y of this type (modulo automorphisms of the domain \mathbb{P}^1). We will now classify lines, first dealing with the bivalent case:

LEMMA 4.17. *Let $(\mathbf{u}, \mathbf{w}) = ((u_1, u_2), (w_1, w_2))$, and let E be the sublattice of dimension one of M generated by u_1 (or equivalently, u_2 , by the balancing condition). There is a natural map $g : Y \rightarrow \mathbb{P}^1$ induced by the map of lattices $M \rightarrow M/E$. Then each $\varphi \in \mathcal{L}_{(\mathbf{u}, \mathbf{w})}$ has image $g^{-1}(p) \subseteq Y$ for some $p \in \mathbb{P}^1 \setminus \partial\mathbb{P}^1$, and the map $\varphi : \mathbb{P}^1 \rightarrow g^{-1}(p)$ is a w_1 -fold branched cover of $g^{-1}(p)$, totally branched precisely at the two points $g^{-1}(p) \cap \partial Y$.*

PROOF. Since $\rho_i = \mathbb{R}_{\geq 0}u_i$ is a cone in the fan Σ_Y defining Y , and $u_2 = -u_1$, every cone in Σ_Y is contained in one of the two half-planes containing $\mathbb{R}u_i$, and hence maps into a ray in $(M/E) \otimes_{\mathbb{Z}} \mathbb{R}$. Thus the projection $M \rightarrow M/E$ defines a map of fans from Σ_Y to the fan in $(M/E) \otimes_{\mathbb{Z}} \mathbb{R}$ defining \mathbb{P}^1 . We can then write the toric boundary of Y as $\partial Y = D_1 \cup D_2 \cup g^{-1}(0) \cup g^{-1}(\infty)$. Since $\varphi(\mathbb{P}^1)$ is disjoint from $g^{-1}(0)$ and $g^{-1}(\infty)$, we see that $\varphi(\mathbb{P}^1)$ must be a fibre of g , but not $g^{-1}(0)$ or $g^{-1}(\infty)$. All fibres of g are isomorphic to \mathbb{P}^1 except $g^{-1}(0)$ and $g^{-1}(\infty)$. In order for $\varphi^{-1}(\partial Y)$ to consist of only two points, φ must be totally branched at ∂Y , giving the result. \square

COROLLARY 4.18. *In the case of Lemma 4.17,*

$$\mathcal{L}(\mathbf{u}, \mathbf{w}) \cong \mathbb{G}(M/E).$$

PROOF. This is just observing that there is a unique cover as described in the Lemma, so $\varphi \in \mathcal{L}(\mathbf{u}, \mathbf{w})$ is completely determined by a point in $p \in \mathbb{G}(M/E) = \mathbb{P}^1 \setminus \partial \mathbb{P}^1$. \square

Next, we consider the trivalent case, so fix $(\mathbf{u}, \mathbf{w}) = ((u_1, u_2, u_3), (w_1, w_2, w_3))$. For convenience, we can assume that Σ_Y , the fan defining Y , only has three rays, defined by $\rho_i = \mathbb{R}_{\geq 0}u_i$. Indeed, there is always a toric blow-down $\alpha : Y \rightarrow Y'$ to such a surface, and by toric transversality, the set of lines of type (\mathbf{u}, \mathbf{w}) on Y and Y' are the same.

LEMMA 4.19. *In the above situation, let Σ_Y be the complete fan with exactly three rays, generated by u_1, u_2 and u_3 . Consider the map*

$$\mathbb{Z}^2 \rightarrow M$$

given by

$$(a, b) \mapsto aw_1u_1 + bw_2u_2.$$

This induces a map of fans from the standard fan for \mathbb{P}^2 as given in Example 1.14, $\Sigma_{\mathbb{P}^2}$, in \mathbb{R}^2 , to Σ_Y , and hence a map

$$f_{(\mathbf{u}, \mathbf{w})} : \mathbb{P}^2 \rightarrow Y.$$

Any line $\varphi : \mathbb{P}^1 \rightarrow Y$ of type (\mathbf{u}, \mathbf{w}) is isomorphic to the composition of a torically transverse linear embedding $\mathbb{P}^1 \rightarrow \mathbb{P}^2$ with $f_{(\mathbf{u}, \mathbf{w})}$.

PROOF. The map $\mathbb{Z}^2 \rightarrow M$ clearly takes the first quadrant to the cone spanned by u_1 and u_2 . Note that this map also takes $(-1, -1)$ to $-w_1u_1 - w_2u_2 = w_3u_3$ by the balancing condition, from which it follows that the cone in $\Sigma_{\mathbb{P}^2}$ generated by $(1, 0)$ and $(-1, -1)$ is taken to the cone in Σ_Y generated by u_1 and u_3 , while the cone in $\Sigma_{\mathbb{P}^2}$ generated by $(0, 1)$ and $(-1, -1)$ is taken to the cone in Σ_Y generated by u_2 and u_3 . This gives the map of fans, hence $f_{(\mathbf{u}, \mathbf{w})} : \mathbb{P}^2 \rightarrow Y$.

Now let C be the normalization of an irreducible component of $\mathbb{P}^1 \times_Y \mathbb{P}^2$, with the map $\mathbb{P}^1 \rightarrow Y$ being φ and the map $\mathbb{P}^2 \rightarrow Y$ being $f_{(\mathbf{u}, \mathbf{w})}$. We have the projection $\tilde{\varphi} : C \rightarrow \mathbb{P}^2$. It will be sufficient to show that first, $\tilde{\varphi}$ is the embedding of a straight line in \mathbb{P}^2 and second, the composition $C \rightarrow \mathbb{P}^1 \times_Y \mathbb{P}^2 \rightarrow \mathbb{P}^1$ is an isomorphism.

First note that the degree of $f_{(\mathbf{u}, \mathbf{w})}$ is the same as the degree of the map of big torus orbits $\mathbb{G}(\mathbb{Z}^2) \rightarrow \mathbb{G}(M)$, and that this is just the index of the image of \mathbb{Z}^2 in M , i.e., the index δ of the sublattice of M generated by w_1u_1 and w_2u_2 .

We also compute the degree of branching of $f_{(\mathbf{u}, \mathbf{w})}$ along the divisor D_i . Without loss of generality, consider the divisor D_1 , and we can assume that $u_1 = (1, 0)$,

$u_2 = (a, b)$ in some basis for M . Let ρ_i be the cone generated by u_i . Then ρ_1^\vee is generated by $(0, \pm 1)$ and $(1, 0)$, and the map $f_{(\mathbf{u}, \mathbf{w})}$ on the affine charts of \mathbb{P}^2 and Y determined by the cones $\mathbb{R}_{\geq 0}(1, 0)$ and ρ_1 respectively is given by the map of rings

$$\mathbb{k}[x, y^{\pm 1}] \rightarrow \mathbb{k}[x, y^{\pm 1}]$$

defined by

$$x \mapsto x^{w_1} y^{w_2 a}, \quad y \mapsto y^{w_2 b}.$$

Since y is a unit, and the divisor D_1 is given by $x = 0$, we see that $f_{(\mathbf{u}, \mathbf{w})}$ has $w_2 b$ branches over D_1 , and each branch is ramified over D_1 with order of ramification w_1 . More generally, since the degree of $f_{(\mathbf{u}, \mathbf{w})}$ is δ (being $bw_1 w_2$ in this case), we see that $f_{(\mathbf{u}, \mathbf{w})}$ has δ/w_i branches over D_i , with ramification order of each branch being w_i . Note that this ramification order agrees with the intersection index of φ with D_i at the unique point of intersection.

Now consider the diagram

$$\begin{array}{ccccc}
 C & & & & \\
 \searrow \psi & & \tilde{\varphi} & & \\
 & \mathbb{P}^1 \times_Y \mathbb{P}^2 & \xrightarrow{\pi_2} & \mathbb{P}^2 & \\
 & \downarrow \pi_1 & & \downarrow f_{(\mathbf{u}, \mathbf{w})} & \\
 & \mathbb{P}^1 & \xrightarrow{\varphi} & Y &
 \end{array}$$

Locally near, say, D_1 , in the above coordinates, we have $\varphi^*(x) = \text{unit} \cdot t^{w_1}$, where t is a local parameter on \mathbb{P}^1 in a neighbourhood of $\varphi^{-1}(D_1)$, with t vanishing at $\varphi^{-1}(D_1)$. On the other hand, $f_{(\mathbf{u}, \mathbf{w})}^*(x) = x^{w_1} \cdot y^{w_2 a} = x^{w_1} \cdot \text{unit}$, so $t^{w_1} = x^{w_1} \cdot \text{unit}$ is a local equation in the ideal of $\mathbb{P}^1 \times_Y \mathbb{P}^2$ inside $\mathbb{P}^1 \times \mathbb{P}^2$ in a neighbourhood of a point p of $\mathbb{P}^1 \times_Y \mathbb{P}^2$ with $\varphi(\pi_1(p)) \in D_1$. From this we see that there are w_1 branches of $\mathbb{P}^1 \times_Y \mathbb{P}^2$ passing through p , each branch mapping locally isomorphically to \mathbb{P}^1 via π_1 . After normalizing, these w_1 branches are separated, so if $p' \in C$ with $\psi(p') = p$, ψ is unramified in a neighbourhood of p' . We also see that if t' is a local parameter for C in a neighbourhood of p' , then $\tilde{\varphi}^*(x) = \text{unit} \cdot t'$. So $\tilde{\varphi}$ is transversal to $f_{(\mathbf{u}, \mathbf{w})}^{-1}(D_1)$. As the same analysis works at D_2 and D_3 , we see that ψ is unramified, hence is an isomorphism as \mathbb{P}^1 is simply connected. Furthermore, $\tilde{\varphi}$ is transversal to each toric divisor of \mathbb{P}^2 , and hence is the embedding of a straight line. \square

COROLLARY 4.20. *In the situation of Lemma 4.19, $\mathbb{G}(M)$ acts simply transitively on $\mathcal{L}_{(\mathbf{u}, \mathbf{w})}$. In particular, $\mathcal{L}_{(\mathbf{u}, \mathbf{w})}$ is (non-canonically) isomorphic to $\mathbb{G}(M)$.*

PROOF. Any line $\varphi: \mathbb{P}^1 \rightarrow Y$ lifts to at most $\delta = \deg f_{(\mathbf{u}, \mathbf{w})}$ distinct lines in \mathbb{P}^2 , and δ is the order of the kernel of the homomorphism $\mathbb{G}(\mathbb{Z}^2) \rightarrow \mathbb{G}(M)$. Now $\mathbb{G}(\mathbb{Z}^2)$ clearly acts simply transitively on the set of torically transverse linear embeddings $\tilde{\varphi}: \mathbb{P}^1 \rightarrow \mathbb{P}^2$, and then $\ker(\mathbb{G}(\mathbb{Z}^2) \rightarrow \mathbb{G}(M))$ acts transitively on the set of lifts $\tilde{\varphi}$ of φ to \mathbb{P}^2 . This shows that the action of $\mathbb{G}(M)$ on $\mathcal{L}_{(\mathbf{u}, \mathbf{w})}$ is simple transitive. \square

We will now put this together to create the underlying morphism of schemes $f: C \rightarrow X_0$ of the desired log curves. We shall call such curves *pre-log curves*:

DEFINITION 4.21. A *torically transverse pre-log curve* in X_0 is a stable map $f : C \rightarrow X_0$ with C a curve, such that for every vertex $v \in \mathcal{P}$, $f^{-1}(D_v) \rightarrow D_v$ is a torically transverse curve, and f satisfies conclusions (1) and (2) of Proposition 4.9.

Since we only used (1) and (2) of Proposition 4.9 to construct the tropical curve associated to a torically transverse log curve, we can in fact associate a tropical curve to any torically transverse pre-log curve.

Now return to the situation at the beginning of the section in which we are given a marked tropical curve

$$h : (\Gamma, x_1, \dots, x_s) \rightarrow M_{\mathbb{R}}$$

with $h(x_i) = P_i$, and associated to this the curve $h : \tilde{\Gamma} \rightarrow M_{\mathbb{R}}$ with extra vertices. We can associate to $h : \Gamma \rightarrow M_{\mathbb{R}}$ another tropical curve $h : \hat{\Gamma} \rightarrow M_{\mathbb{R}}$ obtained by removing all marked edges from Γ and as usual removing any resulting bivalent vertices. The graph $\hat{\Gamma}$ will then have distinguished edges E_1, \dots, E_s with the endpoint of E_{x_i} in Γ now contained in the interior of E_i . Note a priori that the E_i 's need not be distinct (in fact they will be distinct, but we won't need this). If an edge E of $\hat{\Gamma}$ is equal to E_i for some i , we will say that E is a *marked edge* of $\hat{\Gamma}$.

PROPOSITION 4.22. *For each edge E of $\hat{\Gamma}$, choose an orientation. This gives a choice of the endpoints of E , which we denote by $\partial^- E$ and $\partial^+ E$ (if E is non-compact, then by $\partial^- E$ only). After making this choice, let*

$$u_{(\partial^- E, E)} \in M$$

denote a primitive tangent vector to $h(E)$ pointing from $h(\partial^- E)$ to $h(\partial^+ E)$ (or in the unbounded direction of $h(E)$). Then the map

$$(4.6) \quad \Phi : \text{Map}(\hat{\Gamma}^{[0]}, M) \rightarrow \left(\prod_{E \in \hat{\Gamma}^{[1]} \setminus \hat{\Gamma}_{\infty}^{[1]}} M/\mathbb{Z}u_{(\partial^- E, E)} \right) \times \left(\prod_{i=1}^s M/\mathbb{Z}u_{(\partial^- E_i, E_i)} \right)$$

$$H \mapsto \left((H(\partial^+ E) - H(\partial^- E))_E, (H(\partial^- E_i))_i \right).$$

is an inclusion of lattices of finite index. Call this index \mathfrak{D} .

Then \mathfrak{D} is the number of marked torically transverse pre-log curves, up to isomorphism, of the form $f : (C, x_1, \dots, x_s) \rightarrow X_0$ with $f(x_i) = q_i$ and associated tropical curve h .

PROOF. Tensoring Φ with \mathbb{R} , we obtain a description of deformations of the tropical curve $h : \hat{\Gamma} \rightarrow M_{\mathbb{R}}$ preserving the incidence conditions with the points P_i . Indeed, given $H \in \text{Map}(\hat{\Gamma}^{[0]}, M_{\mathbb{R}})$, we can modify h by taking a vertex V to $h(V) + H(V)$. A compact edge E is then mapped to the edge joining $h(\partial^- E) + H(\partial^- E)$ and $h(\partial^+ E) + H(\partial^+ E)$. A non-compact, non-contracted edge E is mapped to the ray $\mathbb{R}_{\geq 0}u_{(\partial^- E, E)} + h(\partial^- E) + H(\partial^- E)$. Now in general, such a deformation is not a tropical curve, and if we want to deform h so it remains a tropical curve without changing the combinatorial type and ensuring that the image of E_i still contains P_i , we need two conditions to be satisfied:

- (1) For each compact edge E , $(h(\partial^+ E) + H(\partial^+ E)) - (h(\partial^- E) + H(\partial^- E))$ is parallel to $u_{(\partial^- E, E)}$, or equivalently,

$$H(\partial^+ E) - H(\partial^- E) = 0 \text{ in } M/\mathbb{Z}u_{(\partial^- E, E)}.$$

- (2) For each i , $H(\partial^- E_i) = 0$ in $M/\mathbb{Z}u_{(\partial^- E_i, E_i)}$. This, along with (1) for $E = E_i$, guarantees that the affine line spanned by the image of E_i still passes through P_i .

Put more succinctly, H provides a deformation of h as a tropical curve only if H is in the kernel of the map Φ tensored with \mathbb{R} . If H is sufficiently close to the origin, then it will produce a deformation of h . Since by assumption h is rigid, the kernel of Φ is in fact zero. A dimension count then shows that the the image of Φ is a sublattice of finite index, say \mathfrak{D} .

Now for every $V \in \tilde{\Gamma}^{[0]}$, write $\mathbf{u}(V)$, $\mathbf{w}(V)$ for the data of the primitive tangent vectors to unmarked edges adjacent to V (pointing away from $h(V)$) and their weights.

Necessarily, the curve C we are trying to build is glued from lines in the toric surfaces $D_{h(V)}$ for V running over $\tilde{\Gamma}^{[0]}$. For a given V , the set of lines of type $(\mathbf{u}(v), \mathbf{w}(v))$ in $D_{h(V)}$ is $\mathcal{L}_{(\mathbf{u}(V), \mathbf{w}(V))}$. Now for $V \in \tilde{\Gamma}^{[0]}$ trivalent with no adjacent marked edge, $\mathcal{L}_{(\mathbf{u}(V), \mathbf{w}(V))}$ is a torsor over $\mathbb{G}(M)$ by Corollary 4.20. For $V \in \tilde{\Gamma}^{[0]}$ bivalent or trivalent with an adjacent marked edge, with adjacent unmarked edges $E^\pm(V)$, $\mathcal{L}_{(\mathbf{u}(V), \mathbf{w}(V))}$ is a torsor over $\mathbb{G}(M/u_{(V, E^\pm(V))})$ by Corollary 4.18. Note that the trivalent vertices with no adjacent marked edges are precisely the vertices of $\hat{\Gamma}$, while the remaining vertices are those vertices of $\tilde{\Gamma}$ which are not vertices of $\hat{\Gamma}$. Thus $\prod_{V \in \tilde{\Gamma}^{[0]}} \mathcal{L}_{(\mathbf{u}(V), \mathbf{w}(V))}$ is a torsor over

$$\mathbb{G}(\text{Map}(\hat{\Gamma}^{[0]}, M)) \times \prod_{V \in \tilde{\Gamma}^{[0]} \setminus \hat{\Gamma}^{[0]}} \mathbb{G}(M/u_{(V, E^-(V))}).$$

In order to get a torically transverse pre-log curve, we need to match up the various lines and marked points. We do this by defining a map

$$(4.7) \quad \Phi' : \prod_{V \in \tilde{\Gamma}^{[0]}} \mathcal{L}_{(\mathbf{u}(V), \mathbf{w}(V))} \rightarrow \prod_{E \in \tilde{\Gamma}^{[1]} \setminus \tilde{\Gamma}_\infty^{[1]}} \mathbb{G}(M/u_{(\partial^- E, E)}) \times \prod_{i=1}^s \mathbb{G}(M/\mathbb{Z}u_{(\partial^- E_i, E_i)})$$

as follows.

(1) Given an element $(\varphi_V)_{V \in \tilde{\Gamma}^{[0]}}$ of the domain of this map, the component of the image of this element under the map Φ' corresponding to $E \in \tilde{\Gamma}^{[1]} \setminus \tilde{\Gamma}_\infty^{[1]}$ is defined as follows. Let $V^\pm = \partial^\pm E$, $v^\pm = h(V^\pm)$, and let ω be the edge of \mathcal{S} joining v^+ and v^- , so that $D_{v^+} \cap D_{v^-} = D_\omega$. We have $\varphi_{V^\pm} : C_\pm \cong \mathbb{P}^1 \rightarrow D_{v^\pm}$, and let $p^\pm = \varphi_{v^\pm}(C_\pm) \cap D_\omega$. The two points p^\pm lie in the big torus orbit in D_ω by transversality, and this orbit is isomorphic to $\mathbb{G}(M/\mathbb{Z}u_{(\partial^- E, E)})$. We then take the component of the image under Φ' of $(\varphi_V)_V$ corresponding to E to be the ratio p^+/p^- . This is 1 precisely if $p^- = p^+$.

(2) Given an element $(\varphi_V)_{V \in \tilde{\Gamma}^{[0]}}$, the component of the image of Φ' corresponding to E_i is defined as follows. Let $V_i \in \tilde{\Gamma}^{[0]}$ be the vertex adjacent to E_{x_i} , so $h(V_i) = P_i$. Since P_i lies in the interior of $h(E_i)$, the projection $M \rightarrow M/\mathbb{Z}u_{(\partial^- E_i, E_i)}$ defines a map of toric varieties $g_i : D_{P_i} \rightarrow \mathbb{P}^1 \supseteq \mathbb{G}(M/\mathbb{Z}u_{(\partial^- E_i, E_i)})$. By Lemma 4.17, the image of φ_{V_i} is a fibre $g_i^{-1}(r_i)$. We then take the component of the image under Φ' of $(\varphi_V)_V$ corresponding to E_i to be $(g_i(q_i))/r_i$. This makes sense. Indeed, by Proposition 4.13, $q_i \in D_{P_i} \setminus (\partial D_{P_i})$ so $g_i(q_i) \in \mathbb{P}^1 \setminus \partial \mathbb{P}^1$. Note that this ratio is 1 precisely when q_i lies in the image of φ_{V_i} .

We can use the map Φ' as follows. Given $\varphi : (C, x_1, \dots, x_s) \rightarrow X$ a torically transverse pre-log curve with $\varphi(x_i) = q_i$ and associated tropical curve h , the elements of $\tilde{\Gamma}^{[0]}$ are in one-to-one correspondence with equivalence classes of irreducible components of C . For $V \in \tilde{\Gamma}^{[0]}$, if C_V denotes the union of irreducible components in the equivalence class corresponding to V , then as h is simple, $\varphi_V := \varphi|_{C_V} : C_V \rightarrow D_{h(V)}$ is a line. Indeed, as h is simple, there are at most 3 points on C_V mapping into $\text{Sing}(X_0)$, as there is by construction a one-to-one correspondence between such points and bounded edges of Γ adjacent to V . Furthermore, again by simplicity, each of these points maps to different irreducible components of $\text{Sing}(X_0)$, since no two edges of Γ adjacent to V can map to the same segment in $M_{\mathbb{R}}$. Since the image of any non-contracted component of C_V intersects $\text{Sing}(X_0)$ in at least two points, C_V can contain only one non-contracted component. Furthermore, since C is rational and φ is stable, in fact C_V cannot contain any contracted component. Hence C_V is irreducible and φ_V is a line. We obtain

$$(\varphi_V)_V \in \prod_{V \in \tilde{\Gamma}^{[0]}} \mathcal{L}_{(\mathbf{u}(V), \mathbf{w}(V))}.$$

Furthermore, the image of this element under Φ' is $(1, \dots, 1)$. Indeed, the fact that the components of type (1) above are 1 is just saying that the lines $\varphi_{\partial^- E}$ and $\varphi_{\partial^+ E}$ pass through the same point in D_{ω} , and the fact that the components of type (2) above are 1 is saying that the line φ_{V_i} passes through the point q_i .

Conversely, given $(\varphi_V)_V \in \prod_{V \in \tilde{\Gamma}^{[0]}} \mathcal{L}_{(\mathbf{u}(V), \mathbf{w}(V))}$ which maps to $(1, \dots, 1)$ under Φ' , it is clear we can reverse the process, gluing the C_V 's to obtain a torically transverse pre-log curve $\varphi : C \rightarrow X_0$. We also need to mark this curve: for $1 \leq i \leq s$, choose a point $x_i \in C_{V_i}$ mapping under φ_{V_i} to q_i . There are $w(E_i)$ choices, but there is an automorphism of C_{V_i} leaving the two points of $\varphi_{V_i}^{-1}(\partial D_{P_i})$ fixed and which takes any one choice of x_i to any other choice. Thus we conclude that up to isomorphism, the set of torically transverse pre-log maps $f : (C, x_1, \dots, x_s) \rightarrow X_0$ passing through q_1, \dots, q_s and with associated tropical curve h is precisely the inverse image of $(1, \dots, 1)$ under Φ' . We wish to compute the cardinality of this set, i.e., the degree of Φ' .

Now Φ' is a map of torsors over algebraic tori whose corresponding map of character lattices is Φ'' , which can be described as follows. For each vertex $V \in \tilde{\Gamma}^{[0]} \setminus \hat{\Gamma}^{[0]}$, let $E(V)$ denote the unique edge of $\hat{\Gamma}$ containing V . For an edge E of $\tilde{\Gamma}$, write \hat{E} for the edge of $\hat{\Gamma}$ containing E . Then we can write

$$\begin{aligned} \Phi'' : \prod_{V \in \tilde{\Gamma}^{[0]}} M \times \prod_{V \in \tilde{\Gamma}^{[0]} \setminus \hat{\Gamma}^{[0]}} M / \mathbb{Z}u_{(\partial^- E(V), E(V))} \\ (4.8) \quad \rightarrow \prod_{E \in \tilde{\Gamma}^{[1]} \setminus \tilde{\Gamma}_{\infty}^{[1]}} M / \mathbb{Z}u_{(\partial^- \hat{E}, \hat{E})} \times \prod_{i=1}^s M / \mathbb{Z}u_{(\partial^- E_i, E_i)}. \end{aligned}$$

Given $(m_V)_{V \in \tilde{\Gamma}^{[0]}}$ with $m_V \in M$ if $V \in \hat{\Gamma}^{[0]}$ and $m_V \in M / \mathbb{Z}u_{(\partial^- E(V), E(V))}$ if $V \in \tilde{\Gamma}^{[0]} \setminus \hat{\Gamma}^{[0]}$, the image of (m_V) under Φ'' has the following components. The component corresponding to an edge $E \in \tilde{\Gamma}^{[1]} \setminus \tilde{\Gamma}_{\infty}^{[1]}$ is $m_{\partial^+ E} - m_{\partial^- E}$. The component corresponding to E_i is m_{V_i} .

We can relate Φ'' to Φ as follows. Let $S \subseteq \tilde{\Gamma}^{[1]} \setminus \tilde{\Gamma}_{\infty}^{[1]}$ be the set of edges E such that \hat{E} is compact and $\partial^+ \hat{E}$ is a vertex of E ; clearly there is a bijection between S

and $\widehat{\Gamma}^{[1]} \setminus \widehat{\Gamma}_{\infty}^{[1]}$. Then we split

$$\prod_{E \in \widehat{\Gamma}^{[1]} \setminus \widehat{\Gamma}_{\infty}^{[1]}} M/\mathbb{Z}u_{(\partial-\widehat{E}, \widehat{E})} = \prod_{E \in S} M/\mathbb{Z}u_{(\partial-\widehat{E}, \widehat{E})} \times \prod_{E \in \widehat{\Gamma}^{[1]} \setminus (\widehat{\Gamma}_{\infty}^{[1]} \cup S)} M/\mathbb{Z}u_{(\partial-\widehat{E}, \widehat{E})}.$$

Now Φ'' induces an isomorphism between

$$\prod_{V \in \widehat{\Gamma}^{[0]} \setminus \widehat{\Gamma}^{[0]}} M/\mathbb{Z}u_{(\partial-E(V), E(V))}$$

and

$$\prod_{E \in \widehat{\Gamma}^{[1]} \setminus (\widehat{\Gamma}_{\infty}^{[1]} \cup S)} M/\mathbb{Z}u_{(\partial-\widehat{E}, \widehat{E})}.$$

We will now show that the index of Φ and the index of Φ'' agree. Since the index of Φ'' is the degree of the map Φ' , the result will follow.

Introducing short-hand notation, let

$$L_1 = \text{Map}(\widehat{\Gamma}^{[0]}, M),$$

$$L_2 = \prod_{E \in \widehat{\Gamma}^{[1]} \setminus \widehat{\Gamma}_{\infty}^{[1]}} M/\mathbb{Z}u_{(\partial-E, E)},$$

$$L_3 = \prod_{i=1}^s M/\mathbb{Z}u_{(\partial-E_i, E_i)},$$

$$M_1 = \prod_{V \in \widehat{\Gamma}^{[0]} \setminus \widehat{\Gamma}^{[0]}} M/\mathbb{Z}u_{(\partial-E(V), E(V))},$$

$$M_2 = \prod_{E \in \widehat{\Gamma}^{[1]} \setminus (\widehat{\Gamma}_{\infty}^{[1]} \cup S)} M/\mathbb{Z}u_{(\partial-\widehat{E}, \widehat{E})},$$

so that we have maps

$$\Phi : L_1 \rightarrow L_2 \oplus L_3$$

$$\Phi'' : L_1 \oplus M_1 \rightarrow L_2 \oplus M_2 \oplus L_3$$

with Φ'' inducing an isomorphism $M_1 \rightarrow M_2$.

Now consider the diagram

$$\begin{array}{ccc} L_1 & \xrightarrow{\Phi} & L_2 \oplus L_3 \\ \downarrow \Psi_1 & & \downarrow \Psi_2 \\ L_1 \oplus M_1 & \xrightarrow{\Phi''} & L_2 \oplus M_2 \oplus L_3 \end{array}$$

where

$$\Psi_1((m_V)_{V \in \widehat{\Gamma}^{[0]}}) = ((m_V)_{V \in \widehat{\Gamma}^{[0]}}, (m_{\partial-E(V)})_{V \in \widehat{\Gamma}^{[0]} \setminus \widehat{\Gamma}^{[0]}})$$

and Ψ_2 is the natural inclusion. Then one checks easily that this diagram is commutative and hence, by the snake lemma, the cokernels of Φ and Φ'' are isomorphic. But the order of the cokernel is the index of the image, hence the result. \square

Given a torically transverse pre-log curve $f : C \rightarrow X_0$ given by Proposition 4.22, we now wish to count the number of non-isomorphic choices of log morphism $f : C^\dagger \rightarrow X_0^\dagger$ with the given underlying morphism of schemes. Now in general, we have a bit too much freedom to put log marked points on C , so we impose

the condition that f should be strict where $X_0^\dagger \rightarrow \text{Spec } \mathbb{k}^\dagger$ is strict (see Definition 3.21). Note that $X_0^\dagger \rightarrow \text{Spec } \mathbb{k}^\dagger$ is strict away from $\text{Sing}(X_0) \cup \partial X_0$. Indeed, in a neighbourhood of a point in $X_0 \setminus (\text{Sing}(X_0) \cup \partial X_0)$, the log structure on X_0 is just $\mathcal{M}_{X_0} = \mathcal{O}_{X_0}^\times \oplus \mathbb{N}$. So this constraint of strictness requires that all points of C mapping to points not in $\text{Sing}(X_0) \cup \partial X_0$ are smooth points, rather than log marked points, of C . Double points of C map to double points of X_0 , in any event, by the construction of $f : C \rightarrow X_0$, and so by Proposition 4.9, (1), any log marked point of C^\dagger must map into ∂X_0 . This is the only real consequence of imposing strictness.

PROPOSITION 4.23. *Let $f : (C, x_1, \dots, x_s) \rightarrow X_0$ be a torically transverse pre-log curve constructed from the simple tropical curve h using Proposition 4.22. Assume further that each edge $h(E)$, for $E \in \widehat{\Gamma}^{[1]} \setminus \widehat{\Gamma}_\infty^{[1]}$, has affine length divisible by $w(E)$. (Note that this can always be achieved by rescaling M .) Then the number of non-isomorphic maps*

$$f^\dagger : (C^\dagger, x_1, \dots, x_s) \rightarrow X_0^\dagger$$

with underlying scheme morphism f which are strict where $X_0^\dagger \rightarrow \text{Spec } \mathbb{k}^\dagger$ is strict is

$$\left(\prod_{E \in \widehat{\Gamma}^{[1]} \setminus \widehat{\Gamma}_\infty^{[1]}} w(E) \right) \cdot \prod_{i=1}^s w(E_i).$$

PROOF. Let $C^\circ \subseteq C$ be the open subset of C given by

$$C^\circ := C \setminus f^{-1}(\text{Sing}(X_0) \cup \partial X_0).$$

Then the requirement is that $f^\dagger : (C^\circ)^\dagger \rightarrow X_0^\dagger$ be strict, so on C° , the log structure is just the pull-back of the log structure on X_0 . Thus there is no choice for the log structure on C° . We need to understand how many ways there are of extending this log structure to C .

We will split the discussion up as follows. We are interested in classifying extensions of $f^\dagger : (C^\circ)^\dagger \rightarrow X_0^\dagger$ to $f^\dagger : C^\dagger \rightarrow X_0^\dagger$ over $\text{Spec } \mathbb{k}^\dagger$, up to isomorphism. More precisely, suppose we have two such extensions, which we write as $f_i^\dagger : C_i^\dagger \rightarrow X_0^\dagger$, for $i = 1, 2$. Here the underlying scheme of C_i^\dagger is C . We say that these extensions are *isomorphic* if there is a log isomorphism $\kappa^\dagger : C_1^\dagger \rightarrow C_2^\dagger$ such that $f_1^\dagger = f_2^\dagger \circ \kappa^\dagger$. We say that such an isomorphism κ^\dagger is *scheme-theoretically trivial* if the underlying scheme morphism of κ^\dagger is the identity.

We will then prove the result in three steps. First, we will show that there is a unique extension of $f^\dagger : (C^\circ)^\dagger \rightarrow X_0^\dagger$ across points $x \in C$ with $f(x) \in \partial X_0$. Second, we will classify extensions of f^\dagger across double points of C up to scheme-theoretically trivial isomorphism. Third, we will take into account isomorphisms which are not scheme-theoretically trivial.

Step 1. Extension across points $x \in C$ with $f(x) \in \partial X_0$. Let $x \in C$ be a point with $f(x) \in \partial X_0$. By toric transversality, $f(x)$ lies in a one-dimensional torus orbit of X . Let $\sigma \in \Sigma_\emptyset$ be the two-dimensional cone corresponding to this orbit, necessarily of the form $\mathbb{R}_{\geq 0}(v, 1) + \mathbb{R}_{\geq 0}(v', 0)$ for $v \in \mathcal{P}$ and some primitive $v' \in M$ generating a ray of Σ . Note that after applying an element of $\text{GL}_3(\mathbb{Z})$, we can assume $v' = (1, 0)$ and $v = (0, 0)$, so that the open affine subset $X_\sigma \subseteq X$ is $X_\sigma = \text{Spec } \mathbb{k}[t, u, y^{\pm 1}]$, with $t = z^{(0,0,1)}$ yielding the regular function π , $u = z^{(1,0,0)}$

and $y = z^{(0,1,0)}$. With these variables, we have $(\partial X) \cap X_\sigma = V(tu)$. The one-dimensional orbit containing $f(x)$ is given by $t = u = 0$. Let s_t and s_u denote the sections of \mathcal{M}_X over X_σ corresponding to t and u respectively. We also denote by s_t and s_u the restriction of these sections to \mathcal{M}_{X_0} . Hence, locally near $f(x)$, \mathcal{M}_{X_0} consists of sections of the form $s_t^a s_u^b \varphi$, where φ is invertible in a neighbourhood of $f(x)$.

How do we put a log structure on C in a neighbourhood of x so that f^\dagger extends to a log morphism across x ? Suppose there is an extension $f^\dagger : C^\dagger \rightarrow X_0^\dagger$ of $f^\dagger : (C^\circ)^\dagger \rightarrow X_0^\dagger$. Then, in a neighbourhood of x , we must have $\alpha_C f^\#(s_u) = f^*(u)$, and $f^*(u)$ is invertible outside of x in a neighbourhood of x . Thus the image of $f^\#(s_u)$ in $\overline{\mathcal{M}}_C$ has support precisely at x in this neighbourhood. This shows that x must be a log marked point.

Given this, there is a unique way of extending $f^\dagger : (C^\circ)^\dagger \rightarrow X_0^\dagger$ to x . In a neighbourhood of x , we can write $\mathcal{M}_C = \mathcal{M}_{(C,x)} \oplus \mathbb{N}\rho_C$, where $\mathcal{M}_{(C,x)}$ is the divisorial log structure defined by $x \in C$ and ρ_C defines the morphism of C^\dagger to $\text{Spec } \mathbb{k}^\dagger$. We then necessarily have $f^\#(s_t) = \rho_C$, $f^\#(s_u) = f^*(u) \in \mathcal{M}_{(C,x)}$. This gives the unique extension of $f^\dagger : (C_0)^\dagger \rightarrow X_0^\dagger$ over $\text{Spec } \mathbb{k}^\dagger$ across x .

Thus the points of C mapping to ∂X_0 do not cause any problems: there are no choices.

Step 2. Extension across nodes $x \in C$. Suppose $x \in C$ is a node. Let E be the corresponding edge of $\tilde{\Gamma}$, ℓ the affine length of $h(E)$, and $e = \ell/w(E)$. We know from Proposition 4.9, (4), that a desired \mathcal{M}_C would necessarily have $\overline{\mathcal{M}}_{C,\bar{x}} = S_e$. We will show that there are in fact precisely $\mu := w(E)$ choices for f^\dagger in a neighbourhood of x up to scheme-theoretically trivial isomorphism.

Write C locally near x as $C = \text{Spec } \mathbb{k}[z, w]/(zw)$, and write X locally near $f(x)$ as

$$X = \text{Spec } \mathbb{k}[u, v, y^{\pm 1}, t]/(uv - t^\ell),$$

so

$$X_0 = \text{Spec } \mathbb{k}[u, v, y^{\pm 1}]/(uv).$$

Here, in the corresponding notation of the proof of Proposition 4.9, $u = z^{\alpha_1}$, $v = z^{\alpha_2}$, $t = z^\rho$. Note that, by Proposition 4.9, (2), we have in a neighbourhood of x ,

$$f^*(u) = z^\mu \cdot \text{unit}, \quad f^*(v) = w^\mu \cdot \text{unit}.$$

With a local change of coordinates on C near x we can assume these units are 1. Now u, v and t are sections of \mathcal{M}_X and to avoid confusion, we write these sections as s_u, s_v, s_t for sections of \mathcal{M}_X and u, v, t for regular functions. We can then restrict s_u, s_v and s_t to sections of \mathcal{M}_{X_0} , which we also call s_u, s_v and s_t . Note that

$$\alpha_{X_0}(s_u) = u, \quad \alpha_{X_0}(s_v) = v, \quad \alpha_{X_0}(s_t) = 0.$$

We are looking for a commutative diagram

$$(4.9) \quad \begin{array}{ccc} f^{-1}\mathcal{M}_{X_0} & \xrightarrow{f^\#} & \mathcal{M}_C \\ \alpha_{X_0} \downarrow & & \downarrow \alpha_C \\ f^{-1}\mathcal{O}_{X_0} & \xrightarrow{f^*} & \mathcal{O}_C \end{array}$$

First, we construct μ such extensions of f^\dagger in a neighbourhood of x . For any $\zeta \in \mathbb{k}^\times$ with $\zeta^\mu = 1$, consider the chart $S_e \rightarrow \mathcal{O}_C$ defined in a neighbourhood of x , defining a log structure \mathcal{M}_C in that neighbourhood, given by

$$(4.10) \quad ((a, b), c) \mapsto \begin{cases} (\zeta^{-1}z)^a w^b & c = 0, \\ 0 & c \neq 0. \end{cases}$$

This describes \mathcal{M}_C locally as a quotient of $\underline{S}_e \oplus \mathcal{O}_C^\times$, so an element of S_e induces an element of \mathcal{M}_C . For each $((a, b), c) \in S_e$, denote by $s_{((a,b),c)}$ the corresponding section of \mathcal{M}_C . Then define $s_z = s_{((1,0),0)}$, $s_w = s_{((0,1),0)}$, so that $\alpha_C(s_z) = \zeta^{-1}z$, $\alpha_C(s_w) = w$. We define the map $C^\dagger \rightarrow \text{Spec } \mathbb{k}^\dagger$ by the section $s_{((0,0),1)}$; this is log smooth by Example 3.26, (2). We define $f^\# : f^{-1}\mathcal{M}_X \rightarrow \mathcal{M}_C$ by

$$(4.11) \quad f^\#(s_u) = s_z^\mu, \quad f^\#(s_v) = s_w^\mu, \quad f^\#(s_t) = s_{((0,0),1)}.$$

This makes (4.9) commutative and hence gives, for each choice of ζ , a choice of log structure on C in a neighbourhood of x and an extension of f^\dagger across x .

There are now two things we need to show. First, we need to show that none of these μ choices are identified via a scheme-theoretically trivial isomorphism. Second, we need to show that all possible extensions are of the above form.

For the first point, suppose that two choices of μ -th root, ζ_1, ζ_2 , give rise to C_1^\dagger, C_2^\dagger , with maps $f_i^\dagger : C_i^\dagger \rightarrow X_0^\dagger$. Suppose there is a scheme-theoretically trivial isomorphism $\kappa^\dagger : C_1^\dagger \rightarrow C_2^\dagger$, i.e., a commutative diagram

$$\begin{array}{ccc} \mathcal{M}_{C_2} & \xrightarrow{\kappa^\#} & \mathcal{M}_{C_1} \\ \alpha_{C_2} \downarrow & & \downarrow \alpha_{C_1} \\ \mathcal{O}_C^\times & \xrightarrow{=} & \mathcal{O}_C^\times \end{array}$$

Then we must have

$$(4.12) \quad \alpha_{C_2}(s_z) = \zeta_2^{-1}z = \alpha_{C_1}(\kappa^\#(s_z))$$

$$(4.13) \quad \alpha_{C_2}(s_w) = w = \alpha_{C_1}(\kappa^\#(s_w))$$

Of course, since κ must be an isomorphism over $\text{Spec } \mathbb{k}^\dagger$, we must have $\kappa^\#(s_t) = s_t$. Now it is clear $\kappa^\#$ must induce the identity on $\overline{\mathcal{M}}_{C_i}$, so we must have $\kappa^\#(s_z) = \varphi_z \cdot s_z$ and $\kappa^\#(s_w) = \varphi_w \cdot s_w$ for φ_z, φ_w invertible functions in a neighbourhood of x . But (4.12) tells us that $\varphi_z = \zeta_1/\zeta_2$ at points where $z \neq 0$, and (4.13) tells us that $\varphi_w = 1$ at points where $w \neq 0$, but applying $\kappa^\#$ to the relation $s_z s_w = s_t^e$ shows that $\varphi_z \varphi_w = 1$. This is only possible if $\zeta_1 = \zeta_2$ and $\varphi_z = \varphi_w = 1$.

For the second point, suppose we are given a diagram (4.9). Consider the induced map $\bar{f}^\# : f^{-1}\overline{\mathcal{M}}_{X_0} \rightarrow \overline{\mathcal{M}}_C$. On stalks at x , we have no choice but for $\bar{f}^\# : S_\ell \rightarrow S_e$ to be given by

$$\bar{f}^\#((1, 0), 0) = ((\mu, 0), 0),$$

$$\bar{f}^\#((0, 1), 0) = ((0, \mu), 0),$$

$$\bar{f}^\#((0, 0), 1) = ((0, 0), 1).$$

Indeed, this was shown in the proof of Proposition 4.9: see (4.5). Thus we can find sections s'_z, s'_w of \mathcal{M}_C in a neighbourhood of x such that $f^\#(s_u) = (s'_z)^\mu$, $f^\#(s_v) = (s'_w)^\mu$. These sections are well-defined up to a μ -th root of unity. Since

$\alpha_C((s'_z)^\mu) = z^\mu$, $\alpha_C((s'_w)^\mu) = w^\mu$, there is a unique choice of s'_z, s'_w subject to the constraint that $\alpha_C(s'_z) = z$, $\alpha_C(s'_w) = w$. Thus, applying $f^\#$ to the relationship $s_u s_v = s_t^\ell$ gives $(s'_z s'_w)^\mu = s_t^\ell$ in \mathcal{M}_C , or $s'_z s'_w = \zeta s_t^e$ for ζ a μ -th root of unity. Now define a map $S_e \rightarrow \mathcal{M}_C$ by

$$((a, b), c) \mapsto (\zeta^{-1} s'_z)^a (s'_w)^b s_t^c.$$

Composing this map with α_C defines a chart for the log structure \mathcal{M}_C as given by (4.10). Given the notation

$$s_z = s_{((1,0),0)}, \quad s_w = s_{((0,1),0)},$$

we have $s'_z = \zeta s_z$, $s'_w = s_w$, and the map $f^\#$ is exactly as given in (4.11).

Step 3. Non-trivial identifications. So far, it appears that we have given $\prod_{E \in \tilde{\Gamma}^{[1]} \setminus \tilde{\Gamma}^{[1]_\infty}} w(E)$ choices of log morphisms. However, we haven't taken into account isomorphisms between these choices which aren't scheme-theoretically trivial. Let $f_i^\dagger : (C_i^\dagger, x_1, \dots, x_s) \rightarrow X_0^\dagger$ be two of the choices constructed above, and suppose $\kappa^\dagger : C_1^\dagger \rightarrow C_2^\dagger$ yields an isomorphism between these choices.

Of course κ^\dagger induces an isomorphism $\kappa^\# : \kappa^{-1} \mathcal{M}_{C_2} \rightarrow \mathcal{M}_{C_1}$, so we can assume κ^\dagger is strict, completely determined by the underlying automorphism $\kappa : C \rightarrow C$ of schemes.

Now since such an automorphism must satisfy $f \circ \kappa = f$, we are somewhat limited as to what κ can do. An automorphism of C induces an automorphism $\tilde{\kappa}$ of $\tilde{\Gamma}$, by $\tilde{\kappa}(V) = V'$ if $\kappa(C_V) = C_{V'}$. We must have $h \circ \tilde{\kappa} = h$. Since h is injective on the set of vertices of $\tilde{\Gamma}$, as h is simple, it follows that $\tilde{\kappa}$ is the identity, and hence for each irreducible component C' of C , $\kappa(C') = C'$.

Let us consider elements of $\text{Aut}(C, x_1, \dots, x_s)$ which preserve components. These elements form the subgroup which is the connected component of the identity of $\text{Aut}(C, x_1, \dots, x_s)$, which we write as $\text{Aut}_0(C, x_1, \dots, x_s)$. Note that all components of C have at least 3 special points (nodes, marked points, or log marked points) except for those components corresponding to bivalent vertices of $\tilde{\Gamma}$, and these have two special points. The automorphism group of a component with more than 2 special points is trivial, and the automorphism group of one with two special points is \mathbb{G}_m . Hence $\text{Aut}_0(C, x_1, \dots, x_s)$ is $(\mathbb{G}_m)^{\#\tilde{\Gamma}_b^{[0]}}$, where $\tilde{\Gamma}_b^{[0]}$ denotes the set of bivalent vertices of $\tilde{\Gamma}$.

We now analyze the effect of these automorphisms on the log structure. Let $E \in \tilde{\Gamma}^{[1]}$ and let $\mu = w(E)$. Assume first that $E \neq E_i$ for any i . In $\tilde{\Gamma}$, we can assume that E splits up into edges F_1, \dots, F_r with $\partial F_i = \{V_{i-1}, V_i\}$ for $1 \leq i \leq r$ if E bounded, and $2 \leq i \leq r$ if E unbounded. If E is unbounded, we also have $\partial F_1 = \{V_1\}$, so that F_1 is unbounded. Let $p_i \in C$ be the special point corresponding to F_i : if F_i is bounded, then p_i is a double point, while if F_i is unbounded, then $f(p_i) \in \partial X_0$ and p_i is a log marked point.

If F_i is a bounded edge, then note that $(C_{V_{i-1}} \setminus \{p_{i-1}\}) \cup (C_{V_i} \setminus \{p_{i+1}\}) \subseteq C$ is isomorphic to $\text{Spec } \mathbb{k}[z_i, w_i]/(z_i w_i)$, with z_i non-zero on $C_{V_{i-1}}$ and w_i non-zero on C_{V_i} . If F_1 is unbounded, then $C_{V_1} \setminus \{p_2\} = \text{Spec } \mathbb{k}[w_1]$. Since each C_{V_i} is a \mathbb{P}^1 , we also have, when both F_i and F_{i+1} are bounded, $z_{i+1} = w_i^{-1}$. For F_i bounded, $h(F_i)$ is an edge $\omega_i \in \mathcal{P}$ of length ℓ_i and $X_{C(\omega_i)} \subseteq X$ can be written as $\text{Spec } \mathbb{k}[u_i, v_i, y_i^{\pm 1}, t]/(u_i v_i - t^{\ell_i})$. In case F_1 is unbounded, $h(F_1) = \omega_1$, we write $X_{C(\omega_1)} = \text{Spec } \mathbb{k}[t, v_1, y_1^{\pm 1}]$. We have on $X_{C(\omega_i)} \cap X_{C(\omega_{i+1})}$, $u_{i+1} = v_i^{-1}$. By the

description of bivalent lines in Proposition 4.17, note $f^*(u_i) = z_i^\mu$ for $2 \leq i \leq r$ and $f^*(v_i) = w_i^\mu$ for $1 \leq i \leq r-1$.

Let $\eta_i : \text{Aut}_0(C, x_1, \dots, x_s) \rightarrow \mathbb{k}^\times$ be the projection to the component corresponding to C_{V_i} . Then for $\kappa \in \text{Aut}_0(C, x_1, \dots, x_s)$, we can describe the action of κ on the component C_{V_i} , which has coordinates $w_i = z_{i+1}^{-1}$, by

$$\kappa^* w_i = \eta_i(\kappa) \cdot w_i, \quad \kappa^* z_{i+1} = \eta_i(\kappa)^{-1} \cdot z_{i+1}.$$

So pulling back by κ is compatible with $f^* u_i = z_i^\mu$, $f^* v_i = w_i^\mu$ if and only if $\eta_i(\kappa)^\mu = 1$ for all i . Now consider the data ζ_1, \dots, ζ_r (or ζ_2, \dots, ζ_r if F_1 is unbounded) of μ -th roots of unity determining the charts for \mathcal{M}_C in neighbourhoods of the double points p_1, \dots, p_r (or p_2, \dots, p_r). So the chart at p_i is given by

$$((a, b), c) \mapsto \begin{cases} (\zeta_i^{-1} z_i)^a w_i^b & c = 0 \\ 0 & c \neq 0 \end{cases}$$

and applying κ replaces this chart with the chart

$$((a, b), c) \mapsto \begin{cases} (\zeta_i^{-1} \eta_{i-1}(\kappa)^{-1} z_i)^a (\eta_i(\kappa) w_i)^b & c = 0 \\ 0 & c \neq 0 \end{cases}$$

This chart is equivalent (see Definition 3.18) to the chart

$$((a, b), c) \mapsto \begin{cases} ((\eta_{i-1}(\kappa) \eta_i(\kappa)^{-1} \zeta_i)^{-1} z_i)^a w_i^b & c = 0 \\ 0 & c \neq 0 \end{cases}$$

via multiplication of the first chart with the function $S_e \rightarrow \mathcal{O}_C^\times$ given by

$$((a, b), c) \mapsto \eta_i(\kappa)^a / \eta_i(\kappa)^b.$$

Thus, taking

$$\eta_i(\kappa) = \prod_{j=i+1}^r \zeta_j^{-1}$$

for $1 \leq i \leq r-1$, we see that applying κ replaces ζ_i with

$$\left(\prod_{j=i}^r \zeta_j^{-1} \right) \left(\prod_{j=i+1}^r \zeta_j \right) \zeta_i = 1$$

for $2 \leq i \leq r$. If F_1 is bounded, then ζ_1 is replaced by $\prod_{i=1}^r \zeta_i$, and this is now a fixed μ -th root of unity, while if F_1 is unbounded, then we have eliminated all of the ζ_i 's. Thus, up to isomorphism, in the bounded case there are only $w(E)$ choices for the log structures at the double points corresponding to F_1, \dots, F_r , while in the unbounded case there is a unique choice for the log structures at the double points corresponding to F_2, \dots, F_r .

Next, suppose $E = E_i$. Using the same notation as in the previous case, this means there is some $1 \leq i \leq r-1$ with V_i the vertex of E_{p_i} in $\tilde{\Gamma}$. Thus C_{V_i} has three special points, so we only can have automorphisms κ with $\eta_i(\kappa) = 1$. From the argument in the unmarked case, it then becomes clear that there are $w(E_i)^2$ possibilities in the bounded case and $w(E_i)$ possibilities in the unbounded case. This now accounts for the total number of isomorphism classes of maps $f^\dagger : (C^\dagger, x_1, \dots, x_s) \rightarrow X_0^\dagger$. \square

Proof of Theorem 4.14: We just need to show the identity

$$(4.14) \quad \mathfrak{D} \cdot \left(\prod_{E \in \widehat{\Gamma}^{[1]} \setminus \widehat{\Gamma}_\infty^{[1]}} w(E) \right) \cdot \left(\prod_{i=1}^s w(E_i) \right) = \text{Mult}(h).$$

First note that $\text{Mult}(h)$ can be computed using Γ or $\widehat{\Gamma}$: the vertices of Γ which become bivalent after removing a marked unbounded edge do not contribute to the multiplicity of h . We go by induction on the number of vertices of $\widehat{\Gamma}$. First, if $\widehat{\Gamma}$ has one vertex, then $\widehat{\Gamma}$ has no compact edges, three unbounded edges, and two of these, E_1 and E_2 , are marked edges, say with tangent directions $u_1, u_2 \in M$ primitive. Then Φ is the obvious projection

$$\Phi : M \rightarrow M/\mathbb{Z}u_1 \times M/\mathbb{Z}u_2;$$

the index of the projection is just $|u_1 \wedge u_2|$, so

$$\mathfrak{D} \cdot \prod_{i=1}^2 w(E_i) = w_1 w_2 |u_1 \wedge u_2| = \text{Mult}(h).$$

In the general case, say with s marked points, $\widehat{\Gamma}$ has $s + 1 = |\Delta|$ unbounded edges, $s - 1$ vertices and $s - 2$ bounded edges. At least one unbounded edge E is unmarked. Let $V = \partial^- E$. If we remove the edge E and vertex V of $\widehat{\Gamma}$, we obtain two connected components $\widehat{\Gamma}_1, \widehat{\Gamma}_2$ with two new non-compact edges (which previously had V as a vertex). There are two cases.

Case 1. Both $\widehat{\Gamma}_1$ and $\widehat{\Gamma}_2$ have vertices. We then get tropical curves $h_i : \widehat{\Gamma}_i \rightarrow M_{\mathbb{R}}$, $i = 1, 2$, obtained by restricting h to $\widehat{\Gamma}_i$, but extending the new non-compact edges so that h_i is proper. Let E' and E'' be these new non-compact edges of $\widehat{\Gamma}_1$ and $\widehat{\Gamma}_2$ respectively. Note that $\widehat{\Gamma}_1$ and $\widehat{\Gamma}_2$ now pass through some subset of the points P_1, \dots, P_s ; in fact, they must split up so that neither h_1 nor h_2 can be deformed preserving this incidence property; otherwise h itself would move in a one-parameter family. Thus we can inductively apply the result to h_1 and h_2 .

Let u_1, \dots, u_{s-2} be primitive tangent vectors to the images under h of the bounded edges of $\widehat{\Gamma}$ ordered so that $u_1, \dots, u_{\ell-2}$ are associated to bounded edges of $\widehat{\Gamma}_1$, $u_{\ell-1}$ to E' , u_ℓ to E'' , and $u_{\ell+1}, \dots, u_{s-2}$ are associated to bounded edges of $\widehat{\Gamma}_2$. Let v_1, \dots, v_ℓ be the primitive tangent vectors associated to marked edges of $\widehat{\Gamma}_1$, and $v_{\ell+1}, \dots, v_s \in M$ the same for $\widehat{\Gamma}_2$. Let Φ_1, Φ_2 be the maps defined in (4.6) for h_1 and h_2 respectively, and \mathfrak{D}_i the order of the cokernel of Φ_i . Note that we have

$$\Phi_1 : \text{Map}(\widehat{\Gamma}_1^{[0]}, M) \rightarrow \prod_{i=1}^{\ell-2} M/\mathbb{Z}u_i \times \prod_{j=1}^{\ell} M/\mathbb{Z}v_j =: B'$$

and

$$\Phi_2 : \text{Map}(\widehat{\Gamma}_2^{[0]}, M) \rightarrow \prod_{i=\ell+1}^{s-2} M/\mathbb{Z}u_i \times \prod_{j=\ell+1}^s M/\mathbb{Z}v_j =: B''$$

while the original map Φ is given by

$\Phi : \text{Map}(\widehat{\Gamma}_1^{[0]}, M) \times \text{Map}(\widehat{\Gamma}_2^{[0]}, M) \times \text{Map}(\{V\}, M) \rightarrow B' \times B'' \times M/\mathbb{Z}u_{\ell-1} \times M/\mathbb{Z}u_\ell$
 given by

$$(H_1, H_2, H') \mapsto (\Phi_1(H_1), \Phi_2(H_2), H_1(V') - H'(V), H_2(V'') - H'(V)),$$

where $V' \in \widehat{\Gamma}_1^{[0]}$, $V'' \in \widehat{\Gamma}_2^{[0]}$ are vertices adjacent to E' and E'' respectively. Then

$$|\operatorname{coker} \Phi| = |\operatorname{coker} \Phi_1| \cdot |\operatorname{coker} \Phi_2| \cdot |\operatorname{coker}(M \rightarrow M/\mathbb{Z}u_{\ell-1} \times M/\mathbb{Z}u_{\ell})|$$

and so

$$\begin{aligned} & \mathfrak{D} \cdot \left(\prod_{E \in \widehat{\Gamma}^{[1]} \setminus \widehat{\Gamma}_{\infty}^{[1]}} w(E) \right) \cdot \left(\prod_{j=1}^s w(E_j) \right) \\ &= \prod_{i=1}^2 \left(\mathfrak{D}_i \cdot \left(\prod_{E \in \widehat{\Gamma}_i^{[1]} \setminus \widehat{\Gamma}_{i,\infty}^{[1]}} w(E) \right) \left(\prod_{E_j \in \widehat{\Gamma}_i^{[1]}} w(E_j) \right) \right) \cdot w(E')w(E'')|u_{\ell-1} \wedge u_{\ell}| \\ &= \operatorname{Mult}(h_1) \operatorname{Mult}(h_2) \operatorname{Mult}_V(h) \\ &= \operatorname{Mult}(h), \end{aligned}$$

the second-to-last line by the induction hypothesis.

Case 2. One of $\widehat{\Gamma}_1, \widehat{\Gamma}_2$ consists just of an unbounded edge, say $\widehat{\Gamma}_2$, in which case, $\widehat{\Gamma}_2$ must pass through a marked point; otherwise, h moves in a one-parameter family. Then Φ takes the form, using the same notation as in Case 1,

$$\begin{aligned} \Phi : \operatorname{Map}(\widehat{\Gamma}_1^{[0]}, M) \times \operatorname{Map}(\{V\}, M) &\rightarrow B' \times M/\mathbb{Z}u_{s-2} \times M/\mathbb{Z}v_s \\ (H_1, H') &\mapsto (\Phi_1(H_1), H_1(V') - H'(V), H'(V)) \end{aligned}$$

and one sees that

$$|\operatorname{coker} \Phi| = |\operatorname{coker} \Phi_1| \cdot |u_{s-2} \wedge v_s|.$$

From this one obtains similarly the desired result. \square

4.4. Classical world \rightarrow log world

The main point of this section, intuitively, is as follows. We are given as usual a degeneration $\pi : X \rightarrow \mathbb{A}_{\mathbb{k}}^1$ of the toric variety X_{Σ} . Suppose we are given a family of curves in the general fibers. We would like to know that in the limit this family gives rise to a torically transverse log curve on X_0 . This may not be true for a general choice of degeneration or general family of curves. However, we shall see that there is always a good choice of toric degeneration which comes close to achieving this goal. With an additional restriction on the family of curves, one can then achieve the goal of obtaining a torically transverse log curve on X_0 .

Let us start to make this more precise, sketching what we shall do. We are interested in the following situation. We are given $\pi : X \rightarrow \mathbb{A}_{\mathbb{k}}^1$ defined as usual by some polyhedral decomposition \mathscr{P} of $M_{\mathbb{R}}$ — in general, we shall assume only that \mathscr{P} satisfies conditions (1) and (2) in Definition 4.5, as we will not choose points $P_1, \dots, P_s \in M_{\mathbb{R}}$ in this section. Suppose furthermore we are given a discrete valuation ring R with residue field \mathbb{k} and quotient field L , and suppose we are given a dominant map $\operatorname{Spec} R \rightarrow \mathbb{A}_{\mathbb{k}}^1$ mapping the closed point p of $\operatorname{Spec} R$ to $0 \in \mathbb{A}_{\mathbb{k}}^1$. We now want to think of a family of curves in the general fibres of $X \rightarrow \mathbb{A}_{\mathbb{k}}^1$ as a

commutative diagram

$$(4.15) \quad \begin{array}{ccc} (C^*, x_1^*, \dots, x_k^*) & \xrightarrow{f^*} & X \setminus X_0 \\ \downarrow & & \downarrow \pi \\ \text{Spec } L & \xrightarrow[\psi]{} & \mathbb{A}_{\mathbb{k}}^1 \setminus \{0\} \end{array}$$

where f^* is a torically transverse stable map.

We would like to fill in the family $C^* \rightarrow \text{Spec } L$ in some nice way. In general, we can't do this, but according to the stable reduction theorem for stable maps (Proposition 2.4), we can do so after a base-change. Specifically, there is some $d \geq 1$ with the following property. If $R' = R[t]/(t^d - u)$ for u a uniformizing parameter of R , then R' is a discrete valuation ring with field of fractions $L' = L[t]/(t^d - u)$, a degree d extension of L . Then, making the base-change $(C')^* = C^* \times_L L'$, we obtain a diagram

$$\begin{array}{ccc} ((C')^*, x_1^*, \dots, x_k^*) & \xrightarrow{f^*} & X \setminus X_0 \\ \downarrow & & \downarrow \pi \\ \text{Spec } L' & \longrightarrow & \mathbb{A}_{\mathbb{k}}^1 \setminus \{0\} \end{array}$$

which can then be extended to a diagram

$$\begin{array}{ccc} (C', x_1, \dots, x_k) & \xrightarrow{f} & X \\ \downarrow & & \downarrow \pi \\ \text{Spec } R' & \longrightarrow & \mathbb{A}_{\mathbb{k}}^1 \end{array}$$

where f is now a stable map. The trouble is that even though f^* is torically transverse, there is no reason to expect that $f_0 : C'_p \rightarrow X_0$ is torically transverse: f_0 could contract some components to one-dimensional strata of X_0 , or map some points of C'_p to zero-dimensional strata. To guarantee that this doesn't happen, we need to blow-up X . We will do this via toric blow-ups, i.e., by subdividing the fan $\Sigma_{\mathcal{P}}$ defining the toric variety X . Because f is already well-behaved outside of X_0 , we will be able to do this so that we don't change $X \setminus X_0$.

In this way, it will be possible to blow-up X to a toric variety \tilde{X} , chosen before we apply stable reduction, to get a diagram

$$(4.16) \quad \begin{array}{ccc} (C', x_1, \dots, x_k) & \xrightarrow{f} & \tilde{X} \\ \downarrow & & \downarrow \pi \\ \text{Spec } R' & \longrightarrow & \mathbb{A}_{\mathbb{k}}^1 \end{array}$$

with f stable, and $f_0 : C'_p \rightarrow \tilde{X}_0$ now torically transverse: the image of f_0 is disjoint from the zero-dimensional strata of \tilde{X}_0 and no irreducible component of C'_p maps into a one-dimensional stratum of \tilde{X}_0 . Achieving this blow-up \tilde{X} will occupy most of the effort in this section.

In fact, we would like to do this so that \tilde{X} comes from a refinement of \mathcal{P} . This is always the case: to obtain \tilde{X} , we refined $\Sigma_{\mathcal{P}}$ to get a fan $\tilde{\Sigma}$ in $\tilde{M}_{\mathbb{R}}$. Since we did

this so that $\tilde{X} \setminus \tilde{X}_0 \cong X \setminus X_0$, in fact we only added cones generated by elements $(m, r) \in \tilde{M}$ with $r > 0$. Thus, by taking

$$\tilde{\mathcal{P}} = \{\sigma \cap (M_{\mathbb{R}} \oplus \{1\}) \mid \sigma \in \tilde{\Sigma}\},$$

we obtain a polyhedral decomposition $\tilde{\mathcal{P}}$ of $M_{\mathbb{R}}$ refining \mathcal{P} . Specifying $\tilde{\mathcal{P}}$ is equivalent to specifying $\tilde{\Sigma}$. However, $\tilde{\mathcal{P}}$ might have non-integral vertices, as some of the rays we added to obtain $\tilde{\Sigma}$ might have primitive generators of the form (m, r) with $r > 1$. We would like to work only with refinements of \mathcal{P} with integral vertices. To rectify this, we can choose e a positive integer such that for any vertex v of $\tilde{\mathcal{P}}$, $ev \in M$. Then we can make a base-change of diagram (4.16) via $\mathbb{A}_{\mathbb{k}}^1 \rightarrow \mathbb{A}_{\mathbb{k}}^1$ given by $t \mapsto t^e$, replacing \tilde{X} with the normalization of $\tilde{X} \times_{\mathbb{A}^1} \mathbb{A}^1$. It is easy to see that this is the toric variety given by the fan $\tilde{\Sigma}$ with the lattice $M \oplus e\mathbb{Z} \subseteq \tilde{M}_{\mathbb{R}}$ instead of \tilde{M} . This procedure preserves the polyhedral decompositions \mathcal{P} and $\tilde{\mathcal{P}}$, but replaces the lattice M with the lattice $\frac{1}{e}M$, making all vertices of $\tilde{\mathcal{P}}$ integral. Once we make this base-change for the map $\tilde{X} \rightarrow \mathbb{A}_{\mathbb{k}}^1$, we replace $\text{Spec } R'$ with an irreducible component of $\text{Spec } R' \times_{\mathbb{A}^1} \mathbb{A}^1$ and C' with the corresponding irreducible component of $C' \times_{\mathbb{A}^1} \mathbb{A}^1$. Thus we can assume we have a diagram (4.16) in which \tilde{X} is specified by a polyhedral decomposition $\tilde{\mathcal{P}}$ of $M_{\mathbb{R}}$ with integral vertices. One can then use results of earlier sections to study the map $f_0 : C'_p \rightarrow \tilde{X}_0$, and understand when this yields a torically transverse stable log curve.

This is the outline of what we are going to do. We can now make a precise statement.

THEOREM 4.24. *Let $\pi : X \rightarrow \mathbb{A}_{\mathbb{k}}^1$ be as usual defined by some polyhedral decomposition \mathcal{P} satisfying conditions (1) and (2) in Definition 4.5. Let R be a discrete valuation ring with residue field \mathbb{k} and quotient field L , and suppose we are given a dominant map $\psi : \text{Spec } R \rightarrow \mathbb{A}_{\mathbb{k}}^1$ mapping the closed point p of $\text{Spec } R$ to $0 \in \mathbb{A}_{\mathbb{k}}^1$. Suppose furthermore we have a commutative diagram (4.15) where f^* is a torically transverse stable map. Then:*

- (1) *Possibly after making a base-change $\mathbb{A}_{\mathbb{k}}^1 \rightarrow \mathbb{A}_{\mathbb{k}}^1$ given by $t \mapsto t^e$ and replacing R by $R[t]/(t^d - u)$ for u a uniformizing parameter of R , there is a refinement $\tilde{\mathcal{P}}$ of \mathcal{P} with integral vertices defining a toric blow-up \tilde{X} of X such that the diagram (4.15) extends to a diagram*

$$(4.17) \quad \begin{array}{ccc} (C, x_1, \dots, x_k) & \xrightarrow{f} & \tilde{X} \\ \downarrow & & \downarrow \pi \\ \text{Spec } R & \xrightarrow{\psi} & \mathbb{A}_{\mathbb{k}}^1 \end{array}$$

such that the restriction f_0 of f to the fibre C_p over the closed point of $\text{Spec } R$ is a torically transverse pre-log curve and $\psi^(t)$ is a uniformizing parameter for R .*

- (2) *Given (1), we obtain an induced diagram of log schemes*

$$(4.18) \quad \begin{array}{ccc} (C^\dagger, x_1, \dots, x_k) & \xrightarrow{f^\dagger} & \tilde{X}^\dagger \\ \downarrow & & \downarrow \pi \\ (\text{Spec } R)^\dagger & \longrightarrow & (\mathbb{A}_{\mathbb{k}}^1)^\dagger \end{array}$$

Here the log structures in (4.18) are the divisorial ones induced by

$$f^{-1}(\partial\tilde{X}) \subseteq C, \partial\tilde{X} \subseteq \tilde{X}, \{p\} \subseteq \text{Spec } R, \text{ and } \{0\} \subseteq \mathbb{A}_{\mathbb{k}}^1.$$

If in addition,

- C_p is genus zero;
- the tropical curve associated to the pre-log curve f_0 is simple;
- $f^{-1}(\overline{\partial(\tilde{X} \setminus \tilde{X}_0)})$ is a disjoint union of sections of $C \rightarrow \text{Spec } R$;

then the induced log map $C_p^\dagger \rightarrow \text{Spec } \mathbb{k}^\dagger$ is log smooth, and in particular f_0^\dagger yields a torically transverse log curve.

Before proving this theorem, we need a few lemmas.

LEMMA 4.25. *Let X be a toric variety and $W \subseteq X$ a proper closed subset with no irreducible component contained in ∂X . Then there exists a toric blow-up $\phi: \tilde{X} \rightarrow X$ such that the proper transform \tilde{W} of W in \tilde{X} does not contain any zero-dimensional stratum.*

PROOF. Any toric blow-up is given by a refinement of the fan Σ defining X . Note that if we obtain the desired result using some refinement of Σ , then a further refinement of Σ will still do the trick. As a consequence, we only need to solve the problem for a single cone σ , since if we have a refinement Σ_σ of each cone σ in Σ which does the trick on the corresponding open affine subset, we can choose a refinement of Σ which restricts to refinements of Σ_σ for each $\sigma \in \Sigma$.

So we can assume, if X is n -dimensional, that σ is an n -dimensional cone and $X = X_\sigma$ (if $\dim \sigma < n$, then X_σ does not have a zero-dimensional stratum). Furthermore, we can replace W with a hypersurface containing W which does not contain a toric stratum of X_σ .

Suppose W is defined by an equation

$$f = \sum_{p \in \sigma^\vee \cap N} a_p z^p \in \mathbb{k}[\sigma^\vee \cap N],$$

and let

$$\Delta_f = \text{Conv} \left(\bigcup_{p \text{ such that } a_p \neq 0} (p + (\sigma^\vee \cap N)) \right) \subseteq N_{\mathbb{R}}.$$

This is the Newton polyhedron of f . Let $\tilde{\Sigma}$ be the normal fan to Δ_f . It clearly is a refinement of Σ , and hence defines a blow-up $\phi: \tilde{X} \rightarrow X$ of X . We need to show that ϕ has the desired properties.

Consider an n -dimensional cone $\tau \in \tilde{\Sigma}$; this will be the normal cone $N_{\Delta_f}(v)$ for some vertex v of Δ_f . So in particular, τ^\vee is the tangent cone $T_v(\Delta_f)$, so $\Delta_f \subseteq v + \tau^\vee$ and $a_v \neq 0$. Now $\phi: X_\tau \rightarrow X_\sigma$ is induced by the inclusion of cones $\sigma^\vee \rightarrow \tau^\vee$, and every monomial in $f \circ \phi$ is divisible by z^v . Thus the proper transform \tilde{W} of W in X_τ is defined by the equation

$$\sum_{p \in \sigma^\vee \cap N} a_p z^{p-v} = 0,$$

and this polynomial has a non-zero constant term. Since all the non-constant monomials vanish at the zero-dimensional stratum of X_τ , \tilde{W} does not contain the zero-dimensional stratum, as desired. \square

We need a slightly stronger version of this:

LEMMA 4.26. *Let X be a toric variety and $W \subseteq X$ a proper closed subset with no irreducible component contained in ∂X . Suppose the codimension of W in X is $> c$. Then there is a toric blow-up $\phi : \tilde{X} \rightarrow X$ such that the proper transform \tilde{W} of W in \tilde{X} is disjoint from any toric stratum of dimension $\leq c$.*

PROOF. Via induction on c , we can assume W is disjoint from toric strata of dimension less than c by performing some toric blow-up. Thus we just need to show that we can find \tilde{X} so that \tilde{W} is disjoint from any c -dimensional torus orbit. Now a c -dimensional torus orbit in X is contained in an affine open subset X_τ for some $\tau \in \Sigma$ with $\text{codim } \tau = c$. We can argue as in the proof of Lemma 4.25 that we just have to deal with the case where $X = X_\tau$.

So assume Σ consists of the faces of τ , and let $\bar{\Sigma}$ be the fan Σ , but viewed as a fan in the smaller subspace $\mathbb{R}\tau$. Any choice of a linear projection $M \rightarrow (\mathbb{R}\tau) \cap M$ defines a map of fans from Σ to $\bar{\Sigma}$, giving a map $\psi : X = X_\Sigma \rightarrow X_{\bar{\Sigma}} =: Y$. Note that $\dim Y = \dim X - c$, and that the unique c -dimensional stratum of X is mapped to the unique zero-dimensional stratum of Y . Also, $\dim \psi(W) \leq \dim W < \dim X - c = \dim Y$, so $\psi(W) \subseteq Y$ is a proper subset with no irreducible component contained in ∂Y . Thus, by Lemma 4.25, there is a subdivision of $\bar{\Sigma}$ giving a toric blow-up $\tilde{Y} \rightarrow Y$ such that the proper transform of $\psi(W)$ in \tilde{Y} is disjoint from the zero-dimensional toric stratum of \tilde{Y} . But the subdivision of $\bar{\Sigma}$ of course also yields a subdivision of Σ , hence a blow-up $\tilde{X} \rightarrow X$. Since the proper transform \tilde{W} of W maps to the proper transform of $\psi(W)$ in \tilde{Y} , \tilde{W} is now disjoint from c -dimensional toric strata in \tilde{X} . \square

Proof of Theorem 4.24. Step 1. We will show that, after suitable base-changes, there is a diagram (4.17) such that $\psi^*(t)$ is a uniformizing parameter of R and for each irreducible component \tilde{D}_v of \tilde{X}_0 , the map $f^{-1}(\tilde{D}_v) \rightarrow \tilde{D}_v$ is a torically transverse stable map.

Let $\tilde{C}^* \rightarrow C^*$ be the normalization of C^* . Let y_1^*, \dots, y_p^* be the points of the conductor locus, i.e., the points in \tilde{C}^* over the double points of C^* . These points need not be defined over the field L , but are defined over a finite extension L' , necessarily the field of fractions of $R[t]/(t^d - u)$ for some d . By replacing L by L' , we can assume that these points are defined over L , so that C^* is obtained by gluing together pairs of points on \tilde{C}^* defined over L . On a given component D^* of \tilde{C}^* , we have some set $\{x_{i_k}^*\}$ of points mapping to marked points of C^* , along with some additional points $\{y_{j_k}^*\}$. It is then enough to show that there exists an \tilde{X} which works for each curve $(D^*, x_{i_1}^*, \dots, x_{i_k}^*, y_{j_1}^*, \dots, y_{j_\ell}^*)$ over $\text{Spec } L$ appearing in \tilde{C}^* , so that we obtain torically transverse stable maps $(D, x_1, \dots, x_{i_k}, y_{j_1}, \dots, y_{j_\ell}) \rightarrow \tilde{X}$. We can then glue these various curves together along pairs of sections labelled by the y 's in the same way C^* is obtained by gluing together pairs of points labelled by the y^* 's. In this way, we reduce to the case that C^* is geometrically irreducible. Indeed, in what follows, we will construct a refinement $\tilde{\mathcal{P}}$ for each irreducible component C^* ; we can then take a common refinement, which by construction will not destroy the desired properties needed for each irreducible component. We can also, by making base-changes $\mathbb{A}_{\mathbb{k}}^1 \rightarrow \mathbb{A}_{\mathbb{k}}^1$, $t \mapsto t^e$, always assume that the refinements only have integral vertices. We shall automatically do this in what follows without comment.

Now let W be the closure of the image of $C^* \rightarrow X$. Then W is a closed subset of dimension at most two, and since the composition $C^* \xrightarrow{f^*} X \xrightarrow{\pi} \mathbb{A}_{\mathbb{k}}^1$ is dominant (as it is also the composition $C^* \rightarrow \text{Spec } L \rightarrow \mathbb{A}_{\mathbb{k}}^1$), the dimension of W must be at least one. We now have two cases.

Case 1. $\dim W = 1$. Note that this can only happen if f^* is constant on C^* , so that W is just the closure of its generic point $f(C^*)$. Now, since $\text{codim } W = 2$, we can use Lemma 4.26 to refine $\tilde{\Sigma}$ to get \tilde{X} and assume that W is disjoint from any 1-dimensional toric stratum of \tilde{X} . Furthermore, \tilde{X} can be chosen so that $\tilde{X} \setminus \tilde{X}_0 \cong X \setminus X_0$, since W is disjoint from $\partial(X \setminus X_0)$: otherwise $f^* : C^* \rightarrow X \setminus X_0$ would not be torically transverse. Thus we can assume \tilde{X} comes from a refinement $\tilde{\mathcal{P}}$ of \mathcal{P} .

Now observe that since π is proper, so is the projection $W \rightarrow \mathbb{A}_{\mathbb{k}}^1$, and this projection is also dominant. The map $f^* : C^* \rightarrow W$ factors through $\text{Spec } L$ as f is constant on C^* . We now appeal to stable reduction for stable curves (a special case of Proposition 2.4, applied in the case where the target space X is a point), obtaining, again possibly after base change, a family $(C, x_1, \dots, x_s) \rightarrow \text{Spec } R$ of stable curves. By the valuative criterion for properness applied to $W \rightarrow \mathbb{A}_{\mathbb{k}}^1$, we obtain a commutative diagram

$$\begin{array}{ccc} \text{Spec } R & \longrightarrow & W \\ \downarrow & & \downarrow \\ \text{Spec } L & \longrightarrow & \mathbb{A}_{\mathbb{k}}^1 \end{array}$$

and the composition $C \rightarrow \text{Spec } R \rightarrow W \hookrightarrow \tilde{X}$ now gives the desired map $f : (C, x_1, \dots, x_s) \rightarrow \tilde{X}$. Since W is disjoint from any one-dimensional toric stratum, $f^{-1}(\tilde{D}_v) \rightarrow \tilde{D}_v$ is a torically transverse stable map for any irreducible component \tilde{D}_v of \tilde{X}_0 intersecting W , and is empty for the other irreducible components of \tilde{X}_0 .

Case 2. $\dim W = 2$. In this case, by Lemma 4.26, we can subdivide Σ to obtain \tilde{X} so that the proper transform of W is disjoint from zero-dimensional strata of \tilde{X} ; again, this can be done so that $\tilde{X} \setminus \tilde{X}_0 = X \setminus X_0$. So we assume this comes from a refinement $\tilde{\mathcal{P}}$ of \mathcal{P} .

Let $\tau \in \tilde{\mathcal{P}}$ be an edge, $D_\tau \subseteq \tilde{X}$ the corresponding one-dimensional stratum, and $X_\tau := X_{C(\tau)}$ the corresponding open affine subset of \tilde{X} . Then $X_\tau \cong \mathbb{G}_m \times V_e$, where

$$V_e = \text{Spec } \mathbb{k}[x, y, t]/(xy - t^e)$$

for e the affine length of τ . Let $C_\tau^* = f^{-1}(X_\tau) \subseteq C^*$, and compose $C_\tau^* \xrightarrow{f} X_\tau$ with the projection $X_\tau \rightarrow V_e$, getting a map $h : C_\tau^* \rightarrow V_e$. There is an open subset $U \subseteq C_\tau^*$ on which h is étale, non-empty if h is dominant (see [83], Proposition 3.8). Let $Z_\tau \subseteq V_e$ be the smallest closed subset containing the image of $C_\tau^* \setminus U$ and the images of the marked points of C_τ^* . Since the image of C_τ^* is disjoint from ∂V_e (given by $t = 0$), no irreducible component of Z_τ is contained in ∂V_e . Thus we can apply Lemma 4.26, and blow-up V_e so that the proper transform of Z_τ is disjoint from the zero-dimensional stratum of V_e . This corresponds to a subdivision of the edge τ , so we can choose a further subdivision of $\tilde{\mathcal{P}}$ which induces this subdivision on τ . After doing so, we can assume that Z_τ is disjoint from the zero-dimensional stratum of V_e . We do this for every $\tau \in \tilde{\mathcal{P}}$ with $\dim \tau = 1$ such that $W \cap D_\tau \neq \emptyset$. This gives \tilde{X} in this case.

Why does this work? First, by the stable reduction theorem for stable maps, Proposition 2.4, we obtain a marked stable map $f : (C, x_1, \dots, x_k) \rightarrow \tilde{X}$ over $\text{Spec } R$ extending $f : (C^*, x_1^*, \dots, x_k^*) \rightarrow \tilde{X}$, possibly after another base-change. Now the image of this extended f is contained in the proper transform \tilde{W} of W in \tilde{X} , and \tilde{W} , by construction, avoids zero-dimensional strata of \tilde{X} . In particular, no irreducible component of $f(C_0)$ can be contained in a one-dimensional toric stratum of \tilde{X} , as it would also then contain a zero-dimensional stratum of \tilde{X} .

This is not quite enough: we still need to show that no irreducible component of C_0 is contracted by f to a point in a one-dimensional stratum of \tilde{X} . So suppose there is some such irreducible component being contracted into a one-dimensional stratum of \tilde{X} indexed by $\tau \in \tilde{\mathcal{P}}$. This defines as before X_τ , $C_\tau = f^{-1}(X_\tau) \subseteq C$, and $X_\tau \cong \mathbb{G}_m \times V_e$ for some e . We have the set $Z_\tau \subseteq V_e$ as above, disjoint from the singular point of V_e . Note that if the projection $h : C_\tau^* \rightarrow V_e$ were not dominant, then Z_τ contains the image of C_τ^* , hence the image of $\tilde{W} \cap X_\tau$ under the projection to V_e . In particular, \tilde{W} is then disjoint from the one-dimensional stratum of X_τ .

Thus we can assume that h is dominant. Let $\tilde{Z}_\tau \subseteq X_\tau$ be the inverse image of Z_τ under the projection $X_\tau \rightarrow V_e$. Set $Z'_\tau = f^{-1}(\tilde{Z}_\tau)$. The map $f : C_\tau \setminus Z'_\tau \rightarrow \tilde{X}$ factors as

$$C_\tau \setminus Z'_\tau \xrightarrow{f'} \text{Spec } R \times_{\mathbb{A}_k^1} (X_\tau \setminus \tilde{Z}_\tau) \rightarrow \tilde{X},$$

where f' is proper. We know that f' is finite except over some finite (possibly empty) subset

$$T \subseteq \text{Spec } R \times_{\mathbb{A}_k^1} (X_\tau \setminus \tilde{Z}_\tau)$$

where the fibres of f' aren't finite, since f' is proper and quasi-finite except over T (because h is dominant). Now consider the Stein factorization of f' as

$$C_\tau \setminus Z'_\tau \xrightarrow{f''} Y_\tau \xrightarrow{g} \text{Spec } R \times_{\mathbb{A}_k^1} (X_\tau \setminus \tilde{Z}_\tau) \rightarrow \tilde{X}$$

where f'' is an isomorphism away from $(f')^{-1}(T)$. Thus $Y_\tau \rightarrow \tilde{X}$ glues to $f|_{C \setminus f^{-1}(T)}$ to give a map $g' : C' \rightarrow \tilde{X}$. This map is marked by compositions $x_i : \text{Spec } R \rightarrow C' \rightarrow C'$, since by construction $C_\tau \setminus Z'_\tau$ does not contain any of the marked points. I now claim that this new map $g' : (C', x_1, \dots, x_k) \rightarrow \tilde{X}$ is a stable map, contradicting stability of $f : (C, x_1, \dots, x_k) \rightarrow \tilde{X}$ unless T was empty.

To do so, we need to show that C'_0 has at worst double points. Consider the composition $h : Y_\tau \rightarrow \text{Spec } R \times_{\mathbb{A}_1^1} (V_e \setminus Z_\tau)$. By construction, this map is finite and is étale outside of the locus given by $t = 0$. On the other hand, if it is branched over a component of $t = 0$, then the fibre over p of $Y_\tau \rightarrow \text{Spec } R$ has non-reduced components, which is impossible as $f : C \rightarrow \tilde{X}$ is a stable map. Thus h is étale except over a finite subset of $\text{Spec } R \times_{\mathbb{A}_1^1} (V_e \setminus Z_\tau)$. In fact, the only possible point where this map can fail to be étale is at the singular point of $\text{Spec } R \times_{\mathbb{A}_1^1} (V_e \setminus Z_\tau)$, by purity of the branch locus (see e.g., [1], X, 3.1). However, étale locally the only finite maps étale over the complement of this singular point look like the canonical covers $V_{e'} \rightarrow V_e$ where $e'|e$. Thus one sees that the fibre over p of Y_τ can only have at worst double points.

This completes the argument for Case 2.

Now consider the map $\psi : \text{Spec } R \rightarrow \mathbb{A}_k^1$ in the diagram (4.17) we have now constructed. Suppose $\psi^*(t) \in \mathfrak{m}_R^d \setminus \mathfrak{m}_R^{d-1}$ for some $d \geq 1$, so $\psi^*(t) = u^d$ for a uniformizing parameter u . If $d = 1$, we are done. Otherwise, make a degree d base extension $\mathbb{A}_k^1 \rightarrow \mathbb{A}_k^1$: then $R \times_{\mathbb{A}_1^1} \mathbb{A}^1 = R[t]/(t^d - u^d)$ consists of d irreducible

components, each isomorphic to $\text{Spec } R$. We may then use any of these irreducible components to get a diagram (4.17) such that $\psi^*(t) = u$.

Step 2. Having obtained the diagram (4.17), we now wish to show that the restriction f_0 of f to C_0 is a torically transverse pre-log curve. In view of Step 1, we just need to show that the conclusions (1) and (2) of Proposition 4.9 hold.

Let $p \in C_0$ be a closed point mapping to the singular locus of \tilde{X}_0 . Consider the induced homomorphism of complete local $\mathbb{k}[[t]]$ -algebras

$$f_0^* : \widehat{\mathcal{O}}_{\tilde{X}, f_0(p)} \rightarrow \widehat{\mathcal{O}}_{C,p}.$$

Now p is either a smooth point of C_0 or a double point of C_0 . In the former case, the map $C \rightarrow \text{Spec } R$ is itself smooth at p , and as a consequence, the log structure on C induced by $f^{-1}(\partial\tilde{X}) \subseteq C$ is log smooth at p . Restricting this log structure to C_0 gives a curve $C_0^\dagger \rightarrow \text{Spec } \mathbb{k}^\dagger$ which is log smooth at p , along with a log morphism $C_0^\dagger \rightarrow \tilde{X}_0^\dagger$ over $\text{Spec } \mathbb{k}^\dagger$. But since $f(p)$ is in the singular locus of \tilde{X}_0 , this contradicts Proposition 4.9 (the argument there was purely local). Thus we can assume p is a double point of C_0 . Note that

$$\widehat{\mathcal{O}}_{C_0,p} = \mathbb{k}[[x, y]]/(xy),$$

from which it follows that

$$\widehat{\mathcal{O}}_{C,p} = \mathbb{k}[[x, y, t]]/(xy - ft^e)$$

for some $e > 0$ and some $f \in \mathbb{k}[[x, y, t]]$. It is easy to check that this $\mathbb{k}[[t]]$ -algebra is then isomorphic to

$$\mathbb{k}[[x, y, t]]/(xy - \lambda t^e)$$

for λ either 0 or 1 and $e > 0$.

If $\lambda = 0$, then C is locally (in the étale topology) reducible in a neighbourhood of p . By restricting f to one of these components, we obtain a contradiction as above in the case that p was a smooth point of C .

If $\lambda = 1$, then again by taking the log structure on C induced by $f^{-1}(\partial\tilde{X}) \subseteq C$, $C^\dagger \rightarrow \text{Spec } R^\dagger$ is log smooth in a neighbourhood of p , so we can apply Proposition 4.9 to conclude that f_0 is a torically transverse pre-log curve. This completes the proof of (1) of Theorem 4.24.

Step 3. All that remains to be shown is that in the situation of (2) of Theorem 4.24, $C^\dagger \rightarrow \text{Spec } R^\dagger$ is in fact log smooth. In fact, the argument of Step 2 showed that this was the case in neighbourhoods of all points $p \in C_0$ mapping into $\text{Sing}(X_0)$. This is also the case for any smooth point $p \in C_0$ not mapping into $\text{Sing}(X_0) \cup \partial X_0$, and the condition that $f^{-1}(\partial(\tilde{X} \setminus \tilde{X}_0))$ is a disjoint union of sections of $C \rightarrow \text{Spec } R$ yields log smoothness at smooth points p mapping into $\partial\tilde{X}_0$.

Finally, the assumption that the tropical curve $h : \Gamma \rightarrow M_{\mathbb{R}}$ associated to the pre-log curve f_0 is simple implies that if $v \in \mathcal{P}$ is a vertex with $h^{-1}(v)$ non-empty, then either $h^{-1}(v)$ is a single trivalent vertex of Γ or $h^{-1}(v)$ consists of one or more points in the interior of edges of Γ . In the former case, $f^{-1}(D_v)$, being rational since C_0 is rational, consists of a single line, while in the latter case, $f^{-1}(D_v)$ consists of a disjoint union of bivalent lines. In particular, C_0 has no double points not mapping into $\text{Sing}(X_0)$. Thus $C_0^\dagger \rightarrow \text{Spec } \mathbb{k}^\dagger$ is log smooth, as desired. \square

4.5. Log world → classical world

We return to our basic situation, with data $P_1, \dots, P_s \in M$, a good lattice polyhedral decomposition \mathcal{P} , and general points $Q_1, \dots, Q_s \in \mathbb{G}(\widehat{M})$ with $s = |\Delta| - 1$. This gives sections $\sigma_1, \dots, \sigma_s : \mathbb{A}_{\mathbb{k}}^1 \rightarrow X$ and points $q_1, \dots, q_s \in X_0$ with $q_i = \sigma_i(0)$.

The main theorem of this section is

THEOREM 4.27. *Let $f_0 : (C_0^\dagger, x_1, \dots, x_s) \rightarrow X_0^\dagger$ be a torically transverse log curve of genus zero with $f(x_i) = q_i$, with associated tropical curve $h : \Gamma \rightarrow M_{\mathbb{R}}$. Suppose that h is simple. Then there exists a unique marked rational curve $(C_\infty, x_1^\infty, \dots, x_s^\infty)$ over $\text{Spec } \mathbb{k}[[t]]$ with a commutative diagram*

$$\begin{array}{ccc} (C_\infty, x_1^\infty, \dots, x_s^\infty) & \xrightarrow{f_\infty} & X \\ \downarrow & & \downarrow \\ \text{Spec } \mathbb{k}[[t]] & \xrightarrow{\psi_\infty} & \mathbb{A}^1 \end{array}$$

such that

- (1) ψ_∞ is induced by the natural inclusion $\mathbb{k}[[t]] \hookrightarrow \mathbb{k}[[t]]$.
- (2) $\sigma_i \circ \psi_\infty = f_\infty \circ x_i^\infty$.
- (3) If C_∞ is given the log structure induced by $f_\infty^{-1}(\partial X) \subseteq C_\infty$ and $\text{Spec } \mathbb{k}[[t]]$ the log structure pulled back from \mathbb{A}^1 , then the induced log morphism on

$$(C_0, x_1, \dots, x_s) \rightarrow X_0$$

over $\text{Spec } \mathbb{k}^\dagger$ coincides with the given f_0 .

PROOF. *Step 1. Finite order deformation theory.*

We will apply the deformation theory of §3.4, in particular Theorem 3.43. We first note that the hypotheses of Theorem 3.43 apply to the data

$$[f_0 : C_0/O_0 \rightarrow X, \mathbf{x}^0],$$

where $f_0 : C_0^\dagger \rightarrow X^\dagger$ is now the composition of the original f_0 and the inclusion $X_0 \hookrightarrow X$, and $x_i^0 = x_i$. Of course C_0 is rational by assumption, and $f_{0*} : \Theta_{C_0^\dagger/\mathbb{k}^\dagger} \rightarrow f_0^* \Theta_{X^\dagger/(\mathbb{A}^1)^\dagger}$ is injective. Indeed, f_{0*} is injective at each generic point of C_0 since φ_0 is étale onto its image in a neighbourhood of each generic point. Since $\Theta_{C_0^\dagger/\mathbb{k}^\dagger}$ is locally free, as f_0 is log smooth, it cannot contain a torsion subsheaf. Hence $\ker f_{0*} = 0$.

We will now show that the map Ξ of Theorem 3.43 is in fact an isomorphism. This then shows that for any lift

$$[f_{k-1} : C_{k-1}/O_{k-1} \rightarrow X, \mathbf{x}^{k-1}]$$

with $\sigma_i \circ \psi_{k-1} = f_{k-1} \circ x_i^{k-1}$, there exists a unique lift

$$[f_k : C_k/O_k \rightarrow X, \mathbf{x}^k]$$

such that $\sigma_i \circ \psi_k = f_k \circ x_i^k$.

First, note that we have a diagram

$$\begin{array}{ccccccc}
 & & & & & & 0 \\
 & & & & & & \downarrow \\
 & & & & & & \Theta_{C_0^\dagger/\mathbb{k}^\dagger}|\{x_1, \dots, x_s\} \\
 & & & & & & \downarrow \\
 0 & \longrightarrow & \Theta_{C_0^\dagger/\mathbb{k}^\dagger}(-\mathbf{x}) & \longrightarrow & f_0^* \Theta_{X^\dagger/(\mathbb{A}^1)^\dagger} & \longrightarrow & \mathcal{N}_{f_0, \mathbf{x}} \longrightarrow 0 \\
 & & \downarrow & & \downarrow = & & \downarrow \\
 0 & \longrightarrow & \Theta_{C_0^\dagger/\mathbb{k}^\dagger} & \longrightarrow & f_0^* \Theta_{X^\dagger/(\mathbb{A}^1)^\dagger} & \longrightarrow & \mathcal{N}_{f_0} \longrightarrow 0 \\
 & & \downarrow & & & & \downarrow \\
 & & \Theta_{C_0^\dagger/\mathbb{k}^\dagger}|\{x_1, \dots, x_s\} & & & & 0 \\
 & & \downarrow & & & & \\
 & & 0 & & & &
 \end{array}$$

by the snake lemma, so in fact $\mathcal{N}_{f_0, \mathbf{x}}$ splits as a sum of the skyscraper sheaf $\Theta_{C_0^\dagger/\mathbb{k}^\dagger}|\{x_1, \dots, x_s\}$ and \mathcal{N}_{f_0} . If we restrict Ξ to $\Theta_{C_0^\dagger/\mathbb{k}^\dagger} \otimes k(x_i)$, it is just given by

$$f_{0*} : \Theta_{C_0^\dagger/\mathbb{k}^\dagger} \otimes k(x_i) \rightarrow T_{X/\mathbb{A}^1, \sigma_i(0)}.$$

Set $T_{C_0/\mathbb{k}, x_i} := \Theta_{C_0^\dagger/\mathbb{k}^\dagger} \otimes k(x_i)$, the Zariski tangent space to C_0 at x_i . Thus to show Ξ is an isomorphism, it is enough to show that

$$\Xi' : H^0(C, \mathcal{N}_{f_0}) \rightarrow \prod_{i=1}^s T_{X/\mathbb{A}^1, \sigma_i(0)} / f_{0*}(T_{C_0/\mathbb{k}, x_i})$$

is an isomorphism. We now compute more explicitly the domain and range of Ξ' .

As usual, let

$$h : \tilde{\Gamma} \rightarrow M_{\mathbb{R}}$$

be obtained from Γ by adding vertices so that for a point $y \in \tilde{\Gamma}$, $h(y)$ is a vertex of \mathcal{P} if and only if y is a vertex of $\tilde{\Gamma}$ or y is contained in a marked unbounded edge. Let

$$h : \hat{\Gamma} \rightarrow M_{\mathbb{R}}$$

be obtained from Γ by removing the unbounded labelled edges, and removing any resulting bivalent vertices.

The range is easy: Since $\Theta_{X^\dagger/(\mathbb{A}^1)^\dagger}$ is canonically $M \otimes \mathcal{O}_X$, we have

$$T_{X/\mathbb{A}^1, \sigma_i(0)} \cong M \otimes_{\mathbb{Z}} \mathbb{k}$$

canonically. On the other hand, let V_i be the vertex of $\tilde{\Gamma}$ which is the boundary of the marked edge E_{x_i} . Then the curve C_{V_i} is a line. Let u_i be a tangent vector to the image under h of either unmarked edge adjacent to V_i . Then we know by Lemma 4.17 that $f_0(C_{V_i})$ is an orbit of $\mathbb{G}(Zu_i) \subseteq \mathbb{G}(M) \subseteq \mathbb{G}(\tilde{M})$. Thus $f_{0*}T_{C_0/\mathbb{k}, x_i} =$

$(\mathbb{Z}u_i) \otimes_{\mathbb{Z}} \mathbb{k} \subseteq M \otimes_{\mathbb{Z}} \mathbb{k}$. So the range of Ξ' is

$$\prod_{i=1}^s (M/\mathbb{Z}u_i) \otimes_{\mathbb{Z}} \mathbb{k} = \prod_{i=1}^s (M/\mathbb{Z}u_{(\partial^- E_i, E_i)}) \otimes_{\mathbb{Z}} \mathbb{k}$$

in the notation of Proposition 4.22.

Next, we describe the domain of Ξ' , $H^0(C_0, \mathcal{N}_{f_0})$. First, we restrict \mathcal{N}_{f_0} to each component C_V of C_0 , noting that the exact sequence

$$0 \rightarrow \Theta_{C_0^\dagger/\mathbb{k}^\dagger} \rightarrow f_0^* \Theta_{X_0^\dagger/\mathbb{k}^\dagger} = M \otimes_{\mathbb{Z}} \mathcal{O}_C \rightarrow \mathcal{N}_{f_0} \rightarrow 0$$

restricts to an exact sequence

$$0 \rightarrow \Theta_{C_0^\dagger/\mathbb{k}^\dagger}|_{C_V} \rightarrow M \otimes_{\mathbb{Z}} \mathcal{O}_{C_V} \rightarrow \mathcal{N}_{f_0}|_{C_V} \rightarrow 0.$$

By Example 3.36, (6), $\Theta_{C_0^\dagger/\mathbb{k}^\dagger}|_{C_V}$ is just $\Theta_{C_V/\mathbb{k}}(-\sum p_i)$, where $\{p_i\}$ is the set of special points on C_V , i.e., points of type (2) and (3) of Example 3.26 on C_0^\dagger . The number of such points is precisely the valency of V in $\tilde{\Gamma}$, *not* counting marked edges. As this valency is either 2 or 3, not counting marked edges, we refer to these two possibilities as the bivalent and trivalent cases respectively. Thus $\Theta_{C_0^\dagger/\mathbb{k}^\dagger}|_{C_V} \cong \mathcal{O}_{\mathbb{P}^1}$ or $\mathcal{O}_{\mathbb{P}^1}(-1)$ in the bivalent and trivalent cases respectively.

Note again in the bivalent case that since $f_0(C_V)$ is an orbit of $\mathbb{G}(\mathbb{Z}u)$, where u is tangent to either edge adjacent to V , the image of $H^0(C_V, \Theta_{C_0^\dagger/\mathbb{k}^\dagger}|_{C_V})$ in $M \otimes_{\mathbb{Z}} \mathbb{k}$ is $(\mathbb{Z}u) \otimes \mathbb{k}$. Thus in the bivalent and trivalent cases,

$$H^0(C_V, \mathcal{N}_{f_0}|_{C_V}) = \begin{cases} (M/\mathbb{Z}u) \otimes_{\mathbb{Z}} \mathbb{k} & V \text{ bivalent} \\ M \otimes_{\mathbb{Z}} \mathbb{k} & V \text{ trivalent.} \end{cases}$$

To understand whether we can glue a collection of sections

$$(s_v) \in \bigoplus_{V \in \tilde{\Gamma}^{[0]}} H^0(C_V, \mathcal{N}_{f_0}|_{C_V})$$

to get a section of $H^0(C_0, \mathcal{N}_{f_0})$, we shall show that at a double point $x \in C_0$ corresponding to an edge E of $\tilde{\Gamma}$, f_{0*} maps the one-dimensional space $\Theta_{C_0^\dagger/\mathbb{k}^\dagger} \otimes k(x)$ isomorphically to the subspace of $\Theta_{X^\dagger/\mathbb{k}^\dagger} \otimes k(f_0(x)) = M \otimes_{\mathbb{Z}} \mathbb{k}$ spanned by the tangent vectors of $h(E)$, which can be described as $\mathbb{Z}u_{(\partial^- E, E)} \otimes_{\mathbb{Z}} \mathbb{k} \subseteq M \otimes_{\mathbb{Z}} \mathbb{k}$.

Indeed, to see this, write C locally as usual as $\text{Spec } \mathbb{k}[z, w]/(zw)$ and X locally in a neighbourhood of $f(x)$ as $\text{Spec } \mathbb{k}[u, v, y^\pm, t]/(uv - t^\ell)$ for some ℓ , with ℓ the affine length of $h(E)$. We can assume $f_0^*u = z^\mu$, $f_0^*v = w^\mu$, where $\mu = w(E)$, $f_0^*(t) = 0$, $f_0^*(y) = g(z, w)$. Then a local generator of $\Theta_{C^\dagger/\mathbb{k}^\dagger} \otimes k(x)$ is $z\partial_z - w\partial_w$ (see Example 3.36, (6)), while local generators of $\Theta_{X^\dagger/(\mathbb{A}^1)^\dagger}$ at $f(x)$ are $u\partial_u - v\partial_v$ and ∂_y . One then sees easily that at x , $f_{0*}(z\partial_z - w\partial_w) = \mu(u\partial_u - v\partial_v)$. As an element of $\Theta_{X^\dagger/\mathbb{k}^\dagger} \otimes k(f_0(x)) = M \otimes_{\mathbb{Z}} \mathbb{k}$, $u\partial_u - v\partial_v$ is seen to be an element of M tangent to the edge $h(E)$, as this derivation vanishes on y .

As a consequence, given $(s_V) \in \bigoplus H^0(C_V, \mathcal{N}_{f_0}|_{C_V})$, s_V and $s_{V'}$ glue at a double point $x \in C_V \cap C_{V'}$ if, as elements of $M \otimes_{\mathbb{Z}} \mathbb{k}$, they differ by an element of

$\mathbb{Z}u_{(\partial^-E, E)}$. So the kernel of the map

$$\begin{aligned} \tilde{\Phi} : \prod_{V \in \tilde{\Gamma}^{[0]}} M \otimes_{\mathbb{Z}} \mathbb{k} &\times \prod_{V \in \tilde{\Gamma}^{[0]} \setminus \tilde{\Gamma}^{[0]}} (M/\mathbb{Z}u_{(\partial^-E(V), E(V))}) \otimes_{\mathbb{Z}} \mathbb{k} \\ &\rightarrow \prod_{E \in \tilde{\Gamma}^{[1]} \setminus \tilde{\Gamma}_{\infty}^{[1]}} (M/\mathbb{Z}u_{(\partial^-E, E)}) \otimes_{\mathbb{Z}} \mathbb{k} \end{aligned}$$

given by

$$(s_V)_{V \in \tilde{\Gamma}^{[0]}} \mapsto (s_{\partial^+E} - s_{\partial^-E})_{E \in \tilde{\Gamma}^{[1]} \setminus \tilde{\Gamma}_{\infty}^{[1]}}$$

is $H^0(C_0, \mathcal{N}_{f_0})$. One easily checks that $\tilde{\Phi}$ is surjective. Thus Ξ' is an isomorphism if and only if

$$\begin{aligned} \Phi'' : \prod_{V \in \tilde{\Gamma}^{[0]}} (M \otimes_{\mathbb{Z}} \mathbb{k}) &\times \prod_{V \in \tilde{\Gamma}^{[0]} \setminus \tilde{\Gamma}^{[0]}} (M/\mathbb{Z}u_{(\partial^-E(V), E(V))}) \otimes_{\mathbb{Z}} \mathbb{k} \\ &\rightarrow \prod_{E \in \tilde{\Gamma}^{[1]} \setminus \tilde{\Gamma}_{\infty}^{[1]}} (M/\mathbb{Z}u_{(\partial^-E, E)}) \otimes_{\mathbb{Z}} \mathbb{k} \times \prod_{i=1}^s (M/\mathbb{Z}u_{(\partial^-E_i, E_i)}) \otimes_{\mathbb{Z}} \mathbb{k} \end{aligned}$$

is an isomorphism. This map is given by

$$(s_V)_{V \in \tilde{\Gamma}^{[0]}} \mapsto ((s_{\partial^+E} - s_{\partial^-E})_{E \in \tilde{\Gamma}^{[1]} \setminus \tilde{\Gamma}_{\infty}^{[1]}}, (s_{V_i})_{1 \leq i \leq s})$$

where V_i is the vertex of E_{x_i} .

This map coincides with the map Φ'' of (4.8), after tensoring the latter map with \mathbb{k} . We saw in the proof of Proposition 4.22 that this was indeed an isomorphism.

Thus Ξ is an isomorphism, so there exists a unique lift

$$[f_k : C_k/O_k \rightarrow X, \mathbf{x}^k]$$

of

$$[f_{k-1} : C_{k-1}/O_{k-1} \rightarrow X, \mathbf{x}^{k-1}]$$

for any choice of log structure O_k^\dagger lifting the log structure on O_{k-1}^\dagger .

Step 2. Taking the limit. Now use the log structure on O_k given by the chart

$$\begin{aligned} \mathbb{N} &\rightarrow \mathbb{k}[t]/(t^{k+1}) \\ n &\mapsto t^n. \end{aligned}$$

This gives us $[f_k : C_k/O_k \rightarrow X, \mathbf{x}^k]$ for all k by Step 1. Forgetting the log structure on $C_k^\dagger \rightarrow O_k^\dagger$, $f_k : C_k/O_k \rightarrow X$ can be viewed as a family of stable maps in X , and hence gives a map from O_k to the stack of stable maps of genus zero with s marked points with target X . These maps are compatible, and hence give a map $O_\infty = \text{Spec } \mathbb{k}[[t]]$ into the stack of stable maps, i.e., give a diagram

$$\begin{array}{ccc} C_\infty & \xrightarrow{f_\infty} & X \\ \downarrow & & \downarrow \\ O_\infty & \longrightarrow & \mathbb{A}^1 \end{array}$$

with f_∞ a stable map of curves. This gives a diagram as desired in Theorem 4.27.

We still need to show that this diagram satisfies condition (3) of the theorem. We first argue that $C_\infty^\dagger \rightarrow O_\infty^\dagger$, with log structure on C_∞ induced by the inclusion $f_\infty^{-1}(\partial X) \subseteq C_\infty$, is in fact log smooth. This just needs to be checked in étale

neighbourhoods of double points and marked points of C_0 . At a double point q , as observed in Step 2 of the proof of Theorem 4.24, we have

$$\widehat{O}_{C_\infty, q} \cong \mathbb{k}[x, y, t]/(xy - \lambda t^e)$$

for $\lambda \in 0, 1$. By applying the characterization of log smoothness of Example 3.26 to $C_k^\dagger \rightarrow O_k^\dagger$ for large k , in fact $\lambda = 1$ and e is determined by the requirement that $\overline{\mathcal{M}}_{C_0, \bar{q}} = S_e$. Thus $C_\infty^\dagger \rightarrow O_\infty^\dagger$ is log smooth at q . For log smoothness at a marked point p , it is enough to observe that $f_\infty^{-1}(\overline{\partial(X \setminus X_0)})$ is a disjoint union of sections as in Step 3 of the proof of Theorem 4.24. This holds because the same is necessarily true for $f_k^{-1}(\overline{\partial(X \setminus X_0)})$ as $f_k : C_k^\dagger \rightarrow O_k^\dagger$ is log smooth. Thus we see f_∞ is log smooth with the induced log structure.

We can now restrict this log structure to C_0 , obtaining another torically transverse log curve $(C_0')^\dagger \rightarrow X_0^\dagger$ over $\text{Spec } \mathbb{k}^\dagger$. This in fact must coincide with the original curve $C_0^\dagger \rightarrow X_0^\dagger$. Indeed, from the proof of Proposition 4.23, the log structure on C_0 is uniquely determined away from the nodes, and at the nodes the local description of Example 3.26 tells us the two log structures coincide. This shows existence.

For uniqueness, given such a $f_\infty : C_\infty \rightarrow X$, we obtain a log map $f_\infty^\dagger : C_\infty^\dagger \rightarrow X^\dagger$ over $O_\infty^\dagger \rightarrow (\mathbb{A}^1)^\dagger$, where the log structure on C_∞ is induced by $f_\infty^{-1}(\partial X) \subseteq C_\infty$ and the log structure on O_∞ is the pull-back of the log structure on \mathbb{A}^1 . This induces a log structure on the fibre over the closed point of O_∞ , C_0^\dagger , and the assumption is that this coincides with the original given $C_0^\dagger \rightarrow X_0^\dagger$ over $\text{Spec } \mathbb{k}^\dagger$. It is then not difficult to see from the description of log smooth curves of Example 3.26 that this implies that $C_\infty^\dagger \rightarrow O_\infty^\dagger$ is log smooth. But then, by the above deformation theory, $C_k^\dagger \cong C_\infty^\dagger \times_{O_\infty^\dagger} O_k^\dagger$ is uniquely determined for each h , so C_∞^\dagger is unique. \square

4.6. The end of the proof

We now complete the proof of Theorem 4.4, i.e., we show that

$$N_{\Delta, \Sigma}^{0, \text{trop}} = N_{\Delta, \Sigma}^{0, \text{hol}}.$$

Here we have fixed points $P_1, \dots, P_s \in M_{\mathbb{Q}}$ so that all tropical rational curves passing through P_1, \dots, P_s are simple. Recall that we have rescaled the lattice M so that $P_i \in M$ for each i . Furthermore, after having chosen the good decomposition \mathcal{P} , we can rescale the lattice again to ensure that for every genus zero tropical curve $h : (\Gamma, x_1, \dots, x_s) \rightarrow M_{\mathbb{R}}$ with $h(x_i) = P_i$, the image of each edge of $\tilde{\Gamma}$ has affine length divisible by its weight, so that the hypotheses of Theorem 4.14 always hold.

Then, after choosing sections $\sigma_1, \dots, \sigma_s : \mathbb{A}^1 \rightarrow X$ as usual, hence points $q_i = \sigma_i(0) \in X_0$, we obtain, by Theorem 4.14, precisely $N_{\Delta, \Sigma}^{0, \text{trop}}$ torically transverse marked log curves

$$f : (C^\dagger, x_1, \dots, x_s) \rightarrow X_0^\dagger$$

of genus zero with $f(x_i) = q_i$. Let $\mathcal{M}_{\Delta, \Sigma}^{0, \text{log}}(\sigma_1, \dots, \sigma_s)$ denote the set of these torically transverse marked log curves.

On the other hand, let $K = \overline{\mathbb{k}((t))}$. There is a canonical map $\text{Spec } K \rightarrow \mathbb{A}_{\mathbb{k}}^1$ coming from the inclusion $\mathbb{k}[[t]] \hookrightarrow K$, and the fibred product $X \times_{\mathbb{A}_{\mathbb{k}}^1} \text{Spec } K$ is isomorphic to $X_\Sigma \times_{\mathbb{k}} K$. Furthermore, each section σ_i then defines a K -valued point of $X_\Sigma \times_{\mathbb{k}} K$, which we also denote by σ_i .

We wish to show that there is a one-to-one correspondence between

$$\mathcal{M}_{\Delta, \Sigma}^{0, \log}(\sigma_1, \dots, \sigma_s)$$

and the set $\mathcal{M}_{\Delta, \Sigma}^{0, \text{hol}}(\sigma_1, \dots, \sigma_s)$ of torically transverse curves of genus zero

$$f : (C, x_1, \dots, x_s) \rightarrow X_{\Sigma} \times_{\mathbb{k}} K$$

over the field K with $f(x_i) = \sigma_i$. This will prove the theorem. Indeed, since K is an algebraically closed field of characteristic zero and the answer should be independent of which algebraically closed field of characteristic zero we use, the number of such torically transverse curves in $X_{\Sigma} \times_{\mathbb{k}} K$ is $N_{\Delta, \Sigma}^{0, \text{hol}}$.

To show this one-to-one correspondence, start with a log curve

$$[f_0 : C_0/O_0 \rightarrow X, \mathbf{x}^0]$$

in $\mathcal{M}_{\Delta, \Sigma}^{0, \log}(\sigma_1, \dots, \sigma_s)$. By Theorem 4.27, this gives a curve

$$[f_{\infty} : C_{\infty}/O_{\infty} \rightarrow X, \mathbf{x}^{\infty}].$$

We have a natural map $\text{Spec } K \rightarrow O_{\infty}$ coming from the inclusion $\mathbb{k}[[t]] \subseteq K$, hence a curve $C = C_{\infty} \times_{O_{\infty}} \text{Spec } K$ over $\text{Spec } K$, along with a map

$$f : C \rightarrow X \times_{\mathbb{A}^1} \text{Spec } K = X_{\Sigma} \times_{\mathbb{k}} K.$$

This is the desired torically transverse curve over K .

Conversely, suppose we are given a curve $f : C \rightarrow X_{\Sigma} \times_{\mathbb{k}} K$ in the moduli space $\mathcal{M}_{\Delta, \Sigma}^{0, \text{hol}}(\sigma_1, \dots, \sigma_s)$. Since

$$K = \bigcup_{d=1}^{\infty} \mathbb{k}((t^{1/d})),$$

and C is finite type over K , in fact there is some d such that C is defined over the field $\mathbb{k}((t^{1/d}))$. This gives us a diagram as in (4.15), with $C^* = C$ and $L = \mathbb{k}((t^{1/d}))$.

Then applying Theorem 4.24, after replacing L with $L = \mathbb{k}((t^{1/de}))$ for some e , making a base-change $\mathbb{A}^1 \rightarrow \mathbb{A}^1$, and blowing up X , we obtain a diagram (4.17). Note that the blow-up $\tilde{X} \rightarrow X$ corresponds to a subdivision of \mathcal{P} . This then gives rise to a torically transverse pre-log curve

$$f_0 : (C_0, x_1^0, \dots, x_s^0) \rightarrow \tilde{X}_0$$

by restricting the map f of (4.17) to the fibre over the closed point of $\text{Spec } R$. This has some associated tropical curve $h : \Gamma \rightarrow M_{\mathbb{R}}$, necessarily passing through P_1, \dots, P_s . Since P_1, \dots, P_s were chosen generically, h must be simple, and so by Theorem 4.24, (2), f_0 is a torically transverse log curve.

But we have already classified all the torically transverse log curves with associated tropical curve passing through P_1, \dots, P_s , and hence this log curve is one of the ones contributing to $N_{\Delta, \Sigma}^{0, \text{trop}}$. In particular, this log curve gives a unique family already defined over $\text{Spec } \mathbb{k}[[t]]$. Thus we in fact have $de = 1$ and no base-change $\mathbb{A}_{\mathbb{k}}^1 \rightarrow \mathbb{A}_{\mathbb{k}}^1$ is required. In particular, all the torically transverse curves defined over K are in fact defined over $\mathbb{k}((t))$, and we have a one-to-one correspondence between these curves and torically transverse log curves.

This completes the proof of Mikhalkin's theorem.

4.7. References and further reading

The exposition in this chapter is based on the paper [86] of Nishinou and Siebert. That paper covers a similar result in all dimensions. The original proof was given by Mikhalkin in [80] after announcing the result in [79]. After Mikhalkin's original announcement, Shustin also provided an independent proof of the result in [105]. In the case of \mathbb{P}^2 , Gathmann and Markwig [31] gave a purely combinatorial argument, by showing that the tropical curve counts satisfy the WDVV equation. There is a growing literature involving other generalizations of Mikhalkin's results which are too numerous to list here.

Period integrals

In Chapter 4, we proved Mikhalkin’s formula for curve-counting in toric surfaces. Specializing to the case of \mathbb{P}^2 , one essentially obtains the A-model for \mathbb{P}^2 , as these formulae compute the Gromov-Witten invariants $\langle T_2^{3d-1} \rangle_{0,d}$ counting the number of rational curves of degree d passing through $3d - 1$ points. We should note that for arbitrary toric surfaces X , Mikhalkin’s formula does not compute Gromov-Witten invariants because there will always be contributions to such Gromov-Witten invariants from curves which are not torically transverse, e.g., have components mapping into the toric boundary of X . However, for \mathbb{P}^2 , we do obtain Gromov-Witten invariants. This allows one to define quantum cohomology for \mathbb{P}^2 purely tropically, and one obtains from this the full A-model variation of semi-infinite Hodge structures.

This then raises the following question: Is it possible to describe the B-model tropically, in such a way that it becomes transparent that the A- and B-models coincide? At first glance, it is not at all obvious how perturbations of a potential function W_0 and oscillatory integrals involving these functions can be computed tropically. However, as we shall demonstrate in this chapter, one can indeed carry out these computations tropically, and in particular give a completely tropical description of the B-model. This chapter is based entirely on the paper [42].

5.1. The perturbed Landau-Ginzburg potential

Our main goal is to find a canonical description of W_0 and its perturbation by giving the “correct” universal unfolding of W_0 . The important clue is the interpretation of Cho and Oh [16] of the potential W_0 in terms of Maslov index two holomorphic disks. We will not explain this here at the level of symplectic geometry, but pass immediately to the tropical version. At the same time, we shall initially work in a broader context of arbitrary toric varieties, putting the description of the mirror to \mathbb{P}^n discussed already in Chapter 2 in the context of a broader setup of mirror symmetry for toric varieties introduced by Givental (see [33]).

5.1.1. Givental’s mirrors of toric varieties. To begin, we fix once and for all a lattice $M = \mathbb{Z}^n$, $N = \text{Hom}_{\mathbb{Z}}(M, \mathbb{Z})$ the dual lattice, $M_{\mathbb{R}} = M \otimes_{\mathbb{Z}} \mathbb{R}$, $N_{\mathbb{R}} = N \otimes_{\mathbb{Z}} \mathbb{R}$. Fix a complete fan Σ in $M_{\mathbb{R}}$. We shall assume that X_{Σ} is a non-singular toric variety. We adopt the same notation used in §1.3: we let T_{Σ} be the free abelian group generated by the elements of $\Sigma^{[1]}$, the one-dimensional cones of Σ , with $\rho \in \Sigma^{[1]}$ corresponding to a generator t_{ρ} . We have the map $r : T_{\Sigma} \rightarrow M$ given by $r(t_{\rho}) = m_{\rho}$, the primitive generator of ρ , and the assumption of non-singularity implies r is surjective. So there is a natural exact sequence

$$0 \longrightarrow K_{\Sigma} \longrightarrow T_{\Sigma} \xrightarrow{r} M \longrightarrow 0$$

defining K_Σ . Dualizing this sequence gives

$$0 \rightarrow N \rightarrow \mathrm{Hom}_{\mathbb{Z}}(T_\Sigma, \mathbb{Z}) \rightarrow \mathrm{Pic} X_\Sigma \rightarrow 0,$$

as discussed in §3.1.2. After tensoring with \mathbb{C}^\times , we get an exact sequence

$$0 \rightarrow N \otimes \mathbb{C}^\times \rightarrow \mathrm{Hom}(T_\Sigma, \mathbb{C}^\times) \xrightarrow{\kappa} \mathrm{Pic} X_\Sigma \otimes \mathbb{C}^\times \rightarrow 0.$$

We should think of κ as providing the family of mirrors to X_Σ . We define the *Kähler moduli space* of X_Σ to be

$$\mathcal{M}_\Sigma := \mathrm{Pic} X_\Sigma \otimes \mathbb{C}^\times = \mathrm{Spec} \mathbb{C}[K_\Sigma],$$

so we have a morphism

$$\kappa : \mathrm{Hom}(T_\Sigma, \mathbb{C}^\times) = \mathrm{Spec} \mathbb{C}[T_\Sigma] \rightarrow \mathcal{M}_\Sigma.$$

Note that a fibre of κ over a closed point of \mathcal{M}_Σ is canonically isomorphic to $\mathrm{Spec} \mathbb{C}[M]$.

In keeping with Chapter 2, after passing to complex manifolds, we can pass to the universal cover of \mathcal{M}_Σ , which we shall write as the vector space

$$\widetilde{\mathcal{M}}_\Sigma := \mathrm{Pic} X_\Sigma \otimes \mathbb{C},$$

with the map $\widetilde{\mathcal{M}}_\Sigma \rightarrow \mathcal{M}_\Sigma$ given by $D \otimes y \mapsto D \otimes e^y$. We then set

$$\check{\mathcal{X}}_\Sigma := \mathrm{Hom}(T_\Sigma, \mathbb{C}^\times) \times_{\mathcal{M}_\Sigma} \widetilde{\mathcal{M}}_\Sigma,$$

pulling back the family given by κ to $\widetilde{\mathcal{M}}_\Sigma$. So far, we have not introduced the formal thickening of this moduli space, nor shall we do so in general.

We shall take W_0 to be the function on $\check{\mathcal{X}}_\Sigma$ defined by

$$(5.1) \quad W_0 := \sum_{\rho \in \Sigma[1]} z^{t_\rho}.$$

This makes sense as a function on $\mathrm{Spec} \mathbb{C}[T_\Sigma]$, hence as a function on $\check{\mathcal{X}}_\Sigma$. We think of $(\check{\mathcal{X}}_\Sigma, W_0)$ as a family of Landau-Ginzburg models which are mirror to X_Σ .

EXAMPLE 5.1. Taking Σ to be the fan for \mathbb{P}^n , this gives precisely the description of the mirror family $\check{\mathcal{X}} \rightarrow \widetilde{\mathcal{M}}$ for \mathbb{P}^n defined in §2.2.3, with the formal coordinates $t_0 = t_2 = \dots = t_n = 0$.

5.1.2. Tropical disks and mirrors to toric varieties. We shall now restrict to the surface case, taking $M = \mathbb{Z}^2$.

We have already defined the notion of a tropical curve in Chapter 1. Here, we define the notion of a *tropical disk* (which was called a *tropical curve with stops* in [85]).

DEFINITION 5.2. Let $\bar{\Gamma}$ be a weighted, connected finite graph without bivalent vertices as in §1.3, with the additional data of a choice of univalent vertex V_{out} , adjacent to a unique edge E_{out} . Let

$$\Gamma' := (\bar{\Gamma} \setminus \bar{\Gamma}_\infty^{[0]}) \cup \{V_{\mathrm{out}}\} \subseteq \bar{\Gamma}.$$

Suppose furthermore that Γ' has first Betti number zero (i.e., Γ' is a tree with one compact external edge and a number of non-compact external edges). Then a *parameterized d -pointed tropical disk in $M_{\mathbb{R}}$ with domain Γ'* is a choice of inclusion $\{p_1, \dots, p_d\} \hookrightarrow \Gamma_\infty^{[1]} \setminus \{E_{\mathrm{out}}\}$ written as $p_i \mapsto E_{p_i}$, with $w(E) = 0$ if and only if

$E = E_{p_i}$ for some i , and a continuous map $h : \Gamma' \rightarrow M_{\mathbb{R}}$ satisfying the same conditions as Definition 1.11, except that there is *no* balancing condition at V_{out} .

An *isomorphism* of tropical disks

$$h_1 : (\Gamma'_1, p_1, \dots, p_d) \rightarrow M_{\mathbb{R}} \text{ and } h_2 : (\Gamma'_2, p_1, \dots, p_d) \rightarrow M_{\mathbb{R}}$$

is a homeomorphism $\Phi : \Gamma'_1 \rightarrow \Gamma'_2$ respecting the marked edges and the weights with $h_1 = h_2 \circ \Phi$. A *tropical disk* is an isomorphism class of parameterized tropical disks.

The *combinatorial type* of a tropical disk $h : (\Gamma', p_1, \dots, p_d) \rightarrow M_{\mathbb{R}}$ is defined to be the homeomorphism class of $\bar{\Gamma}$ with the marked points, weights, and V_{out} , together with, for every vertex V and edge E containing V , the primitive tangent vector to $h(E)$ in M pointing away from V .

While the above definitions make sense in any dimension, the point of restricting to dimension two is that we can make use of Mikhalkin's definition of multiplicity (Definition 1.21), using the same definition for disks, but not counting V_{out} :

DEFINITION 5.3. Suppose $\text{rank } M = 2$. Let $h : \Gamma' \rightarrow M_{\mathbb{R}}$ be a marked tropical disk such that $\bar{\Gamma}$ only has vertices of valency one and three. The *multiplicity* of the disk h is then

$$\text{Mult}(h) := \prod_{V \in \Gamma^{[0]}} \text{Mult}_V(h),$$

where $\text{Mult}_V(h)$ is as defined in Definition 1.21. In particular, there is no contribution from V_{out} .

We now fix $M \cong \mathbb{Z}^2$ and a complete rational polyhedral fan Σ in $M_{\mathbb{R}}$, yielding T_{Σ} and $r : T_{\Sigma} \rightarrow M$ as in §5.1.1.

DEFINITION 5.4. A disk h is a *tropical disk in X_{Σ}* if every $E \in \Gamma_{\infty}^{[1]} \setminus \{E_{\text{out}}\}$ has $h(E)$ either a point or a translate of some $\rho \in \Sigma^{[1]}$.

Analogously to Definition 1.12, if the disk has d_{ρ} unbounded edges which are translates of $\rho \in \Sigma^{[1]}$ (counted with weight), then the *degree* of h is

$$\Delta(h) := \sum_{\rho \in \Sigma^{[1]}} d_{\rho} t_{\rho} \in T_{\Sigma}.$$

Fix general points $P_1, \dots, P_k \in M_{\mathbb{R}}$, and fix a general *base-point* $Q \in M_{\mathbb{R}}$. When we talk about general points in the sequel, we mean that there is an open dense subset (typically the complement of a finite union of polyhedra of codimension at least one) of $M_{\mathbb{R}}^{k+1}$ such that $(P_1, \dots, P_k, Q) \in M_{\mathbb{R}}^{k+1}$ lies in this open subset. This choice of open subset will depend on particular needs.

Associate to the points P_1, \dots, P_k the variables u_1, \dots, u_k in the ring

$$R_k := \frac{\mathbb{C}[u_1, \dots, u_k]}{(u_1^2, \dots, u_k^2)}.$$

DEFINITION 5.5. Let $h : (\Gamma', p_1, \dots, p_d) \rightarrow M_{\mathbb{R}}$ be a tropical disk in X_{Σ} with $h(V_{\text{out}}) = Q$, $h(p_j) = P_{i_j}$, $1 \leq i_1 < \dots < i_d \leq k$. (This ordering removes a $d!$ ambiguity about the labelling of the marked points.) We say h is a *tropical disk in $(X_{\Sigma}, P_1, \dots, P_k)$ with boundary Q* .

The Maslov index of the disk h is

$$MI(h) := 2(|\Delta(h)| - d).$$

The phrase ‘‘Maslov index’’ has a precise definition in the context of pseudo-holomorphic disks with boundaries contained in Lagrangian submanifolds of symplectic manifolds. In particular, the Maslov index enters into a formula for the virtual dimension of a family of such disks. We have given the definition above so that a precisely analogous formula holds, given by the following lemma.

LEMMA 5.6. *If P_1, \dots, P_k, Q are chosen in general position, then the set of Maslov index $2n$ tropical disks in $(X_\Sigma, P_1, \dots, P_k)$ with boundary Q is an $(n - 1)$ -dimensional polyhedral complex. The set of Maslov index $2n$ tropical disks with arbitrary boundary is an $(n + 1)$ -dimensional polyhedral complex.*

PROOF. This is exactly the kind of standard tropical dimension counting argument carried out in Lemma 1.20. We sketch the argument here.

Fix a combinatorial type of tropical disk with d marked points, with degree Δ . If the combinatorial type is general, then the domain Γ' only has trivalent vertices apart from V_{out} . Such a tree has $|\Delta| + d$ unbounded edges and hence $|\Delta| + d - 1$ bounded edges (including E_{out}). A tropical disk $h : \Gamma' \rightarrow M_{\mathbb{R}}$ of this given combinatorial type is then completely determined by the position of $h(V_{\text{out}}) \in M_{\mathbb{R}}$ and the affine lengths of the bounded edges. This produces a cell in the moduli space $\mathcal{M}_{\Delta, d}^{\text{disk}}(X_\Sigma)$ of all d -pointed tropical disks of degree Δ . The closure of this cell is $(\mathbb{R}_{\geq 0})^{|\Delta| + d - 1} \times M_{\mathbb{R}}$. Also, there are only a finite number of combinatorial types of disks of a given degree. Thus $\mathcal{M}_{\Delta, d}^{\text{disk}}(X_\Sigma)$ is a finite $(|\Delta| + d + 1)$ -dimensional polyhedral complex. Furthermore, we have a piecewise linear map $\text{ev} : \mathcal{M}_{\Delta, d}^{\text{disk}}(X_\Sigma) \rightarrow M_{\mathbb{R}}^d$, taking a disk h to the tuple $(h(p_1), \dots, h(p_d))$. Let $E \subseteq \mathcal{M}_{\Delta, d}^{\text{disk}}(X_\Sigma)$ be the union of cells mapping under ev to cells of codimension ≥ 1 in $M_{\mathbb{R}}^d$; then $h(E)$ is a closed subset of $M_{\mathbb{R}}^d$. Thus, if $(P_{i_1}, \dots, P_{i_d}) \in M_{\mathbb{R}}^d$ is not in this closed subset, for $1 \leq i_1 < \dots < i_d \leq k$ distinct indices, then $\text{ev}^{-1}(P_{i_1}, \dots, P_{i_d})$ is a codimension $2d$ subset of $\mathcal{M}_{\Delta, d}^{\text{disk}}(X_\Sigma)$. Thus the dimension of the moduli space of tropical disks of a given degree Δ with arbitrary boundary in $(X_\Sigma, P_1, \dots, P_k)$ is $|\Delta| + 1 - d = MI(h)/2 + 1$. Similarly, if we fix a general boundary point Q , the dimension is $MI(h)/2 - 1$, as claimed. \square

DEFINITION 5.7. Given the data $P_1, \dots, P_k, Q \in M_{\mathbb{R}}$ general, let

$$h : (\Gamma', p_1, \dots, p_d) \rightarrow M_{\mathbb{R}}$$

be a Maslov index two marked tropical disk with boundary Q in $(X_\Sigma, P_1, \dots, P_k)$. Then we can associate to h a monomial in $\mathbb{C}[T_\Sigma] \otimes_{\mathbb{C}} R_k[[y_0]]$,

$$\text{Mono}(h) := \text{Mult}(h) z^{\Delta(h)} u_{I(h)},$$

where $z^{\Delta(h)} \in \mathbb{C}[T_\Sigma]$ is the monomial corresponding to $\Delta(h) \in T_\Sigma$, the subset $I(h) \subseteq \{1, \dots, k\}$ is defined by

$$I(h) := \{i \mid h(p_j) = P_i \text{ for some } j\},$$

and

$$u_{I(h)} = \prod_{i \in I(h)} u_i.$$

Define the k -pointed Landau-Ginzburg potential as

$$W_k(Q) := y_0 + \sum_h \text{Mono}(h) \in \mathbb{C}[T_\Sigma] \otimes_{\mathbb{C}} R_k[[y_0]]$$

where the sum is over all Maslov index two disks h in $(X_\Sigma, P_1, \dots, P_k)$ with boundary Q . By Lemma 5.6, this is a finite sum for P_1, \dots, P_k, Q general.

We define the k -th order thickening of the Kähler moduli space $\widetilde{\mathcal{M}}_\Sigma$ defined in §5.1.1 to be the ringed space

$$\widetilde{\mathcal{M}}_{\Sigma,k} := (\widetilde{\mathcal{M}}_\Sigma, \mathcal{O}_{\widetilde{\mathcal{M}}_{\Sigma,k}}),$$

with elements of $\mathcal{O}_{\widetilde{\mathcal{M}}_{\Sigma,k}}(U)$ for $U \subseteq \widetilde{\mathcal{M}}_\Sigma$ being expressions of the form

$$\sum_{\substack{n=0 \\ I \subseteq \{1, \dots, k\}}}^{\infty} f_{n,I} y_0^n u_I$$

where $f_{n,I}$ is a holomorphic function on U for each n and I .

Similarly, we define the thickened mirror family

$$\check{\mathcal{X}}_{\Sigma,k} := (\check{\mathcal{X}}_\Sigma, \mathcal{O}_{\check{\mathcal{X}}_{\Sigma,k}})$$

in exactly the same way. Thus we have a family

$$\kappa : \check{\mathcal{X}}_{\Sigma,k} \rightarrow \widetilde{\mathcal{M}}_{\Sigma,k}.$$

By construction, $W_k(Q)$ is a regular function on $\check{\mathcal{X}}_{\Sigma,k}$, so we can think of this as providing a family of Landau-Ginzburg potentials.

The sheaf of relative differentials $\Omega_{\check{\mathcal{X}}_{\Sigma,k}/\widetilde{\mathcal{M}}_{\Sigma,k}}^1$ is canonically isomorphic to the trivial locally free sheaf $M \otimes_{\mathbb{Z}} \mathcal{O}_{\check{\mathcal{X}}_{\Sigma,k}}$, with $m \otimes 1$ corresponding to the differential

$$(5.2) \quad \text{dlog } m := \frac{d(z^{\overline{m}})}{z^{\overline{m}}};$$

here $\overline{m} \in T_\Sigma$ is any lift of $m \in M$, and $\text{dlog } m$ is well-defined as a *relative* differential independently of the lift. Thus a choice of generator of $\bigwedge^2 M \cong \mathbb{Z}$ determines a nowhere-vanishing relative holomorphic two-form Ω , canonical up to sign. Explicitly, if $e_1, e_2 \in M$ is a positively oriented basis, then

$$(5.3) \quad \Omega = \text{dlog } e_1 \wedge \text{dlog } e_2.$$

□

REMARK 5.8. Given a fan Σ in $M_{\mathbb{R}} \cong \mathbb{R}^2$, if we take $k = 0$, the Maslov index two disks with boundary Q are precisely the disks of the form $Q + \rho$ for $\rho \in \Sigma^{[1]}$. All these disks have multiplicity one, having no vertices. Thus

$$W_0(Q) = y_0 + \sum_{\rho \in \Sigma^{[1]}} z^{t_\rho}.$$

Other than the term y_0 , whose role in the following discussion will only be necessary to reproduce the correct behaviour of the J -function, this is precisely (5.1). For those familiar with the work of Cho and Oh [16], we point out that these tropical disks coincide precisely with the holomorphic disks classified there, and so we reproduce, tropically, the description of the Landau-Ginzburg potential given in [16].




FIGURE 1. The one additional Maslov index two tropical disk with $k = 1$.

The function $W_k(Q)$ is then intended to be the “correct” perturbation of $W_0(Q)$, in the sense that the parameters appearing in $W_k(Q)$ are closely related to flat coordinates. However, for a general choice of Σ this is not true, the chief problem being that there will be copies of \mathbb{P}^1 in the toric boundary of X_Σ which do not deform to curves intersecting the big torus orbit of X_Σ . This is a standard problem in tropical geometry: tropical geometry cannot “see” these curves. This is not a problem as long as X_Σ is a product of projective spaces, so in particular, we will now restrict to the case of $X_\Sigma = \mathbb{P}^2$. This returns us to the example studied in detail in Chapter 2.

EXAMPLE 5.9. Let Σ be the fan for \mathbb{P}^2 , depicted in Figure 14 of Chapter 1, so that $X_\Sigma = \mathbb{P}^2$. Here $T_\Sigma = \mathbb{Z}^3$ with basis t_0, t_1, t_2 corresponding to ρ_0, ρ_1, ρ_2 , and we write x_i for the monomial $z^{t_i} \in \mathbb{C}[T_\Sigma]$. In this case $\widetilde{\mathcal{M}}_\Sigma = \mathbb{C}$, say with coordinate y_1 , so that the space $\widetilde{\mathcal{M}}_{\Sigma,k}$ has underlying topological space \mathbb{C} . The mirror family $\check{\mathcal{X}}_{\Sigma,k}$ is then defined by the equation $x_0x_1x_2 = e^{y_1}$ in $\mathbb{C}^3 \times \widetilde{\mathcal{M}}_{\Sigma,k}$, and the map

$$\kappa : \check{\mathcal{X}}_{\Sigma,k} \rightarrow \widetilde{\mathcal{M}}_{\Sigma,k}$$

is then given by projection, i.e.,

$$\kappa^*(y_0) = y_0$$

$$\kappa^*(y_1) = y_1$$

$$\kappa^*(u_i) = u_i$$

As discussed in Remark 5.8,

$$W_0(Q) = y_0 + x_0 + x_1 + x_2.$$

If we take $k = 1$, marking one point in \mathbb{P}^2 , we obtain one additional disk, as depicted in Figure 1, and if we take $k = 2$ with P_1 and P_2 chosen as in Figure 2, we have three additional disks. Note that the potential depends on the particular choices of the points P_1, \dots, P_k as well as Q . In the given examples, we have respectively

$$W_1(Q) = y_0 + x_0 + x_1 + x_2 + u_1x_1x_2$$

$$W_2(Q) = y_0 + x_0 + x_1 + x_2 + u_1x_0x_1 + u_2x_0x_1 + u_1u_2x_0x_1^2.$$

□




FIGURE 2. The additional Maslov index two tropical disks with $k = 2$.

5.2. Tropical descendent invariants

As we know from Chapter 2, mirror symmetry for \mathbb{P}^2 gives an expression for the J -function of \mathbb{P}^2 in terms of oscillatory integrals. The J -function in turn involves gravitational descendent Gromov-Witten invariants. On the other hand, the main point of this chapter is to show that the oscillatory integrals can be naturally evaluated in terms of certain tropically described objects. In particular, we will need to define tropical versions of the gravitational descendent invariants, which we shall do in the case of \mathbb{P}^2 .

The tropical count of rational curves of degree d through $3d - 1$ points was ultimately motivated by the actual proof that tropical curves correspond to holomorphic curves: the count arising from Mikhalkin's original proof in [80] is very similar to the way the count arises in Chapter 4. In the case of gravitational descendents, however, it is somewhat more difficult to motivate these formulae. There is no direct proof analogous to the arguments of Chapter 4 that the descendent invariants we define in fact coincide with the holomorphic versions. We will only be able to motivate these definitions by saying that these definitions are what are given by computing oscillatory integrals. The main result of this chapter says that the definition given for tropical descendent invariants is correct if and only if mirror symmetry for \mathbb{P}^2 (Theorem 2.44) holds.

One can also attempt to approach this problem in a purely tropical setting, by considering the moduli spaces of tropical curves, defining the ψ -classes as tropical cycles on these moduli spaces, and applying tropical intersection theory. This

approach was applied by Markwig and Rau in [77], but as yet does not recover all the formulae we give here.

We continue with $M = \mathbb{Z}^2$ and Σ a fan in $M_{\mathbb{R}}$ defining a projective toric variety X_{Σ} .

DEFINITION 5.10. Let $P_1, \dots, P_k \in M_{\mathbb{R}}$ be general. Let $S \subseteq M_{\mathbb{R}}$ be a subset. Define

$$\mathcal{M}_{\Delta, n}^{\text{trop}}(X_{\Sigma}, P_1, \dots, P_k, \psi^{\nu} S)$$

to be the moduli space of rational $(n+1)$ -pointed tropical curves in X_{Σ}

$$h : (\Gamma, p_1, \dots, p_n, x) \rightarrow M_{\mathbb{R}}$$

of degree Δ such that

- (1) $h(p_j) = P_{i_j}$, $1 \leq i_1 < \dots < i_n \leq k$.
- (2) The edge E_x is attached to a vertex V_x of Γ ; the valency of this vertex is denoted $\text{Val}(x)$. Then

$$\text{Val}(x) = \nu + 3.$$

- (3) $h(x) \in S$.
- (4) The weight of each unbounded edge of Γ is either 0 or 1.

LEMMA 5.11. For $P_1, \dots, P_k \in M_{\mathbb{R}}$ general,

- (1) $\mathcal{M}_{\Delta, n}^{\text{trop}}(X_{\Sigma}, P_1, \dots, P_k, \psi^{\nu} M_{\mathbb{R}})$ is a polyhedral complex of dimension $|\Delta| - n - \nu$.
- (2) $\mathcal{M}_{\Delta, n}^{\text{trop}}(X_{\Sigma}, P_1, \dots, P_k, \psi^{\nu} C)$ is a polyhedral complex of dimension $|\Delta| - n - \nu - 1$ for C a general translate of a tropical curve in $M_{\mathbb{R}}$.
- (3) $\mathcal{M}_{\Delta, n}^{\text{trop}}(X_{\Sigma}, P_1, \dots, P_k, \psi^{\nu} Q)$ is a polyhedral complex of dimension $|\Delta| - n - \nu - 2$ for $Q \in M_{\mathbb{R}}$ a general point.

PROOF. This is straightforward, as in Lemma 5.6. The dimension count is as follows. Fix the combinatorial type of the curve to be generic, so that all vertices of Γ are trivalent except for the vertex adjacent to E_x , which is $(\nu+3)$ -valent. Such a tree has $|\Delta| + n + 1$ unbounded edges, and thus has $|\Delta| + n + 1 - (\nu+3)$ bounded edges. The curves of this combinatorial type are then determined by the location of $h(x) \in M_{\mathbb{R}}$ and the lengths of the bounded edges, giving a cell of the form $(\mathbb{R}_{\geq 0})^{|\Delta| + n - \nu - 2} \times M_{\mathbb{R}}$. Fixing $h(p_1), \dots, h(p_n)$ then yields the desired dimension of the moduli space in (1) to be

$$|\Delta| + n - \nu - 2n.$$

This gives (1). For (2) and (3), we consider the map

$$\text{ev}_x : \mathcal{M}_{\Delta, n}^{\text{trop}}(X_{\Sigma}, P_1, \dots, P_k, \psi^{\nu} M_{\mathbb{R}}) \rightarrow M_{\mathbb{R}}$$

given by $\text{ev}_x(h) := h(x)$. Let $E_1 \subseteq \mathcal{M}_{\Delta, n}^{\text{trop}}(X_{\Sigma}, P_1, \dots, P_k, \psi^{\nu} M_{\mathbb{R}})$ be the union of cells which map to codimension ≥ 1 sets in $M_{\mathbb{R}}$, and let E_2 be the union of cells which map to points in $M_{\mathbb{R}}$. Then we need to choose the translate C so that $C \cap h(E_1)$ is zero-dimensional and $C \cap h(E_2) = \emptyset$. Similarly, we need to choose $Q \notin h(E_1)$. Then $\text{ev}_x^{-1}(C)$ or $\text{ev}_x^{-1}(Q)$ are the desired moduli spaces in cases (2) and (3) and are of the desired dimension. \square

We can describe a tropical curve $h : \Gamma \rightarrow M_{\mathbb{R}}$ of genus zero in terms of a collection of disks, by splitting the curve up at a chosen vertex. The following lemma is crucial for describing these curves in terms of disks in what follows.

LEMMA 5.12. *Let $P_1, \dots, P_k \in M_{\mathbb{R}}$ be general and $S \subseteq M_{\mathbb{R}}$ a subset. Let*

$$h \in \mathcal{M}_{\Delta, n}^{\text{trop}}(X_{\Sigma}, P_1, \dots, P_k, \psi^{\nu} S).$$

Let $\Gamma'_1, \dots, \Gamma'_{\nu+2}$ denote the closures of the connected components of $\Gamma \setminus E_x$, with $h_i : \Gamma'_i \rightarrow M_{\mathbb{R}}$ the restrictions of h . Each disk h_i is viewed as being marked by those points $p \in \{p_1, \dots, p_n\}$ with $E_p \subseteq \Gamma'_i$. There is one special case to consider here: if E_x and E_{p_i} share a common vertex V , then we discard the edge E_{p_i} from consideration as well, so we have disks $h_1, \dots, h_{\nu+1}$. (Note that since $h(p_i) \neq h(p_j)$ for $i \neq j$, we never have E_{p_i} and E_{p_j} sharing a common vertex.)

- (1) *If $S = M_{\mathbb{R}}$ and $n = |\Delta| - \nu$, then either*
 - (a) *E_x does not share a vertex with any of the edges E_{p_i} , and then $MI(h_i) = 2$ for all but two choices of i , and for these i , $MI(h_i) = 0$.*
 - (b) *E_x does share a vertex with one of the edges E_{p_i} , and then $MI(h_i) = 2$ for all i .*
- (2) *If $S = C$ is a general translate of a tropical curve in $M_{\mathbb{R}}$ and $n = |\Delta| - \nu - 1$, then $MI(h_i) = 2$ for all but one i , and for this i , $MI(h_i) = 0$.*
- (3) *If $S = \{Q\}$ for a general point Q and $n = |\Delta| - \nu - 2$, then $MI(h_i) = 2$ for all i .*

PROOF. First note that the condition on n and the generality of P_1, \dots, P_k, C , and Q guarantee by the previous lemma that the moduli space under consideration is zero dimensional. If any of the disks h_i can be deformed while keeping its boundary $h_i(x)$ fixed, then this yields a non-trivial deformation of h , which does not exist. Thus by Lemma 5.6 we must have $MI(h_i) \leq 2$ in all cases. Let n_i be the number of marked points on h_i . We note that

$$\begin{aligned} \sum_i \frac{MI(h_i)}{2} &= \sum_i (|\Delta(h_i)| - n_i) \\ &= \begin{cases} |\Delta(h)| - (n - 1) & \text{Case (1) (b)} \\ |\Delta(h)| - n & \text{otherwise} \end{cases} \\ &= \begin{cases} \nu & \text{Case (1) (a)} \\ \nu + 1 & \text{Case (1) (b)} \\ \nu + 1 & \text{Case (2)} \\ \nu + 2 & \text{Case (3)} \end{cases} \end{aligned}$$

Since there are $\nu + 2$ disks except in Case (1) (b), when there are $\nu + 1$ disks, the result follows. \square

We can now define tropical analogues of the descendent Gromov-Witten invariants which appear in the Givental J -function. From now on in this chapter, we only consider the case of $X_{\Sigma} = \mathbb{P}^2$, with Σ the fan with rays generated by $m_0 = (-1, -1)$, $m_1 = (1, 0)$ and $m_2 = (0, 1)$, and t_0, t_1, t_2 the generators of T_{Σ} , with $r(t_i) = m_i$. Let

$$\Delta_d := d(t_0 + t_1 + t_2) \in T_{\Sigma};$$

curves of degree Δ_d should be viewed as usual as degree d tropical curves in \mathbb{P}^2 .

DEFINITION 5.13. Fix general points $Q, P_1, P_2, \dots \in M_{\mathbb{R}}$. Let L be the tropical line (a translate of the union of the one-dimensional cones in the fan of \mathbb{P}^2) with vertex Q .

For a tropical curve h in \mathbb{P}^2 with a marked point x , let $n_0(x), n_1(x)$ and $n_2(x)$ be the number of unbounded rays sharing a common vertex with E_x mapping under h to rays in the directions m_0, m_1 and m_2 respectively. As in Lemma 5.12, we denote by h_1, \dots the tropical disks obtained by removing E_x from Γ , with the outgoing edge of h_i being $E_{i,\text{out}}$. Let $m(h_i) = w(E_{i,\text{out}})m^{\text{prim}}(h_i)$, where $m^{\text{prim}}(h_i) \in M$ is a primitive vector tangent to $h_i(E_{i,\text{out}})$ pointing away from $h(x)$.

Define

$$\begin{aligned} \text{Mult}_x^0(h) &= \frac{1}{n_0(x)!n_1(x)!n_2(x)!} \\ \text{Mult}_x^1(h) &= -\frac{\sum_{k=1}^{n_0(x)} \frac{1}{k} + \sum_{k=1}^{n_1(x)} \frac{1}{k} + \sum_{k=1}^{n_2(x)} \frac{1}{k}}{n_0(x)!n_1(x)!n_2(x)!} \\ \text{Mult}_x^2(h) &= \frac{\left(\sum_{l=0}^2 \sum_{k=1}^{n_l(x)} \frac{1}{k}\right)^2 + \sum_{l=0}^2 \sum_{k=1}^{n_l(x)} \frac{1}{k^2}}{2n_0(x)!n_1(x)!n_2(x)!} \end{aligned}$$

(1) We define

$$\langle P_1, \dots, P_{3d-2-\nu}, \psi^\nu Q \rangle_{0,d}^{\text{trop}}$$

to be

$$\sum_h \text{Mult}(h)$$

where the sum is over all marked tropical rational curves

$$h \in \mathcal{M}_{\Delta_d, 3d-2-\nu}^{\text{trop}}(P_1, \dots, P_{3d-2-\nu}, \psi^\nu Q).$$

We define

$$\text{Mult}(h) := \text{Mult}_x^0(h) \prod_{\substack{V \in \Gamma^{[0]} \\ V \notin E_x}} \text{Mult}_V(h).$$

(2) We define

$$\langle P_1, \dots, P_{3d-1-\nu}, \psi^\nu L \rangle_{0,d}^{\text{trop}}$$

as a sum

$$\sum_h \text{Mult}(h)$$

where the sum is again over all marked tropical rational curves

$$h : (\Gamma, p_1, \dots, p_{3d-1-\nu}, x) \rightarrow M_{\mathbb{R}}$$

with $h(p_i) = P_i$ and satisfying one of the following two conditions.

(a)

$$h \in \mathcal{M}_{\Delta_d, 3d-1-\nu}^{\text{trop}}(P_1, \dots, P_{3d-1-\nu}, \psi^\nu L).$$

Furthermore, no unbounded edge of Γ having a common vertex with E_x other than E_x maps into the connected component of $L \setminus \{Q\}$ containing $h(x)$. By Lemma 5.12, there is precisely one j , $1 \leq j \leq \nu + 2$, with $MI(h_j) = 0$. Suppose also that the connected component of $L \setminus \{Q\}$ containing $h(x)$ is $Q + \mathbb{R}_{>0}m_i$. Then we define

$$\text{Mult}(h) = |m(h_j) \wedge m_i| \text{Mult}_x^0(h) \prod_{\substack{V \in \Gamma^{[0]} \\ V \notin E_x}} \text{Mult}_V(h).$$

Here $m(h_j) \wedge m_i \in \wedge^2 M \cong \mathbb{Z}$, so the absolute value makes sense.

(b) $\nu \geq 1$ and

$$h \in \mathcal{M}_{\Delta_d, 3d-1-\nu}^{\text{trop}}(P_1, \dots, P_{3d-1-\nu}, \psi^{\nu-1}Q).$$

Then

$$\text{Mult}(h) = \text{Mult}_x^1(h) \prod_{\substack{V \in \Gamma^{[0]} \\ V \notin E_x}} \text{Mult}_V(h).$$

(3) We define

$$\langle P_1, \dots, P_{3d-\nu}, \psi^\nu M_{\mathbb{R}} \rangle_{0,d}^{\text{trop}}$$

as a sum

$$\sum_h \text{Mult}(h)$$

where the sum is over all marked tropical rational curves

$$h : (\Gamma, p_1, \dots, p_{3d-\nu}, x) \rightarrow M_{\mathbb{R}}$$

such that $h(p_i) = P_i$ and either

(a)

$$h \in \mathcal{M}_{\Delta_d, 3d-\nu}^{\text{trop}}(P_1, \dots, P_{3d-\nu}, \psi^\nu M_{\mathbb{R}})$$

and E_x does not share a vertex with any of the E_{p_i} 's. Furthermore, no unbounded edge of Γ having a common vertex with E_x other than E_x maps into the connected component of $M_{\mathbb{R}} \setminus L$ containing $h(x)$. By Lemma 5.12, there are precisely two distinct j_1, j_2 with $1 \leq j_1, j_2 \leq \nu + 2$ such that $MI(h_{j_i}) = 0$. Then we define

$$\text{Mult}(h) = |m(h_{j_1}) \wedge m(h_{j_2})| \text{Mult}_x^0(h) \prod_{\substack{V \in \Gamma^{[0]} \\ V \notin E_x}} \text{Mult}_V(h).$$

(b)

$$h \in \mathcal{M}_{\Delta_d, 3d-\nu}^{\text{trop}}(P_1, \dots, P_{3d-\nu}, \psi^\nu M_{\mathbb{R}})$$

and E_x shares a vertex with E_{p_i} . Furthermore, no unbounded edge of Γ having a common vertex with E_x other than E_x and E_{p_i} maps into the connected component of $M_{\mathbb{R}} \setminus L$ containing $h(x)$. Then we define

$$\text{Mult}(h) = \text{Mult}_x^0(h) \prod_{\substack{V \in \Gamma^{[0]} \\ V \notin E_x}} \text{Mult}_V(h).$$

(c) $\nu \geq 1$ and

$$h \in \mathcal{M}_{\Delta_d, 3d-\nu}^{\text{trop}}(P_1, \dots, P_{3d-\nu}, \psi^{\nu-1}L).$$

Furthermore, no unbounded edge of Γ having a common vertex with E_x other than E_x maps into the connected component of $L \setminus \{Q\}$ containing $h(x)$. By Lemma 5.12, there is precisely one j , $1 \leq j \leq \nu + 1$, with $MI(h_j) = 0$. Suppose the connected component of $L \setminus \{Q\}$ containing $h(x)$ is $Q + \mathbb{R}_{>0}m_i$. Then we define

$$\text{Mult}(h) = |m(h_j) \wedge m_i| \text{Mult}_x^1(h) \prod_{\substack{V \in \Gamma^{[0]} \\ V \notin E_x}} \text{Mult}_V(h).$$

(d) $\nu \geq 2$ and

$$h \in \mathcal{M}_{\Delta_d, 3d-\nu}^{\text{trop}}(P_1, \dots, P_{3d-\nu}, \psi^{\nu-2}Q).$$

Then

$$\text{Mult}(h) = \text{Mult}_x^2(h) \prod_{\substack{V \in \Gamma^{[0]} \\ V \notin E_x}} \text{Mult}_V(h).$$

In all cases $S = \{Q\}, L$ or $M_{\mathbb{R}}$, we define for $\sigma \in \Sigma$,

$$(5.4) \quad \langle P_1, \dots, P_{3d-\nu-(2-\dim S)}, \psi^\nu S \rangle_{d,\sigma}^{\text{trop}}$$

to be the contribution to $\langle P_1, \dots, P_{3d-\nu-(2-\dim S)}, \psi^\nu S \rangle_{0,d}^{\text{trop}}$ coming from curves h with $h(x)$ in the interior of $\sigma + Q$. In (1), the only contribution comes from $\sigma = \{0\}$, in (2), the contributions come from the zero- and one-dimensional cones of Σ , and in (3), the contributions come from all cones of Σ . \square

REMARKS 5.14. (1) Note that all moduli spaces involved are zero-dimensional for general choices of Q, P_1, \dots , so the sums make sense.

(2) The formula in Definition 5.13, (1), for $\nu = 0$, gives the tropical curve counting formula for the number of rational curves of degree d passing through $3d - 1$ points, by Mikhalkin's formula (Theorem 4.4). For $\nu > 0$, this coincides with the formula given by Markwig and Rau in [77]. In particular, by the results of that paper,

$$\langle P_1, \dots, P_{3d-2-\nu}, \psi^\nu Q \rangle_{0,d}^{\text{trop}} = \langle T_2^{3d-2-\nu}, \psi^\nu T_2 \rangle_{0,d}.$$

The remaining formulas we give are more mysterious, and have no known justification outside of the mirror symmetry arguments given in this chapter.

(3) It is easy to see that $\langle P_1, \dots, P_{3d-1}, \psi^0 L \rangle_{0,d}^{\text{trop}}$ is d times the number of rational curves through $3d-1$ points. Indeed, the only contribution to this number comes from Definition 5.13, (2) (a). For each tropical rational curve $h : \Gamma \rightarrow M_{\mathbb{R}}$ with $3d-1$ marked points passing through P_1, \dots, P_{3d-1} we obtain a contribution for every point of $h^{-1}(L)$ by marking that point with x . The factor $|m(h_j) \wedge m_i| \text{Mult}_x^0(h)$ for the multiplicity in this case gives the intersection multiplicity of $h(\Gamma)$ with L at each point of $h^{-1}(L)$, as defined in Example 1.6. By the tropical Bézout theorem, the total contribution from h is then $(h(\Gamma) \cdot L) \text{Mult}(h) = d \text{Mult}(h)$.

Thus

$$\langle P_1, \dots, P_{3d-1}, \psi^0 L \rangle_{0,d}^{\text{trop}} = d \langle P_1, \dots, P_{3d-1} \rangle_{0,d}^{\text{trop}} = d \langle T_2^{3d-1} \rangle_{0,d} = \langle T_2^{3d-1}, T_1 \rangle_{0,d},$$

by Mikhalkin's formula and the Divisor Axiom.

(4) $\langle P_1, \dots, P_{3d}, \psi^0 M_{\mathbb{R}} \rangle_{0,d}^{\text{trop}} = 0$. Indeed, the only possible contributions come from Definition 5.13 (3) (a), but there are no rational curves of degree d through $3d$ general points. Thus

$$\langle P_1, \dots, P_{3d}, \psi^0 M_{\mathbb{R}} \rangle_{0,d}^{\text{trop}} = \langle T_2^{3d}, T_0 \rangle_{0,d},$$

as both are zero, the latter by the Fundamental Class Axiom.

We will prove the following in §5.5:

THEOREM 5.15. *The invariants defined in Definition 5.13 are independent of the choice of the P_i 's and Q .*

This allows us to make the following definition.

DEFINITION 5.16. We define

$$\begin{aligned} \langle T_2^{3d-2-\nu}, \psi^\nu T_2 \rangle_{0,d}^{\text{trop}} &:= \langle P_1, \dots, P_{3d-2-\nu}, \psi^\nu Q \rangle_{0,d}^{\text{trop}} \\ \langle T_2^{3d-1-\nu}, \psi^\nu T_1 \rangle_{0,d}^{\text{trop}} &:= \langle P_1, \dots, P_{3d-1-\nu}, \psi^\nu L \rangle_{0,d}^{\text{trop}} \\ \langle T_2^{3d-\nu}, \psi^\nu T_0 \rangle_{0,d}^{\text{trop}} &:= \langle P_1, \dots, P_{3d-\nu}, \psi^\nu M_{\mathbb{R}} \rangle_{0,d}^{\text{trop}} \end{aligned}$$

where the P_i 's and Q have been chosen generally.

We define

$$\langle T_2^m, \psi^\nu T_i \rangle_{0,d}^{\text{trop}} := 0$$

if $m + i + \nu \neq 3d$.

We define the *tropical J-function* for \mathbb{P}^2 by analogy with Example 2.30,

$$\begin{aligned} J_{\mathbb{P}^2}^{\text{trop}} &:= e^{(y_0 T_0 + y_1 T_1)/\hbar} \cup \left(T_0 + \sum_{i=0}^2 \left(\hbar^{-1} y_2 \delta_{2,i} \right. \right. \\ &\quad \left. \left. + \sum_{d \geq 1} \sum_{\nu \geq 0} \langle T_2^{3d+i-2-\nu}, \psi^\nu T_{2-i} \rangle_{0,d}^{\text{trop}} \hbar^{-(\nu+2)} e^{dy_1} \frac{y_2^{3d+i-2-\nu}}{(3d+i-2-\nu)!} \right) T_i \right) \\ &=: \sum_{i=0}^2 J_i^{\text{trop}} T_i. \end{aligned}$$

EXAMPLE 5.17. The tropical descendent invariants are (relatively) easy to compute. For example, consider $\langle \psi^{3d-2} T_2 \rangle_{0,d}^{\text{trop}}$. There is only one tropical curve of degree d with a vertex of valency $3d + 1$ at Q (including the marked edge), namely the curve which has d legs of weight one in each of the three directions $(-1, -1)$, $(1, 0)$, and $(0, 1)$, and hence contributes a multiplicity of $1/(d!)^3$, so

$$\langle \psi^{3d-2} T_2 \rangle_{0,d}^{\text{trop}} = \frac{1}{(d!)^3}.$$

Next, consider $\langle T_2, \psi^{3d-3} T_2 \rangle_{0,d}^{\text{trop}}$. Fixing some point $P_1 \in M_{\mathbb{R}}$, we note that any curve contributing to this invariant will have a $3d$ -valent vertex at Q (including the marked edge), and will decompose into $3d - 1$ Maslov index two disks with boundary Q , with precisely one of them passing through the point P_1 . Taking, say, P_1 as depicted in Figure 1, the only curve contributing to this invariant is then as depicted in Figure 3. Thus

$$\langle T_2, \psi^{3d-3} T_2 \rangle_{0,d}^{\text{trop}} = \frac{1}{d!(d-1)!(d-1)!}.$$

Both these give the correct non-tropical descendent invariants. This will follow from Theorem 5.18, Corollary 5.19 and Theorem 2.44.

We give several additional examples of contributions to tropical descendent Gromov-Witten invariants. In Figure 4, we see a curve which is obtained by gluing together four Maslov index two disks with boundary Q : the vertical line with the attached number 2 means we take that vertical unbounded edge twice, while the 2 attached to the horizontal line segment tells us that that edge has weight two. The contribution to the multiplicity from all vertices except Q is 2×2 . The vertex at Q is 5-valent (remembering that this vertex is also the endpoint of the edge E_x),




FIGURE 3. The one curve contributing to $\langle P_1, \psi^{3d-3}Q \rangle_{0,d}^{\text{trop}}$. The numbers d and $d-1$ indicate not the weight of the edge but the number of copies of the edge.

with $n_0(x) = 1$, $n_1(x) = 0$ and $n_2(x) = 2$. Thus

$$\text{Mult}_x^0(h) = 1/(1!0!2!) = 1/2,$$

$$\text{Mult}_x^1(h) = -(1 + 1 + 1/2)/(1!0!2!) = -5/4,$$

$$\text{Mult}_x^2(h) = ((1 + 1 + 1/2)^2 + 1 + 1 + 1/4)/(2 \cdot 1!0!2!) = 17/8.$$

Thus this curve contributes $4/2 = 2$ to $\langle P_1, \dots, P_5, \psi^2Q \rangle_{0,3}^{\text{trop}}$, $4 \times (-5/4) = -5$ to $\langle P_1, \dots, P_5, \psi^3L \rangle_{0,3}^{\text{trop}}$ and $4 \times (17/8) = 17/2$ to $\langle P_1, \dots, P_5, \psi^4M_{\mathbb{R}} \rangle_{0,3}^{\text{trop}}$.

In Figure 5, we consider another example, with the dotted line being L . In this case, decomposing the curve into a collection of disks by removing E_x , one finds the only one of these disks with Maslov index zero is the diagonal line passing through P_6 . The factor $|m(h_j) \wedge m_1|$ is then 1. We have $n_0(x) = n_1(x) = 0$ and $n_2(x) = 2$. (Note that the diagonal ray passing through P_6 is not the image of an unbounded edge with a common vertex with E_x because of the marked edge mapping to P_6 !) Thus $\text{Mult}_x^0(h) = 1/2$ and $\text{Mult}_x^1(h) = -3/4$. Thus the contribution from this curve to $\langle P_1, \dots, P_6, \psi^2L \rangle_{0,3}^{\text{trop}}$ is 2 and the contribution to $\langle P_1, \dots, P_6, \psi^3M_{\mathbb{R}} \rangle_{0,3}^{\text{trop}}$ is -3 .

Figure 6 shows a curve of degree 4, with the vertex mapping to $h(x)$ having valency 7, and $n_0(x) = n_1(x) = 0$ (necessarily, otherwise this curve would not contribute) and $n_2(x) = 4$. Then $\text{Mult}_x^0(h) = 1/24$, and the two Maslov index zero disks obtained by decomposing this curve have tangent vectors to their outgoing edges given by $(-1, -2)$ and $(1, -2)$. Thus the factor $|m(h_{j_1}) \wedge m(h_{j_2})|$ is 4. All other vertices have multiplicity 1 except for two vertices of multiplicity 2, so the total contribution of this curve to $\langle P_1, \dots, P_8, \psi^4M_{\mathbb{R}} \rangle_{0,4}^{\text{trop}}$ is $2/3$.

We leave it to the reader to give an example of a curve contributing to this descendent invariant of type (3) (b) in the definition of tropical descendent invariants.




FIGURE 4




FIGURE 5

5.3. The main B-model statement

In this section we continue with $M = \mathbb{Z}^2$ and Σ the fan in $M_{\mathbb{R}}$ defining \mathbb{P}^2 . We then have for $k \geq 0$,

$$\kappa : \check{\mathcal{X}}_{\Sigma,k} \rightarrow \widetilde{\mathcal{M}}_{\Sigma,k}$$

as given by Definition 5.7. Given a general choice of points $P_1, \dots, P_k, Q \in M_{\mathbb{R}}$, one obtains the Landau-Ginzburg potential $W_k(Q)$. Note that modulo u_1, \dots, u_k , $W_k(Q) = W_0(Q)$. The space $\check{\mathcal{X}}_{\Sigma,k}$ carries a relative holomorphic two-form defined by (5.3).

Let R be the local system on $\widetilde{\mathcal{M}}_{\Sigma,k} \times \mathbb{C}^\times$ whose fibre over (u, \hbar) is

$$H_2(\kappa^{-1}(u), \text{Re}(W_0(Q)/\hbar) \ll 0).$$

This local system is not concerned about the thickening of the structure sheaf: it carries purely topological information on the underlying topological space of




FIGURE 6

$\widetilde{\mathcal{M}}_{\Sigma,k} \times \mathbb{C}^\times$, which is $\widetilde{\mathcal{M}}_\Sigma \times \mathbb{C}^\times$. In fact, this local system coincides with the one given in §2.2.3 in the case of \mathbb{P}^2 .

From §2.2.4, we know that the local system R is rank 3 and has a multi-valued basis of sections Ξ_0, Ξ_1, Ξ_2 whose integrals over $e^{W_0(Q)/\hbar}\Omega$ are given by (2.38). As a consequence, the integrals

$$\int_{\Xi_i} e^{W_k(Q)/\hbar}\Omega$$

make sense formally, by writing

$$\exp(W_k(Q)/\hbar) = \exp(W_0(Q)/\hbar) \exp((W_k(Q) - W_0(Q))/\hbar)$$

and expanding $\exp((W_k(Q) - W_0(Q))/\hbar)$ as a finite power series, precisely as was done in §2.2.3.

The main theorem of this chapter, to be proved in §5.5, is now:

THEOREM 5.18. *We can write*

$$(5.5) \quad \sum_{i=0}^2 \alpha^i \int_{\Xi_i} e^{W_k(Q)/\hbar} \Omega = \hbar^{-3\alpha} \sum_{i=0}^2 \varphi_i \cdot (\alpha \hbar)^i$$

where

$$\varphi_i(y_0, y_1, u_1, \dots, u_k, \hbar^{-1}) = \delta_{0,i} + \sum_{j=1}^{\infty} \varphi_{i,j}(y_0, y_1, u_1, \dots, u_k) \hbar^{-j},$$

for $0 \leq i \leq 2$, with

$$\begin{aligned} \varphi_{0,1} &= y_0 \\ \varphi_{1,1} &= y_1 \\ \varphi_{2,1} &= y_2 := \sum_{i=1}^k u_i. \end{aligned}$$

Furthermore,

$$\varphi_i = J_i^{\text{trop}}(y_0, y_1, y_2).$$

In fact, this theorem proves the equivalence of tropical curve counting with descendants and mirror symmetry. To make this precise, consider the statement:

STATEMENT 5.19 (Tropical curve counting with descendants).

$$J_{\mathbb{P}^2} = J_{\mathbb{P}^2}^{\text{trop}}.$$

COROLLARY 5.20. *Theorem 2.44 and Statement 5.19 are equivalent. In particular, since Theorem 2.44 is known to be true, Statement 5.19 is true.*

PROOF. We adopt the notation of Chapter 2, with $\widetilde{\mathcal{M}} = \widetilde{\mathcal{M}}^B$ the B-model moduli space, $\pi: \check{\mathcal{X}} \rightarrow \widetilde{\mathcal{M}}$ as given in §2.2.3, with

$$W = t_0 + W_0 + t_2 W_0^2,$$

the potential on $\check{\mathcal{X}}$, universal in a neighbourhood of each point of $\widetilde{\mathcal{M}}$. On the other hand, for a fixed k , we have $\kappa: \check{\mathcal{X}}_{\Sigma,k} \rightarrow \widetilde{\mathcal{M}}_{\Sigma,k}$ constructed in §5.1.2. We replace $\widetilde{\mathcal{M}}$ and $\widetilde{\mathcal{M}}_{\Sigma,k}$ with germs of these spaces at 0, and replace $\check{\mathcal{X}}$ and $\check{\mathcal{X}}_{\Sigma,k}$ with the inverse images of these germs. Then $(\check{\mathcal{X}}, W)$ is a universal unfolding for $(\pi^{-1}(0), W_0)$, which coincides with $(\kappa^{-1}(0), W_k(Q)|_{\kappa^{-1}(0)})$. Thus there is a diagram

$$\begin{array}{ccc} \check{\mathcal{X}}_{\Sigma,k} & \xrightarrow{\eta} & \check{\mathcal{X}} \\ \kappa \downarrow & & \downarrow \pi \\ \widetilde{\mathcal{M}}_{\Sigma,k} & \xrightarrow{\xi} & \widetilde{\mathcal{M}} \end{array}$$

such that $\eta|_{\kappa^{-1}(0)}: \kappa^{-1}(0) \rightarrow \pi^{-1}(0)$ is the identity and $W \circ \eta = W_k(Q)$.

Now both $\widetilde{\mathcal{M}} \times \mathbb{C}^\times$ and $\widetilde{\mathcal{M}}_{\Sigma,k} \times \mathbb{C}^\times$ come with local systems R with stalk at (u, \hbar) given by

$$H_2(\pi^{-1}(u), \text{Re}(W_0/\hbar) \ll 0; \mathbb{C})$$

or

$$H_2(\kappa^{-1}(u), \text{Re}(W_k(Q)/\hbar) \ll 0; \mathbb{C})$$

respectively. Writing the local system on $\widetilde{\mathcal{M}} \times \mathbb{C}^\times$ as R' to distinguish the two cases, we clearly have $\xi^{-1}R' \cong R$.

On $\widetilde{\mathcal{M}}$ we have the B-model variation of semi-infinite Hodge structures, given by \mathcal{E}' , ∇' , $(\cdot, \cdot)_{\mathcal{E}'}$ and Gr' determined by the Euler vector field E' . Similarly, on $\widetilde{\mathcal{M}}_{\Sigma, k}$ we can construct a semi-infinite variation of Hodge structures, using exactly the same procedures: we use unprimed notation. So \mathcal{E} is defined so that $\Gamma(U, \mathcal{E})$ consists of sections of $\mathcal{R}^\vee = (R \otimes_{\mathbb{C}} \mathcal{O}_{\widetilde{\mathcal{M}}_{\Sigma, k} \times \mathbb{C}^\times})^\vee$ over open sets of the form $U \times \{\hbar \mid |\hbar| < \epsilon\}$ given by forms $f\Omega$ on $\kappa^{-1}(U) \times \{\hbar \mid |\hbar| < \epsilon\}$ with f holomorphic, and algebraic when restricted to $\kappa^{-1}(u) \times \{\hbar\}$. The connection ∇ is induced by the Gauss-Manin connection, and the pairing given by the pairing on H_2 . The only thing which is defined slightly differently is the grading operator, with

$$\text{Gr}(s) = \nabla_{\hbar \partial_{\hbar} + E}^{GM}(s) - s,$$

but where

$$E = y_0 \partial_{y_0} + 3 \partial_{y_1} - \sum_{i=1}^k u_i \partial_{u_i}.$$

It is clear that by construction, $\mathcal{E} = \xi^* \mathcal{E}'$, and ∇' , $(\cdot, \cdot)_{\mathcal{E}'}$ pull back to ∇ and $(\cdot, \cdot)_{\mathcal{E}}$. I claim also that Gr' pulls back to Gr . To verify this, note that

$$E = \kappa_* \left(y_0 \partial_{y_0} + \sum_{i=0}^2 x_i \partial_{x_i} - \sum_{i=1}^k u_i \partial_{u_i} \right),$$

and thus thinking of E as a vector field on $\check{\mathcal{X}}_{\Sigma, k}$, $E(W_k(Q)) = W_k(Q)$. Indeed, $E(y_0) = y_0$, and

$$\begin{aligned} E(\text{Mono}(h)) &= E(\text{Mult}(h) z^{\Delta(h)} u_{I(h)}) \\ &= (|\Delta(h)| - \#I(h)) \text{Mult}(h) z^{\Delta(h)} u_{I(h)} \\ &= \text{Mono}(h), \end{aligned}$$

since each h appearing in $W_k(Q)$ is Maslov index two. Thus thinking of E as living on $\check{\mathcal{X}}_{\Sigma, k}$, $\eta_* E$ preserves W , since $(\eta_* E)(W) = E(W \circ \eta)$, and hence $\eta_* E = E'$, as E' is the unique vector field on $\check{\mathcal{X}}$ preserving W . Thus Gr is the pull-back of Gr' .

Now the Frobenius manifold structure on $\widetilde{\mathcal{M}}$ is induced by the splitting of \mathcal{H}' as $\mathcal{H}'_- \oplus \mathcal{E}'_0$, again putting primes on all the notation associated to $\widetilde{\mathcal{M}}$. As flat sections of $\mathcal{E}' \otimes_{\mathcal{O}_{\widetilde{\mathcal{M}}}\{\hbar\}} \mathcal{O}_{\widetilde{\mathcal{M}}}\{\hbar, \hbar^{-1}\}$ pull back to flat sections of $\mathcal{E} \otimes_{\mathcal{O}_{\widetilde{\mathcal{M}}_{\Sigma, k}}\{\hbar\}} \mathcal{O}_{\widetilde{\mathcal{M}}_{\Sigma, k}}\{\hbar, \hbar^{-1}\}$, we see that \mathcal{H} and \mathcal{H}' are canonically isomorphic. Of course \mathcal{E}_0 and \mathcal{E}'_0 are also, so we can use the splitting $\mathcal{H} = \mathcal{H}_- \oplus \mathcal{E}_0$ with $\mathcal{H}_- = \mathcal{H}'_-$.

Let $\Omega'_0 \in \mathcal{H}'$ be the flat section whose value at $0 \in \widetilde{\mathcal{M}}$ is $[\Omega]$; this agrees with the flat section $\Omega_0 \in \mathcal{H}$ whose value at $0 \in \widetilde{\mathcal{M}}_{\Sigma, k}$ is $[\Omega]$. We have maps $\tau' : (\hbar \mathcal{H}'_- / \mathcal{H}'_-) \otimes_{\mathbb{C}} \mathcal{O}_{\widetilde{\mathcal{M}}}\{\hbar\} \rightarrow \mathcal{E}'$ and $\tau : (\hbar \mathcal{H}_- / \mathcal{H}_-) \otimes_{\mathbb{C}} \mathcal{O}_{\widetilde{\mathcal{M}}_{\Sigma, k}}\{\hbar\} \rightarrow \mathcal{E}$ by (2.16); these maps commute with pull-back by ξ . We know from the proof of Proposition 2.43, (3), that $\tau'([\Omega'_0] \otimes 1)$ is represented by some $f'\Omega$ with f' holomorphic on $\check{\mathcal{X}} \times \{\hbar \mid |\hbar| < \epsilon\}$, regular on fibres of π , $f'|_{\pi^{-1}(0)} \equiv 1$, and

$$\sum_{i=0}^2 \alpha^i \int_{\Xi_i} e^{W/\hbar} f' \Omega = \hbar^{-3\alpha} \sum_{i=0}^2 \varphi'_i(\mathbf{t}, \hbar^{-1})(\alpha \hbar)^i$$

such that

$$\varphi'_i(\mathbf{t}, \hbar^{-1}) = \delta_{0,i} + \sum_{j=1}^{\infty} \varphi'_{i,j}(\mathbf{t}) \hbar^{-j}.$$

Furthermore, these two conditions uniquely determine the section of \mathcal{E}' represented by $f'\Omega$. The same is true of $\tau([\Omega_0] \otimes 1)$: it is represented by some $f\Omega$ with f holomorphic on $\check{\mathcal{X}}_{\Sigma,k} \times \{\hbar \mid |\hbar| < \epsilon\}$, regular on fibres of κ , $f|_{\kappa^{-1}(0)} \equiv 1$ and

$$\sum_{i=0}^2 \alpha^i \int_{\Xi_i} e^{W_\kappa(Q)/\hbar} f\Omega = \hbar^{-3\alpha} \sum_{i=0}^2 \varphi_i(y_0, y_1, u_1, \dots, u_k, \hbar^{-1})(\alpha\hbar)^i$$

such that

$$\varphi = \delta_{0,i} + \sum_{j=1}^{\infty} \varphi_{i,j}(y_0, y_1, u_1, \dots, u_k) \hbar^{-j}.$$

Now the main point is that Theorem 5.18 shows that we may take $f \equiv 1$ on $\mathcal{X}_{\Sigma,k} \times \mathbb{C}^\times$. Furthermore, since $[\Omega_0] = \xi^*[\Omega'_0]$, $\tau([\Omega_0] \otimes 1) = \xi^* \tau'([\Omega'_0] \otimes 1)$. Thus if $f'\Omega$ represents $\tau'([\Omega'_0] \otimes 1)$ as above, we see that

$$\varphi_i = \varphi'_i \circ \xi.$$

In particular, if y'_0, y'_1, y'_2 are flat coordinates on $\widetilde{\mathcal{M}}$, i.e., $y'_i = \varphi'_{i,1}$, then

$$y_i = y'_i \circ \xi.$$

Now suppose Theorem 2.44 is true. Then in particular, the conclusions of Proposition 2.45 hold, and $\varphi'_i = J_i$, so $\varphi_i = J_i \circ \xi$. But by Theorem 5.18, $\varphi_i = J_i^{\text{trop}}$. Hence, identifying the coordinates y_i and y'_i , we see $J_i = J_i^{\text{trop}}$, i.e., Statement 5.19 holds.

Conversely, suppose Statement 5.19 holds. We just need to show that the conclusions of Proposition 2.45 hold. Note that we already know (1): this is part of the construction of the Frobenius manifold structure on $\widetilde{\mathcal{M}}$. Then Statement 5.19 shows (2). Since

$$E_A = y_0 \partial_{y_0} + 3\partial_{y_1} - y_2 \partial_{y_2},$$

then, with $y_2 = u_1 + \dots + u_k$,

$$E = y_0 \partial_{y_0} + 3\partial_{y_1} - \sum_{i=1}^k u_i \partial_{u_i} = y_0 \partial_{y_0} + 3\partial_{y_1} - y_2 \partial_{y_2}$$

coincides with E_A . This gives (3). Finally, for (4), we argued in the proof of Proposition 2.43, (2), that $((-)^*(\hbar^{-3\alpha} \alpha^j), \hbar^{-3\alpha} \alpha^k)$ was single-valued, and since the dependence on \hbar was a function of $\log \hbar$, must be independent of \hbar . In particular, by using $\hbar^{-3\alpha} \alpha^j = (1 - 3\alpha \log \hbar + (9/2)(\alpha \log \hbar)^2) \alpha^j$, we see

$$((-)^* \hbar^{-3\alpha} \alpha^0, \hbar^{-3\alpha} \alpha^2) = ((-)^* \alpha^0, \alpha^2).$$

Recall that $1, \alpha, \alpha^2$ is the dual basis to Ξ_0, Ξ_1, Ξ_2 . By the explicit forms for the Ξ_i in Remark 2.42 and the intersection numbers in Example 2.34 we have

$$((-)^* \Xi_0, \Xi_2) = -\frac{1}{(2\pi i)^2},$$

so

$$((-)^* \alpha^0, \alpha^2) = -(2\pi i)^2,$$

and hence

$$(\hbar^{-3\alpha} \alpha^0, \hbar^{-3\alpha} \alpha^2)_{\mathcal{E}} = \hbar^{-2}.$$

Applying (2.39) twice and (2.40), we see that

$$(\hbar^{-3\alpha}\alpha^i, \hbar^{-3\alpha}\alpha^j)_\mathcal{E} = \begin{cases} 0 & i+j > 2 \\ (-1)^i \hbar^{-2} & i+j = 2 \\ (\text{Constant})\hbar^{-2} & i+j < 2. \end{cases}$$

Thus

$$\begin{aligned} (\hbar^{-3\alpha}(\hbar\alpha)^i, \hbar^{-3\alpha}(\hbar\alpha)^j)_{\hbar\mathcal{H}_-/\mathcal{H}_-} &= ((-\hbar)^i \hbar^j (\hbar^{-3\alpha}\alpha^i, \hbar^{-3\alpha}\alpha^j)_\mathcal{E})|_{\hbar=\infty} \\ &= \delta_{i,2-j} \\ &= \int_{\mathbb{P}^2} T_i \cup T_j. \end{aligned}$$

This gives (4) of Proposition 2.45, hence Theorem 2.44. \square

The proofs of Theorems 5.15 and 5.18 will be given in §5.5. These proofs require explicit evaluation of the integrals $\int_{\Xi_i} e^{W_k(Q)/\hbar}\Omega$. While it is not, in general, difficult to get explicit answers for these integrals, it is actually quite difficult to extract a useful combinatorial result from these answers. The calculation for Ξ_0 is, however, quite easy, and to see it carried out explicitly, the reader may look at Proposition 3.13 of [42]. However, the argument even for Ξ_0 will be subsumed in the more general calculation, so we shall not repeat this here.

5.4. Deforming Q and P_1, \dots, P_k

Our ultimate goal in this chapter is to compute the integrals $\int_{\Xi_i} e^{W_k(Q)/\hbar}\Omega$. These integrals involve the explicit choice of Q , as well as the implicit choices of P_1, \dots, P_k . We expect the integrals, however, to be independent of these choices. We need to prove this, and in order to prove this, we need explicit formulas for how $W_k(Q)$ changes as Q or P_1, \dots, P_k varies. In fact, as we shall see, there is a “wall-crossing formula” saying that $W_k(Q)$ changes by an automorphism of the ring in which $W_k(Q)$ lives. These wall-crossings occur when certain moduli spaces of disks are not the correct dimension. These wall-crossings seem to be vital to the theory, and will reoccur under a different disguise in the next chapter.

To motivate the discussion, we will begin with an informal exploration of the role that Maslov index zero disks play.

5.4.1. Crossing Maslov index zero disks. For this discussion, we shall fix P_1, \dots, P_k general and let Q vary. By Lemma 5.6, there are no Maslov index zero disks with boundary Q a general point, as the moduli space of such disks is of dimension -1 . On the other hand, the moduli space of Maslov index zero disks without specified boundary is one-dimensional, so we expect these to exist.

Note that the one-dimensionality of this moduli space arises as follows. If we have a Maslov index zero disk $g : \Gamma' \rightarrow M_{\mathbb{R}}$ with boundary some point Q , with outgoing edge E_{out} , we can simply move Q in the direction parallel to $g(E_{\text{out}})$ so that $g(E_{\text{out}})$ gets shorter or longer. As a consequence, it makes sense to keep track of the entire one-parameter family of Maslov index zero disks obtained in this way by extending $g(E_{\text{out}})$ so it is unbounded; see Figure 7. The left-hand figure shows a Maslov index zero tropical disk, and Q can be varied in a one-parameter family, along the ray labelled as $g(E_{\text{out}})$ in the right-hand picture. Here $g : \Gamma \rightarrow M_{\mathbb{R}}$ is a tropical curve obtained from the tropical disk on the left by extending the outgoing




FIGURE 7

edge of the disk indefinitely. We still call this unbounded edge E_{out} . Note while this is a tropical curve, it is not a curve in \mathbb{P}^2 , because $g(E_{\text{out}})$ is not necessarily parallel to a ray in the fan Σ for \mathbb{P}^2 . However, any $Q \in g(E_{\text{out}})$ is the boundary of a Maslov index zero disk. We shall call such extensions of tropical disks *tropical trees*: we make this precise in §5.4.3.

Now suppose that we have such a tropical tree g , yielding a one-parameter family of Maslov index zero disks. Consider two points Q_1, Q_2 on either side of $g(E_{\text{out}})$. Suppose we have a Maslov index two disk with boundary Q_1 . What might happen to this disk as we move from Q_1 to Q_2 ?

We will describe the possibilities here at an intuitive level; precise results will be given in §5.4.5.

Start with a Maslov index two disk h with boundary Q_1 and outgoing edge $E_{\text{out},h}$. Let $m \in M$ be a primitive tangent vector to $h(E_{\text{out},h})$ pointing away from Q_1 . There are three possibilities:

Case 1. The tangent vector m may point away from $g(E_{\text{out}})$. We can then try to deform the Maslov index two disk so that its boundary moves continuously from Q_1 to Q_2 . In general, as we shall see, it is possible to do this. This gives a Maslov index two disk with boundary Q_2 . But there is another thing we can do: after deforming the disk to one with boundary Q' on $h(E_{\text{out}})$, we obtain a new kind of Maslov index two disk by taking the Maslov index zero disk with boundary at Q' and gluing it to the Maslov index two disk with boundary Q' . One then adds a new outgoing edge, with weight and direction determined by the balancing condition; see Figure 8. Thus the single Maslov index two disk with boundary Q_1 gives rise to two Maslov index two disks with boundary Q_2 : this is the fundamental feature of the wall-crossing phenomenon which we shall study here.

Let h_1, h_2 be the two Maslov index two disks with boundary Q_2 constructed in this fashion, with h_2 obtained via the gluing procedure. Then clearly $\text{Mono}(h_1) = \text{Mono}(h)$, but h_2 has one additional vertex V over and above the vertices coming from h and g , whose multiplicity is

$$w(E_{\text{out}})w(E_{\text{out},h})|m \wedge m'|,$$

where m' is a primitive tangent vector to $g(E_{\text{out}})$. Note that in terms of the degree $\Delta(h)$ of the disk h , we in fact have $w(E_{\text{out},h})m = r(\Delta(h))$, as can be seen simply by summing the balancing conditions at each vertex. Similarly, if we write $\Delta(g)$ for the degree of the Maslov index zero disks which yield g , then we can write




FIGURE 8

$w(E_{out})m' = r(\Delta(g))$. Thus

$$\begin{aligned}
 (5.6) \quad \text{Mono}(h_1) + \text{Mono}(h_2) &= \text{Mono}(h)(1 + |r(\Delta(g)) \wedge r(\Delta(h))| \text{Mono}(g)) \\
 &= \text{Mono}(h)(1 + \text{Mono}(g))^{|r(\Delta(g)) \wedge r(\Delta(h))|}.
 \end{aligned}$$

The second equality follows from the fact that $u_{I(g)}^2 = 0$.

Case 2. The second case is that m points towards $g(E_{out})$. In this case, we can reverse the role of Q_1 and Q_2 , so that by using the same argument, $W_k(Q_1)$ contains two monomials $\text{Mono}(h_1) + \text{Mono}(h_2)$ coming from a term $\text{Mono}(h)$ in $W_k(Q_2)$. We can rewrite (5.6), in this case, as

$$(5.7) \quad \text{Mono}(h) = (\text{Mono}(h_1) + \text{Mono}(h_2))(1 + \text{Mono}(g))^{-|r(\Delta(g)) \wedge r(\Delta(h))|}.$$

Case 3. If m is parallel to $g(E_{out})$, then the expectation is that it is impossible to glue on a copy of the Maslov index zero disk as we cross $g(E_{out})$, so that there is no change of contribution to $W_k(Q)$ as Q crosses $g(E_{out})$ from such an h .

We shall prove rigorously in §5.4.5 that the above discussion is accurate; a priori, it is not clear that Maslov index two disks can be deformed at will. But let us consider the implications of the formulas (5.6) and (5.7). Define an element n_g

of $N = \text{Hom}_{\mathbb{Z}}(M, \mathbb{Z})$ by

$$m \mapsto \begin{cases} |r(\Delta(g)) \wedge m| & \text{if } m \text{ points away from } g(E_{\text{out}}), \\ -|r(\Delta(g)) \wedge m| & \text{if } m \text{ points towards } g(E_{\text{out}}), \\ 0 & \text{if } m \text{ is parallel to } g(E_{\text{out}}). \end{cases}$$

It is easy to see that this is linear: in fact, this can be written as a composition

$$M \rightarrow \bigwedge^2 M \xrightarrow{\cong} \mathbb{Z}$$

where the first map is

$$m \mapsto r(\Delta(g)) \wedge m$$

and the second is a choice of one of two isomorphisms so that $r(\Delta(g)) \wedge m$ is mapped to a positive integer if m points away from $g(E_{\text{out}})$.

We then define an automorphism θ_g of the ring $\mathbb{C}[T_{\Sigma}] \otimes_{\mathbb{C}} R_k[[y_0]]$ by

$$\theta_g(z^m) = z^m(1 + \text{Mono}(g))^{(n_g, r(m))}$$

for $m \in T_{\Sigma}$, and

$$\theta_g(a) = a$$

for $a \in R_k[[y_0]]$. Note that θ_g has inverse given by

$$\theta_g^{-1}(z^m) = z^m(1 + \text{Mono}(g))^{-(n_g, r(m))},$$

hence is an automorphism. Then putting (5.6) and (5.7) together, along with the discussion of the third case, we get

$$(5.8) \quad W_k(Q_2) = \theta_g(W_k(Q_1)).$$

We have not derived a rigorous proof of this yet, nor indeed a rigorous statement. We shall do this completely in Theorem 5.35. The main point to absorb now is that Maslov index zero disks are responsible for changes in $W_k(Q)$ as Q varies, and $W_k(Q)$ changes in a very controlled way which can be expressed using automorphisms of $\mathbb{C}[T_{\Sigma}] \otimes_{\mathbb{C}} R_k[[y_0]]$ determined by the Maslov index zero disks.

We shall find a similar, but more complex, behaviour as the points P_1, \dots, P_k move.

We shall now explore these ideas more rigorously, beginning by describing the group of automorphisms that arise in wall crossings.

5.4.2. Automorphisms. Our goal now is to specify a subgroup of the group of automorphisms of $\mathbb{C}[T_{\Sigma}] \otimes_{\mathbb{C}} R_k[[y_0]]$. These automorphisms will, in particular, be *symplectomorphisms*, in the sense that they preserve the relative holomorphic two-form Ω defined in (5.3). However, we will not consider all symplectomorphisms, but rather a subgroup with special properties, a variant of a group first defined by Kontsevich and Soibelman in [70]. These properties will not be explored thoroughly until the next chapter: the discussion here can be viewed as a warm-up.

In this subsection, we have $M \cong \mathbb{Z}^2$, and Σ an arbitrary complete fan in $M_{\mathbb{R}}$ defining a non-singular toric surface. We have the standard exact sequence

$$0 \rightarrow K_{\Sigma} \rightarrow T_{\Sigma} \xrightarrow{r} M \rightarrow 0.$$

To begin, we define the *module of log derivations* of $\mathbb{C}[T_{\Sigma}] \otimes_{\mathbb{C}} R_k[[y_0]]$ to be the module

$$\Theta(\mathbb{C}[T_{\Sigma}] \otimes_{\mathbb{C}} R_k[[y_0]]) := \text{Hom}_{\mathbb{Z}}(M, \mathbb{C}[T_{\Sigma}] \otimes_{\mathbb{C}} R_k[[y_0]]) = (\mathbb{C}[T_{\Sigma}] \otimes_{\mathbb{C}} R_k[[y_0]]) \otimes_{\mathbb{Z}} N.$$

An element $f \otimes n$ is written as $f\partial_n$, and acts as a derivation on $\mathbb{C}[T_\Sigma] \otimes_{\mathbb{C}} R_k[[y_0]]$ over $\mathbb{C}[K_\Sigma] \otimes_{\mathbb{C}} R_k[[y_0]]$ via

$$f\partial_n(z^m) = f\langle n, r(m) \rangle z^m.$$

Given $\xi \in \mathfrak{m}_{R_k} \Theta(\mathbb{C}[T_\Sigma] \otimes_{\mathbb{C}} R_k[[y_0]])$, where $\mathfrak{m}_{R_k} = (u_1, \dots, u_k)$ is the maximal ideal of R_k , we define

$$\exp(\xi) \in \text{Aut}(\mathbb{C}[T_\Sigma] \otimes_{\mathbb{C}} R_k[[y_0]])$$

by

$$\exp(\xi)(a) = a + \sum_{i=1}^{\infty} \frac{\xi^i(a)}{i!}.$$

This is a finite sum given the assumption on ξ , as $\mathfrak{m}_{R_k}^{k+1} = 0$. It is easy to check that $\exp(\xi)$ is indeed a ring automorphism.

For a general log derivation, it can be hard to describe this automorphism. Here, however, is a simple example: Take $m \in T_\Sigma$ with $r(m) \neq 0$, $n \in N$ with $\langle n, r(m) \rangle = 0$, $I \subseteq \{1, \dots, k\}$ non-empty, and $c \in \mathbb{C}$. Then

$$(5.9) \quad \exp(cu_I z^m \partial_n)(z^{m'}) = z^{m'} + cu_I z^m \partial_n z^{m'} = z^{m'}(1 + cu_I \langle n, r(m') \rangle z^m).$$

This is precisely the form of the automorphism θ_g appearing in §5.4.1.

There is a natural Lie bracket defined on the module of log derivations:

$$\begin{aligned} [z^m \partial_n, z^{m'} \partial_{n'}] &:= z^m \partial_n(z^{m'}) \partial_{n'} - z^{m'} \partial_{n'}(z^m) \partial_n \\ &= z^{m+m'} (\langle n, r(m') \rangle \partial_{n'} - \langle n', r(m) \rangle \partial_n). \end{aligned}$$

This is just the ordinary Lie bracket

Now let

$$\mathfrak{v}_{\Sigma, k} := \bigoplus_{\substack{m \in T_\Sigma \\ r(m) \neq 0}} z^m (\mathfrak{m}_{R_k} \otimes r(m)^\perp) \subseteq \Theta(\mathbb{C}[T_\Sigma] \otimes_{\mathbb{C}} R_k[[y_0]]).$$

In other words, $\mathfrak{v}_{\Sigma, k}$ consists of log derivations which are \mathfrak{m}_{R_k} -linear combinations of log derivations of the form $z^m \partial_n$ where $r(m) \neq 0$ and $\langle n, r(m) \rangle = 0$.

Here is the punch-line of this discussion: $\mathfrak{v}_{\Sigma, k}$ is closed under the Lie bracket. Indeed, given m, n and m', n' satisfying the above two requirements, we note that

$$[z^m \partial_n, z^{m'} \partial_{n'}] = z^{m+m'} \partial_{\langle n, r(m') \rangle n' - \langle n', r(m) \rangle n}.$$

Now

$$\begin{aligned} &\langle \langle n, r(m') \rangle n' - \langle n', r(m) \rangle n, r(m+m') \rangle \\ &= \langle n, r(m') \rangle \langle n', r(m) \rangle - \langle n', r(m) \rangle \langle n, r(m') \rangle \\ &= 0 \end{aligned}$$

because $\langle n, r(m) \rangle = \langle n', r(m') \rangle = 0$ by assumption. On the other hand, in the case that $r(m+m') = 0$, then $r(m) = -r(m')$ and $\langle n, r(m') \rangle = \langle n', r(m) \rangle = 0$. Thus in this case $[z^m \partial_n, z^{m'} \partial_{n'}] = 0$ anyway. This demonstrates that $\mathfrak{v}_{\Sigma, k}$ is a Lie subalgebra of the module of log derivations.

Now set

$$\mathbb{V}_{\Sigma, k} = \{\exp(\xi) \mid \xi \in \mathfrak{v}_{\Sigma, k}\}.$$

As a consequence of the observation that $\mathfrak{v}_{\Sigma, k}$ is closed under Lie bracket, so that $\mathfrak{v}_{\Sigma, k}$ is a Lie algebra, $\mathbb{V}_{\Sigma, k}$ is a group, with multiplication given by the Baker-Campbell-Hausdorff formula. Note that as $\mathfrak{v}_{\Sigma, k}$ is generated by derivations $cu_I z^m \partial_n$

with $r(m) \neq 0$, $\langle n, r(m) \rangle = 0$, $\mathbb{V}_{\Sigma, k}$ is generated by automorphisms of the form $\exp(cu_I z^m \partial_n)$, as described in (5.9).

It is not difficult to check that $\exp(cu_I z^m \partial_n)$ preserves Ω ; while this can be done directly, it will follow shortly anyway. In fact, the original version of this group introduced in [70] was defined as a group of Hamiltonian symplectomorphisms.

This description is as follows. We fix the relative holomorphic two-form Ω , as in (5.3). This is determined by a choice of generator of $\bigwedge^2 M$, i.e., an identification $\bigwedge^2 M \cong \mathbb{Z}$. Given an element $f \in \mathbb{C}[T_\Sigma] \otimes_{\mathbb{C}} R_k[[y_0]]$, we denote by X_f the corresponding Hamiltonian vector field with respect to the symplectic form Ω . In other words, this is the vector field X_f such that $\iota(X_f)\Omega = df$. Here, we are considering the relative differential with respect to the map κ . So df is an element of the dual to the module of log derivations, i.e., an element of $(\mathbb{C}[T_\Sigma] \otimes_{\mathbb{C}} R_k[[y_0]]) \otimes_{\mathbb{Z}} M$. We write $f \otimes m$ as $f \operatorname{dlog} m$, and $d(z^m) = z^m \operatorname{dlog} r(m)$.

To describe X_{z^m} , we introduce an identification of N with M induced by the chosen identification $\bigwedge^2 M \cong \mathbb{Z}$ given by Ω . An element $m \in M$ induces an element $X_m \in N$ given by the composition

$$\begin{aligned} X_m : M &\rightarrow \bigwedge^2 M \xrightarrow{\cong} \mathbb{Z} \\ m' &\mapsto m \wedge m' \end{aligned}$$

So if $f = z^m$, then X_f is the vector field $-z^m X_{r(m)}$. Indeed, writing $\Omega = \operatorname{dlog} e_1 \wedge \operatorname{dlog} e_2$ for e_1, e_2 an oriented basis of M , and $r(m) = a_1 e_1 + a_2 e_2$, we see

$$\begin{aligned} -z^m \iota(X_{r(m)})\Omega &= -z^m (\langle X_{r(m)}, e_1 \rangle \operatorname{dlog} e_2 - \langle X_{r(m)}, e_2 \rangle \operatorname{dlog} e_1) \\ &= -z^m (-a_2 \operatorname{dlog} e_2 - a_1 \operatorname{dlog} e_1) \\ &= z^m \operatorname{dlog} r(m) \\ &= d(z^m). \end{aligned}$$

Note that $\langle X_m, m \rangle = 0$ trivially. Thus one sees that in fact $\mathfrak{v}_{\Sigma, k}$ can be described as the set of Hamiltonian vector fields associated to functions f of the form $\sum_i c_i u_{I_i} z^{m_i}$, where $c_i \in \mathbb{C}$, $I_i \subseteq \{1, \dots, k\}$ is a non-empty index set, and $m_i \in T_\Sigma$ with $r(m_i) \neq 0$. Of course, the exponential of a Hamiltonian vector field preserves Ω , and hence $\mathbb{V}_{\Sigma, k}$ is a group of symplectomorphisms.

For future use, we have the following standard lemma:

LEMMA 5.21. *If $f \in \mathfrak{m}_{R_k}(\mathbb{C}[T_\Sigma] \otimes_{\mathbb{C}} R_k[[y_0]])$ and $\theta \in \mathbb{V}_{\Sigma, k}$, then*

$$\theta \circ X_f \circ \theta^{-1} = X_{\theta(f)}.$$

PROOF. We can assume θ is a generator of $\mathbb{V}_{\Sigma, k}$, i.e.,

$$\theta(z^{m''}) = z^{m''} (1 + r(m) \wedge r(m'')) cu_I z^m$$

for some $m \in T_\Sigma$, $r(m) \neq 0$. By linearity, we can assume that $f = z^{m'}$ for some m' . Then computing, we see

$$\begin{aligned}
& (\theta \circ X_f \circ \theta^{-1})(z^{m''}) \\
&= -(\theta \circ (z^{m'} X_{r(m')})) (z^{m''} - cu_I r(m) \wedge r(m'') z^{m+m''}) \\
&= -\theta(r(m') \wedge r(m'')) z^{m'+m''} \\
&\quad - cu_I[r(m') \wedge (r(m) + r(m''))][r(m) \wedge r(m'')] z^{m+m'+m''} \\
&= -r(m') \wedge r(m'') z^{m'+m''} \\
&\quad - cu_I[r(m') \wedge r(m'')][r(m) \wedge (r(m') + r(m''))] z^{m+m'+m''} \\
&\quad + cu_I[r(m') \wedge (r(m) + r(m''))][r(m) \wedge r(m'')] z^{m+m'+m''} \\
&= -r(m') \wedge r(m'') z^{m'+m''} \\
&\quad - cu_I[r(m) \wedge r(m')][r(m + m') \wedge r(m'')] z^{m+m'+m''}.
\end{aligned}$$

On the other hand,

$$\begin{aligned}
X_{\theta(f)}(z^{m''}) &= X_{z^{m'} + r(m) \wedge r(m') cu_I z^{m+m'}}(z^{m''}) \\
&= -r(m') \wedge r(m'') z^{m'+m''} \\
&\quad - cu_I[r(m) \wedge r(m')][r(m + m') \wedge r(m'')] z^{m+m'+m''}.
\end{aligned}$$

□

5.4.3. Scattering diagrams. We shall now introduce a pictorial method, which we call the method of *scattering diagrams*, of keeping track of calculations in the group $\mathbb{V}_{\Sigma,k}$. As usual, we fix $M \cong \mathbb{Z}^2$, Σ a complete fan in $M_{\mathbb{R}}$ defining a non-singular toric surface.

DEFINITION 5.22. Fix $k \geq 0$.

(1) A *ray* or *line* is a pair $(\mathfrak{d}, f_{\mathfrak{d}})$ such that

- $\mathfrak{d} \subseteq M_{\mathbb{R}}$ is given by

$$\mathfrak{d} = m'_0 - \mathbb{R}_{\geq 0} r(m_0)$$

if \mathfrak{d} is a ray and

$$\mathfrak{d} = m'_0 - \mathbb{R} r(m_0)$$

if \mathfrak{d} is a line, for some $m'_0 \in M_{\mathbb{R}}$ and $m_0 \in T_\Sigma$ with $r(m_0) \neq 0$. The set \mathfrak{d} is called the *support* of the line or ray. If \mathfrak{d} is a ray, m'_0 is called the *initial point* of the ray, written as $\text{Init}(\mathfrak{d})$.

- $f_{\mathfrak{d}} \in \mathbb{C}[z^{m_0}] \otimes_{\mathbb{C}} R_k \subseteq \mathbb{C}[T_\Sigma] \otimes_{\mathbb{C}} R_k[[y_0]]$.
- $f_{\mathfrak{d}} \equiv 1 \pmod{(u_1, \dots, u_k) z^{m_0}}$.

(2) A *scattering diagram* \mathfrak{D} is a finite collection of lines and rays.

If \mathfrak{D} is a scattering diagram, we write

$$\text{Supp}(\mathfrak{D}) := \bigcup_{\mathfrak{d} \in \mathfrak{D}} \mathfrak{d} \subseteq M_{\mathbb{R}}$$

and

$$\text{Sing}(\mathfrak{D}) := \bigcup_{\mathfrak{d} \in \mathfrak{D}} \partial \mathfrak{d} \cup \bigcup_{\substack{\mathfrak{d}_1, \mathfrak{d}_2 \\ \dim \mathfrak{d}_1 \cap \mathfrak{d}_2 = 0}} \mathfrak{d}_1 \cap \mathfrak{d}_2.$$




FIGURE 9. The path on the left is homotopic to the path on the right in $M_{\mathbb{R}} \setminus \text{Sing}(\mathfrak{D})$. The automorphism $\theta_{\gamma, \mathfrak{D}}$ is the same on the left and right because the double crossing of the same line cancels.

Here $\partial \mathfrak{d} = \{\text{Init}(\mathfrak{d})\}$ if \mathfrak{d} is a ray, and is empty if \mathfrak{d} is a line.

CONSTRUCTION 5.23. Given a smooth immersion $\gamma : [0, 1] \rightarrow M_{\mathbb{R}} \setminus \text{Sing}(\mathfrak{D})$ with endpoints not in $\text{Supp}(\mathfrak{D})$, such that γ intersects elements of \mathfrak{D} transversally, we can define a ring automorphism $\theta_{\gamma, \mathfrak{D}} \in \mathbb{V}_{\Sigma, k}$, the γ -ordered product of \mathfrak{D} . Explicitly, we can find numbers

$$0 < t_1 \leq t_2 \leq \dots \leq t_s < 1$$

and elements $\mathfrak{d}_i \in \mathfrak{D}$ such that $\gamma(t_i) \in \mathfrak{d}_i$ and $\mathfrak{d}_i \neq \mathfrak{d}_j$ if $t_i = t_j$, $i \neq j$, and s taken as large as possible. Then for each i , define $\theta_{\gamma, \mathfrak{d}_i} \in \mathbb{V}_{\Sigma, k}$ to be

$$\begin{aligned} \theta_{\gamma, \mathfrak{d}_i}(z^m) &= z^m f_{\mathfrak{d}_i}^{\langle n_0, r(m) \rangle} \\ \theta_{\gamma, \mathfrak{d}_i}(a) &= a \end{aligned}$$

for $m \in T_{\Sigma}$, $a \in R_k[[y_0]]$, where $n_0 \in N$ is chosen to be primitive, annihilates the tangent space to \mathfrak{d}_i , and is finally completely determined by the sign convention that

$$\langle n_0, \gamma'(t_i) \rangle < 0.$$

We then define

$$\theta_{\gamma, \mathfrak{D}} = \theta_{\gamma, \mathfrak{d}_s} \circ \dots \circ \theta_{\gamma, \mathfrak{d}_1}.$$

Note there is still some ambiguity to the ordering if γ crosses several overlapping rays. However, an easy check shows that automorphisms associated to parallel rays commute, so the order is irrelevant. Automorphisms associated with non-parallel rays do not necessarily commute, hence the need for γ to avoid points of $\text{Sing}(\mathfrak{D})$.

We will also allow the possibility that γ is piecewise linear so that γ' may not be defined at t_i . In this case, we insist that if the path γ intersects \mathfrak{d} , then γ passes from one side of \mathfrak{d} to the other. We then take n_0 so that γ passes from the side of \mathfrak{d} where n_0 is larger to the side where it is smaller.

It is easy to check that $\theta_{\gamma, \mathfrak{D}}$ only depends on the homotopy class of the path γ inside $M_{\mathbb{R}} \setminus \text{Sing}(\mathfrak{D})$. The key point is the simple observation that if γ is a path which crosses a ray in \mathfrak{D} twice consecutively and in opposite directions, then the automorphisms induced by the two crossings are inverse to each other, and hence cancel, see Figure 9.

EXAMPLE 5.24. Let $\mathfrak{D} = \{(\mathfrak{d}_1, f_{\mathfrak{d}_1}), (\mathfrak{d}_2, f_{\mathfrak{d}_2}), (\mathfrak{d}_3, f_{\mathfrak{d}_3})\}$ with

$$\begin{aligned} \mathfrak{d}_1 &= \mathbb{R}r(m_1), & f_{\mathfrak{d}_1} &= 1 + c_1 w_1 z^{m_1}, \\ \mathfrak{d}_2 &= \mathbb{R}r(m_2), & f_{\mathfrak{d}_2} &= 1 + c_2 w_2 z^{m_2}, \\ \mathfrak{d}_3 &= -\mathbb{R}_{\geq 0}(r(m_1 + m_2)), & f_{\mathfrak{d}_3} &= 1 + c_1 c_2 w_{\text{out}} |r(m_1) \wedge r(m_2)| z^{m_1+m_2}, \end{aligned}$$

where $m_1, m_2 \in T_\Sigma \setminus K_\Sigma$, and w_1, w_2 and w_{out} are the indices of $r(m_1)$, $r(m_2)$ and $r(m_1 + m_2)$ respectively. Suppose $c_1, c_2 \in R_k$ satisfy $c_1^2 = c_2^2 = 0$. Then we can calculate $\theta_{\gamma, \mathfrak{D}}$ if γ is a loop around the origin, as in Figure 10, where $\theta_{\gamma, \mathfrak{D}} = \theta_2^{-1} \theta_3 \theta_1^{-1} \theta_2 \theta_1$, with θ_i the automorphism coming from the first crossing of \mathfrak{d}_i . Note that if we use the standard orientation on \mathbb{R}^2 , then to define θ_i , for $i = 1, 2$, we can take $n_i = \frac{-1}{w_i} X_{r(m_i)}$: this is a cotangent vector which annihilates the tangent space to \mathfrak{d}_i and is negative on γ' when γ crosses \mathfrak{d}_1 and \mathfrak{d}_2 for the first time. On the other hand, when γ crosses \mathfrak{d}_3 , θ_3 is defined using

$$n_3 = \frac{1}{w_{\text{out}}} X_{r(m_1)+r(m_2)}.$$

Thus applying successively $\theta_1, \theta_2, \theta_1^{-1}, \theta_3$ and θ_2^{-1} , we get, keeping in mind that $c_1^2 = c_2^2 = 0$,

$$\begin{aligned} z^m &\mapsto z^m - c_1 \langle X_{r(m_1)}, r(m) \rangle z^{m+m_1} \\ &\mapsto z^m - c_2 \langle X_{r(m_2)}, r(m) \rangle z^{m+m_2} - c_1 \langle X_{r(m_1)}, r(m) \rangle z^{m+m_1} \\ &\quad + c_1 c_2 \langle X_{r(m_1)}, r(m) \rangle \langle X_{r(m_2)}, r(m) + r(m_1) \rangle z^{m+m_1+m_2} \\ &\mapsto z^m - c_2 \langle X_{r(m_2)}, r(m) \rangle z^{m+m_2} \\ &\quad + c_1 c_2 (- \langle X_{r(m_2)}, r(m) \rangle \langle X_{r(m_1)}, r(m) + r(m_2) \rangle \\ &\quad + \langle X_{r(m_1)}, r(m) \rangle \langle X_{r(m_2)}, r(m) + r(m_1) \rangle) z^{m+m_1+m_2} \\ &\mapsto z^m - c_2 \langle X_{r(m_2)}, r(m) \rangle z^{m+m_2} \\ &\quad + c_1 c_2 (\langle X_{r(m_1)}, r(m) \rangle \langle X_{r(m_2)}, r(m_1) \rangle \\ &\quad - \langle X_{r(m_1)}, r(m_2) \rangle \langle X_{r(m_2)}, r(m) \rangle \\ &\quad + \langle X_{r(m_1)+r(m_2)}, r(m) \rangle |r(m_1) \wedge r(m_2)|) z^{m+m_1+m_2} \\ &\mapsto z^m + c_1 c_2 (\langle X_{r(m_1)}, r(m) \rangle \langle X_{r(m_2)}, r(m_1) \rangle \\ &\quad - \langle X_{r(m_1)}, r(m_2) \rangle \langle X_{r(m_2)}, r(m) \rangle \\ &\quad + \langle X_{r(m_1)+r(m_2)}, r(m) \rangle |r(m_1) \wedge r(m_2)|) z^{m+m_1+m_2}. \end{aligned}$$

Note that from our choice of orientation, $|r(m_1) \wedge r(m_2)| = \langle X_{r(m_1)}, r(m_2) \rangle = -\langle X_{r(m_2)}, r(m_1) \rangle$, from which it follows that the coefficient of $z^{m+m_1+m_2}$ above is zero, so that $\theta_2^{-1} \theta_3 \theta_1^{-1} \theta_2 \theta_1$ is the identity. \square

We shall now construct a scattering diagram from tropical trees, essentially encoding all Maslov index zero disks.

DEFINITION 5.25. A *tropical tree* in $(X_\Sigma, P_1, \dots, P_k)$ is a d -pointed tropical curve $h : (\Gamma, p_1, \dots, p_d) \rightarrow M_{\mathbb{R}}$ with $h(p_j) = P_{i_j}$ $1 \leq i_1 < \dots < i_d \leq k$, along with the additional data of a choice of an unmarked unbounded edge $E_{\text{out}} \in \Gamma_{\infty}^{[1]}$ such that for any $E \in \Gamma_{\infty}^{[1]} \setminus \{E_{\text{out}}\}$, $h(E)$ is a point or is a translate of some $\rho \in \Sigma^{[1]}$. The degree of h , $\Delta(h)$, is defined *without* counting the edge $h(E_{\text{out}})$, which need not be a translate of any $\rho \in \Sigma^{[1]}$.




FIGURE 10

The Maslov index of h is

$$MI(h) := 2(|\Delta(h)| - d).$$

Given h and a point V_{out} in the interior of E_{out} , we can remove the unbounded component of $E_{\text{out}} \setminus \{V_{\text{out}}\}$ from Γ to obtain Γ' . Note V_{out} is a univalent vertex of Γ' . Take $h' : \Gamma' \rightarrow M_{\mathbb{R}}$ with $h' = h|_{\Gamma'}$. Then h' is a tropical disk with boundary $h(V_{\text{out}})$ and Maslov index $MI(h') = MI(h)$, since $|\Delta(h')| = |\Delta(h)|$.

Conversely, given a tropical disk $h' : \Gamma' \rightarrow M_{\mathbb{R}}$, one can extend E_{out} to an unbounded edge and obtain a tropical tree $h : \Gamma \rightarrow M_{\mathbb{R}}$.

As in Lemma 5.6, standard tropical dimension counting arguments show that, for general choice of P_1, \dots, P_k , a tropical tree h moves in a family of dimension $MI(h)/2$. In particular, the set of Maslov index zero trees is a finite set, which we denote by $\text{Trees}(\Sigma, P_1, \dots, P_k)$. As usual, with general choice of P_1, \dots, P_k , we can assume all these trees are trivalent.

DEFINITION 5.26. We define $\mathfrak{D}(\Sigma, P_1, \dots, P_k)$ to be the scattering diagram which contains one ray for each element h of $\text{Trees}(\Sigma, P_1, \dots, P_k)$. The ray corresponding to h is of the form $(\mathfrak{d}, f_{\mathfrak{d}})$, where

- $\mathfrak{d} = h(E_{\text{out}})$.
- $f_{\mathfrak{d}} = 1 + w_{\Gamma}(E_{\text{out}}) \text{Mult}(h) z^{\Delta(h)} u_{I(h)}$, where $u_{I(h)} = \prod_{i \in I(h)} u_i$ and $I(h) \subseteq \{1, \dots, k\}$ is defined by

$$I(h) := \{i \mid h(p_j) = P_i \text{ for some } j\}.$$


EXAMPLE 5.27. $\mathfrak{D}(\Sigma, P_1, P_2)$ is illustrated in Figure 11, where Σ is the fan for \mathbb{P}^2 .

PROPOSITION 5.28. Let P_1, \dots, P_k be generally chosen. If

$$P \in \text{Sing}(\mathfrak{D}(\Sigma, P_1, \dots, P_k))$$

is a singular point with $P \notin \{P_1, \dots, P_k\}$, and γ_P is a small loop around P , then $\theta_{\gamma_P, \mathfrak{D}(\Sigma, P_1, \dots, P_k)} = \text{Id}$.

PROOF. Let P be such a singular point. Suppose that $\mathfrak{d} \in \mathfrak{D}(\Sigma, P_1, \dots, P_k)$ has $\text{Init}(\mathfrak{d}) = P$, and let h be the corresponding tree. Then the unique vertex V of Γ on E_{out} satisfies $h(V) = P$. Note that no edge with vertex V is contracted by h since $P \notin \{P_1, \dots, P_k\}$. In addition, by the generality assumption, V is trivalent,

FIGURE 11. The scattering diagram $\mathfrak{D}(\Sigma, P_1, P_2)$.

so if we cut Γ at V , we obtain two tropical disks $h'_1 : \Gamma'_1 \rightarrow M_{\mathbb{R}}$ and $h'_2 : \Gamma'_2 \rightarrow M_{\mathbb{R}}$ with boundary P and $V_{\text{out}} = V$ in both cases. Now $MI(h) = MI(h'_1) + MI(h'_2)$, so $MI(h'_1) = MI(h'_2) = 0$ is the only possibility. Thus h'_1, h'_2 extend to tropical trees $h_i : \Gamma_i \rightarrow M_{\mathbb{R}}$, with corresponding rays $\mathfrak{d}_1, \mathfrak{d}_2$. Note that $P \neq \text{Init}(\mathfrak{d}_1), \text{Init}(\mathfrak{d}_2)$ and $I(h_1) \cap I(h_2) = \emptyset$. So every ray \mathfrak{d} with $P = \text{Init}(\mathfrak{d})$ arises from the collision of two rays $\mathfrak{d}_1, \mathfrak{d}_2$ with $P \neq \text{Init}(\mathfrak{d}_i)$.

Conversely, if we are given two such rays $\mathfrak{d}_1, \mathfrak{d}_2$ passing through P corresponding to trees h_1 and h_2 with $I(h_1) \cap I(h_2) = \emptyset$, we obtain a new tree by cutting h_1 and h_2 at P to get Maslov index zero disks $h'_i : \Gamma'_i \rightarrow M_{\mathbb{R}}$ with boundary at P . Next glue Γ'_1 and Γ'_2 at the outgoing vertex V , and add an additional unbounded edge E_{out} with endpoint V to get a graph Γ . If $E_{\text{out},1}, E_{\text{out},2}$ are the two outgoing edges of Γ'_1 and Γ'_2 respectively, with primitive tangent vector to $h'_i(E_{\text{out},i})$ pointing away from P being m'_i , then we define $h : \Gamma \rightarrow M_{\mathbb{R}}$ to restrict to h'_i on Γ'_i and to take E_{out} to the ray $P - \mathbb{R}_{\geq 0}(w_{\Gamma}(E_{\text{out},1})m'_1 + w_{\Gamma}(E_{\text{out},2})m'_2)$. By taking $w_{\Gamma}(E_{\text{out}})$ to be the index of $w_{\Gamma}(E_{\text{out},1})m'_1 + w_{\Gamma}(E_{\text{out},2})m'_2$, we find that h is balanced at V . Thus h is a tropical tree, whose Maslov index is zero.

After making these observations, to prove the proposition, define a new scattering diagram \mathfrak{D}_P , whose elements are in one-to-one correspondence with elements of $\mathfrak{D}(\Sigma, P_1, \dots, P_k)$ containing P . If $(\mathfrak{d}, f_{\mathfrak{d}}) \in \mathfrak{D}(\Sigma, P_1, \dots, P_k)$ is a ray containing P then the corresponding element of \mathfrak{D}_P will be $(\mathfrak{d}', f_{\mathfrak{d}})$, where \mathfrak{d}' is the tangent line (through the origin) of \mathfrak{d} if $P \neq \text{Init}(\mathfrak{d})$ and is the ray $\mathfrak{d} - P$ with endpoint the origin otherwise. If γ_0 is a loop around the origin with the same orientation as γ_P , then $\theta_{\gamma_0, \mathfrak{D}_P} = \theta_{\gamma_P, \mathfrak{D}(\Sigma, P_1, \dots, P_k)}$.

First consider the simplest case, when \mathfrak{D}_P contains two lines and at most one ray. If the two lines correspond to trees h_1 and h_2 , and $I(h_1) \cap I(h_2) \neq \emptyset$, then h_1 and h_2 cannot be glued as above since they pass through some common marked point P_i . Thus \mathfrak{D}_P contains no rays. In this case, the automorphisms associated

to \mathfrak{d}_1 and \mathfrak{d}_2 commute by Example 5.24 since $u_{I(h_1)}u_{I(h_2)} = 0$, and so $\theta_{\gamma_0, \mathfrak{D}_P}$ is the identity.

If, on the other hand, $I(h_1) \cap I(h_2) = \emptyset$, then h_1 and h_2 can be glued to obtain a new tree h , and \mathfrak{D}_P consists of three elements \mathfrak{d}_1 , \mathfrak{d}_2 and \mathfrak{d} , corresponding to h_1, h_2 and h respectively. Now

$$f_{\mathfrak{d}_i} = 1 + w_\Gamma(E_{\text{out}, i}) \text{Mult}(h_i) z^{\Delta(h_i)} u_{I(h_i)}$$

for $i = 1, 2$ and

$$\begin{aligned} f_{\mathfrak{d}} &= 1 + w_\Gamma(E_{\text{out}}) \text{Mult}(h) z^{\Delta(h)} u_{I(h)} \\ &= 1 + w_\Gamma(E_{\text{out}}) \text{Mult}(h_1) \text{Mult}(h_2) \text{Mult}_V(h) z^{\Delta(h_1) + \Delta(h_2)} u_{I(h_1)} u_{I(h_2)} \\ &= 1 + w_\Gamma(E_{\text{out}}) \text{Mult}(h_1) \text{Mult}(h_2) w_{\Gamma_1}(E_{\text{out}, 1}) w_{\Gamma_2}(E_{\text{out}, 2}) \\ &\quad \cdot |m'_1 \wedge m'_2| z^{\Delta(h_1) + \Delta(h_2)} u_{I(h_1)} u_{I(h_2)} \\ &= 1 + w_\Gamma(E_{\text{out}}) \text{Mult}(h_1) \text{Mult}(h_2) \\ &\quad \cdot |r(\Delta(h_1)) \wedge r(\Delta(h_2))| z^{\Delta(h_1) + \Delta(h_2)} u_{I(h_1)} u_{I(h_2)}. \end{aligned}$$

Here we are using the fact that if we sum the balancing condition over all vertices of Γ'_i , we get

$$w_{\Gamma'_i}(E_{\text{out}, i}) m'_i = r(\Delta(h_i)).$$

Thus from Example 5.24, $\theta_{\gamma_0, \mathfrak{D}_P}$ is the identity.

For the general case, we have some finite set of lines in \mathfrak{D}_P , along with some rays. Suppose that there are three lines in \mathfrak{D}_P corresponding to trees h_1, h_2 and h_3 with $I(h_1), I(h_2)$ and $I(h_3)$ mutually disjoint. Then as in the case of two lines above, these trees can be glued at P , obtaining a Maslov index zero tree with a quadrivalent vertex. However, since P_1, \dots, P_k are in general position, no Maslov index zero tree has a vertex with valence > 3 . Thus this case does not occur. On the other hand, given two lines corresponding to trees h_1, h_2 with $I(h_1) \cap I(h_2) = \emptyset$, these two trees can be glued as above at P to obtain a new Maslov index zero tree. Thus the rays in \mathfrak{D}_P are in one-to-one correspondence with pairs of lines $\mathfrak{d}_1, \mathfrak{d}_2 \in \mathfrak{D}_P$ corresponding to trees h_1 and h_2 with $I(h_1) \cap I(h_2) = \emptyset$. So we can write

$$\mathfrak{D}_P = \{\mathfrak{d}_1, \dots, \mathfrak{d}_n\} \cup \bigcup_{j=1}^m \mathfrak{D}_j$$

where $\mathfrak{d}_1, \dots, \mathfrak{d}_n$ are lines corresponding to trees h such that $I(h) \cap I(h') \neq \emptyset$ for any Maslov index zero tree h' with outgoing edge passing through P , and $\mathfrak{D}_1, \dots, \mathfrak{D}_m$ are scattering diagrams each consisting of two lines and one ray, with the lines corresponding to trees h_1 and h_2 with $I(h_1) \cap I(h_2) = \emptyset$ and the ray corresponding to the tree obtained by gluing h_1 and h_2 at P .

Now computing $\theta_{\gamma_0, \mathfrak{D}_P}$ is an exercise in commutators. Note that if $\mathfrak{d}_1, \mathfrak{d}_2 \in \mathfrak{D}_P$ correspond to two trees h_1, h_2 with $I(h_1) \cap I(h_2) \neq \emptyset$, then as already observed, $\theta_{\gamma_0, \mathfrak{d}_1}$ and $\theta_{\gamma_0, \mathfrak{d}_2}$ commute. Thus, after using this commutation, one can write

$$\theta_{\gamma_0, \mathfrak{D}_P} = \left(\prod_{i=1}^n \theta_{\gamma_0, \mathfrak{d}_i} \circ \theta_{\gamma_0, \mathfrak{d}_i}^{-1} \right) \circ \prod_{j=1}^m \theta_{\gamma_0, \mathfrak{D}_j}.$$

Of course $\theta_{\gamma_0, \mathfrak{d}_i} \circ \theta_{\gamma_0, \mathfrak{d}_i}^{-1} = \text{Id}$ and $\theta_{\gamma_0, \mathfrak{D}_j} = \text{Id}$ by the special case already carried out. Thus $\theta_{\gamma_0, \mathfrak{D}_P} = \text{Id}$ in this general case. \square

REMARK 5.29. Note that the rays in $\mathfrak{D} = \mathfrak{D}(\Sigma, P_1, \dots, P_k)$ with endpoint P_i are in one-to-one correspondence with Maslov index two disks in

$$(X_\Sigma, P_1, \dots, P_{i-1}, P_{i+1}, \dots, P_k)$$

with boundary P_i . Indeed, taking any such Maslov index two disk, extending the outgoing edge to get a tropical tree, we can mark the point on this outgoing edge which maps to P_i , thus getting a tropical tree with Maslov index zero in $(X_\Sigma, P_1, \dots, P_k)$. The corresponding ray in \mathfrak{D} has endpoint P_i . Conversely, given a ray in \mathfrak{D} with endpoint P_i , this corresponds to a Maslov index zero tree such that the vertex V adjacent to E_{out} is the vertex of a marked edge E_x mapping to P_i . By cutting this tree at V and removing the marked edge mapping to P_i , we get a Maslov index two disk with boundary P_i .

Furthermore, by the general position of the P_j , there are no rays in \mathfrak{D} containing P_i but which don't have P_i as an endpoint.

5.4.4. Broken lines. One benefit of this scattering diagram approach is that it is easy to describe the Maslov index two disks with boundary a general point Q , using what we call *broken lines*:

DEFINITION 5.30. A *broken line* is a continuous proper piecewise linear map

$$\beta : (-\infty, 0] \rightarrow M_{\mathbb{R}}$$

with endpoint $Q = \beta(0)$, along with some additional data described as follows. Let

$$-\infty = t_0 < t_1 < \dots < t_n = 0$$

be the smallest set of real numbers such that $\beta|_{(t_{i-1}, t_i)}$ is linear. Then for each $1 \leq i \leq n$, we are given the additional data of a monomial $c_i z^{m_i^\beta} \in \mathbb{C}[T_\Sigma] \otimes_{\mathbb{C}} R_k[[y_0]]$ with $m_i^\beta \in T_\Sigma \setminus K_\Sigma$ and $0 \neq c_i \in R_k$. Furthermore, this data should satisfy the following properties:

- (1) For each i , $r(m_i^\beta) = -\beta'(t)$ for $t \in (t_{i-1}, t_i)$.
- (2) $m_1^\beta = t_\rho$ for some $\rho \in \Sigma^{[1]}$ and $c_1 = 1$.
- (3) $\beta(t_i) \in \text{Supp}(\mathfrak{D}(\Sigma, P_1, \dots, P_k)) \setminus \text{Sing}(\mathfrak{D}(\Sigma, P_1, \dots, P_k))$ for $1 \leq i \leq n$.
- (4) If $\beta(t_i) \in \mathfrak{d}_1 \cap \dots \cap \mathfrak{d}_s$ (necessarily this intersection is one-dimensional), then $c_{i+1} z^{m_{i+1}^\beta}$ is a term in

$$(\theta_{\beta, \mathfrak{d}_1} \circ \dots \circ \theta_{\beta, \mathfrak{d}_s})(c_i z^{m_i^\beta}).$$

By this, we mean the following. Suppose $f_{\mathfrak{d}_j} = 1 + c_{\mathfrak{d}_j} z^{m_{\mathfrak{d}_j}}$, $1 \leq j \leq s$, with $c_{\mathfrak{d}_j}^2 = 0$, and $n \in N$ is primitive, orthogonal to all the \mathfrak{d}_j 's, chosen so that

$$\begin{aligned} (\theta_{\beta, \mathfrak{d}_1} \circ \dots \circ \theta_{\beta, \mathfrak{d}_s})(c_i z^{m_i^\beta}) &= c_i z^{m_i^\beta} \prod_{j=1}^s (1 + c_{\mathfrak{d}_j} z^{m_{\mathfrak{d}_j}})^{\langle n, r(m_i^\beta) \rangle} \\ &= c_i z^{m_i^\beta} \prod_{j=1}^s (1 + \langle n, r(m_i^\beta) \rangle c_{\mathfrak{d}_j} z^{m_{\mathfrak{d}_j}}). \end{aligned}$$

Then we must have

$$c_{i+1} z^{m_{i+1}^\beta} = c_i z^{m_i^\beta} \prod_{j \in J} \langle n, r(m_i^\beta) \rangle c_{\mathfrak{d}_j} z^{m_{\mathfrak{d}_j}}$$




FIGURE 12. The broken lines ending at Q .

for some index set $J \subseteq \{1, \dots, s\}$. We think of this as β being bent at time t_i by the rays $\{\mathfrak{D}_j \mid j \in J\}$.

EXAMPLE 5.31. Again, in the case of \mathbb{P}^2 , $k = 2$, Figure 12 shows the broken lines with $\beta(0)$ the given point Q . The segments of each broken line are labelled with their corresponding monomial.

PROPOSITION 5.32. *If $Q \notin \text{Supp}(\mathfrak{D}(\Sigma, P_1, \dots, P_k))$ is general, then there is a one-to-one correspondence between broken lines with endpoint Q and Maslov index two disks with boundary Q . In addition, if β is a broken line corresponding to a disk h , and cz^m is the monomial associated to the last segment of β (the one whose endpoint is $t_n = 0$), then*

$$cz^m = \text{Mono}(h).$$

PROOF. We first prove the following claim:

Claim: Let $h : \Gamma' \rightarrow M_{\mathbb{R}}$ be a Maslov index two disk in $(X_{\Sigma}, P_1, \dots, P_k)$ with boundary $Q' \in M_{\mathbb{R}}$. Suppose furthermore that all vertices of Γ' except V_{out} are trivalent and h cannot be deformed continuously in a family of Maslov index two disks with boundary Q' . Then there is a uniquely determined subset $\Xi = \Xi(h) \subseteq \overline{\Gamma}$ which is a union of edges of $\overline{\Gamma}$ and is homeomorphic to $[-\infty, 0]$, connecting some point in $\overline{\Gamma}_{\infty}^{[0]} \setminus \{V_{\text{out}}\}$ to V_{out} , satisfying:

- (1) Ξ is disjoint from ∂E_{p_i} for all i .
- (2) The restriction of h to the closure of any connected component of $\Gamma' \setminus \Xi$ is a Maslov index zero disk.

PROOF. We proceed inductively on the number of vertices of Γ' . If Γ' has only one vertex, V_{out} , then Γ' has only one edge and no marked edges. We simply take $\Xi(h)$ to consist of this edge.

For the induction step, let Γ' have outgoing edge E_{out} with vertices V_{out} and V , and $h(V_{\text{out}}) = Q'$.

First, we will show that V cannot be a vertex of some marked edge E_{p_i} . If it were, so that $h(V) = h(p_i) = P_j$ for some j , then we can cut h at V and remove the marked edge E_{p_i} . This gives a disk h' with boundary P_j but with one less marked point than h . Hence h' is Maslov index four, and thus by Lemma 5.6, h' can be deformed in a one-parameter family while keeping the endpoint P_j fixed. Note that for small deformations of h' , the edge $h'(E'_{\text{out}})$ does not change its tangent direction. Thus a deformation of h' can be extended to a deformation of h . This contradicts the assumption that h cannot be deformed. Thus V cannot be a vertex of some E_{p_i} .

Now split h at V , letting Γ'_1 and Γ'_2 be the closures of the two connected components of $\Gamma' \setminus \{V\}$ not containing V_{out} . Let $h_i = h|_{\Gamma'_i}$. This gives two disks h_1, h_2 with boundary $h(V)$. We have $MI(h) = MI(h_1) + MI(h_2)$. Suppose $MI(h_1) \geq 4$. Then h_1 can be deformed leaving the endpoint $h(V)$ fixed, and by gluing such a deformation to h_2 , we obtain a deformation of h , again a contradiction. Thus $MI(h_1), MI(h_2) \leq 2$, so we must have $MI(h_1) = 0$ and $MI(h_2) = 2$ or vice versa.

Without loss of generality, assume $MI(h_2) = 2$. Note that h_2 is now a Maslov index two disk with boundary $Q'' = h_2(V)$. If h_2 could be deformed in a family of disks with boundary Q'' , then by gluing these deformations to h_1 , we obtain a deformation of h , a contradiction. Thus h_2 satisfies the hypotheses of the Claim, but Γ'_2 has fewer vertices than Γ' . So $\Xi(h_2)$ exists by the induction hypothesis, and we can take $\Xi(h)$ to be

$$\Xi(h) = E_{\text{out}} \cup \Xi(h_2).$$

$\Xi(h)$ satisfies the two desired properties because $\Xi(h_2)$ does, E_{out} is disjoint from ∂E_{p_i} for all i , and h_1 is a Maslov index zero disk. \square

Now fix a Maslov index two disk $h : \Gamma' \rightarrow M_{\mathbb{R}}$ with boundary Q . By the generality of Q, P_1, \dots, P_k , h satisfies the hypotheses of the Claim. Taking $\beta = h|_{\Xi(h)}$, we see that β is piecewise linear. Let $-\infty = t_0 < \dots < t_n = 0$ be chosen as in the definition of broken line. Each t_i corresponds to a vertex V_i of Γ' . Of course $\Gamma' \setminus \{V_i\}$ for $i \neq n$ has two connected components not containing V_{out} , and the proof of the claim shows that restricting h to the closure of one of these two connected components yields a Maslov index two disk with boundary $h(V_i)$ which we now call h_i . The other component similarly yields a Maslov index zero disk. Hence $\beta(t_i) \in \text{Supp}(\mathfrak{D}(\Sigma, P_1, \dots, P_k))$ for $1 \leq i \leq n$. We take the monomial $c_i z^{m_i}$ to be $\text{Mono}(h_i)$, and need to check that with this data β is now a broken line.

We have just shown condition (3), and condition (2) is obvious. So for Q sufficiently general, we only need to verify conditions (1) and (4). We need to show that the monomial $\text{Mono}(h_{i+1})$ attached to the edge joining V_i and V_{i+1} arises from the monomial $\text{Mono}(h_i)$ attached to the edge joining V_{i-1} and V_i , as in Condition (4). Suppose that the two subtrees in Γ' rooted at V_i are g and h_i , with $MI(g) = 0$. Now

$$\text{Mono}(h_i) = \text{Mult}(h_i) z^{\Delta(h_i)} u_{I(h_i)},$$

and if \mathfrak{d} is the ray corresponding to the tropical tree obtained from g , then

$$f_{\mathfrak{d}} = 1 + w_{\text{out}}(g) \text{Mult}(g) z^{\Delta(g)} u_{I(g)}.$$

Here $w_{\text{out}}(g)$ denotes the weight of the outgoing edge of g . By summing the balancing condition at each vertex of either of the two disks, we can write $-w_{\text{out}}(g)m'_1 = r(\Delta(g))$, and $-w_{\text{out}}(h_i)m'_2 = r(\Delta(h_i))$ with $m'_i \in M$ primitive, with m'_1 tangent

to the outgoing edge of g and m'_2 tangent to the outgoing edge of h_i , both pointing towards $h(V_i)$. In particular, this shows condition (1), assuming the correct parameterization of $\Xi(h)$ has been chosen.

Now the automorphism $\theta_{\beta, \mathfrak{d}}$ can be defined by taking $n_0 = \pm X_{m'_1}$, with the choice of sign being given by the requirement that $\langle n_0, -m'_2 \rangle > 0$, as the convention on n_0 says that n_0 should be negative on vectors pointing in the direction we cross \mathfrak{d} ; but m'_2 is such a vector so n_0 is positive on $-m'_2$. Thus $\langle n_0, -m'_2 \rangle = |m'_1 \wedge m'_2|$, and

$$\begin{aligned} & \theta_{\beta, \mathfrak{d}}(\text{Mono}(h_i)) \\ &= \text{Mult}(h_i) z^{\Delta(h_i)} u_{I(h_i)} \\ & \quad + \text{Mult}(h_i) \text{Mult}(g) |m'_1 \wedge m'_2| w_{\text{out}}(g) w_{\text{out}}(h_i) z^{\Delta(h_i) + \Delta(g)} u_{I(h_i)} u_{I(g)}. \end{aligned}$$

Now it is the second term we are interested in, and this is

$$\begin{aligned} \text{Mult}(h_i) \text{Mult}(g) \text{Mult}_{V_i}(h_{i+1}) z^{\Delta(h_{i+1})} u_{I(h_{i+1})} &= \text{Mult}(h_{i+1}) z^{\Delta(h_{i+1})} u_{I(h_{i+1})} \\ &= \text{Mono}(h_{i+1}) \end{aligned}$$

as desired.

Conversely, given a broken line β , it is easy to construct the corresponding Maslov index two disk, by attaching Maslov index zero disks to the domain $(-\infty, 0]$ of β at each bending point. In particular, if $\beta(t_i)$ lies in rays $\mathfrak{d}_1, \dots, \mathfrak{d}_s \in \mathfrak{D}(\Sigma, P_1, \dots, P_k)$, and β is bent at time t_i by a subset $\{\mathfrak{d}_j \mid j \in J\}$ of these rays, then for each $j \in J$ we attach the Maslov index zero disk with endpoint $\beta(t_i)$ corresponding to \mathfrak{d}_j to $t_i \in (-\infty, 0]$. (Note that by general position of the P_i 's and Q , in fact we can assume that $\#J = 1$.) It is clear that this reverses the above process of passing from a Maslov index two disk to a broken line. \square

5.4.5. Wall-crossing for Q . We will now prove a wall-crossing formula as Q varies inside $M_{\mathbb{R}}$. To do so, we will understand how broken lines with endpoint Q vary as Q varies. For this purpose, we introduce the notion of a deformation of a broken line and a degenerate broken line:

DEFINITION 5.33. A *family of broken lines* consists of the data:

- A continuous map $B : (-\infty, 0] \times I \rightarrow M_{\mathbb{R}}$ with $I \subseteq \mathbb{R}$ an interval.
- Continuous functions $t_0, \dots, t_n : I \rightarrow [-\infty, 0]$ such that $-\infty = t_0(s) < t_1(s) < \dots < t_n(s) = 0$ for $s \in I$.
- Monomials $c_i z^{m_i^B}$, $1 \leq i \leq n$.

This data satisfies the condition that $B_s := B|_{(-\infty, 0] \times \{s\}}$ is a broken line in the usual sense for all $s \in I$, with the data $t_0(s) < \dots < t_n(s)$ and monomials $c_i z^{m_i^B}$, $1 \leq i \leq n$.

We say $B_{s'}$ is a *deformation* of B_s for $s, s' \in I$.

DEFINITION 5.34. A *degenerate broken line* is a limit of broken lines which bends at a point of $\text{Sing}(\mathfrak{D}(\Sigma, P_1, \dots, P_k))$. More precisely, a degenerate broken line is a continuous proper map $\beta : (-\infty, 0] \rightarrow M_{\mathbb{R}}$ along with data $-\infty = t_0 < t_1 \leq \dots \leq t_n = 0$ and monomials $c_i z^{m_i^{\beta}}$ such that there is:

- A continuous map $B : (-\infty, 0] \times [0, 1] \rightarrow M_{\mathbb{R}}$.
- Continuous functions $\bar{t}_0, \dots, \bar{t}_n : [0, 1] \rightarrow [-\infty, 0]$ such that

$$-\infty = \bar{t}_0(s) \leq \bar{t}_1(s) \leq \dots \leq \bar{t}_n(s) = 0$$

for $s \in [0, 1]$, with strict inequality for $s < 1$.

This data satisfies

- $B_s := B|_{(-\infty, 0] \times \{s\}}$ (with the data $c_i z^{m_i}$) is a broken line in the usual sense for $s < 1$;
- $\beta = B_1$ and $\bar{t}_i(1) = t_i$ for all i ;
- $\beta(t_i) \in \text{Sing}(\mathfrak{D}(\Sigma, P_1, \dots, P_k))$ for some i .

Note that in taking such a limit, we might have \bar{t}_{i-1} and \bar{t}_i coming together for various i , so the limit might have fewer linear segments.

We now have the wall-crossing theorem:

THEOREM 5.35. *If $Q, Q' \in M_{\mathbb{R}} \setminus \text{Supp}(\mathfrak{D}(\Sigma, P_1, \dots, P_k))$ are general, and γ is a path connecting Q and Q' for which $\theta_{\gamma, \mathfrak{D}(\Sigma, P_1, \dots, P_k)}$ is defined, then*

$$W_k(Q') = \theta_{\gamma, \mathfrak{D}(\Sigma, P_1, \dots, P_k)}(W_k(Q)).$$

PROOF. Let $\mathfrak{D} = \mathfrak{D}(\Sigma, P_1, \dots, P_k)$. Let \mathfrak{U} be the union of $\text{Supp}(\mathfrak{D})$ and the union of images of all degenerate broken lines, with arbitrary endpoint. It is clear that $\dim \mathfrak{U} \leq 1$ (of course equal to 1 provided $k \geq 1$).

Note that a broken line β can always be deformed continuously. This can be done as follows. We translate the initial ray $\beta((-\infty, t_1])$ of β . Inductively, this deforms all the remaining segments of β . As long as one of the bending points does not reach a singular point of \mathfrak{D} , each bending point remains inside exactly the same set of rays in \mathfrak{D} , and therefore the deformed broken line can bend in exactly the same way as β . Thus we run into trouble building this deformation only when this deformation of β converges to a degenerate broken line bending at a singular point of $\text{Sing}(\mathfrak{D})$, as then the set of rays containing a bending point may jump.

From this it is clear that as long as the endpoint of β stays within one connected component of $M_{\mathbb{R}} \setminus \mathfrak{U}$, β can be deformed continuously. More precisely, if we consider a path $\gamma : [0, 1] \rightarrow \mathfrak{u}$, for \mathfrak{u} a connected component of $M_{\mathbb{R}} \setminus \mathfrak{U}$, and if β is a broken line with endpoint $\gamma(0)$, then there is a continuous deformation B with $\beta = B_0$ and with $B_s(0) = \gamma(s)$, $0 \leq s \leq 1$.

By Proposition 5.32, the Maslov index two disks with boundary Q are in one-to-one correspondence with the broken lines with endpoint Q for Q general. Thus, by the above discussion, $W_k(Q)$ is constant for general Q inside a connected component of $M_{\mathbb{R}} \setminus \mathfrak{U}$.

We will now analyze carefully how broken lines change if their endpoint passes in between different connected components of $M_{\mathbb{R}} \setminus \mathfrak{U}$. So now consider two connected components \mathfrak{u}_1 and \mathfrak{u}_2 of $M_{\mathbb{R}} \setminus \mathfrak{U}$. Let $L = \overline{\mathfrak{u}_1} \cap \overline{\mathfrak{u}_2}$, and assume $\dim L = 1$. Let Q_1 and Q_2 be general points in \mathfrak{u}_1 and \mathfrak{u}_2 , near L , positioned on opposite sides of L . Let $\gamma : [0, 1] \rightarrow M_{\mathbb{R}}$ be a short general path connecting Q_1 and Q_2 crossing L precisely once. Let s_0 be the only time at which $\gamma(s_0) \in L$. By choosing γ sufficiently generally, we can assume that $\gamma(s_0)$ is a point in a neighbourhood of which \mathfrak{U} is a manifold.

Let $\mathfrak{B}(Q_i)$ be the set of broken lines with endpoint Q_i . Let $n_0 \in N$ be a primitive vector annihilating the tangent space to L and taking a smaller value on Q_1 than Q_2 . We can decompose $\mathfrak{B}(Q_i)$ into three sets $\mathfrak{B}^+(Q_i)$, $\mathfrak{B}^-(Q_i)$, and $\mathfrak{B}^0(Q_i)$ as follows. For $\beta \in \mathfrak{B}(Q_i)$, let $m_\beta = \beta_*(-\partial/\partial t|_{t=0})$. Then $\beta \in \mathfrak{B}^+(Q_i)$, $\mathfrak{B}^-(Q_i)$, or $\mathfrak{B}^0(Q_i)$ depending on whether $\langle n_0, m_\beta \rangle > 0$, $\langle n_0, m_\beta \rangle < 0$, or $\langle n_0, m_\beta \rangle = 0$. This

gives decompositions

$$\begin{aligned} W_k(Q_1) &= W_k^-(Q_1) + W_k^0(Q_1) + W_k^+(Q_1) \\ W_k(Q_2) &= W_k^-(Q_2) + W_k^0(Q_2) + W_k^+(Q_2). \end{aligned}$$

We will show

$$(5.10) \quad \theta_{\gamma, \mathfrak{D}}(W_k^-(Q_1)) = W_k^-(Q_2),$$

$$(5.11) \quad \theta_{\gamma, \mathfrak{D}}^{-1}(W_k^+(Q_2)) = W_k^+(Q_1),$$

$$(5.12) \quad W_k^0(Q_2) = W_k^0(Q_1).$$

From this follows the desired identity

$$\theta_{\gamma, \mathfrak{D}}(W_k(Q_1)) = W_k(Q_2),$$

as $\theta_{\gamma, \mathfrak{D}}$ is necessarily the identity on $W_k^0(Q_1)$. One then uses this inductively to see that this holds for any path γ with endpoints in $M_{\mathbb{R}} \setminus \mathfrak{U}$ for which $\theta_{\gamma, \mathfrak{D}}$ is defined.

Proof of (5.10) and (5.11). If β is a broken line with endpoint Q_1 , then $\beta([t_{n-1}, 0]) \cap L = \emptyset$ if $\langle n_0, m_\beta \rangle \leq 0$, while $\beta([t_{n-1}, 0]) \cap L \neq \emptyset$ if $\langle n_0, m_\beta \rangle > 0$. (Here we are using Q_1 very close to L .) On the other hand, if β has endpoint Q_2 , then $\beta([t_{n-1}, 0]) \cap L = \emptyset$ if $\langle n_0, m_\beta \rangle \geq 0$ and $\beta([t_{n-1}, 0]) \cap L \neq \emptyset$ if $\langle n_0, m_\beta \rangle < 0$.

To see, say, (5.10), we proceed as follows. Let $\beta \in \mathfrak{B}^-(Q_1)$. By the previous paragraph, $\beta([t_{n-1}, 0]) \cap L = \emptyset$. Let $c_n z^{m_n^\beta}$ be the monomial associated to the last segment of β , and write $\theta_{\gamma, \mathfrak{D}}(c_n z^{m_n^\beta})$ as a sum of monomials $\sum_{i=1}^s d_i z^{m_i}$ as in Definition 5.30, (4). We can then deform β continuously along γ to time s_0 . Indeed, by the definition of \mathfrak{U} , if β converged to a degenerate broken line through $\text{Sing}(\mathfrak{D})$, the image of this broken line would be contained in \mathfrak{U} , and then \mathfrak{U} , already containing L , would not be a manifold in a neighbourhood of $\gamma(s_0)$.

Let β' be the deformation of β with endpoint $\gamma(s_0)$. For $1 \leq i \leq s$, we then get a broken line β'_i by adding a short line segment to β' in the direction $-r(m_i)$, with attached monomial $d_i z^{m_i}$. This new broken line has endpoint in \mathfrak{u}_2 , and hence can be deformed to a broken line $\beta''_i \in \mathfrak{B}^-(Q_2)$. We note that the line may not actually bend at L if $d_i z^{m_i}$ is the term $c_n z^{m_n^\beta}$ appearing in $\theta_{\gamma, \mathfrak{D}}(c_n z^{m_n^\beta})$. See Figure 13.

Conversely, any broken line $\beta \in \mathfrak{B}^-(Q_2)$ clearly arises in this way.

From this, (5.10) becomes clear. (5.11) is identical. \square

Proof of (5.12). We will show that there are partitions $\mathfrak{B}^0(Q_1) = \coprod_{i=1}^s \mathfrak{B}_i^1$ and $\mathfrak{B}^0(Q_2) = \coprod_{i=1}^s \mathfrak{B}_i^2$ such that for each i , the contributions to $W_k(Q_1)$ and $W_k(Q_2)$ from \mathfrak{B}_i^1 and \mathfrak{B}_i^2 are the same.

For simplicity of exposition, we will describe this in the case that any degenerate broken line with endpoint $\gamma(s_0)$ passes through at most one point of $\text{Sing}(\mathfrak{D})$; we leave it to the reader to deal with the general case: this is notationally, but not conceptually, more complicated.

Let $\beta_1 \in \mathfrak{B}^0(Q_1)$. If β_1 deforms continuously to a broken line β_2 in $\mathfrak{B}^0(Q_2)$, then β_1 and β_2 will each appear in one-element sets in the partition, say $\beta_1 \in \mathfrak{B}_i^1$, $\beta_2 \in \mathfrak{B}_i^2$, and clearly both these sets contribute the same term to $W_k(Q_1)$ and $W_k(Q_2)$.

Now suppose $\beta_1 \in \mathfrak{B}^0(Q_1)$ cannot be deformed continuously to any $\beta_2 \in \mathfrak{B}^0(Q_2)$. This means that β_1 must deform to a degenerate broken line at time s_0 : i.e., there is a $B : (-\infty, 0] \times [0, s_0] \rightarrow M_{\mathbb{R}}$ as in Definition 5.34 such that $B|_{(-\infty, 0] \times [0, s_0]}$ is a continuous deformation and B_{s_0} is a degenerate broken line




FIGURE 13. Broken lines with endpoints Q_1 and Q_2 .

bending at a point P of $\text{Sing}(\mathfrak{D})$. So there must be some j such that $B(\bar{t}_j(s_0), s_0) = P \in \text{Sing}(\mathfrak{D})$.

There are two cases we need to analyze: either $P \in \{P_1, \dots, P_k\}$ or $P \notin \{P_1, \dots, P_k\}$.

Case 1. $P = P_i$ for some i . Because all rays with endpoint P_i involve the monomial u_i , a broken line can only bend along at most one ray with endpoint P_i , and, as observed above, β_1 bends along at least one such ray. So call the ray with endpoint P_i along which β_1 bends \mathfrak{d}_1 , corresponding to a Maslov index zero tree $\tilde{h}_1 : \tilde{\Gamma}_1 \rightarrow M_{\mathbb{R}}$. This tree passes through P_i , and by cutting this tree at P_i and removing the marked edge mapping to P_i , we obtain a Maslov index two disk $h_1 : \Gamma'_1 \rightarrow M_{\mathbb{R}}$ with boundary P_i . Let $\bar{\beta}_2$ be the broken line with endpoint P_i corresponding to this Maslov index 2 disk. See Figure 14.

Next, recalling that $B(\bar{t}_j(s_0), s_0) = P_i$, let $\bar{\beta}'_2 : [\bar{t}_j(s_0), 0] \rightarrow M_{\mathbb{R}}$ be the restriction of B to $[\bar{t}_j(s_0), 0] \times \{s_0\}$: this is a piece of a broken line starting at P_i . We can then concatenate $\bar{\beta}_2$ with $\bar{\beta}'_2$ by identifying 0 in the domain of $\bar{\beta}_2$ with $\bar{t}_j(s_0)$ in the domain of $\bar{\beta}'_2$, obtaining what we hope will be a degenerate broken line β'_2 passing through P_i .

Note that $B|_{(-\infty, \bar{t}_j(s_0)] \times \{s_0\}}$ is a broken line with endpoint P_i , and hence corresponds to a Maslov index two disk $h_2 : \Gamma'_2 \rightarrow M_{\mathbb{R}}$ with endpoint P_i . By extending the edge E'_{out} of Γ'_2 to an unbounded edge, we get a tropical tree $\tilde{h}_2 : \tilde{\Gamma}_2 \rightarrow M_{\mathbb{R}}$, and once we mark the point on $\tilde{\Gamma}_2$ which maps to P_i , it becomes a Maslov index zero tree and hence corresponds to a ray $\mathfrak{d}_2 \in \mathfrak{D}$ with endpoint P_i .

Note that the function attached to \mathfrak{d}_i is $1 + w_{\Gamma_i}(E_{\text{out}, i})u_i \text{Mono}(h_i)$. On the other hand, the monomial attached to the last segment of $B|_{(-\infty, \bar{t}_j(s_0)] \times \{s_0\}}$, i.e., $c_j z^{m_j^{\beta_1}}$, is $\text{Mono}(h_2)$, while the monomial attached to the last segment of $\bar{\beta}_2$ is $\text{Mono}(h_1)$. Thus, in particular, the monomial $c_{j+1} z^{m_{j+1}^{\beta_1}}$ is obtained from the bend of β_1 at \mathfrak{d}_1 , and hence is

$$(5.13) \quad w_{\Gamma_1}(E_{\text{out}, 1}) \langle n_1, r(\Delta(h_2)) \rangle u_i \text{Mono}(h_1) \text{Mono}(h_2).$$




FIGURE 14. The change in a broken line as it passes through a singular point $P \in \{P_1, \dots, P_k\}$.

Here $n_1 \in N$ is primitive, orthogonal to \mathfrak{d}_1 , and positive on $r(\Delta(h_2))$.

We can now deform β'_2 by moving the endpoint of $\bar{\beta}_2$ along \mathfrak{d}_2 away from P_i , moving $\bar{\beta}'_2$ along with it. However, we also need to keep track of monomials: we have to make sure that the monomial on the first segment of $\bar{\beta}'_2$ is the one which would arise when β'_2 bends along \mathfrak{d}_2 . However, this latter monomial is a term obtained by applying the automorphism associated to crossing \mathfrak{d}_2 to $\text{Mono}(h_1)$, and is thus precisely

$$(5.14) \quad w_{\Gamma_2}(E_{\text{out},2}) \langle n_2, r(\Delta(h_1)) \rangle u_i \text{Mono}(h_1) \text{Mono}(h_2).$$

Again, $n_2 \in N$ is primitive, orthogonal to \mathfrak{d}_2 , and positive on $r(\Delta(h_1))$. In fact, for $i = 1, 2$,

$$w_{\Gamma_i}(E_{\text{out},i}) n_i = \pm X_{r(\Delta(h_i))},$$

with the sign chosen so that both (5.13) and (5.14) coincide with

$$|r(\Delta(h_1)) \wedge r(\Delta(h_2))| u_i \text{Mono}(h_1) \text{Mono}(h_2).$$

As a result, β'_2 can now be deformed away from the singular point P_i , giving a broken line β_2 with endpoint Q_2 . Note that in no way does this represent a continuous deformation: the broken line really jumps as it passes through P_i .

Note that this process is reversible. If we start with β_2 and try to deform it through P_i as above, we obtain β_1 .

To conclude, in this case, we can take one-element sets in the partition of the form $\beta_1 \in \mathfrak{B}_i^1$ and $\beta_2 \in \mathfrak{B}_i^2$ for some i . They both give the same contribution to $W_k(Q)$.

Case 2. $P \notin \{P_1, \dots, P_k\}$. Let \mathfrak{D}_P be the scattering diagram constructed in the proof of Proposition 5.28, consisting of a ray with initial point the origin for

each ray in $\mathfrak{D}(\Sigma, P_1, \dots, P_k)$ with initial point P , and a line through the origin for each ray in $\mathfrak{D}(\Sigma, P_1, \dots, P_k)$ through P with initial point not equal to P . By Proposition 5.28,

$$(5.15) \quad \theta_{\gamma, \mathfrak{D}_P} = \text{Id}$$

for γ a loop around P .

Since the behaviour of \mathfrak{D}_P is relatively simple, one can analyze this situation by a careful case-by-case analysis: this was carried out in [42]. Here, we will give a simpler and more general approach, due to Carl, Pumperla and Siebert in work in progress [13].

Define a *trajectory* in $M_{\mathbb{R}}$ to be a pair (\mathfrak{t}, cz^m) where $cz^m \in \mathbb{C}[T_{\Sigma}] \otimes R_k$ is a monomial and either

- (1) $\mathfrak{t} = \mathbb{R}_{\geq 0}r(m)$, in which case we say (\mathfrak{t}, cz^m) is an *incoming trajectory*, or
- (2) $\mathfrak{t} = -\mathbb{R}_{\geq 0}r(m)$, in which case we say (\mathfrak{t}, cz^m) is an *outgoing trajectory*.

As usual, \mathfrak{t} is called the *support* of (\mathfrak{t}, cz^m) .

CLAIM 5.36. *Given an incoming trajectory $\mathfrak{t}_0 = (\mathbb{R}_{\geq 0}r(m_0), c_0z^{m_0})$, there is a set of outgoing trajectories $\{\mathfrak{t}_i\}$ with $\mathfrak{t}_i = (-\mathbb{R}_{\geq 0}r(m_i), c_iz^{m_i})$ with the following property. For each \mathfrak{t}_i (including $i = 0$), let x_i be a point in a connected component of $M_{\mathbb{R}} \setminus \text{Supp}(\mathfrak{D}_P)$ whose closure contains the support of \mathfrak{t}_i . Let γ_i be a path from x_i to x_0 not passing through P ; by (5.15), $\theta_i := \theta_{\gamma_i, \mathfrak{D}_P}$ is independent of this choice of path. Then*

$$(5.16) \quad c_0z^{m_0} = \sum_i \theta_i(c_iz^{m_i}).$$

Furthermore, the set of trajectories is unique in the sense that given a ray $\mathbb{R}_{\geq 0}m \subseteq M_{\mathbb{R}}$, the sum $\sum_i c_iz^{m_i}$ over all $i \neq 0$ with the support of \mathfrak{t}_i being $\mathbb{R}_{\geq 0}m$ is uniquely determined by \mathfrak{t}_0 .

PROOF. First note that (5.16) does not depend on the precise choice of x_i or x_0 : if \mathfrak{t}_i (including the case $i = 0$) is contained in a ray or line of \mathfrak{D}_P , then $c_iz^{m_i}$ is invariant under the corresponding automorphism.

We now construct the set $\{\mathfrak{t}_i\}$ by induction. At the p th step, we will find a set \mathfrak{F}_p of outgoing trajectories for which (5.16) holds modulo the ideal \mathfrak{m}^p , where \mathfrak{m} is the ideal generated by u_1, \dots, u_k . For $p = 1$, we take

$$\mathfrak{F}_1 = \{(-\mathbb{R}_{\geq 0}r(m_0), c_0z^{m_0})\},$$

which works since $\theta_i \equiv \text{Id} \pmod{\mathfrak{m}}$.

Assume now that we have constructed $\mathfrak{F}_p = \{\mathfrak{t}_i \mid i \in I\}$. Then, by the induction hypothesis,

$$c_0z^{m_0} - \sum_i \theta_i(c_iz^{m_i}) = \sum c_jz^{m_j} \pmod{\mathfrak{m}^{p+1}}$$

with $c_j \in \mathfrak{m}^p$. Then take

$$\mathfrak{F}_{p+1} = \mathfrak{F}_p \cup \{(-\mathbb{R}_{\geq 0}r(m_j), c_jz^{m_j})\};$$

(5.16) now holds modulo \mathfrak{m}^{p+1} .

Since $\mathfrak{m}^{k+1} = 0$, this process terminates. It is also clear that at each step there are no choices to be made, hence the uniqueness. \square

The point of this set of outgoing trajectories is that it tells us precisely what terms broken lines coming near P can produce as they bend near P . More precisely,




FIGURE 15. The set \mathfrak{B} of local broken lines with a given initial monomial

consider a kind of local form of broken line: these will be continuous maps $\beta : \mathbb{R} \rightarrow M_{\mathbb{R}}$ with image disjoint from $\text{Sing}(\mathfrak{D}_P)$ along with data

$$-\infty = t_0 < \dots < t_n = \infty$$

and monomials $c_i z^{m_i^\beta}$ as in Definition 5.30, satisfying conditions (1), (3) and (4) of that definition. In other words, we do not restrict the initial monomial, and we have no endpoint. Other than this, β is a broken line. We call this a *local broken line*.

Fix a monomial $c_0 z^{m_0}$ and a general point $x \in M_{\mathbb{R}}$ contained in a connected component of $M_{\mathbb{R}} \setminus \text{Sing}(\mathfrak{D}_P)$ whose closure contains $\mathbb{R}_{\geq 0}r(m_0)$. Consider the (necessarily finite) set \mathfrak{B} of local broken lines β such that the monomial associated with the first segment (t_0, t_1) is $c_0 z^{m_0}$ and there is a $t \in (t_0, t_1)$ with $\beta(t) = x$. (We consider two local broken lines to be the same if they differ just by a time translation). See Figure 15. For $\beta_i \in \mathfrak{B}$, let $c_i z^{m_i}$ be the monomial attached to the last line segment of β_i , and let γ_i be a path from a point $x_i = \beta_i(t)$ for $t \gg 0$ to x . Let $\theta_i = \theta_{\gamma_i, \mathfrak{D}_P}$.

CLAIM 5.37. $\sum_i \theta_i(c_i z^{m_i}) = c_0 z^{m_0}$.

PROOF. Note that if \mathfrak{D}_P just consisted of one line, then the broken lines in \mathfrak{B} would bend at at most one point, the intersection of $x - \mathbb{R}_{\geq 0}r(m_0)$ with the unique element of \mathfrak{D}_P . It then follows immediately from condition (4) of Definition 5.30 that the claim is true in this case.

Using this, the general case then follows easily if we take, for the path γ_i , the segment of β_i running from x_i backwards to x . □

Fix an outgoing direction $-\mathbb{R}_{\geq 0}m$, and consider the sum $\sum_i c_i z^{m_i}$ where i runs over all indices such that $-r(m_i) \in -\mathbb{R}_{\geq 0}m$. As a consequence of the uniqueness statement of Claim 5.36, this sum is well-defined, irrespective of the choice of the point x .

We can now complete the proof in this case. Return to the situation at the beginning of the proof of (5.12): we have $\beta_1 \in \mathfrak{B}^0(Q_1)$ which deforms to a degenerate broken line B_{s_0} passing through P .

We will take a set $\mathfrak{B}_i^1 \subseteq \mathfrak{B}^0(Q_1)$ to be the set of all broken lines in $\mathfrak{B}^0(Q_1)$ which deform to the same limiting map as B_{s_0} (but may involve different monomials). We define \mathfrak{B}_i^2 similarly to be the set of all broken lines in $\mathfrak{B}^0(Q_2)$ which deform to the same limiting map. These two sets give contributions to $W_k(Q_1)$ and $W_k(Q_2)$ respectively; we will be done once we show these contributions are the same.

Recalling that we are deforming the endpoints of the broken lines in $\mathfrak{B}^0(Q_1)$ and $\mathfrak{B}^0(Q_2)$ along a path γ , we define for $s \neq s_0$ the set \mathfrak{B}_s of broken lines with endpoint $\gamma(s)$ which again deform to the limiting map B_{s_0} , so that each broken line in \mathfrak{B}_s for $s < s_0$ is a deformation of a unique broken line in \mathfrak{B}_i^1 and each broken line in \mathfrak{B}_s for $s > s_0$ is a deformation of a unique broken line in \mathfrak{B}_i^2 .

Let $x_0 := B_{s_0}(t)$ for t smaller but close to the \bar{t} such that $B_{s_0}(\bar{t}) = P$. Locally near x_0 , the image of B_{s_0} splits $M_{\mathbb{R}}$ into two connected components, and broken lines in \mathfrak{B}_s for $s < s_0$ are locally contained in one of these connected components, and broken lines in \mathfrak{B}_s for $s > s_0$ are locally contained in the other. Fix points x, x' on either side of the image of B_{s_0} near x_0 , chosen so that some element of \mathfrak{B}_s for some $s < s_0$ passes through, say, x . For each broken line $\beta \in \mathfrak{B}_i^1$, one can find at least one $s < s_0$ such that β can be deformed to a $\beta_s \in \mathfrak{B}_s$ which passes through x . Similarly, for each broken line $\beta \in \mathfrak{B}_i^2$, one can find at least one $s > s_0$ such that β can be deformed to a $\beta_s \in \mathfrak{B}_s$ which passes through x' . For each $\beta \in \mathfrak{B}_i^1$, we make a choice of one of these β_s 's passing through x , and let \mathfrak{B}_x be the set of these choices. We define $\mathfrak{B}_{x'}$ similarly, so that \mathfrak{B}_x is in one-to-one correspondence with \mathfrak{B}_i^1 and $\mathfrak{B}_{x'}$ is in one-to-one correspondence with \mathfrak{B}_i^2 . In particular, if $c_{\beta} z^{m_{\beta}}$ denotes the monomial attached to the last segment of a broken line β , we just need to show that

$$\sum_{\beta \in \mathfrak{B}_x} c_{\beta} z^{m_{\beta}} = \sum_{\beta \in \mathfrak{B}_{x'}} c_{\beta} z^{m_{\beta}}.$$

To show this, it is enough to show the same statement where one replaces $c_{\beta} z^{m_{\beta}}$ with the monomial attached to β right after β finishes bending near P . The result now follows from Claim 5.37 and the uniqueness statement of Claim 5.36. \square

5.4.6. Varying P_1, \dots, P_k . We now turn our attention to studying the dependence of $W_k(Q)$ on the points P_1, \dots, P_k . For this discussion, we will need a three-dimensional version of scattering diagrams.

DEFINITION 5.38. Let $L \subseteq \mathbb{R}$ be a closed interval. Let π_1 and π_2 be the projections of $M_{\mathbb{R}} \times L$ onto $M_{\mathbb{R}}$ and L respectively. A *scattering diagram* in $M_{\mathbb{R}} \times L$ is a finite set \mathfrak{D} consisting of pairs $(\mathfrak{d}, f_{\mathfrak{d}})$ such that

- $\mathfrak{d} \subseteq M_{\mathbb{R}} \times L$ is a polyhedron of dimension two such that $\pi_2(\mathfrak{d})$ is one-dimensional. Furthermore there is a one-dimensional subset $\mathfrak{b} \subseteq M_{\mathbb{R}} \times L$ and an element $m_0 \in T_{\Sigma}$ with $r(m_0) \neq 0$ such that

$$\mathfrak{d} = \mathfrak{b} - \mathbb{R}_{\geq 0}(r(m_0), 0).$$

- $f_{\mathfrak{d}} \in \mathbb{C}[z^{m_0}] \otimes_{\mathbb{C}} R_k \subseteq \mathbb{C}[T_{\Sigma}] \otimes_{\mathbb{C}} R_k[[y_0]]$.
- $f_{\mathfrak{d}} \equiv 1 \pmod{(u_1, \dots, u_k)z^{m_0}}$.

We define

$$\text{Sing}(\mathfrak{D}) = \bigcup_{\mathfrak{d} \in \mathfrak{D}} \partial \mathfrak{d} \cup \bigcup_{\substack{\mathfrak{d}_1, \mathfrak{d}_2 \in \mathfrak{D} \\ \dim \mathfrak{d}_1 \cap \mathfrak{d}_2 = 1}} \mathfrak{d}_1 \cap \mathfrak{d}_2.$$

This is a one-dimensional subset of $M_{\mathbb{R}} \times L$. Let $\text{Interstices}(\mathfrak{D})$ be the finite set of points where $\text{Sing}(\mathfrak{D})$ is not a manifold. We will denote by $\text{Joints}(\mathfrak{D})$ the set of

closures of the connected components of $\text{Sing}(\mathfrak{D}) \setminus \text{Interstices}(\mathfrak{D})$, calling elements of $\text{Joints}(\mathfrak{D})$ and $\text{Interstices}(\mathfrak{D})$ *joints* and *interstices* respectively.¹ We call a joint *horizontal* if its image under π_2 is a point; otherwise we call a joint *vertical*.

For a path γ in $(M_{\mathbb{R}} \times L) \setminus \text{Sing}(\mathfrak{D})$, one can define an element

$$\theta_{\gamma, \mathfrak{D}} \in \mathbb{V}_{\Sigma, k}$$

exactly as in the case of a scattering diagram in $M_{\mathbb{R}}$. Indeed, we just need to define the automorphism $\theta_{\gamma, \mathfrak{d}}$ when γ crosses $(\mathfrak{d}, f_{\mathfrak{d}})$ at time t_i . Assuming γ passes from one side of \mathfrak{d} to the other, we choose $n_0 \in N$ primitive with $\langle n_0, r(m_0) \rangle = 0$ and n_0 smaller on the side of \mathfrak{d} that γ passes into, as usual. We can then also define as usual

$$\theta_{\gamma, \mathfrak{d}}(z^m) = z^m f_{\mathfrak{d}}^{\langle n_0, r(m) \rangle}.$$

Again, it is easy to check that $\theta_{\gamma, \mathfrak{D}}$ only depends on the homotopy type of the path γ inside $(M_{\mathbb{R}} \times L) \setminus \text{Sing}(\mathfrak{D})$.

A *broken line* in $M_{\mathbb{R}} \times L$ is a map $\beta : (-\infty, 0] \rightarrow M_{\mathbb{R}} \times L$, along with data $t_0 < \dots < t_n$ and monomials $c_i z^{m_i^{\beta}}$, such that

- (1) $\pi_2 \circ \beta$ is constant, say with image $P \in L$.
- (2) $\pi_1 \circ \beta$ is a broken line in the sense of Definition 5.30 with respect to the scattering diagram \mathfrak{D}_P in $M_{\mathbb{R}}$ given, after identifying $M_{\mathbb{R}} \times \{P\}$ with $M_{\mathbb{R}}$, by

$$\mathfrak{D}_P := \{(\mathfrak{d} \cap (M_{\mathbb{R}} \times \{P\}), f_{\mathfrak{d}}) \mid (\mathfrak{d}, f_{\mathfrak{d}}) \in \mathfrak{D} \text{ such that } \mathfrak{d} \cap (M_{\mathbb{R}} \times \{P\}) \neq \emptyset\}.$$

We can now describe how $W_k(Q)$ varies as the points P_1, \dots, P_k are varied.

THEOREM 5.39. *Let W and W' be $W_k(Q)$ for two different choices of general points P_1, \dots, P_k and P'_1, \dots, P'_k . Then*

$$W' = \theta(W)$$

for some $\theta \in \mathbb{V}_{\Sigma, k}$.

PROOF. We shall show this result by induction on k , noting that the base case $k = 1$ is obvious, as moving P_1 and keeping Q fixed is the same thing as moving Q and keeping P_1 fixed.

It is clearly enough to show this result in the case that only P_1 changes, for then we can successively use the same result for P_2, \dots, P_k . So consider a choice of general points P_1, \dots, P_k and P'_1 . Consider the line segment L joining P_1 and P'_1 . For all but a finite number of points $P \in L$, we can assume P, P_2, \dots, P_k will be sufficiently general so that $\text{Trees}(\Sigma, P, P_2, \dots, P_k)$ is finite, and all elements of this set are trivalent. This gives rise to a family of scattering diagrams $\mathfrak{D}(\Sigma, P, P_2, \dots, P_k)$ for $P \in L$ general. We can put these scattering diagrams together into a scattering diagram $\tilde{\mathfrak{D}} = \tilde{\mathfrak{D}}(\Sigma, L, P_2, \dots, P_k)$ in $M_{\mathbb{R}} \times L$. $\tilde{\mathfrak{D}}$ is determined by the requirement that for $P \in L$ general,

$$\begin{aligned} \mathfrak{D}(\Sigma, P, P_2, \dots, P_k) \\ = \{(\tilde{\mathfrak{d}} \cap (M_{\mathbb{R}} \times \{P\}), f_{\tilde{\mathfrak{d}}}) \mid (\tilde{\mathfrak{d}}, f_{\tilde{\mathfrak{d}}}) \in \tilde{\mathfrak{D}} \text{ such that } \tilde{\mathfrak{d}} \cap (M_{\mathbb{R}} \times \{P\}) \neq \emptyset\}. \end{aligned}$$

¹This terminology is adopted from [49].

To keep track of the dependence of $W_k(Q)$ on the point $P \in L$, we write $W_k(Q; P)$. We wish now to show that if γ is a general path in $M_{\mathbb{R}} \times L$ joining (Q, P_1) to (Q, P'_1) , then

$$(5.17) \quad W_k(Q; P'_1) = \theta_{\gamma, \tilde{\mathfrak{D}}}(W_k(Q; P_1)).$$

It is enough to show

- (1) $W_k(Q; P)$ is constant for $(Q, P) \in M_{\mathbb{R}} \times L$ varying within a connected component of $(M_{\mathbb{R}} \times L) \setminus \text{Supp}(\tilde{\mathfrak{D}})$.
- (2) For two such connected components separated by a wall $(\tilde{\mathfrak{d}}, f_{\tilde{\mathfrak{d}}})$ and points (Q, P) , (Q', P') on either side of the wall, we have

$$W_k(Q'; P') = \theta_{\gamma, \mathfrak{D}}(W_k(Q; P))$$

for γ a short path joining (Q, P) with (Q', P') .

Once we show (1), Theorem 5.35 already shows (2): as there are no walls in $\tilde{\mathfrak{D}}$ projecting to points in L , we can always choose points (Q, P) , (Q', P') on opposite sides of a wall with $P = P'$, and then we are in the case already shown in Theorem 5.35. So we only need to show (1).

To show (1), we use the same technique we used for the variation of Q , deforming broken lines. Take (Q, P) and (Q', P') general within a connected component of $(M_{\mathbb{R}} \times L) \setminus \text{Supp}(\tilde{\mathfrak{D}})$ and move from (Q, P) to (Q', P') via a general path γ . Consider broken lines in $M_{\mathbb{R}} \times L$ with endpoint $\gamma(t)$. As t varies, we can continuously deform a broken line with endpoint $\gamma(t)$ unless the broken line converges to one passing through a singular point of $\tilde{\mathfrak{D}}$. However, since such a family of broken lines traces out a two-dimensional subset of $M_{\mathbb{R}} \times L$, by choosing γ sufficiently general we can be sure that none of these broken lines converge to broken lines passing through interstices of $\tilde{\mathfrak{D}}$, as interstices are codimension three. However, they can pass through joints, and this requires some care.

The first observation is that we have already analyzed in the proof of Theorem 5.35 what happens if a broken line passes through a vertical joint. Indeed, we can just as well assume that γ has been chosen so that at a time t_0 when a broken line passes through a vertical joint, $\pi_2(\gamma(t))$ remains constant for t in a neighbourhood of t_0 . Then we are in precisely the situation analyzed in Theorem 5.35.

So we only need to see what happens if a broken line passes through a horizontal joint. Note that horizontal joints occur when two or more parallel rays in a scattering diagram come together as the point P varies; this can typically lead to values of P with families of Maslov index zero disks or the existence of Maslov index -2 disks: see Figures 16 and 17.

In fact, it is enough to show that if j is a horizontal joint and γ_j is a small loop in $M_{\mathbb{R}} \times L$ around the joint, then $\theta_{\gamma_j, \tilde{\mathfrak{D}}} = \text{Id}$. Indeed, if j projects to $P \in L$, j is contained in some polygons $\tilde{\mathfrak{d}}_1, \dots, \tilde{\mathfrak{d}}_n \in \tilde{\mathfrak{D}}$, and necessarily for $P' \in L$ near P , $\tilde{\mathfrak{d}}_i \cap (M_{\mathbb{R}} \times \{P'\})$ is either a ray parallel to j or is empty. Thus, as $P' \in L$ moves from one side of P to the other, some parallel rays $\mathfrak{d}_1, \dots, \mathfrak{d}_p$ in $\mathfrak{D}(\Sigma, P', P_2, \dots, P_k)$ come together to yield the joint and then turn into parallel rays $\mathfrak{d}'_1, \dots, \mathfrak{d}'_{p'}$ on the other side of P . Let $\mathfrak{D}_1, \mathfrak{D}_2$ be the scattering diagrams in $M_{\mathbb{R}}$ given by $\mathfrak{D}(\Sigma, P', P_2, \dots, P_k)$ for P' very close to P , but on opposite sides of P . Let γ




FIGURE 16. As P_1 moves down, we suddenly get a Maslov index -2 tree in the second picture. In addition, a family of Maslov index zero trees appear, as shown in Figure 17.




FIGURE 17

be a path which is a short line segment crossing j , so that we can write

$$\begin{aligned} \theta_{\gamma, \mathfrak{D}_1} &= \theta_{\gamma, \mathfrak{d}_1} \circ \dots \circ \theta_{\gamma, \mathfrak{d}_p}, \\ \theta_{\gamma, \mathfrak{D}_2} &= \theta_{\gamma, \mathfrak{d}'_1} \circ \dots \circ \theta_{\gamma, \mathfrak{d}'_p}. \end{aligned}$$

(Note that the ordering is immaterial as all these automorphisms commute). But $\theta_{\gamma_j, \tilde{\mathfrak{D}}} = \theta_{\gamma, \mathfrak{D}_1}^{-1} \circ \theta_{\gamma, \mathfrak{D}_2}$, so if $\theta_{\gamma_j, \tilde{\mathfrak{D}}} = \text{Id}$, we have $\theta_{\gamma, \mathfrak{D}_1} = \theta_{\gamma, \mathfrak{D}_2}$. This means that, by Definition 5.30, (4), broken lines will behave in the same way on either side of P near the joint j . Note that the actual set of broken lines on either side may be different, because we are not claiming that the set $\{\theta_{\gamma, \mathfrak{d}_1}, \dots, \theta_{\gamma, \mathfrak{d}_p}\}$ coincides with

$\{\theta_{\gamma, \partial'_1}, \dots, \theta_{\gamma, \partial'_p}\}$, but rather the total contribution from bends along the two sets of broken lines remains the same.

To show that $\theta_{\gamma_j, \tilde{\mathfrak{D}}} = \text{Id}$ for each horizontal joint, we proceed as follows. For $I \subseteq \{1, \dots, k\}$, define

$$\text{Ideal}(I) := \langle u_i | i \notin I \rangle \subseteq \mathbb{C}[T_\Sigma] \otimes_{\mathbb{C}} R_k[[y_0]].$$

We use induction, showing

Claim. For $k' \geq 0$ and $\#I = k'$, we have

$$\theta_{\gamma_j, \tilde{\mathfrak{D}}} \equiv \text{Id} \pmod{\text{Ideal}(I)}$$

for every horizontal joint j .

The base case with $k' = 0$ is trivial, because all automorphisms are trivial modulo the ideal (u_1, \dots, u_k) . So assume the claim for all $k'' < k'$. Fix a set I with $\#I = k'$. Fix an orientation on $M_{\mathbb{R}} \times L$, so that if any joint j is given an orientation, this determines the orientation of a loop γ_j around j . We wish to study $\theta_{\gamma_j, \tilde{\mathfrak{D}}}$ for j horizontal.

Note that as $\theta_{\gamma_j, \tilde{\mathfrak{D}}}$ for j horizontal only involves a composition of automorphisms associated to parallel rays, we can in fact write

$$\theta_{\gamma_j, \tilde{\mathfrak{D}}}(z^{m'}) = f_j^{(n_j, r(m'))} z^{m'}$$

for some $n_j \in N$ primitive and zero on the tangent space to j . Also,

$$f_j \in \mathbb{C}[\{m \in T_\Sigma | r(m) \text{ is tangent to } j\}] \otimes_{\mathbb{C}} R_k[[y_0]].$$

Note that f_j depends on the choice of sign of n_j . Assume we have chosen these consistently, in the sense that if any two joints j and j' have the same tangent space, then $n_j = n_{j'}$.

We need to show $f_j \equiv 1 \pmod{\text{Ideal}(I)}$. Fix some $m \in T_\Sigma$. For each horizontal joint j , let the term in $f_j \pmod{\text{Ideal}(I)}$ involving z^m be $c_{m,j} z^m$. Thus $c_{m,j} = \bar{c}_{m,j} \prod_{i \in I} u_i$ for some $\bar{c}_{m,j} \in \mathbb{C}$ since, by the induction hypothesis, $f_j \equiv 1 \pmod{\text{Ideal}(I')}$ for any $I' \subsetneq I$.

Let us first observe that if $\bar{c}_{m,j} \neq 0$, then $r(m) \neq 0$. Indeed, note that f_j is a product of factors of the form $(1 + c_{m'} z^{m'})^{\pm 1}$ with $r(m') \neq 0$ by construction of $\theta_{\gamma_j, \tilde{\mathfrak{D}}}$. Then $\log f_j$ is a sum of expressions of the form $\pm \log(1 + c_{m'} z^{m'})$. After expanding this out using the (finite) Taylor series, we see that $\log f_j = \sum \pm c_{m'} z^{m'}$ with $r(m') \neq 0$ for every m' appearing in this sum. On the other hand, modulo $\text{Ideal}(I)$, $\log f_j = \sum_{m \in T_\Sigma} \bar{c}_{m,j} z^m$. Hence $c_{m,j} = 0$ if $r(m) = 0$, so $\bar{c}_{m,j} = 0$ if $r(m) = 0$.

So we will fix $m \in T_\Sigma$ with $r(m) \neq 0$, and focus on showing $\bar{c}_{m,j} = 0$ for all horizontal joints j . We will also include here the case that j is a vertical joint, by setting $\bar{c}_{m,j} = 0$ for vertical joints. Note that $\bar{c}_{m,j}$ depends on the orientation of j . A change of orientation of j changes the direction of γ_j , replacing f_j with f_j^{-1} . This changes the sign of $\bar{c}_{m,j}$. As a result, we can view

$$j \mapsto \bar{c}_{m,j}$$

as a 1-chain for the one-dimensional simplicial complex $\text{Sing}(\tilde{\mathfrak{D}})$. Here the choice of orientation on j is implicit.

Subclaim. $j \mapsto \bar{c}_{m,j}$ is a 1-cycle.




FIGURE 18. The loop depicted, going successively around x_1, x_2, x_3 , and x_4 , is contractible in $S \setminus \{x_1, \dots, x_4\}$

PROOF. We need to check the 1-cycle condition at each interstice of $\tilde{\mathfrak{D}}$, so let $(Q, P) \in \text{Interstices}(\tilde{\mathfrak{D}})$. Consider a small two-sphere S in $M_{\mathbb{R}} \times L$ with center (Q, P) . Then suppose that $x_1, \dots, x_s \in S$ are distinct points such that

$$\{x_1, \dots, x_s\} = \bigcup_{j \in \text{Joints}(\tilde{\mathfrak{D}})} j \cap S.$$

Choose a base-point $y \in S, y \notin \text{Supp}(\tilde{\mathfrak{D}})$. We can choose small counterclockwise loops $\gamma_1, \dots, \gamma_s$ in S around x_1, \dots, x_s and paths β_i joining y with the base-point of γ_i in such a way so that

$$\beta_1 \gamma_1 \beta_1^{-1} \dots \beta_s \gamma_s \beta_s^{-1} = 1$$

in $\pi_1(S \setminus \{x_1, \dots, x_s, y\})$; see Figure 18. Because $\theta_{\gamma, \tilde{\mathfrak{D}}}$ only depends on the homotopy type of γ in $(M_{\mathbb{R}} \times L) \setminus \text{Sing}(\tilde{\mathfrak{D}})$, we obtain the equality

$$(5.18) \quad \theta_{\beta_s}^{-1} \circ \theta_{\gamma_s} \circ \theta_{\beta_s} \circ \dots \circ \theta_{\beta_1}^{-1} \circ \theta_{\gamma_1} \circ \theta_{\beta_1} = \text{Id}.$$

Here, we have dropped the $\tilde{\mathfrak{D}}$'s in the subscripts.

We now distinguish between two cases.

Case 1. The interstice (Q, P) does not satisfy $Q \in \{P, P_2, \dots, P_k\}$. Then by Proposition 5.28, $\theta_{\gamma_i} = \text{Id}$ for each γ_i which is a loop around a vertical joint containing (Q, P) . On the other hand, modulo $\text{Ideal}(I)$, for γ_i around a horizontal joint j_i , by the induction hypothesis, f_{j_i} is of the form $1 + (\dots) \prod_{i \in I} u_i$. One then checks θ_{γ_i} necessarily commutes, modulo $\text{Ideal}(I)$, with any element of $\mathbb{V}_{\Sigma, k}$. This can easily be seen as in Example 5.24, using the fact that $u_j \prod_{i \in I} u_i \equiv 0$

mod $\text{Ideal}(I)$ for any j . Thus, in particular, θ_{γ_i} commutes with θ_{β_i} . So (5.18) becomes

$$\prod \theta_{\gamma_i} \equiv \text{Id} \quad \text{mod Ideal}(I)$$

where the product is over all γ_i around horizontal joints. Applying this identity to a monomial $z^{m'}$, we obtain

$$\prod f_{j_i}^{\langle n_{j_i}, r(m') \rangle} z^{m'} = z^{m'} \quad \text{mod Ideal}(I),$$

which, after expansion and reading off the coefficient of $z^{m+m'}$, gives the identity

$$(5.19) \quad \sum \langle n_{j_i}, r(m') \rangle \bar{c}_{m, j_i} = 0$$

for any $m' \in T_\Sigma$. Now a monomial z^m can only appear in f_{j_i} if $r(m)$ is in fact tangent to j_i , so the only horizontal joints containing (Q, P) with $\bar{c}_{m, j_i} \neq 0$ are the joints contained in the affine line $(Q, P) + \mathbb{R}(r(m), 0)$. Let s be the number of joints contained in this line and containing (Q, P) . Then either $s = 0, 1$ or 2 . If $s = 0$, there is nothing to prove. If $s = 1$, with j_i the only such joint, it follows from (5.19) that $\bar{c}_{m, j_i} = 0$. If $s = 2$, let j_{i_1}, j_{i_2} be the two such joints. Then (5.19) implies that $\bar{c}_{m, j_{i_1}} = \bar{c}_{m, j_{i_2}}$, assuming j_{i_1} and j_{i_2} are oriented in the same direction. This shows that the 1-cycle condition holds at (Q, P) .

Case 2. The interstice (Q, P) satisfies $Q \in \{P, P_2, \dots, P_k\}$, say $Q = P_i$. We'll write $i = 1$ if $Q = P$. The argument is almost the same, but now there are two vertical joints, say j_1 and j_2 with endpoint (Q, P) , with $j_1, j_2 \subseteq \{P_i\} \times L$ if $i > 1$ and $j_1, j_2 \subseteq \{(P', P') | P' \in L\} \subseteq M_{\mathbb{R}} \times L$ if $i = 1$. Without loss of generality we can take the base-point y near x_1 and assume β_1 is a constant path so that $\theta_{\beta_1} = \text{Id}$. The argument will be the same as in Case 1 once we show that

$$(5.20) \quad \theta_{\beta_2}^{-1} \circ \theta_{\gamma_2} \circ \theta_{\beta_2} \circ \theta_{\gamma_1} = \text{Id}.$$

To do so, consider the scattering diagram $\mathfrak{D}(\Sigma, P', P_2, \dots, P_k)$ for $P' \in L$, P' near P but $P' \neq P$. By Remark 5.29, the rays emanating from P_i (P' if $i = 1$) in $\mathfrak{D}(\Sigma, P', P_2, \dots, P_k)$ are in one-to-one correspondence with the terms in $W_{k-1}(P_i; P') - y_0$, where $W_{k-1}(P_i; P')$ denotes $W_{k-1}(P_i)$ computed using the marked points $P', \dots, P_{i-1}, P_{i+1}, \dots, P_k$ (or P_2, \dots, P_k if $i = 1$). In particular, given a term cz^m in $W_{k-1}(P_i; P') - y_0$, the corresponding ray carries the function $1 + u_i cw(m)z^m$, where $w(m)$ is the index of $r(m)$. Note that if γ is a simple loop around P_i , then the contribution to $\theta_{\gamma, \mathfrak{D}(\Sigma, P', P_2, \dots, P_k)}$ from such a ray is $\exp(\pm X_{u_i cz^m})$. Here the sign only depends on the orientation of γ and the chosen identification of $\bigwedge^2 M$ with \mathbb{Z} . All automorphisms attached to the rays emanating from P_i commute by Example 5.24 because $u_i^2 = 0$, so

$$\theta_{\gamma, \mathfrak{D}(\Sigma, P', P_2, \dots, P_k)} = \prod \exp(\pm X_{u_i cz^m}) = \exp(\pm X_{u_i (W_{k-1}(P_i; P') - y_0)}).$$

Here the product is over all terms cz^m appearing in $W_{k-1}(P_i; P') - y_0$. Furthermore, if $P' \in \pi_2(j_1) \setminus \{P\}$ and $P'' \in \pi_2(j_2) \setminus \{P\}$, then by (5.17) applied inductively to $k-1$ points if $i \neq 1$, and by Theorem 5.35, if $i = 1$,

$$u_i W_{k-1}(P_i; P'') = u_i \theta_{\beta_2} (W_{k-1}(P_i; P')).$$

It then follows from Lemma 5.21 that

$$\theta_{\gamma_2} = (\theta_{\beta_2} \circ \theta_{\gamma_1} \circ \theta_{\beta_2}^{-1})^{-1},$$

the last inverse on the right since γ_1 and γ_2 are homotopic to loops in $M_{\mathbb{R}} \times \{P'\}$ and $M_{\mathbb{R}} \times \{P''\}$ respectively with *opposite orientations*. This shows (5.20). We can then finish as in Case 1.

This completes the proof of the subclaim. \square

To complete the proof of the claim, hence the theorem, we now note that the cycle σ given by $j \mapsto \bar{c}_{m,j}$ is in fact zero. Indeed, picking a given joint with $\bar{c}_{m,j} \neq 0$, the fact that σ is a cycle implies that the line containing j can be written as a union of joints j' with orientation compatible with that on j , with $\bar{c}_{m,j'} = \bar{c}_{m,j}$. However, there must be one joint j' contained in this line which is unbounded in the direction $r(m)$. But none of the polyhedra of $\tilde{\mathfrak{D}}$ containing j' can involve a monomial of the form z^m , since a ray carrying a monomial z^m is unbounded only in the direction $-r(m)$. Thus $0 = \bar{c}_{m,j'} = \bar{c}_{m,j}$ as desired. \square

5.4.7. Independence of the integrals.

LEMMA 5.40. *Let $\theta \in \mathbb{V}_{\Sigma,k}$, $(u, \hbar) \in \widetilde{\mathcal{M}}_{\Sigma,k} \times \mathbb{C}^\times$ and suppose f is in the ideal generated by u_1, \dots, u_k in $\mathbb{C}[T_\Sigma] \otimes_{\mathbb{C}} R_k[[y_0]]$. Then for any cycle*

$$\Xi \in H_2(\kappa^{-1}(u), \operatorname{Re}(W_0(Q)/\hbar) \ll 0; \mathbb{C}),$$

we have

$$\int_{\Xi} e^{(W_0(Q)+f)/\hbar} \Omega = \int_{\Xi} e^{\theta(W_0(Q)+f)/\hbar} \Omega.$$

PROOF. It is enough to show the lemma for $\theta = \exp(cz^{m_0} X_{r(m_0)})$ with $m_0 \in T_\Sigma$, $r(m_0) \neq 0$ and $c^2 = 0$, as such elements generate $\mathbb{V}_{\Sigma,k}$. Note that if $W_0(Q) + f = \sum_m c_m z^m$, then

$$\theta(W_0(Q) + f) = \sum_m c_m (z^m + \langle X_{r(m_0)}, r(m) \rangle c z^{m_0+m})$$

and

$$e^{\theta(W_0(Q)+f)/\hbar} = e^{(W_0(Q)+f)/\hbar} \left(1 + \sum_m \hbar^{-1} c c_m \langle X_{r(m_0)}, r(m) \rangle z^{m_0+m} \right).$$

Furthermore, $d(z^m \operatorname{dlog}(z^{m_0})) = -\langle X_{r(m_0)}, r(m) \rangle z^m \Omega$. Thus

$$\begin{aligned} & (e^{\theta(W_0(Q)+f)/\hbar} - e^{(W_0(Q)+f)/\hbar}) \Omega \\ &= e^{(W_0(Q)+f)/\hbar} \left(\hbar^{-1} \sum_m c c_m \langle X_{r(m_0)}, r(m) \rangle z^{m_0+m} \right) \Omega \\ &= -d(cz^{m_0} e^{(W_0(Q)+f)/\hbar} \operatorname{dlog}(z^{m_0})). \end{aligned}$$

The result then follows from Stokes' theorem and the fact that $e^{(W_0(Q)+f)/\hbar}$ goes to zero rapidly on the unbounded part of Ξ . \square

LEMMA 5.41. *For $\Xi \in H_2((\check{\mathcal{X}}_{\Sigma,k})_\kappa, \operatorname{Re}(W_0(Q)/\hbar) \ll 0)$, the integral*

$$\int_{\Xi} e^{W_k(Q)/\hbar} \Omega$$

is independent of the choice of Q and P_1, \dots, P_k .

PROOF. This follows immediately from Theorems 5.35, 5.39, and Lemma 5.40. \square

5.5. Evaluation of the period integrals

Our main goal in this section is the computation of the integrals

$$(5.21) \quad \int_{\Xi_i} e^{W_k(Q)/\hbar} \Omega$$

in the case of \mathbb{P}^2 . In doing so, we will prove Theorems 5.15 and 5.18. We continue with the notation $\Sigma, T_\Sigma, t_i, \rho_i$ of Example 5.9 as well as $m_i = r(t_i)$.

5.5.1. The raw integrals. As we shall see in this subsection, it is not at all difficult to produce an actual expression for (5.21), provided one knows what $W_k(Q)$ is. This is only the first step, as it will take some effort to convert this explicit expression into something which is useful for us. That will take the remainder of the section to do.

LEMMA 5.42. *Restricting to $x_0 x_1 x_2 = 1$, we have*

$$\begin{aligned} & \sum_{i=0}^2 \alpha^i \int_{\Xi_i} e^{(x_0+x_1+x_2)/\hbar} \Omega \\ &= \hbar^{-3\alpha} \left(\sum_{d=0}^{\infty} \frac{\hbar^{-3d}}{(d!)^3} - 3\alpha \sum_{d=1}^{\infty} \frac{\hbar^{-3d}}{(d!)^3} \sum_{k=1}^d \frac{1}{k} \right. \\ & \quad \left. + \frac{9}{2} \alpha^2 \sum_{d=1}^{\infty} \frac{\hbar^{-3d}}{(d!)^3} \left(\left(\sum_{k=1}^d \frac{1}{k} \right)^2 + \frac{1}{3} \sum_{k=1}^d \frac{1}{k^2} \right) \right) \\ &=: \hbar^{-3\alpha} \sum_{d=0}^{\infty} \hbar^{-3d} (B_0(d) + \alpha B_1(d) + \alpha^2 B_2(d)), \end{aligned}$$

where the last equality defines the numbers $B_0(d), B_1(d), B_2(d)$.

PROOF. This is the expansion of the explicit expression given in Proposition 2.40. \square

We can use this to compute (5.21) by writing

$$\int_{\Xi_i} e^{W_k(Q)/\hbar} \Omega = e^{y_0/\hbar} \int_{\Xi_i} e^{(x_0+x_1+x_2)/\hbar} e^{(W_k(Q)-W_0(Q))/\hbar} \Omega.$$

The factor $e^{(W_k(Q)-W_0(Q))/\hbar}$ can then be expanded in a Taylor series, noting that in any term, each monomial in $W_k(Q) - W_0(Q)$ can appear at most once, because it has a coefficient of square zero; thus this expansion is quite easy and is finite. Thus we only need to calculate, with $x_0 x_1 x_2 = e^{y_1}$,

$$\sum_{i=0}^2 \alpha^i \int_{\Xi_i} e^{(x_0+x_1+x_2)/\hbar} x_0^{n_0} x_1^{n_1} x_2^{n_2} \Omega.$$

LEMMA 5.43. *With $x_0 x_1 x_2 = e^{y_1}$,*

$$\sum_{i=0}^2 \alpha^i \int_{\Xi_i} e^{(x_0+x_1+x_2)/\hbar} x_0^{n_0} x_1^{n_1} x_2^{n_2} \Omega = \hbar^{-3\alpha} e^{\alpha y_1} \sum_{i=0}^2 \psi_i(n_0, n_1, n_2) \alpha^i,$$

where

$$\psi_i(n_0, n_1, n_2) = \sum_{d=0}^{\infty} D_i(d, n_0, n_1, n_2) \hbar^{-(3d-n_0-n_1-n_2)} e^{d y_1}$$

with D_i given as follows. First,

$$D_0(d, n_0, n_1, n_2) = \begin{cases} \frac{1}{(d-n_0)!(d-n_1)!(d-n_2)!} & \text{if } d \geq n_0, n_1, n_2 \\ 0 & \text{otherwise.} \end{cases}$$

Second, if $d \geq n_0, n_1, n_2$, then

$$D_1(d, n_0, n_1, n_2) = -\frac{\sum_{k=1}^{d-n_0} \frac{1}{k} + \sum_{k=1}^{d-n_1} \frac{1}{k} + \sum_{k=1}^{d-n_2} \frac{1}{k}}{(d-n_0)!(d-n_1)!(d-n_2)!}$$

while if $n_0, n_1 \leq d < n_2$, then

$$D_1(d, n_0, n_1, n_2) = \frac{(-1)^{n_2-d-1}(n_2-d-1)!}{(d-n_0)!(d-n_1)!},$$

with similar expressions if instead $d < n_0$ or $d < n_1$. If d is smaller than two of n_0, n_1, n_2 , then

$$D_1(d, n_0, n_1, n_2) = 0.$$

Third, if $d \geq n_0, n_1, n_2$, then

$$D_2(d, n_0, n_1, n_2) = \frac{\left(\sum_{l=0}^2 \sum_{k=1}^{d-n_l} \frac{1}{k}\right)^2 + \sum_{l=0}^2 \sum_{k=1}^{d-n_l} \frac{1}{k^2}}{2(d-n_0)!(d-n_1)!(d-n_2)!}$$

while if $n_0, n_1 \leq d < n_2$,

$$D_2(d, n_0, n_1, n_2) = \frac{(-1)^{d-n_2}(n_2-d-1)!}{(d-n_0)!(d-n_1)!} \left(\sum_{k=1}^{d-n_0} \frac{1}{k} + \sum_{k=1}^{d-n_1} \frac{1}{k} + \sum_{k=1}^{n_2-d-1} \frac{1}{k} \right),$$

with similar expressions if instead $d < n_0$ or $d < n_1$. If $n_0 \leq d < n_1, n_2$, then

$$D_2(d, n_0, n_1, n_2) = \frac{(-1)^{n_1+n_2}(n_1-d-1)!(n_2-d-1)!}{(d-n_0)!},$$

with similar expressions if instead $n_1 \leq d < n_0, n_2$ or $n_2 \leq d < n_0, n_1$. Finally, if $d < n_0, n_1, n_2$, then

$$D_2(d, n_0, n_1, n_2) = 0.$$

PROOF. Consider the integral

$$I_i(a_0, a_1, a_2) = \int_{\Xi_i} e^{a_0 x_0 + a_1 x_1 + a_2 x_2} \Omega,$$

with $a_0, a_1, a_2 \in \mathbb{C}^\times$ and $x_0 x_1 x_2 = 1$. Then

$$\frac{\partial^{n_0+n_1+n_2}}{\partial a_0^{n_0} \partial a_1^{n_1} \partial a_2^{n_2}} I_i = \int_{\Xi_i} e^{a_0 x_0 + a_1 x_1 + a_2 x_2} x_0^{n_0} x_1^{n_1} x_2^{n_2} \Omega.$$

Evaluate this at $a_0 = a_1 = a_2 = e^{y_1/3}/\hbar$ and make the change of variables $x_i \mapsto x_i e^{-y_1/3}$ in the integral. Note that as $\Omega = \frac{dx_1 \wedge dx_2}{x_1 x_2}$, such a change of variables does not affect Ω . Then with $x_0 x_1 x_2 = e^{y_1}$ we obtain

$$\frac{\partial^{n_0+n_1+n_2}}{\partial a_0^{n_0} \partial a_1^{n_1} \partial a_2^{n_2}} I_i \Big|_{a_i=e^{y_1/3}/\hbar} = \int_{\Xi_i} e^{(x_0+x_1+x_2)/\hbar} e^{-(n_0+n_1+n_2)y_1/3} x_0^{n_0} x_1^{n_1} x_2^{n_2} \Omega.$$

On the other hand, I_i can be calculated by making the substitution

$$\begin{aligned} x_0 &\mapsto (a_1 a_2 / a_0^2)^{1/3} x_0 \\ x_1 &\mapsto (a_0 a_2 / a_1^2)^{1/3} x_1 \\ x_2 &\mapsto (a_0 a_1 / a_2^2)^{1/3} x_2 \end{aligned}$$

in I_i which gives

$$I_i(a_0, a_1, a_2) = \int_{\Xi_i} e^{(a_0 a_1 a_2)^{1/3} (x_0 + x_1 + x_2)} \Omega.$$

Thus we can compute $\sum_{i=0}^2 \alpha^i I_i(a_0, a_1, a_2)$ by substituting in $\hbar = (a_0 a_1 a_2)^{-1/3}$ in the formula of Lemma 5.42. To differentiate the resulting expression, note that under this substitution, $\hbar^{-(3\alpha+3d)}$ becomes $(a_0 a_1 a_2)^{\alpha+d}$ and

$$\begin{aligned} &\frac{\partial^{n_0+n_1+n_2}}{\partial a_0^{n_0} \partial a_1^{n_1} \partial a_2^{n_2}} (a_0 a_1 a_2)^{\alpha+d} \Big|_{a_i=e^{y_i/3}/\hbar} \\ &= \hbar^{-(3\alpha+3d-n_0-n_1-n_2)} e^{(\alpha+d-(n_0+n_1+n_2)/3)y_1} \\ &\quad \cdot \prod_{k=1}^{n_0} (\alpha+d-k+1) \prod_{k=1}^{n_1} (\alpha+d-k+1) \prod_{k=1}^{n_2} (\alpha+d-k+1) \\ &= \hbar^{-(3\alpha+3d-n_0-n_1-n_2)} e^{(\alpha+d-(n_0+n_1+n_2)/3)y_1} \\ &\quad \cdot (C_0(d, n_0, n_1, n_2) + \alpha C_1(d, n_0, n_1, n_2) + \alpha^2 C_2(d, n_0, n_1, n_2)), \end{aligned}$$

where the last equality defines C_0, C_1 and C_2 . One then sees that

$$\psi_i(n_0, n_1, n_2) = \sum_{d=0}^{\infty} \sum_{k=0}^i B_k(d) C_{i-k}(d, n_0, n_1, n_2) \hbar^{-(3d-n_0-n_1-n_2)} e^{dy_1}$$

with the B_i 's defined in Lemma 5.42. Furthermore, computing the C_i 's, we see

$$C_0(d, n_0, n_1, n_2) = \begin{cases} \frac{(d!)^3}{(d-n_0)!(d-n_1)!(d-n_2)!} & \text{if } d \geq n_0, n_1, n_2 \\ 0 & \text{otherwise.} \end{cases}$$

If $d \geq n_0, n_1, n_2$, then

$$C_1(d, n_0, n_1, n_2) = \frac{(d!)^3}{(d-n_0)!(d-n_1)!(d-n_2)!} \left(\sum_{k=d-n_0+1}^d \frac{1}{k} + \sum_{k=d-n_1+1}^d \frac{1}{k} + \sum_{k=d-n_2+1}^d \frac{1}{k} \right),$$

while if *one* of n_0, n_1, n_2 is larger than d , we have

$$C_1(d, n_0, n_1, n_2) = \prod_{\substack{k=d-n_0+1 \\ k \neq 0}}^d k \prod_{\substack{k=d-n_1+1 \\ k \neq 0}}^d k \prod_{\substack{k=d-n_2+1 \\ k \neq 0}}^d k.$$

Otherwise

$$C_1(d, n_0, n_1, n_2) = 0.$$

If $d \geq n_0, n_1, n_2$ then

$$C_2(d, n_0, n_1, n_2) = \frac{(d!)^3}{2(d-n_0)!(d-n_1)!(d-n_2)!} \left(\left(\sum_{k=d-n_0+1}^d \frac{1}{k} + \sum_{k=d-n_1+1}^d \frac{1}{k} + \sum_{k=d-n_2+1}^d \frac{1}{k} \right)^2 - \left(\sum_{k=d-n_0+1}^d \frac{1}{k^2} + \sum_{k=d-n_1+1}^d \frac{1}{k^2} + \sum_{k=d-n_2+1}^d \frac{1}{k^2} \right) \right).$$

If $n_1, n_2 \leq d < n_0$, then

$$C_2(d, n_0, n_1, n_2) = \left(\prod_{\substack{k=d-n_0+1 \\ k \neq 0}}^d k \right) \frac{(d!)^2}{(d-n_1)!(d-n_2)!} \left(\sum_{\substack{k=d-n_0+1 \\ k \neq 0}}^d \frac{1}{k} + \sum_{k=d-n_1+1}^d \frac{1}{k} + \sum_{k=d-n_2+1}^d \frac{1}{k} \right).$$

We have similar expressions if $d < n_1$ or $d < n_2$. If two of n_0, n_1 and n_2 are larger than d , then

$$C_2(d, n_0, n_1, n_2) = \prod_{\substack{k=d-n_0+1 \\ k \neq 0}}^d k \prod_{\substack{k=d-n_1+1 \\ k \neq 0}}^d k \prod_{\substack{k=d-n_2+1 \\ k \neq 0}}^d k.$$

Finally, if $n_0, n_1, n_2 > d$, then

$$C_2(d, n_0, n_1, n_2) = 0.$$

A laborious calculation now gives the forms given in the Lemma for the coefficients D_i . \square

DEFINITION 5.44. For $m \in T_\Sigma$, $m = \sum_{i=0}^2 n_i t_i$ with $n_i \geq 0$ for all i , define

$$\begin{aligned} \psi_i(m) &:= \psi_i(n_0, n_1, n_2) \\ D_i(d, m) &:= D_i(d, n_0, n_1, n_2) \end{aligned}$$

and

$$|m| := n_0 + n_1 + n_2.$$

5.5.2. What we need to show. If we have an explicit expression for $W_k(Q)$, the results of the previous subsection allow us to write down (5.21). In general, the result (except for the integral over Ξ_0 , as we shall see), bears no immediate resemblance to the desired form predicted by Theorem 5.18. Some massaging is necessary before the integral begins to look correct.

First, though, we will start on our proof of Theorems 5.15 and 5.18 by making precise which equality needs to be shown in terms of the integrals computed in the previous subsection.

We make the following definition for keeping track of the terms which will appear in (5.21).

DEFINITION 5.45. Fix P_1, \dots, P_k general. For Q general, let S_k (or $S_k(Q)$ if the dependence on Q needs to be emphasized) be a finite set of triples (c, ν, m) with $c \in R_k$ a monomial, $\nu \geq 0$ an integer, and $m \in T_\Sigma$ such that

$$(5.22) \quad e^{(W_k(Q) - W_0(Q))/\hbar} = \sum_{(c, \nu, m) \in S_k} c \hbar^{-\nu} z^m,$$

with each term $c\hbar^{-\nu}z^m$ of the form $\hbar^{-\nu}\prod_{i=1}^{\nu}\text{Mono}(h_i)$ for h_1, \dots, h_{ν} distinct Maslov index two tropical disks with boundary Q .

Let

$$L_i^d = L_i^d(Q) := \sum_{(c, \nu, m) \in S_k} c\hbar^{-(3d+\nu-|m|)} D_i(d, m).$$

We can now clarify what needs to be proved. The following lemma reduces Theorems 5.15 and 5.18 to three equalities.

LEMMA 5.46. *Let Q be chosen generally, and let L be the tropical line with vertex Q . The three equalities*

$$(5.23) \quad L_0^d = \delta_{0,d} + \sum_{\nu \geq 0} \sum_{\substack{I \subseteq \{1, \dots, k\} \\ I = \{i_1, \dots, i_{3d-\nu-2}\} \\ i_1 < \dots < i_{3d-\nu-2}}} \langle P_{i_1}, \dots, P_{i_{3d-\nu-2}}, \psi^{\nu} Q \rangle_{0,d}^{\text{trop}} u_I \hbar^{-(\nu+2)}$$

$$(5.24) \quad L_1^d = \sum_{\nu \geq 0} \sum_{\substack{I \subseteq \{1, \dots, k\} \\ I = \{i_1, \dots, i_{3d-\nu-1}\} \\ i_1 < \dots < i_{3d-\nu-1}}} \langle P_{i_1}, \dots, P_{i_{3d-\nu-1}}, \psi^{\nu} L \rangle_{0,d}^{\text{trop}} u_I \hbar^{-(\nu+1)}$$

$$(5.25) \quad L_2^d = y_2 \hbar \delta_{0,d} + \sum_{\nu \geq 0} \sum_{\substack{I \subseteq \{1, \dots, k\} \\ I = \{i_1, \dots, i_{3d-\nu}\} \\ i_1 < \dots < i_{3d-\nu}}} \langle P_{i_1}, \dots, P_{i_{3d-\nu}}, \psi^{\nu} M_{\mathbb{R}} \rangle_{0,d}^{\text{trop}} u_I \hbar^{-\nu}$$

imply Theorems 5.15 and 5.18.

PROOF. Let us be precise about what needs to be shown to prove Theorems 5.15 and 5.18. If we write, for $0 \leq i \leq 2$ and $y_2 = \sum_{i=1}^k u_i$,

$$K_i^{\text{trop}} = \sum_{d \geq 1} \sum_{\nu \geq 0} \langle T_2^{3d+i-2-\nu}, \psi^{\nu} T_{2-i} \rangle_{0,d}^{\text{trop}} \hbar^{-(\nu+2)} e^{dy_1} \frac{y_2^{3d+i-2-\nu}}{(3d+i-2-\nu)!},$$

then

$$\begin{aligned} J_0^{\text{trop}} &= e^{y_0/\hbar} (1 + K_0^{\text{trop}}), \\ J_1^{\text{trop}} &= e^{y_0/\hbar} (\hbar^{-1} y_1 (1 + K_0^{\text{trop}}) + K_1^{\text{trop}}), \\ J_2^{\text{trop}} &= e^{y_0/\hbar} \left(\frac{\hbar^{-2} y_1^2}{2} (1 + K_0^{\text{trop}}) + \hbar^{-1} y_1 K_1^{\text{trop}} + \hbar^{-1} y_2 + K_2^{\text{trop}} \right). \end{aligned}$$

We wish to compare these expressions with the expressions obtained via period integrals over Ξ_0 , Ξ_1 and Ξ_2 . Consider the φ_i 's defined in (5.5). Expanding the integral in (5.5) by using Lemma 5.43 and $e^{\alpha y_1} = 1 + y_1 \alpha + y_1^2 \alpha^2 / 2$, the left-hand side of (5.5) is

$$\begin{aligned} \sum_{(c, \nu, m) \in S_k} c e^{y_0/\hbar} \hbar^{-(3\alpha+\nu)} e^{\alpha y_1} \sum_{i=0}^2 \psi_i(m) \alpha^i \\ = \hbar^{-3\alpha} e^{y_0/\hbar} \sum_{(c, \nu, m) \in S_k} c \hbar^{-\nu} \sum_{i=0}^2 \sum_{k=0}^i \frac{y_1^k}{k!} \psi_{i-k}(m) \alpha^i. \end{aligned}$$

Comparing this with the right-hand side of (5.5), we get

$$\begin{aligned}\varphi_0 &= e^{y_0/\hbar} \sum_{(c,\nu,m) \in S_k} c \hbar^{-\nu} \psi_0(m), \\ \varphi_1 &= e^{y_0/\hbar} \sum_{(c,\nu,m) \in S_k} c \hbar^{-(\nu+1)} (y_1 \psi_0(m) + \psi_1(m)), \\ \varphi_2 &= e^{y_0/\hbar} \sum_{(c,\nu,m) \in S_k} c \hbar^{-(\nu+2)} \left(\frac{y_1^2}{2} \psi_0(m) + y_1 \psi_1(m) + \psi_2(m) \right).\end{aligned}$$

Thus to show $\varphi_i = J_i^{\text{trop}}$, we need to show the following three equalities:

$$(5.26) \quad \sum_{(c,\nu,m) \in S_k} c \hbar^{-\nu} \psi_0(m) = 1 + K_0^{\text{trop}},$$

$$(5.27) \quad \sum_{(c,\nu,m) \in S_k} c \hbar^{-\nu} \psi_1(m) = \hbar K_1^{\text{trop}},$$

$$(5.28) \quad \sum_{(c,\nu,m) \in S_k} c \hbar^{-\nu} \psi_2(m) = \hbar^2 (\hbar^{-1} y_2 + K_2^{\text{trop}}).$$

Then, using the expansion for ψ_i in Lemma 5.43, (5.26), (5.27) and (5.28) are equivalent, if we compare the coefficients of e^{dy_1} on both sides, to:

$$(5.29) \quad L_0^d = \delta_{0,d} + \sum_{\nu \geq 0} \langle T_2^{3d-\nu-2}, \psi^\nu T_2 \rangle_{0,d}^{\text{trop}} \frac{y_2^{3d-\nu-2}}{(3d-\nu-2)!} \hbar^{-(\nu+2)},$$

$$(5.30) \quad L_1^d = \sum_{\nu \geq 0} \langle T_2^{3d-\nu-1}, \psi^\nu T_1 \rangle_{0,d}^{\text{trop}} \frac{y_2^{3d-\nu-1}}{(3d-\nu-1)!} \hbar^{-(\nu+1)},$$

$$(5.31) \quad L_2^d = y_2 \hbar \delta_{0,d} + \sum_{\nu \geq 0} \langle T_2^{3d-\nu}, \psi^\nu T_0 \rangle_{0,d}^{\text{trop}} \frac{y_2^{3d-\nu}}{(3d-\nu)!} \hbar^{-\nu}.$$

Now suppose we have shown (5.23), (5.24) and (5.25). The left-hand sides of these equations come from the period integrals, and hence are independent of the locations of Q and P_1, \dots, P_k by Lemma 5.41. So the right-hand side is also independent of the locations of Q and P_1, \dots, P_k . So in particular, once we show (5.23), (5.24) and (5.25), we find that the invariants $\langle T_2^{3d+i-2-\nu}, \psi^\nu T_{2-i} \rangle_{0,d}^{\text{trop}}$ are well-defined, showing Theorem 5.15, and also showing (5.29), (5.30) and (5.31), as

$$\frac{y_2^n}{n!} = \sum_{\substack{I \subseteq \{1, \dots, k\} \\ \#I = n}} u_I.$$

Also, this shows $\varphi_i = J_i^{\text{trop}}$. In particular, $\varphi_{i,1} = y_i$ for $0 \leq i \leq 2$. This gives Theorem 5.18. \square

We will now refine the expressions L_i^d which we need to compute, splitting it up into a number of terms which will yield contributions of different flavours.

DEFINITION 5.47. For each cone $\sigma \in \Sigma$, σ is the image under r of a proper face $\tilde{\sigma}$ of the cone $K \subseteq T_\Sigma \otimes \mathbb{R}$ generated by t_0, t_1, t_2 (i.e., the first octant). For $d \geq 0$,

denote by $K_d \subseteq K$ the cube

$$K_d = \left\{ \sum_{i=0}^2 n_i t_i \mid 0 \leq n_i \leq d \right\}$$

and for $\sigma \in \Sigma$, define

$$\tilde{\sigma}_d := (\tilde{\sigma} + K_d) \setminus \bigcup_{\substack{\tau \subseteq \sigma \\ \tau \in \Sigma}} (\tilde{\tau} + K_d).$$

Here $+$ denotes Minkowski sum.

EXAMPLE 5.48. We have the following examples of $\tilde{\sigma}_d$. Let $m = \sum_i n_i t_i \in K$.

- If $\sigma = \{0\}$, then $m \in \tilde{\sigma}_d$ if and only if $d \geq \max\{n_0, n_1, n_2\}$.
- If $\sigma = \rho_0$, then $m \in \tilde{\sigma}_d$ if and only if $n_1, n_2 \leq d < n_0$.
- If $\sigma = \rho_1 + \rho_2$, then $m \in \tilde{\sigma}_d$ if and only if $n_0 \leq d < n_1, n_2$.

DEFINITION 5.49. For $\sigma \in \Sigma$, define

$$L_{i,\sigma}^d = L_{i,\sigma}^d(Q) := \sum_{\substack{(c,\nu,m) \in S_k \\ m \in \tilde{\sigma}_d}} c \hbar^{-(3d+\nu-|m|)} D_i(d, m).$$

We then have the obvious

$$\text{LEMMA 5.50. } L_i^d = \sum_{\sigma \in \Sigma} L_{i,\sigma}^d.$$

5.5.3. The asymptotic behaviour of $L_{i,\omega}^d(Q)$. The main trick for understanding the terms $L_{i,\omega}^d(Q)$ for $\omega \neq \{0\}$ (the case of $\{0\}$ will be elementary) will be to understand the asymptotic behaviour of $L_{i,\omega}^d(Q)$ as Q heads off to infinity along a ray inside ω . In fact, the following lemma will show that this term is eventually zero. We will then be able to compute $L_{i,\omega}^d(Q)$ for arbitrary Q by understanding how it changes as Q heads out along various rays.

LEMMA 5.51. *Let $\omega \in \Sigma$, and let $v \in \omega$ be non-zero (hence ruling out $\omega = \{0\}$). Then*

$$(5.32) \quad \lim_{s \rightarrow \infty} L_{i,\omega}^d(Q + sv) = 0.$$

PROOF. We first note that with $\omega \neq \{0\}$,

$$(5.33) \quad \text{if } m \in \tilde{\omega}_d, \text{ then } r(m) \in \bigcup_{\substack{\sigma \supseteq \omega \\ \sigma \in \Sigma}} \text{Int}(\sigma).$$

Next, for sufficiently large s , $Q + sv$ lies in an unbounded connected component \mathcal{C} of $M_{\mathbb{R}} \setminus \text{Supp}(\mathfrak{D})$, where $\mathfrak{D} = \mathfrak{D}(\Sigma, P_1, \dots, P_k)$. By taking s sufficiently large, we can assume \mathcal{C} is the last component entered as $s \rightarrow \infty$. To show (5.32), it will be enough to show that if $Q + sv \in \mathcal{C}$, there exists a convex cone $K' \subseteq M_{\mathbb{R}}$ with

$$K' \cap \bigcup_{\substack{\sigma \supseteq \omega \\ \sigma \in \Sigma}} \text{Int}(\sigma) = \emptyset$$

such that $W_k(Q + sv) - W_0(Q + sv)$ only contains monomials z^m with $r(m) \in K'$. It then follows that all monomials z^m in

$$\exp((W_k(Q + sv) - W_0(Q + sv))/\hbar)$$

satisfy $r(m) \in K'$, and hence by (5.33), $m \notin \tilde{\omega}_d$. This implies (5.32).

So we study monomials z^m appearing in $W_k(Q+sv) - W_0(Q+sv)$ and construct a cone K' with the desired properties. We will make use of the asymptotic cone to the closure $\overline{\mathcal{C}}$ of \mathcal{C} , $\text{Asym}(\overline{\mathcal{C}})$, which is defined to be the Hausdorff limit $\lim_{\epsilon \rightarrow 0} \epsilon \overline{\mathcal{C}}$. Note that the connected components of $M_{\mathbb{R}} \setminus \mathfrak{D}(\Sigma, P_1)$ are $P_1 - \text{Int}(\sigma)$ where σ runs over the maximal cones of Σ . Since $\text{Supp}(\mathfrak{D}(\Sigma, P_1)) \subseteq \text{Supp} \mathfrak{D}$, one sees that $\text{Asym}(\overline{\mathcal{C}})$ is contained in some cone $-\sigma$ with $\sigma \in \Sigma$ maximal and $(-\sigma) \cap \omega \neq \{0\}$. Note also that $\text{Asym}(\overline{\mathcal{C}})$ is a ray if the unbounded edges of $\overline{\mathcal{C}}$ are parallel. Let $\mathfrak{d}_1, \mathfrak{d}_2$ denote the two unbounded edges of $\overline{\mathcal{C}}$.

Now for general s , a term cz^m in $W_k(Q+sv)$ corresponds to a broken line β with given data $-\infty = t_0 < \dots < t_p = 0$, $m_i^\beta \in T_\Sigma$ as in Definition 5.30, and $m = m_p^\beta$. If $-r(m) \notin \mathbb{R}_{>0}v$, then for s sufficiently large, with $+$ denoting Minkowski sum,

$$Q + sv \notin \mathbb{R}_{\geq 0}(-r(m)) + (\partial \overline{\mathcal{C}} \setminus (\mathfrak{d}_1 \cup \mathfrak{d}_2)).$$

Indeed, $\partial \overline{\mathcal{C}} \setminus (\mathfrak{d}_1 \cup \mathfrak{d}_2)$ is bounded, so the asymptotic cone of the right-hand side is $\mathbb{R}_{\geq 0}(-r(m))$, which does not contain v by assumption. Thus, taking a sufficiently large s , we note that β cannot last enter \mathcal{C} via $\partial \overline{\mathcal{C}} \setminus (\mathfrak{d}_1 \cup \mathfrak{d}_2)$ since the last line segment of β is in the direction $-r(m)$. So for sufficiently large s , β must enter $\overline{\mathcal{C}}$ by crossing one of \mathfrak{d}_1 or \mathfrak{d}_2 . In what follows, we will not need to study the case $-r(m) \in \mathbb{R}_{>0}v$ as the cone K' we construct will always contain $-v$.

We can now assume that for large s , β enters $\mathcal{C} = \mathcal{C}_n$ from another unbounded connected component \mathcal{C}_{n-1} of $M_{\mathbb{R}} \setminus \text{Supp}(\mathfrak{D})$. Necessarily, the m_i^β attached to β while β passes through \mathcal{C}_{n-1} satisfies $-r(m_i^\beta) \notin \text{Asym}(\overline{\mathcal{C}_{n-1}})$. Indeed, otherwise β could not hit an unbounded edge of $\overline{\mathcal{C}_{n-1}}$. Again, for large enough s , one sees similarly that β must enter \mathcal{C}_{n-1} through the other unbounded edge of $\overline{\mathcal{C}_{n-1}}$, and we can then continue this process inductively, with β passing only through unbounded edges via a sequence of unbounded components $\mathcal{C}_0, \dots, \mathcal{C}_n$. When β bends, it then always bends outward, as depicted in Figure 19. From this we make the following two observations:

- (C1) If the edges corresponding to \mathfrak{d}_1 and \mathfrak{d}_2 of $\text{Asym}(\overline{\mathcal{C}})$ are generated by v_1, v_2 respectively (possibly $v_1 = v_2$) and β enters \mathcal{C} by crossing \mathfrak{d}_i , then $-r(m)$ lies in a half-plane with boundary $\mathbb{R}v_i$ containing $\text{Asym}(\overline{\mathcal{C}})$; otherwise, β cannot reach the interior of $\overline{\mathcal{C}}$.
- (C2) For any j , $1 \leq j \leq p$, $-r(m)$ lies in the half-plane with boundary $\mathbb{R}r(m_j^\beta)$ containing v_i corresponding to the edge \mathfrak{d}_i that β crosses to enter \mathcal{C} . This follows from the behaviour described above about how β bends.

Without loss of generality, let us assume for the ease of drawing pictures that $\omega = \rho_2$ or $\rho_1 + \rho_2$ and $\text{Asym}(\overline{\mathcal{C}}) \subseteq -(\rho_0 + \rho_1)$. See Figure 20. Note that, as depicted there, we must have $v_2 \in \rho_1 + \rho_2$.

We analyze the possibilities for β : we have three cases, based on whether the initial direction of β is $-m_1, -m_2$, or $-m_0$.

Case 1. $r(m_0^\beta) = m_1$. Then β must enter \mathcal{C} via \mathfrak{d}_2 . By (C1), $-r(m)$ lies in the half-plane with boundary $\mathbb{R}v_2$ containing $\text{Asym}(\overline{\mathcal{C}})$, and by (C2), $-r(m)$ lies in the half-plane with boundary $\mathbb{R}m_1$ containing $\text{Asym}(\overline{\mathcal{C}})$. Thus

$$-r(m) \in (-\mathbb{R}_{\geq 0}m_1 + \mathbb{R}_{\geq 0}v_2).$$

Case 2. $r(m_0^\beta) = m_2$. Then either $\mathbb{R}_{\geq 0}m_2 \subseteq \text{Asym}(\overline{\mathcal{C}})$ or $\text{Asym}(\overline{\mathcal{C}}) \subseteq \rho_1 + \rho_2$ since $v \in \text{Asym}(\overline{\mathcal{C}})$. In the first case, β has no opportunity to bend, and so corresponds to the monomial x_2 , which doesn't appear in $W_k(Q+sv) - W_0(Q+sv)$.




FIGURE 19




FIGURE 20

In the second case, β bends at time t_1 as it crosses a ray $\mathfrak{d} \in \mathfrak{D}$ with $f_{\mathfrak{d}} = 1 + c_{\mathfrak{d}}z^{m_{\mathfrak{d}}}$ with $-r(m_{\mathfrak{d}}) \in \text{Int}(\rho_1 + \rho_2)$. Now $r(m_1^{\beta}) = m_2 + r(m_{\mathfrak{d}})$, so it follows that $-r(m_1^{\beta}) \in \rho_1 + \rho_2$. (Here we use integrality of $m_{\mathfrak{d}}$ and $m_2 = (0, 1)$.) Thus by (C1) and (C2),

$$-r(m) \in (\mathbb{R}_{\geq 0}m_1 + \mathbb{R}_{\geq 0}v_1).$$

Case 3. $r(m_0^\beta) = m_0$. In this case β must enter \mathcal{C} through the edge \mathfrak{d}_1 since $\text{Asym}(\overline{\mathcal{C}}) \subseteq -(\rho_0 + \rho_1)$. Then one sees from (C1) and (C2) that

$$-r(m) \in (\mathbb{R}_{\geq 0}(-m_0) + \mathbb{R}_{\geq 0}v_1).$$

We now see that if $\mathbb{R}_{\geq 0}m_2 \subseteq \text{Asym}(\overline{\mathcal{C}})$, (which always happens if v is proportional to m_2 , in particular when $\omega = \rho_2$), then of these three cases, only cases 1 and 3 can occur, and in fact $r(m), -v \in \rho_0 + \rho_1$. Thus $K' = \rho_0 + \rho_1$ is the desired cone, proving the claim in this case.

If $\mathbb{R}_{\geq 0}m_2 \not\subseteq \text{Asym}(\mathcal{C})$, then v is not proportional to m_2 and $\omega = \rho_1 + \rho_2$. In this case, the above three cases show that $-r(m)$ is always contained in the upper half-plane. Thus K' the lower half-plane is the desired cone, proving the claim in this case. \square

5.5.4. Wall crossing for $L_{i,\omega}^d(Q)$. The next step is to explain how $L_{i,\omega}^d(Q)$ depends on Q via a wall-crossing formula. Combining this formula with Lemma 5.51 allows us to describe $L_{i,\omega}^d(Q)$ as a sum of contributions from various wall-crossings.

DEFINITION 5.52. Let $\mathfrak{D} = \mathfrak{D}(\Sigma, P_1, \dots, P_k)$. Let $\mathcal{C}_1, \mathcal{C}_2$ be two connected components of $M_{\mathbb{R}} \setminus \text{Supp}(\mathfrak{D})$ with $\dim \overline{\mathcal{C}}_1 \cap \overline{\mathcal{C}}_2 = 1$. Let $Q_i \in \mathcal{C}_i$ be general points, and let γ be a path from Q_1 to Q_2 , passing through $\text{Supp}(\mathfrak{D})$ only at one time t_0 , with $\gamma(t_0) \notin \text{Sing}(\mathfrak{D})$. Let $\mathfrak{d} \in \mathfrak{D}$ be a ray with $\gamma(t_0) \in \mathfrak{d}$, and let $n_{\mathfrak{d}} \in N$ be a primitive vector which is orthogonal to \mathfrak{d} and satisfies $\langle n_{\mathfrak{d}}, \gamma'(t_0) \rangle < 0$. Writing $f_{\mathfrak{d}} = 1 + c_{\mathfrak{d}}z^{m_{\mathfrak{d}}}$, note that

$$\theta_{\gamma, \mathfrak{d}}(z^m) = z^m + c_{\mathfrak{d}} \langle n_{\mathfrak{d}}, r(m) \rangle z^{m+m_{\mathfrak{d}}}.$$

Now take a pair $\omega \subsetneq \tau$ with $\omega, \tau \in \Sigma$ and $\dim \tau = \dim \omega + 1$. Note that there is a unique index $j \in \{0, 1, 2\}$ such that $m_j \notin \omega$ but $m_j \in \tau$; call this index $j(\omega, \tau)$. Then define

$$L_{i, \mathfrak{d}, \gamma, \omega \rightarrow \tau}^d := \sum_{(c, \nu, m)} \langle n_{\mathfrak{d}}, m_{j(\omega, \tau)} \rangle c_{\mathfrak{d}} c D_i(d, m + m_{\mathfrak{d}} + t_{j(\omega, \tau)}) \hbar^{-(\nu + 3d - |m + m_{\mathfrak{d}}|)},$$

where the sum is over all $(c, \nu, m) \in S_k(Q_1)$ such that $m + m_{\mathfrak{d}} \in \tilde{\omega}_d$ but $m + m_{\mathfrak{d}} + t_{j(\omega, \tau)} \in \tilde{\tau}_d$. If $(c, \nu, m) \in S_k(Q_1)$ satisfies this condition, then we say *the term $c \hbar^{-\nu} z^m$ contributes to $L_{i, \mathfrak{d}, \gamma, \omega \rightarrow \tau}^d$* .

Define

$$L_{i, \gamma, \omega \rightarrow \tau}^d := \sum_{\mathfrak{d}} L_{i, \mathfrak{d}, \gamma, \omega \rightarrow \tau}^d,$$

where the sum is over all $\mathfrak{d} \in \mathfrak{D}$ with $\gamma(t_0) \in \mathfrak{d}$.

For an arbitrary path γ in $M_{\mathbb{R}} \setminus \text{Sing}(\mathfrak{D})$ with $\gamma(0) = Q, \gamma(1) = Q'$, choose a partition of $[0, 1]$, $0 = t_0 < t_1 < \dots < t_n = 1$, such that $\gamma|_{[t_{j-1}, t_j]}$ is a path of the sort considered above, connecting endpoints in adjacent connected components. Then define

$$L_{i, \gamma, \omega \rightarrow \tau}^d := \sum_{j=1}^n L_{i, \gamma|_{[t_{j-1}, t_j]}, \omega \rightarrow \tau}^d.$$

\square

Here then are the relevant wall-crossing formulas.

LEMMA 5.53. *Let P_1, \dots, P_k be general. Let γ be a path in $M_{\mathbb{R}} \setminus \text{Sing}(\mathfrak{D})$ with $\gamma(0) = Q$, $\gamma(1) = Q'$. Then for $\dim \rho = 1$, $\rho \in \Sigma$,*

$$(5.34) \quad L_{i,\rho}^d(Q') - L_{i,\rho}^d(Q) = L_{i,\gamma,\{0\} \rightarrow \rho}^d - \sum_{\substack{\sigma \in \Sigma \\ \rho \subsetneq \sigma}} L_{i,\gamma,\rho \rightarrow \sigma}^d$$

while for $\dim \sigma = 2$, $\sigma \in \Sigma$,

$$(5.35) \quad L_{i,\sigma}^d(Q') - L_{i,\sigma}^d(Q) = \sum_{\substack{\rho \in \Sigma \\ \dim \rho = 1 \\ \rho \subsetneq \sigma}} L_{i,\gamma,\rho \rightarrow \sigma}^d.$$

PROOF. It is enough to show this for γ a short path connecting Q and Q' in two adjacent components \mathcal{C}_1 and \mathcal{C}_2 of $M_{\mathbb{R}} \setminus \text{Supp}(\mathfrak{D})$ as in Definition 5.52. Suppose that at time t_0 , $\gamma(t_0) \in \mathfrak{d}_1 \cap \dots \cap \mathfrak{d}_s$ for rays $\mathfrak{d}_1, \dots, \mathfrak{d}_s \in \mathfrak{D}$. Of course, $\dim \mathfrak{d}_i \cap \mathfrak{d}_j = 1$. We can then write, for $n_{\mathfrak{d}} = n_{\mathfrak{d}_i}$ for any i ,

$$\begin{aligned} \theta_{\gamma,\mathfrak{D}}(z^m) &= z^m \prod_{i=1}^s f_{\mathfrak{d}_i}^{\langle n_{\mathfrak{d}}, r(m) \rangle} \\ &= z^m \prod_{i=1}^s (1 + c_{\mathfrak{d}_i} \langle n_{\mathfrak{d}}, r(m) \rangle z^{m_{\mathfrak{d}_i}}) \\ &= z^m + \sum_{i=1}^s c_{\mathfrak{d}_i} \langle n_{\mathfrak{d}}, r(m) \rangle z^{m+m_{\mathfrak{d}_i}}. \end{aligned}$$

Here the last equality follows from $c_{\mathfrak{d}_i} c_{\mathfrak{d}_j} = 0$ for $i \neq j$. This is the case by the assumption that P_1, \dots, P_k are general. Indeed, if $c_{\mathfrak{d}_i} c_{\mathfrak{d}_j} \neq 0$, then the Maslov index zero trees h_i and h_j corresponding to \mathfrak{d}_i and \mathfrak{d}_j would have $I(h_i) \cap I(h_j) = \emptyset$. However, a generic perturbation of the marked points with indices in $I(h_i)$ would deform \mathfrak{d}_i without deforming \mathfrak{d}_j , so that $\mathfrak{d}_i \cap \mathfrak{d}_j = \emptyset$.

Now

$$W_k(Q') = \theta_{\gamma,\mathfrak{D}}(W_k(Q))$$

by Theorem 5.35. Using the expansion (5.22) and $W_0(Q) = y_0 + \sum_{j=0}^2 z^{t_j}$,

$$\begin{aligned} &\exp((W_k(Q') - W_0(Q'))/\hbar) \\ &= \exp((\theta_{\gamma,\mathfrak{D}}(W_k(Q)) - W_0(Q))/\hbar) \\ &= \theta_{\gamma,\mathfrak{D}}(\exp((W_k(Q) - W_0(Q))/\hbar)) \cdot \exp((\theta_{\gamma,\mathfrak{D}}(W_0(Q)) - W_0(Q))/\hbar) \\ &= \theta_{\gamma,\mathfrak{D}}\left(\sum_{(c,m,\nu) \in S_k(Q)} c \hbar^{-\nu} z^m\right) \left(1 + \hbar^{-1} \sum_{\ell=1}^s \sum_{j=0}^2 c_{\mathfrak{d}_\ell} \langle n_{\mathfrak{d}}, m_j \rangle z^{m_{\mathfrak{d}_\ell} + t_j}\right) \\ &= \exp((W_k(Q) - W_0(Q))/\hbar) \\ &\quad + \sum_{(c,\nu,m) \in S_k(Q)} \sum_{\ell=1}^s \left(c_{\mathfrak{d}_\ell} c \hbar^{-\nu} (\langle n_{\mathfrak{d}}, r(m) \rangle z^{m+m_{\mathfrak{d}_\ell}} \right. \\ &\quad \quad \left. + \hbar^{-1} \sum_{j=0}^2 \langle n_{\mathfrak{d}}, m_j \rangle z^{m+m_{\mathfrak{d}_\ell} + t_j} \right). \end{aligned}$$

We interpret this as follows. For each $(c, \nu, m) \in S_k(Q)$ and each ℓ , look at the four terms

$$c_{\partial_\ell} c \hbar^{-\nu} (\langle n_{\partial}, r(m) \rangle z^{m+m_{\partial_\ell}} + \hbar^{-1} \sum_{j=0}^2 \langle n_{\partial}, m_j \rangle z^{m+m_{\partial_\ell}+t_j}).$$

These four terms contribute the expression

$$c_{\partial_\ell} c \hbar^{-(3d+\nu-|m+m_{\partial_\ell}|)} (\langle n_{\partial}, r(m) \rangle D_i(d, m+m_{\partial_\ell}) + \sum_{j=0}^2 \langle n_{\partial}, m_j \rangle D_i(d, m+m_{\partial_\ell}+t_j))$$

to $L_i^d(Q')$. One can check that in fact this total contribution is zero, either by direct but tedious checking from the formulas for D_i , or by applying Lemma 5.40 with $f = c \hbar^{-(\nu-1)} z^m$ and $\theta = \theta_{\gamma, \partial}$.

Now if $m+m_{\partial_\ell}$ and $m+m_{\partial_\ell}+t_j$, $0 \leq j \leq 2$, all lie in the same $\tilde{\omega}_d$, then these terms produce no total contribution to $L_{i,\tau}^d(Q')$ for any $\tau \in \Sigma$, including $\tau = \omega$. On the other hand, these four terms can contribute to different $L_{i,\omega}^d(Q')$'s if $m+m_{\partial_\ell}$ and $m+m_{\partial_\ell}+t_j$, $j=0,1,2$, don't all lie in $\tilde{\omega}_d$ for the same $\omega \in \Sigma$. This can happen only if $m+m_{\partial} \in \tilde{\omega}_d$ but $m+m_{\partial}+t_j \in \tilde{\tau}_d$ for some j with $\omega \subsetneq \tau \in \Sigma$ with $\dim \tau = \dim \omega + 1$ and $m_j \in \tau$, $m_j \notin \omega$. In this case, $L_{i,\tau}^d(Q') - L_{i,\tau}^d(Q)$ has a contribution of the form $cc_{\partial_\ell} \langle n_{\partial}, m_j \rangle \hbar^{-(3d+\nu-|m+m_{\partial_\ell}|)} D_i(d, m+m_{\partial_\ell}+t_j)$. Thus $L_{i,\omega}^d(Q') - L_{i,\omega}^d(Q)$ must have a contribution coming from the same term, but with opposite sign. This gives the lemma. \square

We can now use the asymptotic behaviour of the expressions $L_{i,\omega}^d(Q)$ and the above wall-crossing formula to rewrite the needed expressions:

LEMMA 5.54. *Let γ_j be the straight line path joining Q with $Q+sm_j$ for $s \gg 0$. Let $\gamma_{j,j+1}$ be the loop based at Q which passes linearly from Q to $Q+sm_j$, then takes a large circular arc to $Q+sm_{j+1}$, and then proceeds linearly from $Q+sm_{j+1}$ to Q . Here we take j modulo 3, and $\gamma_{j,j+1}$ is always a counterclockwise loop. Let $\sigma_{j,j+1} = \rho_j + \rho_{j+1}$, a two-dimensional cone in Σ . Then*

$$(5.36) \quad L_i^d(Q) = L_{i,\{0\}}^d(Q) - \sum_{j=0}^2 L_{i,\gamma_j,\{0\} \rightarrow \rho_j}^d - \sum_{j=0}^2 L_{i,\gamma_j,j+1,\rho_{j+1} \rightarrow \sigma_{j,j+1}}^d.$$

PROOF. By Lemma 5.51, $L_{i,\sigma}^d(Q+sm_j) = 0$ for any $\sigma \in \Sigma$ with $\rho_j \subseteq \sigma$. Thus by (5.34) and (5.35), we have

$$\begin{aligned} L_{i,\rho_j}^d(Q) &= -L_{i,\gamma_j,\{0\} \rightarrow \rho_j}^d + \sum_{\substack{\sigma \in \Sigma \\ \rho_j \subsetneq \sigma}} L_{i,\gamma_j,\rho_j \rightarrow \sigma}^d, \\ L_{i,\sigma_{j,j+1}}^d(Q) &= - \sum_{\substack{\rho \in \Sigma \\ \dim \rho = 1 \\ \rho \subsetneq \sigma_{j,j+1}}} L_{i,\gamma_j,\rho \rightarrow \sigma_{j,j+1}}^d. \end{aligned}$$

Note we have broken symmetry for the second equation.

Adding together contributions from the ρ_j 's and σ_{j_1, j_2} 's, we see from Lemma 5.50 that

$$\begin{aligned} L_i^d(Q) - L_{i, \{0\}}^d(Q) &= - \sum_{j=0}^2 L_{i, \gamma_j, \{0\} \rightarrow \rho_j}^d \\ &\quad - (L_{i, \gamma_0, \rho_1 \rightarrow \sigma_{0,1}}^d - L_{i, \gamma_1, \rho_1 \rightarrow \sigma_{0,1}}^d) \\ &\quad - (L_{i, \gamma_1, \rho_2 \rightarrow \sigma_{1,2}}^d - L_{i, \gamma_2, \rho_2 \rightarrow \sigma_{1,2}}^d) \\ &\quad - (L_{i, \gamma_2, \rho_0 \rightarrow \sigma_{2,0}}^d - L_{i, \gamma_0, \rho_0 \rightarrow \sigma_{2,0}}^d). \end{aligned}$$

Again by Lemma 5.51, it follows that the contribution to $L_{i, \gamma_j, j+1, \rho_{j+1} \rightarrow \sigma_{j, j+1}}^d$ from the large circular arc is zero. Hence

$$L_i^d(Q) - L_{i, \{0\}}^d(Q) = - \sum_{j=0}^2 L_{i, \gamma_j, \{0\} \rightarrow \rho_j}^d - \sum_{j=0}^2 L_{i, \gamma_j, j+1, \rho_{j+1} \rightarrow \sigma_{j, j+1}}^d,$$

the desired result. \square

5.5.5. The final steps. We have now finished massaging $L_i^d(Q)$ into a usable form. We will find that each term in (5.36) has a geometric interpretation in the equations (5.23)-(5.25): $L_{i, \{0\}}^d(Q)$ yields the contributions to the right-hand sides of these equations coming from expressions of the form $\langle \dots \rangle_{d, \{0\}}^{\text{trop}}$ in the sense of (5.4), $-L_{i, \gamma_j, \{0\} \rightarrow \rho_j}^d$ will yield contributions coming from $\langle \dots \rangle_{d, \rho_j}^{\text{trop}}$, and $-L_{i, \gamma_j, j+1, \rho_{j+1} \rightarrow \sigma_{j, j+1}}^d$ will yield contributions coming from $\langle \dots \rangle_{d, \sigma_{j, j+1}}^{\text{trop}}$.

We first deal with the simplest term, $L_{i, \{0\}}^d$. The idea in this case is very simple: the period integrals simply tell us all the ways of gluing together tropical disks with boundary Q to obtain a tropical curve with a vertex of high valency mapping to Q .

LEMMA 5.55.

$$\begin{aligned} L_{i, \{0\}}^d(Q) &= \delta_{0,d} \delta_{0,i} + \sum_{\nu \geq i} \sum_{\substack{I \subseteq \{1, \dots, k\} \\ I = \{i_1, \dots, i_{3d-2+i-\nu}\} \\ i_1 < \dots < i_{3d-2+i-\nu}}} \langle P_{i_1}, \dots, P_{i_{3d-2+i-\nu}}, \psi^\nu S \rangle_{d, \{0\}}^{\text{trop}} u_I \hbar^{-(\nu+2-i)} \end{aligned}$$

where $S = Q, L$ the tropical line with vertex Q , or $M_{\mathbb{R}}$ in the cases $i = 0, 1$ and 2.

PROOF. If $d = 0$, the only element $(c, \nu, m) \in S_k$ which contributes to $L_{i, \{0\}}^d$ is $(1, 0, 0)$, corresponding to the constant term in $e^{(W_k(Q) - W_0(Q))/\hbar}$. This contributes 1 if $i = 0$ and 0 otherwise, hence the term $\delta_{0,d} \delta_{0,i}$.

Now assume $d \neq 0$. Let $(c, \nu, m) \in S_k$ with $m = \sum_{i=0}^2 n_i t_i$. Then (c, ν, m) contributes to $L_{i, \{0\}}^d$ only if $n_0, n_1, n_2 \leq d$. Write

$$c \hbar^{-\nu} z^m = \hbar^{-\nu} \prod_{i=1}^{\nu} \text{Mono}(h_i) = \hbar^{-\nu} \prod_{i=1}^{\nu} \text{Mult}(h_i) z^{\Delta(h_i)} u_{I(h_i)}$$

for $h_i : \Gamma'_i \rightarrow M_{\mathbb{R}}$ Maslov index two disks with boundary Q for $1 \leq i \leq \nu$. Let Γ be the graph obtained by identifying the outgoing vertices $V_{\text{out}, i}$ of $\Gamma'_1, \dots, \Gamma'_\nu$ to get a single vertex V_{out} and then adding $(d - n_0) + (d - n_1) + (d - n_2) + 1$ additional unbounded edges with vertex V_{out} . We define $h : \Gamma \rightarrow M_{\mathbb{R}}$ to be h_i on each subgraph $\Gamma'_i \subseteq \Gamma$. Furthermore, for $0 \leq i \leq 2$, h maps $d - n_i$ of the new

unbounded edges to the ray $Q + \mathbb{R}_{\geq 0}m_i$. Finally, the last unbounded ray is labelled with an x and is contracted by h .

I claim that h is now a balanced tropical curve. Indeed, the balancing condition clearly holds at all vertices of Γ except possibly at V_{out} . Consider the tangent vectors to the various edges with endpoint V_{out} . If $m^{\text{prim}}(h_i)$ is a primitive tangent vector to the outgoing edge $h(E_{\text{out},i})$ of Γ'_i pointing away from Q , then by summing the balancing condition over all vertices of Γ'_i not including $V_{\text{out},i}$, one obtains

$$r(\Delta(h_i)) = w(E_{\text{out},i})m^{\text{prim}}(h_i).$$

There are then in addition $d - n_i$ edges with weight one attached to V_{out} in the direction m_i . Now by assumption,

$$\sum_{i=1}^{\nu} \Delta(h_i) + \sum_{j=0}^2 (d - n_j)t_j = d(t_0 + t_1 + t_2),$$

so applying r we get

$$\sum_{i=1}^{\nu} r(\Delta(h_i)) + \sum_{j=0}^2 (d - n_j)m_j = 0.$$

But this is precisely the balancing condition at V_{out} .

The contribution of this term $c\hbar^{-\nu}z^m$ to $L_{i,\{0\}}^d$ is then

$$\begin{aligned} & \hbar^{-(3d+\nu-n_0-n_1-n_2)} D_i(d, n_0, n_1, n_2) \prod_{i=1}^{\nu} \text{Mult}(h_i)u_{I(h_i)} \\ &= \hbar^{-(3d+\nu-n_0-n_1-n_2)} u_{I(h)} \text{Mult}_x^i(h) \prod_{\substack{V \in \Gamma^{[0]} \\ V \notin E_x}} \text{Mult}_V(h), \end{aligned}$$

comparing the definitions of $\text{Mult}_x^i(h)$ and $D_i(d, n_0, n_1, n_2)$. Note that the valency $\text{Val}(V_{\text{out}})$ of the vertex V_{out} in h is $\nu + 3d - (n_0 + n_1 + n_2) + 1$. Suppose that $I(h) = \{i_1, \dots, i_{3d-2+i-\nu'}\}$ for some ν' . Noting that h is obtained by gluing $\text{Val}(V_{\text{out}}) - 1$ Maslov index two disks, we see that

$$\begin{aligned} \text{Val}(V_{\text{out}}) - 1 &= \sum_{i=1}^{\nu} (|\Delta(h_i)| - \#I(h_i)) + (d - n_1) + (d - n_2) + (d - n_3) \\ &= 3d - (3d - 2 + i - \nu') = \nu' + 2 - i. \end{aligned}$$

Then the curve h contributes precisely the correct contribution, as given by Definition 5.13, (1) (a), (2) (b), or (3) (d), to

$$(5.37) \quad \langle P_{i_1}, \dots, P_{i_{3d-2+i-\nu'}}, \psi^{\nu'} S \rangle_{d,\{0\}}^{\text{trop}} u_{I(h)} \hbar^{-(\nu'+2-i)}.$$

Conversely, given any curve h contributing to (5.37), it follows from Lemma 5.12, (3), that h is obtained by gluing together some collection of Maslov index two disks with endpoint Q in the manner described above. Now some of these Maslov index two disks may not have any marked points: these are rays in the directions m_0, m_1 and m_2 . There are at most d of these, so define the non-negative numbers n_0, n_1 and n_2 so that $d - n_i$ is the number of rays appearing in the direction m_i . The remaining Maslov index two disks, h_1, \dots, h_{ν} , all have marked points, and hence there is a term $\hbar^{-\nu} \prod_{i=1}^{\nu} \text{Mono}(h_i)$ appearing in $\exp((W_k(Q) - W_0(Q))/\hbar)$. This

term will make the same contribution to $L_{i,\{0\}}^d$ that the curve h makes to (5.37), by the above argument. \square

Next, the interpretation of $-L_{i,\gamma_j,\{0\}\rightarrow\rho_j}^d$:

LEMMA 5.56.

$$-L_{i,\gamma_j,\{0\}\rightarrow\rho_j}^d = \sum_{\nu \geq i-1} \sum_{\substack{I \subseteq \{1, \dots, k\} \\ I = \{i_1, \dots, i_{3d-2+i-\nu}\} \\ i_1 < \dots < i_{3d-2+i-\nu}}} \langle P_{i_1}, \dots, P_{i_{3d-2+i-\nu}}, \psi^\nu S \rangle_{d,\rho_j}^{\text{trop}} u_I \hbar^{-(\nu+2-i)}$$

with $S = Q, L$ or $M_{\mathbb{R}}$ for $i = 0, 1$ and $i = 2$ respectively. Here, as usual, L is a tropical line with vertex Q .

PROOF. This is vacuous for $i = 0$, as both sides are zero, so we assume $i \geq 1$. Without loss of generality, consider $L_{i,\gamma_0,\{0\}\rightarrow\rho_0}^d$. This quantity is a sum of contributions from each point $P \in Q + (\rho_0 \setminus \{0\})$ which is the intersection of $Q + \rho_0$ with a ray $\mathfrak{d} \in \mathfrak{D} = \mathfrak{D}(\Sigma, P_1, \dots, P_k)$. Write

$$f_{\mathfrak{d}} = 1 + c_{\mathfrak{d}} z^{m_{\mathfrak{d}}}.$$

Let us consider the contribution to $L_{i,\gamma_0,\{0\}\rightarrow\rho_0}^d$ from a small segment γ of γ_0 which only crosses \mathfrak{d} . Let γ run from Q_1 to Q_2 . Now \mathfrak{d} corresponds to a Maslov index zero tree passing through P , and by cutting it at P , we obtain a Maslov index zero disk $h_1 : \Gamma'_1 \rightarrow M_{\mathbb{R}}$ with boundary P . Then

$$f_{\mathfrak{d}} = 1 + w_{\Gamma'_1}(E_{\text{out},1}) \text{Mult}(h_1) z^{\Delta(h_1)} u_{I(h_1)}.$$

Furthermore a term $cz^m \hbar^{-\nu}$ in $\exp((W_k(Q_1) - W_0(Q_1))/\hbar)$ arises from ν distinct Maslov index two disks with boundary Q_1 , say $h_2, \dots, h_{\nu+1}$ (each with at least one marked point), and

$$cz^m \hbar^{-\nu} = \hbar^{-\nu} \prod_{i=2}^{\nu+1} \text{Mult}(h_i) z^{\Delta(h_i)} u_{I(h_i)}.$$

In order for this term to contribute to $L_{i,\gamma,\{0\}\rightarrow\rho_0}^d$, $m + m_{\mathfrak{d}} = \sum_{i=1}^{\nu+1} \Delta(h_i)$ must be of the form $dt_0 + n_1 t_1 + n_2 t_2$ with $n_1, n_2 \leq d$. Assume this is the case. The disks $h_2, \dots, h_{\nu+1}$ deform to disks with boundary at P , which we also call $h_2, \dots, h_{\nu+1}$. Write these disks as $h_i : \Gamma'_i \rightarrow M_{\mathbb{R}}$. Each Γ'_i , $1 \leq i \leq \nu+1$, has a vertex $V_{\text{out},i}$.

Using this data, we can construct an actual tropical curve as follows. Let Γ be the graph obtained by identifying all the outgoing vertices $V_{\text{out},i}$ in $\Gamma'_1, \dots, \Gamma'_{\nu+1}$, to obtain a graph with a distinguished vertex V_{out} , and then attaching

$$(d - n_1) + (d - n_2) + 1$$

additional unbounded edges with vertex V_{out} . We then define $h : \Gamma \rightarrow M_{\mathbb{R}}$ to agree with h_i on $\Gamma'_i \subseteq \Gamma$. In addition, h takes the first $d - n_1$ new unbounded edges to $P + \mathbb{R}_{\geq 0} m_1$; the second $d - n_2$ new unbounded edges to $P + \mathbb{R}_{\geq 0} m_2$; and the last unbounded edge is contracted and marked with the label x . Note Γ has valency at V_{out} given by $\text{Val}(V_{\text{out}}) = \nu + 1 + (d - n_1) + (d - n_2) + 1$. Thus we obtain a parameterized curve $h : \Gamma \rightarrow M_{\mathbb{R}}$ with $h(x) = P$. The balancing condition needs to be checked at V_{out} , but as in the proof of Lemma 5.55, the fact that

$$\sum_{i=1}^{\nu+1} \Delta(h_i) + (d - n_1)t_1 + (d - n_2)t_2 = d(t_0 + t_1 + t_2)$$

shows that the balancing condition indeed holds at V_{out} .

The contribution of this term to $-L_{i,\gamma,\{0\}\rightarrow\rho_0}^d$ is

$$(5.38) \quad -\langle n_{\mathfrak{d}}, m_0 \rangle u_{I(h)} w_{\Gamma_1'}(E_{\text{out},1}) \left(\prod_{i=1}^{\nu+1} \text{Mult}(h_i) \right) D_i(d, d+1, n_1, n_2) \hbar^{-(\nu+3d-(d+n_1+n_2))}.$$

Note that $n_{\mathfrak{d}}$ is primitive, annihilates $r(m_{\mathfrak{d}})$, and must be positive on $-m_0$. Furthermore, the chosen isomorphism $\bigwedge^2 M \cong \mathbb{Z}$ identifies $w(E_{1,\text{out}})n_{\mathfrak{d}}$, up to sign, with $X_{r(m_{\mathfrak{d}})}$. Thus setting $m(h_1) = r(m_{\mathfrak{d}})$ as in Definition 5.13, we see that

$$-\langle n_{\mathfrak{d}}, m_0 \rangle w_{\Gamma_1'}(E_{\text{out},1}) = |m(h_1) \wedge m_0|.$$

Thus (5.38) coincides with

$$|m(h_1) \wedge m_0| u_{I(h)} D_i(d, d+1, n_1, n_2) \left(\prod_{\substack{V \in \Gamma^{[0]} \\ V \not\subseteq E_x}} \text{Mult}_V(h) \right) \hbar^{-(\text{Val}(V_{\text{out}})-2)}.$$

Now $D_i(d, d+1, n_1, n_2) = \text{Mult}_x^{i-1}(h)$ as defined in Definition 5.13 via direct comparison with the definitions of the D_i 's. Furthermore, if $I(h) = \{i_1, \dots, i_{3d-2+i-\nu'}\}$ for some ν' , we see that, as h is obtained by gluing one Maslov index zero disk to $\text{Val}(V_{\text{out}}) - 2$ Maslov index two disks, we have

$$\begin{aligned} \text{Val}(V_{\text{out}}) - 2 &= \sum_{i=1}^{\nu+1} (|\Delta(h_i)| - \#I(h_i)) + (d - n_1) + (d - n_2) \\ &= 3d - (3d - 2 + i - \nu') = \nu' + 2 - i. \end{aligned}$$

Thus, by Definition 5.13, the term under consideration contributes to $-L_{i,\gamma_0,\{0\}\rightarrow\rho_0}^d$ by exactly the same amount that the curve h contributes to

$$(5.39) \quad \langle P_{i_1}, \dots, P_{i_{3d-2+i-\nu'}}, \psi^{\nu'} S_{d,\rho_0}^{\text{trop}} u_{I(h)} \hbar^{-(\nu'+2-i)},$$

as desired.

Conversely, given any curve h contributing to $\langle P_{i_1}, \dots, P_{i_{3d-2+i-\nu'}}, \psi^{\nu'} S_{d,\rho_0}^{\text{trop}} \rangle$ with $h(E_x) = P \in Q + (\rho_0 \setminus \{0\})$, the procedure of Lemma 5.12, (2), shows that h must arise in precisely the way described above. Indeed, that lemma shows that h can be obtained by gluing together a number of Maslov index two disks and one Maslov index zero disk with boundary P . By the condition of Definition 5.13, (2), (a) or (3), (c), none of the unmarked Maslov index two disks can be parallel to m_0 . Thus we can assume that h is obtained by gluing together Maslov index two disks $h_2, \dots, h_{\nu+1}$, a Maslov index zero disk h_1 , $d - n_1$ rays in the direction m_1 and $d - n_2$ rays in the direction m_2 . Then $\hbar^{-\nu} \prod_{i=2}^{\nu+1} \text{Mono}(h_i)$ appears in $\exp((W_k(Q_1) - W_0(Q_1))/\hbar)$ for Q_1 a point near P on $Q + \rho_0$, and so we see this term makes the same contribution to $-L_{i,\gamma_j,\{0\}\rightarrow\rho_j}^d$ that the curve h makes to (5.39), as desired. \square

We now come to the most difficult case, interpreting $-L_{i,\gamma_j,j+1,\rho_{j+1}\rightarrow\sigma_{j,j+1}}^d$. The first step is to localize the calculation to singular points of $\mathfrak{D} = \mathfrak{D}(\Sigma, P_1, \dots, P_k)$.

LEMMA 5.57. *For each point $P \in \text{Sing}(\mathfrak{D})$, let γ_P be a small counterclockwise loop around P , small enough so it doesn't go around any other point of $\text{Sing}(\mathfrak{D})$.*

Then

$$L_{i,\gamma_{j,j+1},\rho_{j+1}\rightarrow\sigma_{j,j+1}}^d = \sum_{P \in \text{Sing}(\mathfrak{D}) \cap (Q + \sigma_{j,j+1})} L_{i,\gamma_P,\rho_{j+1}\rightarrow\sigma_{j,j+1}}^d.$$

PROOF. If \mathcal{C}_1 and \mathcal{C}_2 are closures of two connected components of $M_{\mathbb{R}} \setminus \text{Supp } \mathfrak{D}$ with $\dim \mathcal{C}_1 \cap \mathcal{C}_2 = 1$, and γ is a short path from $\text{Int}(\mathcal{C}_1)$ into $\text{Int}(\mathcal{C}_2)$ just crossing $\text{Int}(\mathcal{C}_1 \cap \mathcal{C}_2)$ once, then $L_{i,\gamma,\rho_{j+1}\rightarrow\sigma_{j,j+1}}^d$ is independent of γ and its endpoints. Furthermore, reversing the direction of γ changes the sign of $L_{i,\gamma_P,\rho_{j+1}\rightarrow\sigma_{j,j+1}}^d$. Let \mathcal{P} be the polyhedral decomposition of $M_{\mathbb{R}}$ induced by $\text{Supp}(\mathfrak{D})$, and let $\check{\mathcal{P}}$ be the dual cell decomposition. This is obtained by taking a point v_σ to be the barycentre of a maximal cell σ of \mathcal{P} . Then $\check{\mathcal{P}}$ consists of the zero-cells $\{v_\sigma\}$, the one-cells with endpoints v_σ and $v_{\sigma'}$ if $\dim \sigma \cap \sigma' = 1$, and two-cells having vertices $\{v_\sigma \mid P \in \sigma\}$ for $P \in \text{Sing}(\mathfrak{D})$. We can take γ_P to be the counterclockwise boundary of the two-cell in $\check{\mathcal{P}}$ corresponding to P . It then becomes clear that if we sum $L_{i,\gamma_P,\rho_{j+1}\rightarrow\sigma_{j,j+1}}^d$ over all singular points P in $Q + \sigma_{j,j+1}$, the contribution from traversing an edge with endpoints $v_\sigma, v_{\sigma'}$ of $\check{\mathcal{P}}$ appears twice with opposite signs provided $\sigma \cap \sigma'$ is contained in $\sigma_{j,j+1}$. The only surviving contributions come from edges of a path homotopic in $M_{\mathbb{R}} \setminus \text{Sing}(\mathfrak{D})$ to $\gamma_{j,j+1}$, hence the result. \square

LEMMA 5.58. *Let $P \in \text{Sing}(\mathfrak{D}) \cap (Q + \sigma_{j,j+1})$, and suppose that*

$$P \notin \{P_1, \dots, P_k\}.$$

Then

$$-L_{i,\gamma_P,\rho_{j+1}\rightarrow\sigma_{j,j+1}}^d = \sum_{\nu \geq 0} \sum_h \text{Mult}(h) u_{I(h)} \hbar^{-(\nu+2-i)},$$

where the sum is over all curves h contributing to $\langle P_{i_1}, \dots, P_{i_{3d-2+i-\nu}}, \psi^\nu S \rangle_{d,\sigma_{j,j+1}}^{\text{trop}}$ for various $i_1 < \dots < i_{3d-2+i-\nu}$ such that $h(E_x) = P$.

PROOF. Note that this is vacuous for $i = 0$ or 1 as both sides are zero, so we can assume $i = 2$. To save on typing, we set

$$L_{P,j} := L_{2,\gamma_P,\rho_{j+1}\rightarrow\sigma_{j,j+1}}^d.$$

Fix a base-point Q' near P . Consider a term $c\hbar^{-\nu} z^m$ in

$$\exp((W_k(Q') - W_0(Q'))/\hbar)$$

of the form

$$(5.40) \quad c\hbar^{-\nu} z^m = \hbar^{-\nu} \prod_{i=3}^{\nu+2} (\text{Mult}(h_i) z^{\Delta(h_i)} u_{I(h_i)})$$

where the h_i 's are Maslov index two disks with boundary Q' , but none of the h_i 's come from broken lines which bend near P . As a result, this term appears in $\exp((W_k(Q'') - W_0(Q''))/\hbar)$ for all Q'' general in a small open neighbourhood of P .

Suppose that such a term $c\hbar^{-\nu} z^m$ contributes to $L_{P,j}$ when γ_P crosses a ray $\mathfrak{d} \in \mathfrak{D}$ with $P \in \mathfrak{d}$, $P \neq \text{Init}(\mathfrak{d})$. But γ_P crosses this ray \mathfrak{d} twice, in opposite directions, so $c\hbar^{-\nu} z^m$ will contribute to $L_{P,j}$ twice, but with opposite signs. Thus these contributions cancel, and don't contribute to the total in $L_{P,j}$.

Thus we only need analyze contributions arising when γ_P crosses a ray \mathfrak{d} with $\text{Init}(\mathfrak{d}) = P$ or contributions from monomials as in (5.40) where some of the h_i 's come from broken lines which do bend near P .

Recall from the argument of Proposition 5.28 that we can assume we can decompose the set of rays of \mathfrak{D} passing through P as

$$\{\mathfrak{d}_1, \dots, \mathfrak{d}_n\} \cup \bigcup_{j=1}^m \mathfrak{D}_j$$

where $\mathfrak{d}_1, \dots, \mathfrak{d}_n$ do not have P as an initial point and correspond to trees h such that $I(h) \cap I(h') \neq \emptyset$ for any Maslov index zero tree h' with outgoing edge passing through P . On the other hand, \mathfrak{D}_i has three elements consisting of two rays which do not have P as an initial point and one ray which does, which corresponds to the tree obtained by gluing the two trees corresponding to the other two rays. Finally, if h corresponds to a ray in \mathfrak{D}_i and h' corresponds to a ray in \mathfrak{D}_j for $i \neq j$, then $I(h) \cap I(h') \neq \emptyset$.

As a consequence, to analyze the contributions to $L_{i, \gamma_P, \rho_{j+1} \rightarrow \sigma_{j, j+1}}$, the above discussion shows we can assume that there are precisely three rays, $\mathfrak{d}_1, \mathfrak{d}_2, \mathfrak{d}_3$ passing through P , with $\text{Init}(\mathfrak{d}_1), \text{Init}(\mathfrak{d}_2) \neq P$ and $\text{Init}(\mathfrak{d}_3) = P$. Now $\mathfrak{d}_1, \mathfrak{d}_2$ correspond to Maslov index zero trees passing through P , and by cutting them, we obtain Maslov index zero disks $h_i : \Gamma'_i \rightarrow M_{\mathbb{R}}$, $i = 1, 2$ with boundary P , and for $i = 1, 2$,

$$f_{\mathfrak{d}_i} = 1 + w_{\Gamma'_i}(E_{\text{out}, i}) \text{Mult}(h_i) z^{\Delta(h_i)} u_{I(h_i)}.$$

We now analyze how additional terms $c\hbar^{-\nu} z^m$ which can contribute to $L_{P, j}$ may arise. In what follows, assume that $c\hbar^{-\nu} z^m$ is as in (5.40) in which none of the broken lines corresponding to $h_3, \dots, h_{\nu+2}$ bend at $\mathfrak{d}_1, \mathfrak{d}_2$ or \mathfrak{d}_3 .

Write

$$m + \Delta(h_1) + \Delta(h_2) = \sum_{j=0}^2 n_j t_j.$$

We have the following possibilities of additional contributions:

- (I) $c\hbar^{-\nu} z^m$ may contribute to $L_{P, j}$ when γ_P crosses \mathfrak{d}_3 . This contribution can only occur if $n_{j+2} \leq n_j = d < n_{j+1}$.
- (II) After crossing \mathfrak{d}_1 , new terms of the form (leaving off the coefficients) $z^{m+\Delta(h_1)}$ and $z^{m+\Delta(h_1)+t_\ell}$, $\ell = 0, 1, 2$, may appear in $\exp((W_k - W_0)/\hbar)$. Thus, when we cross \mathfrak{d}_2 , these new terms may contribute to $L_{P, j}$. Note that $z^{m+\Delta(h_1)}$ only contributes when crossing \mathfrak{d}_2 if $n_{j+2} \leq n_j = d < n_{j+1}$. The term $z^{m+\Delta(h_1)+t_j}$ only contributes if $n_{j+2} \leq d$, $n_j = d - 1$ and $d < n_{j+1}$. The term $z^{m+\Delta(h_1)+t_{j+1}}$ only contributes if $n_{j+2} \leq n_j = d \leq n_{j+1}$. The term $z^{m+\Delta(h_1)+t_{j+2}}$ only contributes if $n_{j+2} < n_j = d < n_{j+1}$.
- (III) After crossing \mathfrak{d}_2 , new terms of the form (leaving off the coefficients) $z^{m+\Delta(h_2)}$ and $z^{m+\Delta(h_2)+t_\ell}$, $\ell = 0, 1, 2$ may appear in $\exp((W_k - W_0)/\hbar)$. Thus, when we cross \mathfrak{d}_1 , these new terms may contribute to $L_{P, j}$. Note that $z^{m+\Delta(h_2)}$ only contributes when crossing \mathfrak{d}_1 if $n_{j+2} \leq n_j = d < n_{j+1}$. The term $z^{m+\Delta(h_2)+t_j}$ only contributes if $n_{j+2} \leq d$, $n_j = d - 1$ and $d < n_{j+1}$. The term $z^{m+\Delta(h_2)+t_{j+1}}$ only contributes if $n_{j+2} \leq n_j = d \leq n_{j+1}$. The term $z^{m+\Delta(h_2)+t_{j+2}}$ only contributes if $n_{j+2} < n_j = d < n_{j+1}$.

There are now three cases when these additional contributions to $L_{P, j}$ occur.

Case (a). $n_{j+2} \leq n_j = d < n_{j+1}$. In this case, (leaving off the coefficients), z^m gives a contribution to $L_{P, j}$ of type (I) when γ_P crosses \mathfrak{d}_3 , and $z^{m+\Delta(h_i)}$, $z^{m+\Delta(h_i)+t_{j+1}}$, or $z^{m+\Delta(h_i)+t_{j+2}}$ (if $n_{j+2} < d$) may give contributions of type (II) or (III) when γ_P crosses \mathfrak{d}_1 and \mathfrak{d}_2 . Now note that the total change to $L_{i, \sigma_{j, j+1}}^d$ due to

these monomials as we traverse the loop γ_P is the sum of the contributions of these monomials to $L_{i,\gamma_P,\rho_j \rightarrow \sigma_{j,j+1}}^d$ and $L_{i,\gamma_P,\rho_{j+1} \rightarrow \sigma_{j,j+1}}^d$. However, the total contribution to the change of $L_{i,\sigma_{j,j+1}}^d$ is necessarily zero, as γ_P is a loop, and because $d < n_{j+1}$, none of these monomials contribute to any change of $L_{i,\gamma_P,\rho_j \rightarrow \sigma_{j,j+1}}^d$. Hence the total contribution of these monomials to $L_{i,\gamma_P,\rho_{j+1} \rightarrow \sigma_{j,j+1}}^d = L_{P,j}$ is also zero.

Case (b). $n_{j+2} \leq d$, $n_j = d - 1$, $d < n_{j+1}$. In this case only the terms $z^{m+\Delta(h_i)+t_j}$, $i = 1, 2$, may contribute. However, the same argument as in Case (a) shows that the total contribution from these terms is zero.

Case (c). $n_{j+2} \leq n_j = n_{j+1} = d$. In this case, contributions to $L_{P,j}$ only arise from terms of the form $z^{m+\Delta(h_i)+t_{j+1}}$. Choose $n_{\mathfrak{d}_i}$ so that at the first time τ_i when γ_P passes through \mathfrak{d}_i , $\langle n_{\mathfrak{d}_i}, \gamma'_P(\tau_i) \rangle < 0$. By interchanging the labelling of \mathfrak{d}_1 and \mathfrak{d}_2 and choosing the base-point Q' appropriately, we can assume firstly that γ_P passes initially through \mathfrak{d}_1 and then \mathfrak{d}_2 , and secondly that $\langle n_{\mathfrak{d}_i}, m_{j+1} \rangle \geq 0$ for $i = 1, 2$. Write $f_{\mathfrak{d}_i} = 1 + c_{\mathfrak{d}_i} z^{m_{\mathfrak{d}_i}}$ for $i = 1, 2$.

Then the term

$$\langle n_{\mathfrak{d}_1}, m_{j+1} \rangle c_{\mathfrak{d}_1} \hbar^{-(\nu+1)} z^{m+\Delta(h_1)+t_{j+1}}$$

appears in $\exp((W_k(\gamma_P(t)) - W_0(\gamma_P(t)))/\hbar)$ right after γ_P crosses \mathfrak{d}_1 the first time (and disappears when we cross \mathfrak{d}_1 for the second time), and hence, when γ_P crosses \mathfrak{d}_2 for the first time, we obtain a contribution to $L_{P,j}$ of

$$\begin{aligned} & \langle n_{\mathfrak{d}_2}, m_j \rangle \langle n_{\mathfrak{d}_1}, m_{j+1} \rangle c_{\mathfrak{d}_1} c_{\mathfrak{d}_2} \\ & \cdot D_2(d, m + \Delta(h_1) + \Delta(h_2) + t_j + t_{j+1}) \hbar^{-(\nu+3d-|m+\Delta(h_1)+\Delta(h_2)|)}. \end{aligned}$$

On the other hand, the term $\langle n_{\mathfrak{d}_2}, m_{j+1} \rangle c_{\mathfrak{d}_2} \hbar^{-(\nu+1)} z^{m+\Delta(h_2)+t_{j+1}}$ appears after γ_P crosses \mathfrak{d}_2 for the first time (and disappears when we cross \mathfrak{d}_2 for the second time), and hence, when γ_P crosses \mathfrak{d}_1 for the second time, we obtain a contribution to $L_{P,j}$ of

$$\begin{aligned} & \langle -n_{\mathfrak{d}_1}, m_j \rangle \langle n_{\mathfrak{d}_2}, m_{j+1} \rangle c_{\mathfrak{d}_1} c_{\mathfrak{d}_2} \\ & \cdot D_2(d, m + \Delta(h_1) + \Delta(h_2) + t_j + t_{j+1}) \hbar^{-(\nu+3d-|m+\Delta(h_1)+\Delta(h_2)|)}. \end{aligned}$$

Note that

$$\begin{aligned} \langle n_{\mathfrak{d}_2}, m_j \rangle \langle n_{\mathfrak{d}_1}, m_{j+1} \rangle - \langle n_{\mathfrak{d}_1}, m_j \rangle \langle n_{\mathfrak{d}_2}, m_{j+1} \rangle &= -|n_{\mathfrak{d}_1} \wedge n_{\mathfrak{d}_2}| \\ &= -|m^{\text{prim}}(h_1) \wedge m^{\text{prim}}(h_2)| \end{aligned}$$

as $n_{\mathfrak{d}_1}, n_{\mathfrak{d}_2}$ form a positively oriented basis of $N_{\mathbb{R}}$, and m_j, m_{j+1} form a positively oriented basis of $M_{\mathbb{R}}$.

Now the Maslov index two disks $h_3, \dots, h_{\nu+2}$ deform to disks with boundary P , which we also call $h_3, \dots, h_{\nu+2}$. We can then glue together the disks $h_1, \dots, h_{\nu+2}$ along with $d - n_{j+2}$ copies of the Maslov index two disk with no marked points in the direction m_{j+2} . These are glued at their respective outgoing vertices, yielding a vertex V_{out} , and we add one additional unbounded edge E_x with the label x , also attached to the vertex V_{out} . This yields a graph Γ , whose valency at V_{out} is $\text{Val}(V_{\text{out}}) = \nu + 3 + d - n_{j+2}$. Thus we obtain a parameterized curve $h : \Gamma \rightarrow M_{\mathbb{R}}$ with $h(x) = P$. Again one easily checks the balancing condition at V_{out} .

Thus the total contribution arising in the ways analyzed from $c\hbar^{-\nu}z^m$ to $-L_{P,j}$ is

$$\begin{aligned} & |m^{\text{prim}}(h_1) \wedge m^{\text{prim}}(h_2)| w_{\Gamma'_1}(E_{\text{out},1}) w_{\Gamma'_2}(E_{\text{out},2}) \cdot \\ & \cdot \text{Mult}_x^0(h) \left(\prod_{\substack{V \in \Gamma^{[0]} \\ V \not\subseteq E_x}} \text{Mult}_V(h) \right) \hbar^{-(\nu+3d-(2d+n_{j+2}))} u_{I(h)} \\ = & |m(h_1) \wedge m(h_2)| \text{Mult}_x^0(h) \left(\prod_{\substack{V \in \Gamma^{[0]} \\ V \not\subseteq E_x}} \text{Mult}_V(h) \right) \hbar^{-(\text{Val}(V_{\text{out}})-3)} u_{I(h)}. \end{aligned}$$

One sees that if $I(h) = \{i_1, \dots, i_{3d-\nu'}\}$ for some ν' , then since h is obtained by gluing two Maslov index zero disks with $\text{Val}(V_{\text{out}}) - 3$ Maslov index two disks, we have

$$\begin{aligned} \text{Val}(V_{\text{out}}) - 3 &= \sum_{i=1}^{\nu+2} (|\Delta(h_i)| - \#I(h_i)) + d - n_{j+2} \\ &= 3d - (3d - \nu') = \nu'. \end{aligned}$$

Thus we see that the coefficient of the contributions analyzed above from $c\hbar^{-\nu}z^m$ to $-L_{P,j}$ is precisely the contribution of h to

$$(5.41) \quad \langle P_{i_1}, \dots, P_{i_{3d-\nu'}}, \psi^{\nu'} M_{\mathbb{R}} \rangle_{d, \sigma_{j,j+1}}^{\text{trop}} u_{I(h)} \hbar^{-\nu'}$$

as desired.

Conversely, given an h contributing to (5.41) with $h(x) = P$, one can cut it at P , using Lemma 5.12, (1), decomposing it into tropical disks. Then we see that h arises precisely as above, exactly as in the proofs of Lemmas 5.55 and 5.56. Thus we see that $-L_{P,j}$ is the contribution to (5.41) from maps with $h(E_x) = P$. \square

LEMMA 5.59. *Let $P \in \sigma_{j,j+1} \cap \text{Sing}(\mathfrak{D})$, and suppose that $P = P_\ell$ for some ℓ . Then*

$$-L_{i, \gamma_P, \rho_{j+1} \rightarrow \sigma_{j,j+1}}^d = u_\ell \hbar \delta_{d,0} \delta_{2,i} + \sum_{\nu \geq 0} \sum_h \text{Mult}(h) u_{I(h)} \hbar^{-(\nu+2-i)},$$

where the sum is over all curves h contributing to $\langle P_{i_1}, \dots, P_{i_{3d-2+i-\nu}}, \psi^\nu S \rangle_{d, \sigma_{j,j+1}}^{\text{trop}}$ for various $i_1 < \dots < i_{3d-2+i-\nu}$ such that $h(E_x) = P$.

PROOF. Again, we may assume $i = 2$, and write

$$L_{P,j} = L_{2, \gamma_P, \rho_{j+1} \rightarrow \sigma_{j,j+1}}^d.$$

Choose a basepoint Q' near P_ℓ . By Remark 5.29, there is a one-to-one correspondence between rays in \mathfrak{D} containing P_ℓ and Maslov index two disks with boundary P_ℓ not having P_ℓ as a marked point. With Q' sufficiently near P_ℓ , these Maslov index two disks deform to ones with boundary at Q' , so the Maslov index two disks with boundary P_ℓ not having P_ℓ as a marked point are in one-to-one correspondence with the Maslov index two disks with boundary Q' not having P_ℓ as a marked point.

If we are interested in terms in $\exp((W_k(Q') - W_0(Q'))/\hbar)$ which may contribute to $L_{P_\ell,j}$, we only need to look at those terms in $\exp((W_k(Q') - W_0(Q'))/\hbar)$ which

do not have u_ℓ as a factor, as any term that does will not produce any new terms as we cross a ray through P_ℓ . So consider a term $c\hbar^{-\nu}z^m$ of the form

$$(5.42) \quad c\hbar^{-\nu}z^m = \hbar^{-\nu} \prod_{p=1}^{\nu} \text{Mult}(h_p) z^{\Delta(h_p)} u_{I(h_p)},$$

where each of these disks h_p with boundary Q' does not pass through P_ℓ , and hence corresponds to a disk with boundary P_ℓ , which we also write as $h_p : \Gamma'_p \rightarrow M_{\mathbb{R}}$. By extending these disks to trees and marking P_ℓ , we obtain Maslov index zero trees, corresponding to rays \mathfrak{d}_p in \mathfrak{D} with initial point P_ℓ . In addition, we have rays $\mathfrak{c}_p \in \mathfrak{D}$, $p = 0, 1, 2$, with initial point P_ℓ , corresponding to the three Maslov index two disks with boundary Q' with no marked points. These do not appear in $W_k(Q') - W_0(Q')$, so are distinct from the \mathfrak{d}_p 's.

In what follows, we write $m = \sum_{p=1}^{\nu} \Delta(h_p) = \sum_{j=0}^2 n_j t_j$, and take $n_{\mathfrak{d}_p}$ and $n_{\mathfrak{c}_p}$ to have their sign chosen so that they are negative on γ'_P when γ_P crosses the corresponding ray. Note that as γ_P is counterclockwise, if we use the identification $\bigwedge^2 M \cong \mathbb{Z}$ given by the standard orientation, i.e., $m_1 \wedge m_2 \mapsto 1$, then $w_{\Gamma'_p}(E_{\text{out},p}) n_{\mathfrak{d}_p} = X_{r(\Delta(h_p))}$. So

$$(5.43) \quad \sum_{p=1}^{\nu} w_{\Gamma'_p}(E_{\text{out},p}) n_{\mathfrak{d}_p} = X_{r(m)}.$$

On the other hand, $\langle n_{\mathfrak{c}_{j+1}}, m_j \rangle = -1$ and $\langle n_{\mathfrak{c}_{j+2}}, m_j \rangle = 1$.

We can now view this term $c\hbar^{-\nu}z^m$ as giving rise to contributions to $L_{P,j}$ in the following four ways:

(I) γ_P crosses \mathfrak{d}_l for some $1 \leq l \leq \nu$. Then the term

$$\hbar^{-(\nu-1)} \prod_{\substack{p=1 \\ p \neq l}}^{\nu} \text{Mult}(h_p) z^{\Delta(h_p)} u_{I(h_p)}$$

in $\exp((W_k(Q') - W_0(Q'))/\hbar)$ contributes to $L_{P,j}$ if $n_{j+2} \leq n_j = d < n_{j+1}$, in which case the contribution is

$$\langle n_{\mathfrak{d}_l}, m_j \rangle \left(\prod_{p=1}^{\nu} \text{Mult}(h_p) u_{I(h_p)} \right) u_\ell w_{\Gamma'_l}(E_{\text{out},l}) D_2(d, m + t_j) \hbar^{-(\nu+3d-|m|-1)}.$$

Note that such a contribution requires $\nu > 0$.

(II) γ_P crosses \mathfrak{c}_j . If $c\hbar^{-\nu}z^m$ contributes to $L_{P,j}$ when γ_P crosses \mathfrak{c}_j , its contribution would involve a factor of $\langle n_{\mathfrak{c}_j}, m_j \rangle = 0$, hence there is no contribution.

(III) γ_P crosses \mathfrak{c}_{j+1} . We get a contribution from $c\hbar^{-\nu}z^m$ if $n_{j+2} \leq n_j = d \leq n_{j+1}$, in which case the contribution is

$$\langle n_{\mathfrak{c}_{j+1}}, m_j \rangle \left(\prod_{p=1}^{\nu} \text{Mult}(h_p) u_{I(h_p)} \right) u_\ell D_2(d, m + t_j + t_{j+1}) \hbar^{-(\nu+3d-|m|-1)}.$$

(IV) γ_P crosses \mathfrak{c}_{j+2} . We get a contribution from $c\hbar^{-\nu}z^m$ if $n_{j+2} < n_j = d < n_{j+1}$, in which case we get

$$\langle n_{\mathfrak{c}_{j+2}}, m_j \rangle \left(\prod_{p=1}^{\nu} \text{Mult}(h_p) u_{I(h_p)} \right) u_\ell D_2(d, m + t_j + t_{j+2}) \hbar^{-(\nu+3d-|m|-1)}.$$

We now consider three cases.

Case (a). $n_{j+2} \leq n_j = d < n_{j+1}$, $\nu > 0$. In this case, ignoring the common factors

$$\hbar^{-(\nu+3d-|m|-1)} u_\ell \prod_{p=1}^{\nu} \text{Mult}(h_p) u_{I(h_p)},$$

the total contribution is, using (5.43) and the formula for D_2 in Lemma 5.43,

$$\begin{aligned} & \left\langle \sum_{l=1}^{\nu} w_{\Gamma_l}(E_{\text{out},l}) n_{\mathfrak{d}_l}, m_j \right\rangle D_2(d, m + t_j) - D_2(d, m + t_j + t_{j+1}) \\ & + \begin{cases} D_2(d, m + t_j + t_{j+2}) & n_{j+2} < d \\ 0 & n_{j+2} = d \end{cases} \\ = & r(m) \wedge m_j (-1)^{n_j+n_{j+1}+1} \frac{(n_j - d)!(n_{j+1} - d - 1)!}{(d - n_{j+2})!} \\ & - (-1)^{n_j+n_{j+1}+2} \frac{(n_j - d)!(n_{j+1} - d)!}{(d - n_{j+2})!} \\ & + (-1)^{n_j+n_{j+1}+1} \frac{(n_j - d)!(n_{j+1} - d - 1)!}{(d - n_{j+2})!} (d - n_{j+2}) \\ = & ((n_{j+2} - n_{j+1}) + (n_{j+1} - d) + (d - n_{j+2})) \cdot \\ & \cdot (-1)^{n_j+n_{j+1}+1} \frac{(n_j - d)!(n_{j+1} - d - 1)!}{(d - n_{j+2})!} \\ = & 0. \end{aligned}$$

So there is no contribution to $L_{P,j}$ from this case.

Case (b). $n_{j+2} \leq n_j = d = n_{j+1}$, $\nu > 0$. In this case we only get a contribution from (III). In this case, we can glue together the disks h_1, \dots, h_ν along with $d - n_{j+2}$ copies of the Maslov index two disk with no marked points in the direction m_{j+2} . These are glued at their respective outgoing vertices, yielding a vertex V_{out} , and we add two additional marked unbounded edges E_x and E_{p_l} for some l attached to V_{out} . This yields a graph Γ , whose valency at V_{out} is $\text{Val}(V_{\text{out}}) = \nu + (d - n_{j+2}) + 2 = \nu + 3d - |m| + 2$. Thus we obtain a parameterized curve $h : \Gamma \rightarrow M_{\mathbb{R}}$ with $h(V_{\text{out}}) = h(x) = h(p_l) = P_\ell$. The contribution to $-L_{P,j}$ from (III) in this case is then easily seen by inspection to be

$$(5.44) \quad \text{Mult}_x^0(h) \left(\prod_{\substack{V \in \Gamma^{[0]} \\ V \notin E_x}} \text{Mult}_V(h) \right) u_{I(h)} \hbar^{-(\text{Val}(V_{\text{out}})-3)}.$$

Suppose that $I(h) = \{i_1, \dots, i_{3d-\nu'}\}$ for some ν' , recalling that $\ell \in I(h)$ since we added the marked edge E_{p_l} mapping to P_ℓ . Since h is obtained by gluing $\text{Val}(V_{\text{out}}) - 2$ Maslov index two disks, we have

$$\begin{aligned} \text{Val}(V_{\text{out}}) - 2 &= \sum_{i=1}^{\nu} (|\Delta(h_i)| - \#I(h_i)) + d - n_{j+2} \\ &= 3d - (3d - \nu' - 1) = \nu' + 1. \end{aligned}$$

Thus we see that (5.44) is precisely the contribution of h to

$$\langle P_{i_1}, \dots, P_{i_{3d-\nu'}}, \psi^{\nu'} M_{\mathbb{R}} \rangle_{d, \sigma_{j, j+1}}^{\text{trop}} u_{I(h)} \hbar^{-\nu'}$$

from Definition 5.13 (3) (b). As in the other cases we have considered, conversely any such curve h will give rise to the correct monomial $c\hbar^{-\nu} z^m$ by cutting the curve at P .

Case (c). $\nu = 0$. There is only one element $(c, \nu, m) \in S_k$ with $\nu = 0$, namely $(1, 0, 0)$ corresponding to the constant monomial 1. So $n_0 = n_1 = n_2 = 0$ and we have no contribution unless $d = 0$. Again, this contribution to $L_{P, j}$ only arises from (III), and is

$$\langle n_{c_{j+1}}, m_j \rangle u_{\ell} D_2(0, t_j + t_{j+1}) \hbar = -u_{\ell} \hbar.$$

This gives the remaining claimed terms in $-L_{P, j}$. □

LEMMA 5.60.

$$\begin{aligned} -L_{i, \gamma_{j, j+1}, \rho_{j+1} \rightarrow \sigma_{j, j+1}}^d &= \sum_{\substack{\ell \text{ s.t.} \\ P_{\ell} \in Q + \sigma_{j, j+1}}} u_{\ell} \hbar \delta_{0, d} \delta_{2, i} + \\ &+ \sum_{\nu \geq 0} \sum_{\substack{I \subseteq \{1, \dots, k\} \\ I = \{i_1, \dots, i_{3d-2+i-\nu}\} \\ i_1 < \dots < i_{3d-2+i-\nu}}} \langle P_{i_1}, \dots, P_{i_{3d-2+i-\nu}}, \psi^{\nu} S \rangle_{d, \sigma_{j, j+1}}^{\text{trop}} u_I \hbar^{-(\nu+2-i)} \end{aligned}$$

for $S = Q, L$ or $M_{\mathbb{R}}$ for $i = 0, 1$ and $i = 2$ respectively.

PROOF. This is just putting together the results of Lemmas 5.57, 5.58, and 5.59. □

We finally have:

Proof of Theorems 5.15 and 5.18. By Lemma 5.46, it is enough to prove (5.23), (5.24) and (5.25). However, these now follow from Lemmas 5.54, 5.55, 5.56, and 5.60. □

5.6. References and further reading

The material of this section is drawn entirely from the paper [42]. The original suggestion of using Maslov index two holomorphic disks to define the Landau-Ginzburg potential was due to Cho and Oh in [16]; however, they only defined the non-perturbed potential in this way. The work of Fukaya, Oh, Ohta and Ono [26] and [23] pursued ideas analogous to the ideas presented here, using holomorphic disks rather than tropical disks. The advantage of using holomorphic disks is that one can apply this to all toric varieties, whereas tropical geometry has difficulty seeing holomorphic curves on the boundary of a toric variety. On the other hand, the cost is that one must make use of the immense work [24], [25] and it is very difficult to make any computations. Chan and Leung in [14] introduced the idea of using tropical geometry to study the Landau-Ginzburg potential; again, they only worked with the unperturbed version. Finally, for a purely tropical, non-mirror-symmetric, approach to gravitational descendent invariants for \mathbb{P}^2 , see [77].

Part 3

The Gross-Siebert program

The program and two-dimensional results

We have now seen how, at least for \mathbb{P}^2 , tropical geometry gives a very satisfactory explanation for the A- and B-models of mirror symmetry. In this chapter, we will explore how similar ideas may be used to explain mirror symmetry for more complicated varieties, especially for Calabi-Yau manifolds.

The ideas described here represent joint work with Bernd Siebert, in a program we began in 2001 (see [47], [48], [40], [51], [49]). There are a number of expository articles explaining aspects of this program already: see especially [41], which explores how this program grew out of a study of the Strominger-Yau-Zaslow (SYZ) conjecture, and [3], Chapters 6 and 7, for a more thorough description of aspects of the SYZ conjecture.

Here, we shall take the more ahistorical point of view, and describe the program by way of analogy with the ideas already covered in this book. We begin by explaining roughly the correspondence between the tropical manifolds introduced in Chapter 1 and nice degenerating families of varieties with effective anti-canonical class — this includes the Calabi-Yau case, the best known example of mirror symmetry. This brings us to the fundamental problem of reconstructing such a family of varieties from a tropical manifold.

The next two sections of this chapter explain the solution to this problem given by myself and Siebert in [49] for the simplest, two-dimensional case. A complete argument is given, modulo results of [45]. The proof I give here is more of a hybrid of the approach of Kontsevich and Soibelman [70], which only deals with the two-dimensional case, and the approach of [49], which deals with all dimensions, but is significantly harder.

6.1. The program

The main content of Chapter 4 was a correspondence theorem between tropical curves with vertices in $M_{\mathbb{Q}}$ and algebraic curves defined over the field $\overline{\mathbb{k}((t))}$. We saw there that a type of strong integrality for the tropical curves allowed us to conclude that the algebraic curves were defined over $\mathbb{k}((t))$, and not just over $\overline{\mathbb{k}((t))}$.

In general, one may not wish to work with tropical curves whose vertices lie in $M_{\mathbb{Q}}$, in which case one should work over the field

$$K = \left\{ \sum_i a_i t^{r_i} \mid a_i \in \mathbb{k}, r_i \in \mathbb{R} \text{ a sequence with } r_i \rightarrow \infty \right\}.$$

This field carries a valuation $\nu : K^{\times} \rightarrow \mathbb{R}$ given by

$$\nu \left(\sum_i a_i t^{r_i} \right) = \min\{r_i\}.$$

Given a variety $X \subseteq (K^\times)^n$, we can then define the *tropicalization* of X , $\text{Trop}(X)$, as

$$\text{Trop}(X) = \{(\nu(x_1), \dots, \nu(x_n)) \mid (x_1, \dots, x_n) \in X\} \subseteq \mathbb{R}^n.$$

It turns out that $\text{Trop}(X)$ is a type of tropical variety. We will not, however, define precisely what we mean by this, but there is much discussion of the definition of tropical varieties in the literature; see, for example, [96], [81]. Certainly tropical varieties include the tropical curves already studied in this book.

One then asks the following natural question: when can a tropical variety in \mathbb{R}^n be realised as the tropicalization of a variety in $(K^\times)^n$? In Chapter 4, we essentially answered this question for tropical curves of genus zero in \mathbb{R}^2 ; in fact, the same results are true for curves of any genus in \mathbb{R}^2 , as shown by Mikhalkin in [80]. However, in general, in \mathbb{R}^n for $n \geq 3$, only tropical curves of genus zero come from actual curves; in higher genus this need not be the case: see [106] for some results in this direction. Very little is known for tropical varieties of higher dimension.

Let us change our point of view a bit. In the previous paragraph, \mathbb{R}^n was the ambient space, the “tropicalization” of $(K^\times)^n$, and we were interested in which tropical subvarieties of \mathbb{R}^n are tropicalizations of varieties in $(K^\times)^n$. Instead, let us try to change the ambient space. Perhaps there are more interesting choices which correspond to more interesting varieties. In particular, if we replace \mathbb{R}^n with a tropical affine manifold B in the sense of Definition 1.22, we obtain a larger set of ambient varieties. As we saw in Chapter 1, tropical affine manifolds are precisely the manifolds where one can still talk about tropical curves. The trouble is that there are few interesting examples of compact tropical affine manifolds (\mathbb{R}^n/Γ for a lattice Γ being one such example, corresponding to a complex torus), so we need to allow tropical affine manifolds with singularities, or more precisely, what we call a tropical manifold in Definition 1.27.

So let B be a tropical manifold. It now becomes natural to try to associate to such a manifold a variety defined over the field K . It is not immediately clear what the connection between this variety and the tropical object should be, and we will in fact avoid this question, for to give a proper answer requires working in the category of rigid analytic spaces (see [70]). However, we will be able to give a sensibly motivated suggestion if we assume furthermore that B is integral. The advantage of working with integral tropical manifolds is that we can hope that the corresponding variety is then defined over the subfield $\mathbb{k}((t))$, as was the case in Chapter 4.

Using the hints from Chapter 4, what we should in fact try to associate to B is a scheme \mathcal{X} defined over $\text{Spec } \mathbb{k}[[t]]$. The generic fibre \mathcal{X}_η will be the desired variety over $\mathbb{k}((t))$, while the fibre \mathcal{X}_0 over the closed point will be a degenerate variety. This should be of a similar flavour to the situation in Chapter 4, when we considered degenerations of \mathbb{P}^2 or other toric surfaces. These degenerations arose from polyhedral decompositions of \mathbb{R}^2 . So, we should expect some choice involved in \mathcal{X}_0 . By analogy, it is reasonable that we need to make use of the polyhedral decomposition \mathcal{P} of B , as in Definition 1.27. This is the basic context of the Gross-Siebert program.

In the remainder of this section, we will sketch some basic ideas of the Gross-Siebert program without giving any technical details.

6.1.1. The fan picture, or the A-model. So, given (B, \mathcal{P}) , let's first try to guess the form of the central fibre X_0 of a degeneration $\mathcal{X} \rightarrow \text{Spec } \mathbb{k}[[t]]$ we may wish to associate to (B, \mathcal{P}) . Again, taking guidance from Chapter 4, recall that we considered pairs $(M_{\mathbb{R}}, \mathcal{P})$. Each vertex v of \mathcal{P} defined a fan Σ_v , defining a toric variety D_v , and X_0 was a union of irreducible components $X_0 = \bigcup_v D_v$ with v running over all vertices of \mathcal{P} . The irreducible components of X_0 are glued together in a way dictated by the combinatorics of \mathcal{P} .

The point is that we have exactly the same structure from an integral tropical manifold (B, \mathcal{P}) . For each cell $\tau \in \mathcal{P}$, we obtain a fan Σ_τ coming from the fan structure on B . Corresponding to Σ_τ is a toric variety X_τ (we shall use X_τ now instead of the D_τ we used in Chapter 4). In particular, for each vertex $v \in \mathcal{P}$, we have a toric variety X_v . Furthermore, if v, w are vertices of $\tau \in \mathcal{P}$, then Σ_τ can be described as a quotient fan of both Σ_v and Σ_w , so X_τ is naturally a toric stratum of both X_v and X_w . If τ is the smallest cell of \mathcal{P} containing v and w , then we glue X_v and X_w along the two copies of $X_\tau \subseteq X_v, X_\tau \subseteq X_w$.

In this way, we build a variety $X_0 = X_0(B, \mathcal{P})$ which is a union of toric varieties. For example, applying this to Example 1.28, (3), with (B, \mathcal{P}) given by $\partial\Xi_3$, one obtains a union of eight \mathbb{P}^2 's, glued together to form an octahedron.

In general, the situation is actually a bit more complex, as one has a choice of identifications of $X_\tau \subseteq X_v$ and $X_\tau \subseteq X_w$ preserving the toric strata. This gives some moduli of possible gluings, but there is always a canonical choice which identifies the identity element of the big tori in the two copies of X_τ . We shall always assume in this chapter that we have made this canonical choice of gluings to avoid having to keep track of a lot of extra data.

Note that, so far, we have only made use of the fan structure on B , and made no use of the fact that the cells of \mathcal{P} themselves have the structure of lattice polytope. On the other hand, in the situation in Chapter 4, it was very important to consider X_0 along with a log structure X_0^\dagger and a log morphism $X_0^\dagger \rightarrow \text{Spec } \mathbb{k}^\dagger$, preferably log smooth. In fact, the lattice polytope structure gives us a hint as to what the log structure should be, by determining the ghost sheaf $\overline{\mathcal{M}}_{X_0}$.

It is not difficult to see from the example of Chapter 4 what this ghost sheaf should be. It is essentially entirely defined by its stalks at generic points of strata of X_0 . If η_τ is the generic point of $X_\tau \subseteq X_0$, then in the case of the degeneration in Chapter 4, built from $(M_{\mathbb{R}}, \mathcal{P})$, one has

$$(6.1) \quad \overline{\mathcal{M}}_{X_0, \eta_\tau} = C(\tau)^\vee \cap (N \oplus \mathbb{Z}).$$

In fact, in the general case of (B, \mathcal{P}) , we can do the same thing. Given $\tau \in \mathcal{P}$, τ is a lattice polytope, so we can identify it with a lattice polytope $\tau \subseteq M_{\mathbb{R}}$. Then $C(\tau)^\vee \cap (N \oplus \mathbb{Z})$ still makes sense, and we require that (6.1) holds. We also require that $\overline{\mathcal{M}}_{X_0}$ is constant on the big torus orbit of X_τ .

It turns out this information is enough to determine $\overline{\mathcal{M}}_{X_0}$, but it is much harder to get an actual log structure. One does this by attempting to classify extensions

$$1 \rightarrow \mathcal{O}_{X_0}^\times \rightarrow \mathcal{M}_{X_0} \rightarrow \overline{\mathcal{M}}_{X_0} \rightarrow 0$$

which yield log structures, along with sections ρ of \mathcal{M}_{X_0} defining a morphism to $\text{Spec } \mathbb{k}^\dagger$. This classification was the main result of [48]. The chief difficulty is that in fact such an extension does not exist globally, but only away from a codimension two closed subset $Z \subseteq X_0$. As a consequence, [48] gives a description of log morphisms

$X_0^\dagger \rightarrow \text{Spec } \mathbb{k}^\dagger$ which are only log smooth off of a closed codimension two subset $Z \subseteq X_0$.

EXAMPLE 6.1. The existence of the bad set Z is fundamental to the theory. Consider, for example, a variety \mathcal{X} given by $tf_4 + x_0 \cdots x_3 = 0$ in $\mathbb{P}^3 \times \mathbb{A}^1$, where \mathbb{P}^3 has coordinates x_0, \dots, x_3 , \mathbb{A}^1 has coordinate t , and f_4 is a general homogeneous polynomial of degree 4 in the variables x_0, \dots, x_3 . Then the projection $f : \mathcal{X} \rightarrow \mathbb{A}^1$ is a degeneration of a quartic K3 surface to the union of coordinate planes in \mathbb{P}^3 . Set $X_0 = f^{-1}(0)$.

Now note that \mathcal{X} has singularities precisely at the 24 points

$$\{t = f_4 = 0\} \cap \text{Sing}(X_0).$$

For a general choice of f_4 these points are ordinary double points. Étale locally near these double points, $X_0 \subseteq \mathcal{X}$ looks like

$$V(t) \subseteq \text{Spec } \mathbb{k}[x, y, w, t]/(xy - wt).$$

This inclusion was studied in Example 3.20. So if we take the log structure on X_0 induced by the divisorial log structure $X_0 \subseteq \mathcal{X}$, the log structure fails even to be fine at the 24 points. These 24 points yield the bad set Z ; clearly the induced map $X_0^\dagger \rightarrow \text{Spec } \mathbb{k}^\dagger$ is not log smooth here.

One can see that the integral tropical manifold giving rise to such an X_0^\dagger is given by the construction of Example 1.28, (3), applied to the reflexive polytope

$$\Xi = \text{Conv}\{(-1, -1, -1), (1, 0, 0), (0, 1, 0), (0, 0, 1)\}.$$

□

The final step is to realise X_0^\dagger as the central fibre of a family $\mathcal{X} \rightarrow \text{Spec } \mathbb{k}[t]$. The log structure on X_0^\dagger should be the pull-back of a divisorial log structure on \mathcal{X} induced by $D \cup X_0$, where $D \subseteq \mathcal{X}$ is some divisor whose restriction to each fibre is an anti-canonical divisor (so in the Calabi-Yau case, D is empty).

This is essentially the problem solved in [49], and whose solution will be explained, in the two-dimensional case, in §6.2.

However, without having solved this problem, we can still ask about a correspondence between tropical curves in B and curves in \mathcal{X} , or more specifically, curves in $\mathcal{X}_{\bar{\eta}}$, where η is the generic point of $\text{Spec } \mathbb{k}[t]$. As suggested in Chapter 4, to explore this correspondence, one should first consider a correspondence between tropical curves in B and log curves in X_0^\dagger . A properly defined notion of Gromov-Witten invariants for X_0^\dagger should then give the same Gromov-Witten invariants for $\mathcal{X}_{\bar{\eta}}$. In this way, Ξ hopes to compute Gromov-Witten invariants of $\mathcal{X}_{\bar{\eta}}$ via tropical geometry.

This is currently the most undeveloped aspect of the program, though work in progress with Siebert [46], as well as work of Parker [91], [92], [93], is developing the theory of log Gromov-Witten invariants. We will only exhibit a very simple example, demonstrating the role that the singularities of the affine and log structures play.

EXAMPLE 6.2. Let us consider the surface $(\check{B}, \check{\mathcal{P}})$ discussed in Example 1.31, (5) and Example 1.33. I claim first that this surface corresponds to a degeneration of a cubic surface of the form $\mathcal{X} = V(tf_3 + x_1x_2x_3) \subseteq \mathbb{P}^3 \times \mathbb{A}^1$, where f_3 is a general homogeneous polynomial of degree 3. We have $\pi : \mathcal{X} \rightarrow \mathbb{A}^1$, the projection,

and the fibre $X_0 = \pi^{-1}(0)$. The log structure on \mathcal{X} is the divisorial one given by $V(x_0) \cup X_0 \subseteq \mathcal{X}$, and this induces a log structure X_0^\dagger .

To see that \check{B} does indeed correspond to X_0^\dagger , let $X'_0 = X_0(\check{B}, \check{\mathcal{P}})$ be the variety obtained from \check{B} by the construction described in this section. We first note that $\check{\mathcal{P}}$ has three vertices, and for each $v \in \check{\mathcal{P}}$, Σ_v is the fan for \mathbb{P}^2 . So \check{B} gives rise to three copies of \mathbb{P}^2 glued together in precisely the way the three coordinate planes in X_0 are glued together. So $X_0 = X'_0$. To make sure that \check{B} also gives the correct information about $\overline{\mathcal{M}}_{X_0}$, we look at generic points of strata of X_0 . For example, at the stratum $x = (1 : 0 : 0 : 0)$ of X_0 , $\overline{\mathcal{M}}_{X_0, x} = \mathbb{N}^3$, as $V(x_0) \cup X_0$ is a divisor with normal crossings at x . But the corresponding maximal cell of $\check{\mathcal{P}}$ is the two-dimensional standard simplex σ , and

$$C(\sigma)^\vee \cap \mathbb{Z}^3 \cong \mathbb{N}^3.$$

Similarly, the point $x = (0 : 1 : 0 : 0)$ is another stratum, which corresponds to an infinite rectangle σ in $\check{\mathcal{P}}$. Again, $\overline{\mathcal{M}}_{X_0, x} = \mathbb{N}^3$ as $V(x_0) \cup X_0$ is a normal crossings divisor at x with the irreducible components $V(x_0)$, $V(x_2)$ and $V(x_3)$ of $V(x_0) \cup X_0$ passing through x . On the other hand, writing, say,

$$\sigma = \text{Conv}\{(0, 0), (1, 0)\} + \mathbb{R}_{\geq 0}(0, 1),$$

then $C(\sigma)$ is the cone generated by $(0, 0, 1)$, $(1, 0, 1)$ and $(0, 1, 0)$. Thus

$$C(\sigma)^\vee \cap \mathbb{Z}^3 \cong \mathbb{N}^3.$$

Similarly, one checks that the stalks of $\overline{\mathcal{M}}_{X_0}$ along one- and two-dimensional strata agree with the monoids defined using one- or zero-dimensional cells of $\check{\mathcal{P}}$, respectively. This shows that the structure of the stalks of $\overline{\mathcal{M}}_{X_0}$ are indeed the stalks specified by the data $(\check{B}, \check{\mathcal{P}})$. Note also that the total space of \mathcal{X} has, for general choice of f_3 , nine singularities at $V(f_3) \cap \text{Sing}(\mathcal{X}_0)$. These nine singular points form the bad set Z .

Next consider one of the tropical curves illustrated in Example 1.33. Based on our experience in Chapter 4, we should look for a log curve mapping to a line in one of the \mathbb{P}^2 components of X_0 , giving $f : \mathbb{P}^1 \rightarrow X_0$. However, we also know from Proposition 4.9 that there will be problems when a point of this \mathbb{P}^1 is mapped into $\text{Sing}(X_0)$: Proposition 4.9 rules this out. However, there is one crucial difference between the current situation and the situation in Chapter 4: namely, there is the locus $Z \subseteq X_0$ where the log structure X_0^\dagger is not fine. The nine points of Z correspond to the nine singular points in \check{B} . It turns out that the argument of Proposition 4.9 breaks down precisely if every point $p \in \mathbb{P}^1$ with $f(p) \in \text{Sing}(X_0)$ satisfies $f(p) \in Z$.

It is easy to count such curves. In each plane, there are precisely $3 \times 3 = 9$ lines satisfying this property. Since there are three planes, this gives 27 possible choices. Each of these choices deforms to give an actual line in $\pi^{-1}(t)$ for $t \neq 0$, and this accounts for the 27 lines on the cubic surface. Morally, these correspond to the 27 tropical curves described in Example 1.33. These tropical curves are allowed to terminate at singular points of \check{B} , and this represents the fact that we can have log curves $f : C^\dagger \rightarrow X_0^\dagger$ with irreducible components passing through points of Z with no matching irreducible component on another irreducible component of X_0 . \square

While this description is very vague, it demonstrates the basic strategy for A-model calculations: (B, \mathcal{P}) gives rise to X_0^\dagger , and tropical curve counting on

(B, \mathcal{P}) should be equivalent to computing log Gromov-Witten invariants on X_0^\dagger , which in turn should be equivalent to Gromov-Witten invariants on a smoothing of X_0^\dagger . Note, however, that one feature of this approach is that if the theory of log Gromov-Witten invariants is properly set up, the comparison of log Gromov-Witten invariants on X_0^\dagger with Gromov-Witten invariants on a smoothing should be automatic, and we never need to know anything about the smoothing.

It is also natural to consider ample line bundles \mathcal{L} on X_0 . In particular, when we compute Gromov-Witten invariants, we might focus on curves of a given degree on X_0 . A choice of such a line bundle gives rise to extra data on B . Indeed, if we restrict \mathcal{L} to any toric stratum X_τ of X_0 , we obtain an ample line bundle on the toric variety X_τ , which is specified by a strictly convex integral PL function $\varphi_\tau : |\Sigma_\tau| \rightarrow \mathbb{R}$, well-defined up to a linear function. It is then easy to check that if $\tau_1 \subseteq \tau_2$, so that $X_{\tau_2} \subseteq X_{\tau_1}$ naturally, then φ_{τ_2} differs by a linear function from $\varphi_{\tau_1}(\tau_2)$ as defined in Definition 1.10. (See the discussion in §3.1.2 concerning restriction of line bundles to toric strata.)

If $S_\tau : U_\tau \rightarrow \mathbb{R}^k$ defines the fan structure along τ , then the collection $\varphi = \{\varphi_\tau \circ S_\tau\}$ defines a strictly convex multi-valued integral PL function on (B, \mathcal{P}) . This gives a triple $(B, \mathcal{P}, \varphi)$.

To summarize, the data $(B, \mathcal{P}, \varphi)$ corresponds, in the fan picture, to a pair $(X_0^\dagger, \mathcal{L})$, where X_0^\dagger determines \mathcal{P} and \mathcal{L} determines φ . As mentioned previously, there will in general be a whole moduli space of such pairs arising from choices of gluing, but we shall ignore this here.

6.1.2. The cone picture, or the B-model. There is in fact another way to interpret the pair (B, \mathcal{P}) , inspired by Example 3.6, the Mumford degeneration. The Mumford degeneration coming from a lattice polytope $\Delta \subseteq \mathbb{N}_\mathbb{R}$, with a polyhedral decomposition \mathcal{P} and strictly convex PL function φ , is a degeneration of \mathbb{P}_Δ whose central fibre is $\bigcup_{\sigma \in \mathcal{P}_{\max}} \mathbb{P}_\sigma$.

Let us generalise this, starting with the pair (B, \mathcal{P}) . Each $\sigma \in \mathcal{P}$ defines a projective toric variety \mathbb{P}_σ . In particular, if $\sigma_1, \sigma_2 \in \mathcal{P}_{\max}$ with $\tau = \sigma_1 \cap \sigma_2$, we can glue together \mathbb{P}_{σ_1} and \mathbb{P}_{σ_2} along the common toric stratum \mathbb{P}_τ . In this way one builds a scheme $\check{X}_0 = \check{X}_0(B, \mathcal{P})$. Again, as in the fan picture, there is a choice to this gluing, but we shall choose the canonical one, which identifies the identity elements in the big tori of \mathbb{P}_τ .

Each \mathbb{P}_σ also carries an ample line bundle $\mathcal{O}_{\mathbb{P}_\sigma}(1)$. Again, one can glue together these line bundles (though not necessarily for every choice of gluing of the components of \check{X}_0), to obtain an ample line bundle \mathcal{L} on \check{X}_0 . So essentially the data (B, \mathcal{P}) determines a pair $(\check{X}_0, \mathcal{L})$. It does not, however, determine a log structure, which requires the additional choice of a multi-valued strictly convex PL function φ on B , in analogy with the Mumford degeneration.

Given such a function, it determines the monoids $\overline{\mathcal{M}}_{\check{X}_0, \bar{\eta}_\tau}$ for η_τ the generic point of a stratum \mathbb{P}_τ for $\tau \in \mathcal{P}$ as follows. Up to affine linear functions, φ is determined in a neighbourhood of τ by a PL function $\varphi_\tau : |\Sigma_\tau| \rightarrow \mathbb{R}$, where Σ_τ is the fan in \mathbb{R}^k determined by the fan structure along τ , with k the codimension of τ . Then

$$(6.2) \quad \overline{\mathcal{M}}_{\check{X}_0, \bar{\eta}_\tau} = \{(m, r) \mid m \in |\Sigma_\tau|, r \geq \varphi_\tau(m)\} \cap (\mathbb{Z}^k \oplus \mathbb{Z}).$$

As in the fan picture, these stalks determine $\overline{\mathcal{M}}_{\check{X}_0}$, and then one follows the same procedure for classifying extensions $\mathcal{M}_{\check{X}_0}$ of $\overline{\mathcal{M}}_{\check{X}_0}$ by $\mathcal{O}_{\check{X}_0}^\times$ yielding log structures. So one sees the symmetry between the fan and cone pictures. In the fan picture, \mathcal{P} determines $\overline{\mathcal{M}}$ and φ determines \mathcal{L} , and in the cone picture, \mathcal{P} determines \mathcal{L} and φ determines $\overline{\mathcal{M}}$.

Now in the Calabi-Yau case, the B-model involves period integrals measuring variation of complex structure. To even talk about this, we need a family $\check{\mathcal{X}} \rightarrow \text{Spec } \mathbb{k}[[t]]$ which has \check{X}_0^\dagger as the central fibre. Thus it is not enough to obtain \check{X}_0^\dagger , but we need the entire family. This is the context in which we need to work:

QUESTION 6.3. *Given $(B, \mathcal{P}, \varphi)$ yielding \check{X}_0^\dagger , is there a family $\check{\mathcal{X}} \rightarrow \text{Spec } \mathbb{k}[[t]]$ with central fibre \check{X}_0 , along with a family $D \subseteq \check{\mathcal{X}}$ of anti-canonical divisors inducing a divisorial log structure $D \cup \check{X}_0 \subseteq \check{\mathcal{X}}$ which restricts to the log structure \check{X}_0^\dagger on the central fibre?*

The main result of [49] gives sufficient conditions on \check{X}_0^\dagger for a positive solution to this problem. Furthermore, this solution, as we shall see, has a very tropical nature.

In §6.2, we will give most of the details of the argument for the construction of the family $\check{\mathcal{X}} \rightarrow \text{Spec } \mathbb{k}[[t]]$ in the case $\dim B = 2$, the higher-dimensional case being much harder.

6.1.3. Mirror symmetry and the discrete Legendre transform. To first approximation, the previous two sections argue that to the data $(B, \mathcal{P}, \varphi)$ one can associate the log spaces $X_0(B, \mathcal{P})^\dagger$, coming from the construction in the fan picture, and $\check{X}_0(B, \mathcal{P})^\dagger$, coming from the construction in the cone picture. As we saw, there are choices of gluing and log structure, so there are really moduli spaces of such structures, but for simplicity, we ignore this here.

The above discussion now suggests the following procedure for mirror symmetry:

- (1) Start with a degeneration $\pi : \mathcal{X} \rightarrow \text{Spec } \mathbb{k}[[t]]$ which is “nice,” i.e., the central fibre is a union of toric varieties and π is log smooth away from some set Z . This set should be codimension two and not contain any toric strata. There are a number of more technical conditions. We won’t give the precise set of conditions here, but see [48], Definition 4.1 or [41], Definition 7.1 for the full definition. In short, we want π to be what we call a *toric degeneration*. Also, choose a relatively ample line bundle \mathcal{L} on \mathcal{X} .
- (2) The log structure on $X_0 = \pi^{-1}(0)$ determines an integral tropical manifold (B, \mathcal{P}) with singularities with $X_0 = X_0(B, \mathcal{P})$ (in the fan picture), and \mathcal{L} determines a multi-valued integral PL strictly convex function φ , giving a triple $(B, \mathcal{P}, \varphi)$.
- (3) Now view $(B, \mathcal{P}, \varphi)$ in the cone picture, determining $\check{X}_0(B, \mathcal{P})$ with a log structure. Construct a family $\check{\mathcal{X}} \rightarrow \text{Spec } \mathbb{k}[[t]]$ with central fibre $\check{X}_0(B, \mathcal{P})$. The families $\mathcal{X}, \check{\mathcal{X}} \rightarrow \text{Spec } \mathbb{k}[[t]]$ should be mirror.

Sometimes, we may wish to consider the A- and B-models on one side, say for $\mathcal{X} \rightarrow \text{Spec } \mathbb{k}[[t]]$. As well as the triple $(B, \mathcal{P}, \varphi)$ with $X_0 = X_0(B, \mathcal{P})$ and \mathcal{L} determining φ , there is another triple $(\check{B}, \check{\mathcal{P}}, \check{\varphi})$, for which $X_0^\dagger = \check{X}_0(\check{B}, \check{\mathcal{P}})^\dagger$.

What is the relationship between $(B, \mathcal{P}, \varphi)$ and $(\check{B}, \check{\mathcal{P}}, \check{\varphi})$? The notation should give this away: they are related by the discrete Legendre transform of §1.5.

Indeed, there is an inclusion reversing correspondence between \mathcal{P} and $\check{\mathcal{P}}$, written as $\tau \mapsto \check{\tau}$, since both \mathcal{P} and $\check{\mathcal{P}}$ are in a one-to-one correspondence with strata of X_0 , i.e., $X_\tau = \mathbb{P}_{\check{\tau}}$. The fan Σ_τ is the normal fan to $\check{\tau}$ precisely because $X_\tau = \mathbb{P}_{\check{\tau}}$. Furthermore, $\check{\tau}$ is the Newton polyhedron of φ_τ . Also, $\Sigma_{\check{\tau}}$ supports the function $\check{\varphi}_{\check{\tau}}$ which determines the log structure on $\check{X}_0(\check{B}, \check{\mathcal{P}})$ via (6.2). Comparing (6.2) with (6.1), we see that we need the two cones $C(\tau)^\vee$ and $\{(m, r) \mid m \in |\Sigma_{\check{\tau}}|, r \geq \check{\varphi}_{\check{\tau}}(m)\}$ to be isomorphic. It is an easy exercise to check that this is the case if and only if τ is a translate of the Newton polyhedron of $\check{\varphi}_{\check{\tau}}$. Comparing with the construction of §1.5, one sees that $(\check{B}, \check{\mathcal{P}}, \check{\varphi})$ is indeed the discrete Legendre transform of $(B, \mathcal{P}, \varphi)$.

So, at one level, the discrete Legendre transform realises mirror symmetry.

EXAMPLE 6.4. Let us consider the mirror to \mathbb{P}^2 from this point of view. $\mathbb{P}^2 \times \mathbb{A}_{\mathbb{k}}^1 \rightarrow \mathbb{A}_{\mathbb{k}}^1$ is a (trivial!) degeneration of \mathbb{P}^2 to a single toric variety. (B, \mathcal{P}) is just $(\mathbb{R}^2, \Sigma_{\mathbb{P}^2})$, with $\Sigma_{\mathbb{P}^2}$ the fan for \mathbb{P}^2 . Choosing an ample line bundle $\mathcal{O}_{\mathbb{P}^2}(n)$ on \mathbb{P}^2 determines (up to linear functions) an integral PL function $\varphi : B \rightarrow \mathbb{R}$ whose Newton polyhedron is, up to translation, $\text{Conv}\{(0, 0), (n, 0), (0, n)\}$. This Newton polyhedron is \check{B} with $\check{\mathcal{P}}$ being the faces of \check{B} and $\check{\varphi}$ being linear. Let us take $n = 1$.

Viewing $(B, \mathcal{P}, \varphi)$ in the cone picture, we can simply construct the Mumford degeneration. Here, we have

$$\tilde{\Delta} = \{(m, r) \mid m \in \mathbb{R}^2, r \geq \varphi(m)\},$$

and

$$\begin{aligned} C(\tilde{\Delta}) &= \overline{\{(sm, sr, s) \mid s \geq 0, (m, r) \in \tilde{\Delta}\}} \\ &= \{(m, r, s) \mid s \geq 0, (m, r) \in \tilde{\Delta}\}, \end{aligned}$$

since $\tilde{\Delta}$ is invariant under rescaling. Thus

$$C(\tilde{\Delta}) \cap (\mathbb{Z}^2 \oplus \mathbb{Z} \oplus \mathbb{Z}) = (\tilde{\Delta} \cap (\mathbb{Z}^2 \oplus \mathbb{Z})) \oplus \mathbb{N}$$

and

$$\text{Proj } \mathbb{k}[(\tilde{\Delta} \cap (\mathbb{Z}^2 \oplus \mathbb{Z})) \oplus \mathbb{N}] \cong \text{Spec } \mathbb{k}[\tilde{\Delta} \cap (\mathbb{Z}^2 \oplus \mathbb{Z})].$$

Thus setting

$$P = \{(m, r) \mid m \in \mathbb{Z}^2, r \geq \varphi(m)\} \subseteq \mathbb{Z}^2 \oplus \mathbb{Z},$$

we get a natural map

$$\text{Spec } \mathbb{k}[P] \rightarrow \text{Spec } \mathbb{k}[\mathbb{N}] = \mathbb{A}_{\mathbb{k}}^1$$

induced by the natural inclusion $\mathbb{N} \hookrightarrow P$ given by $n \mapsto (0, 0, n)$. Note that $P \cong \mathbb{N}^3$: with the correct choice of φ , we can take the generators of P to be $(1, 0, 0)$, $(0, 1, 0)$ and $(-1, -1, 1)$, and the map can be written as

$$\begin{aligned} \text{Spec } \mathbb{k}[x_0, x_1, x_2] &\rightarrow \text{Spec } \mathbb{k}[\kappa], \\ \kappa &\mapsto x_1 x_1 x_2. \end{aligned}$$

This gives the mirror family to \mathbb{P}^2 given in §2.2. This works for other toric varieties too, but a variant of this construction to get multi-dimensional families is necessary to get the mirror construction given in §5.1.1.




FIGURE 1

6.2. From integral tropical manifolds to degenerations in dimension two

We will now fix for this section a triple $(B, \mathcal{P}, \varphi)$ where

- (B, \mathcal{P}) is an integral tropical manifold of dimension two, possibly with boundary, and possibly non-compact.
- φ is a piecewise linear multi-valued strictly convex function with integral slopes.

We need to make two further assumptions about the singularities of B , which concern the monodromy of the local system Λ around each singular point. Recall that Λ is the local system on $B_0 \subseteq B$ of integral vector fields determined by the integral affine structure on B_0 .

By construction, these singularities only occur on compact edges of \mathcal{P} not contained in the boundary of B . Let $\omega \in \mathcal{P}$ be such an edge, contained in two two-dimensional cells $\sigma_{\pm} \in \mathcal{P}$. Label the vertices of ω as v_+ and v_- . Let $d_{\omega} \in \Lambda_{v_+}$ be a primitive tangent vector to ω at v_+ , pointing towards v_- . Let $\check{d}_{\omega} \in \check{\Lambda}_{v_+}$ be the unique integral cotangent vector which is primitive, annihilates d_{ω} , and takes positive values on tangent vectors pointing into σ_+ . We can also view d_{ω} as selecting which of the two vertices of ω is v_+ and which is v_- , and similarly can view \check{d}_{ω} as specifying which cell is σ_+ and which cell is σ_- .

Now consider a loop γ which is based at v_+ and successively passes through σ_+ , v_- , σ_- , and back to v_+ , as in Figure 1. Then it is not difficult to see that the monodromy of the local system Λ around γ , which is a linear transformation

$$T_{\gamma} : \Lambda_{v_+} \rightarrow \Lambda_{v_+},$$

takes the form

$$T_{\gamma}(m) = m + \kappa_{\omega} \langle \check{d}_{\omega}, m \rangle d_{\omega}$$

for some integer κ_{ω} . (See Example 1.28, (4) where this is worked out in an explicit case.) Note that κ_{ω} is independent of choices: interchanging v_+ and v_- or σ_+ and σ_- changes the sign of d_{ω} or \check{d}_{ω} , but also reverses the direction of γ .

DEFINITION 6.5. B is *positive* if $\kappa_{\omega} \geq 0$ for all edges $\omega \in \mathcal{P}$ which are compact and not contained in ∂B . We say B is *simple* if $\kappa_{\omega} \in \{0, 1\}$ for all such edges ω .

Positivity is a necessity even for the existence of a log structure on $\check{X}_0(B, \mathcal{P})$; if positivity fails, the log structure will be bad along an entire stratum of $\check{X}_0(B, \mathcal{P})$, and then there is no hope of a smoothing. Simplicity is not necessary, but will make our arguments, well, simpler.

We shall henceforth assume (B, \mathcal{P}) is positive and simple. Our goal now is to construct an explicit smoothing of $\check{X}_0(B, \mathcal{P})$. In particular, we shall give a proof of the following formal statement:

THEOREM 6.6. *Given (B, \mathcal{P}) positive and simple with $\dim B = 2$, and given a multi-valued strictly convex integral PL function φ on (B, \mathcal{P}) with integral slopes, one can construct a formal flat smoothing $\check{\mathfrak{X}}(B, \mathcal{P}, \varphi) \rightarrow \mathrm{Spf} \mathbb{k}[[t]]$.*

By formal, we mean that $\check{\mathfrak{X}}(B, \mathcal{P}, \varphi)$ is a formal scheme, with underlying topological space $\check{X}_0(B, \mathcal{P})$. The formal scheme $\mathrm{Spf} \mathbb{k}[[t]]$ is the completion of the affine line at the origin.

To get an actual flat deformation $\check{\mathcal{X}}(B, \mathcal{P}, \varphi)$ over $\mathrm{Spec} \mathbb{k}[[t]]$, we need one more ingredient. Provided that $\check{\mathcal{X}}(B, \mathcal{P}, \varphi) \rightarrow \mathrm{Spf} \mathbb{k}[[t]]$ is proper (which is the case if B is compact) and there is a relatively ample line bundle \mathcal{L} on $\check{\mathfrak{X}}(B, \mathcal{P}, \varphi)$, then the Grothendieck existence theorem ([52], 5.4.5) tells us that $\check{\mathfrak{X}}(B, \mathcal{P}, \varphi)$ is obtained by the formal completion of some scheme $\check{\mathcal{X}}(B, \mathcal{P}, \varphi)$ over $\mathrm{Spec} \mathbb{k}[[t]]$ along the central fibre, isomorphic to $\check{X}_0(B, \mathcal{P})$. The flat family $\check{\mathcal{X}}(B, \mathcal{P}, \varphi) \rightarrow \mathrm{Spec} \mathbb{k}[[t]]$ is the desired one.

As we mentioned earlier, $\check{X}_0(B, \mathcal{P})$ carries a natural ample line bundle, obtained by gluing together the line bundles $\mathcal{O}_{\mathbb{P}^\sigma}(1)$ on the irreducible components \mathbb{P}^σ of $\check{X}_0(B, \mathcal{P})$, for $\sigma \in \mathcal{P}_{\max}$. One can in fact check that this extends to an ample line bundle on $\check{\mathfrak{X}}(B, \mathcal{P}, \varphi)$; this is really just a technical point, but we shall not do this here. Instead, we shall focus on the proof of Theorem 6.6, which is a special case of the main result of [49].

6.2.1. Warmup: the Mumford degeneration. To explain the basic strategy, we will first study the case of the Mumford degeneration in greater detail. We start with $B \subseteq N_{\mathbb{R}}$ a lattice polyhedron, \mathcal{P} a polyhedral decomposition of B into lattice polytopes, and $\varphi : B \rightarrow \mathbb{R}$ a convex piecewise linear function with integral slopes. As we know from Example 3.6, this determines a polyhedron

$$\tilde{\Delta} = \{(n, r) \in N_{\mathbb{R}} \oplus \mathbb{R} \mid n \in B, r \geq \varphi(n)\}$$

which in turn defines a toric variety $\mathbb{P}_{\tilde{\Delta}}$ with a map

$$\pi : \mathbb{P}_{\tilde{\Delta}} \rightarrow \mathbb{A}_{\mathbb{k}}^1.$$

Now the basic reason that this approach can't work immediately for more interesting B is that this construction is too global, and in particular requires an embedding of B into $N_{\mathbb{R}}$. In general, we only have local embeddings of B into $N_{\mathbb{R}}$, away from the singular points of B . Thus, before we can generalize this idea, we first need to make this construction more local.

The first step is to consider infinitesimal versions of this. Letting

$$O_k := \mathrm{Spec} \mathbb{k}[t]/(t^{k+1}),$$

consider

$$\pi_k : X_k := \mathbb{P}_{\tilde{\Delta}} \times_{\mathbb{A}^1} O_k \rightarrow O_k.$$

One reason this might be a reasonable thing to do is that we are only asking, in general, for a degeneration over $\text{Spec } \mathbb{k}[[t]]$, not over $\mathbb{A}_{\mathbb{k}}^1$, and in general we don't expect any better (see Example 6.15). Thus we can try to build this deformation order by order, and take a limit. This is precisely what we will do in the general case.

The next step is to describe in detail an affine cover of X_k . Let $\check{\Sigma}$ be the normal fan to $\tilde{\Delta}$, so that $\mathbb{P}_{\tilde{\Delta}} = X_{\check{\Sigma}}$. Of course, $X_{\check{\Sigma}}$ has an affine cover indexed by cones of $\check{\Sigma}$. In particular, if $\tau \in \mathcal{P}$ and

$$\tilde{\tau} := \{(n, r) \mid n \in \tau, r = \varphi(n)\} \subseteq \tilde{\Delta}$$

is the corresponding horizontal face of $\tilde{\Delta}$ projecting to τ , denote by $\tilde{\tau}$ the normal cone $N_{\tilde{\Delta}}(\tilde{\tau})$. Note that $\tilde{\tau}$ determines an affine open subset of $X_{\check{\Sigma}}$, which we shall write as U_{τ} . Then $\{U_{\tau} \mid \tau \in \mathcal{P}\}$ form an open cover of $\mathbb{P}_{\tilde{\Delta}}$. Indeed, while $\tilde{\tau}$ does not run over all cones in $\check{\Sigma}$ as τ runs over cells of \mathcal{P} , all maximal cones of $\check{\Sigma}$ will appear, being dual to zero-dimensional faces of $\tilde{\Delta}$. Since the affine subsets of $X_{\check{\Sigma}}$ corresponding to maximal cones of a fan Σ cover X_{Σ} , we see $\{U_{\tau} \mid \tau \in \mathcal{P}\}$ covers $\mathbb{P}_{\tilde{\Delta}}$.

The open set U_{τ} can be described as follows, using the fact (1.3) that

$$(\tilde{\tau})^{\vee} = (N_{\tilde{\Delta}}(\tilde{\tau}))^{\vee} = T_{\tilde{\tau}}\tilde{\Delta}.$$

(See also Remark 3.5.) We can describe $T_{\tilde{\tau}}\tilde{\Delta}$ as follows. For $\sigma \in \mathcal{P}_{\max}$, let $m_{\sigma} \in M$ be the slope of $\varphi|_{\sigma}$, i.e.,

$$m_{\sigma} = d(\varphi|_{\sigma}) \in M.$$

Let $\varphi_{\tau} : N_{\mathbb{R}} \rightarrow \mathbb{R}$ be defined by

$$\varphi_{\tau}(n) := \max\{\langle m_{\sigma}, n \rangle \mid \tau \subseteq \sigma \in \mathcal{P}_{\max}\}.$$

Then one easily checks that

$$T_{\tilde{\tau}}\tilde{\Delta} = \{(n, r) \in N_{\mathbb{R}} \oplus \mathbb{R} \mid r \geq \varphi_{\tau}(n)\}.$$

Let

$$\begin{aligned} P_{\tau} &:= T_{\tilde{\tau}}\tilde{\Delta} \cap (N \oplus \mathbb{Z}) \\ &= \{(n, r) \in N \oplus \mathbb{Z} \mid r \geq \varphi_{\tau}(n)\}. \end{aligned}$$

Then by definition,

$$U_{\tau} = \text{Spec } \mathbb{k}[P_{\tau}].$$

Of course, π restricts to a map $\pi : U_{\tau} \rightarrow \mathbb{A}_{\mathbb{k}}^1$ given by the regular function $t = z^{(0,1)}$, and we then make a base-change,

$$\pi_k : U_{\tau}^k = U_{\tau} \times_{\mathbb{A}^1} O_k \rightarrow O_k.$$

Describing the sets U_{τ}^k is now a more local problem on B . However, it is still insufficiently local, as we shall see. To further refine this description, we note that U_{τ}^k and $U_{\tau}^0 = \pi^{-1}(0) \cap U_{\tau}$ have the same underlying topological spaces, and hence the same irreducible components. We would like to describe U_{τ}^k by gluing together the irreducible components of U_{τ}^k .

Explicitly, this is done as follows. We first make the definition

DEFINITION 6.7. Let $p = (n, r) \in P_\tau$, and $\tau \subseteq \sigma \in \mathcal{P}_{\max}$. Then the order of p along σ is

$$\text{ord}_\sigma(p) = r - \langle m_\sigma, n \rangle.$$

This is precisely the order of vanishing of the monomial $z^p \in \mathbb{k}[P_\tau]$ along the toric divisor of U_τ specified by $\tilde{\sigma}$.

Denote by $\mathcal{P}_{\max}^\partial$ the set of cells of \mathcal{P} contained in ∂B of dimension $\dim B - 1$. For $\tau \subseteq \sigma \in \mathcal{P}_{\max}^\partial$, let $m_\sigma \in M$ be a primitive generator of the ray $N_\Delta(\sigma)$; note that this is a ray both in the normal fan for Δ and in the normal fan for $\tilde{\Delta}$. Then define

$$\text{ord}_\sigma(p) = \langle m_\sigma, n \rangle.$$

This is precisely the order of vanishing of the monomial $z^p \in \mathbb{k}[P_\tau]$ along the toric divisor specified by the maximal face $\{(n, r) \in \tilde{\Delta} \mid n \in \sigma\}$ of $\tilde{\Delta}$.

Given a pair $\omega \subseteq \tau$ with $\omega, \tau \in \mathcal{P}$, set

$$R_\omega := \mathbb{k}[P_\omega].$$

Consider the monomial ideal $I_{\omega, \tau}^k \subseteq P_\omega$ defined by

$$I_{\omega, \tau}^k := \{p \in P_\omega \mid \exists \sigma \in \mathcal{P}_{\max} \cup \mathcal{P}_{\max}^\partial \text{ with } \tau \subseteq \sigma \text{ s.t. } \text{ord}_\sigma(p) > k\}.$$

$I_{\omega, \tau}^k$ then generates an ideal in R_ω , which we also write as $I_{\omega, \tau}^k$, and define

$$R_{\omega, \tau}^k := R_\omega / I_{\omega, \tau}^k.$$

□

Note that by convexity of φ_ω ,

$$I_{\omega, \tau}^0 = P_\omega \setminus \{(n, r) \mid n \in T_\omega \tau, r = \varphi_\omega(n)\}.$$

So

$$R_{\omega, \tau}^0 \cong \mathbb{k}[(T_\omega \tau) \cap N].$$

As $T_\omega \tau = (N_\tau(\omega))^\vee$ by (1.3) and $N_\tau(\omega)$ is a cone in the normal fan $\check{\Sigma}_\tau$ of τ , in fact $\text{Spec } R_{\omega, \tau}^0$ can be viewed as the affine subset of \mathbb{P}_τ (a toric stratum of $\pi^{-1}(0)$) determined by the face $\omega \subseteq \tau$. Thus $\text{Spec } R_{\omega, \tau}^k$ can be viewed as a kind of thickening of the closed subset $\text{Spec } R_{\omega, \tau}^0$ inside U_ω .

Given $\omega \subseteq \tau_1 \subseteq \tau_2$ with $\omega, \tau_1, \tau_2 \in \mathcal{P}$, we have an inclusion $I_{\omega, \tau_2}^k \subseteq I_{\omega, \tau_1}^k$, and hence a surjection

$$\psi_{\tau_1, \tau_2} : R_{\omega, \tau_2}^k \twoheadrightarrow R_{\omega, \tau_1}^k.$$

EXAMPLE 6.8. In Figure 2, we depict a one-dimensional example. Here, $B = [0, 4] \subseteq N_{\mathbb{R}} = \mathbb{R}$, and the graph of φ is the lower part of $\tilde{\Delta}$ as depicted. We depict $\check{\Sigma}$, the normal fan to $\tilde{\Delta}$, on the right. Figure 3 shows φ_ω and P_ω for various choices of ω . Figure 4 shows the ideals $I_{\omega, \tau}^k$ for various choices. □

Fixing ω , the set of rings

$$\{R_{\omega, \tau}^k \mid \omega \subseteq \tau\},$$

along with the maps ψ_{τ_1, τ_2} , now form an inverse system. We have

LEMMA 6.9.

$$\varprojlim R_{\omega, \tau}^k \cong R_\omega / (t^{k+1})$$

where $t = z^\rho$ for $\rho = (0, 1) \in P_\omega$.




FIGURE 2

PROOF. An element of the inverse limit is a collection of elements $(f_\tau)_{\omega \subseteq \tau}$ with $f_\tau \in R_{\omega, \tau}^k$ such that, whenever $\tau_1 \subseteq \tau_2$,

$$\psi_{\tau_1, \tau_2}(f_{\tau_2}) = f_{\tau_1}.$$

Since $t^{k+1} \in I_{\omega, \tau}^k$ for all $\tau \supseteq \omega$, we have natural surjections

$$\psi_\tau : R_\omega / (t^{k+1}) \rightarrow R_{\omega, \tau}^k.$$

Then clearly the map

$$f \mapsto (\psi_\tau(f))_{\omega \subseteq \tau}$$

defines a homomorphism

$$\psi : R_\omega / (t^{k+1}) \rightarrow \varprojlim R_{\omega, \tau}^k,$$

and we just need to show it is an isomorphism.

For injectivity, suppose $\psi(\sum_p c_p z^p) = 0$, where p runs over elements of

$$P_\omega \setminus ((k+1)\rho + P_\omega).$$

This can only happen if $\psi(z^p) = 0$ whenever $c_p \neq 0$. But if $\psi(z^p) = 0$, we have $p \in I_{\omega, \tau}^k$ for all $\tau \supseteq \omega$. In particular, if we take $\tau = \sigma \in \mathcal{P}_{\max}$, we must have $r > \langle m_\sigma, n \rangle + k$ if we write $p = (n, r)$. Thus for all $\sigma \in \mathcal{P}_{\max}$ containing ω , $r - (k+1) \geq \langle m_\sigma, n \rangle$, so $r - (k+1) \geq \varphi_\omega(n)$. Hence $p - (k+1)\rho \in P_\omega$, i.e., $z^p \in (t^{k+1})$. This shows injectivity.

For surjectivity, consider $(f_\tau)_{\omega \subseteq \tau}$ in the inverse limit. Then we can write

$$f_\tau = \sum_{p \in P_\omega \setminus I_{\omega, \tau}^k} c_{p, \tau} z^p.$$

Suppose that $p \in P_\omega \setminus I_{\omega, \tau}^k$ and $p \in P_\omega \setminus I_{\omega, \tau'}^k$ for two distinct $\tau, \tau' \supseteq \omega$. If we show that $c_{p, \tau} = c_{p, \tau'}$, so that we can define $c_p = c_{p, \tau}$ independently of τ , we can then define $f \in R_\omega / (t^{k+1})$ by

$$f = \sum_{p \in P_\omega \setminus \bigcap_\tau I_{\omega, \tau}^k} c_p z^p,$$

so that $\psi(f) = (f_\tau)_{\omega \subseteq \tau}$, showing surjectivity.







FIGURE 3




FIGURE 4. The shaded areas indicate the points of the monoids in the ideal.

To show this independence, first note that if $\tau \subseteq \sigma \in \mathcal{P}_{\max}$, $\tau' \subseteq \sigma' \in \mathcal{P}_{\max}$, then $I_{\omega,\sigma}^k \subseteq I_{\omega,\tau}^k$, $I_{\omega,\sigma'}^k \subseteq I_{\omega,\tau'}^k$, so $c_{p,\sigma} = c_{p,\tau}$ and $c_{p,\sigma'} = c_{p,\tau'}$. Thus it is enough to show $c_{p,\sigma} = c_{p,\sigma'}$ provided $p \notin I_{\omega,\sigma}^k, I_{\omega,\sigma'}^k$.

First consider the case that there is a codimension one cell $\rho \in \mathcal{P}$ with $\sigma \cap \sigma' = \rho$. Then $I_{\omega,\rho}^k = I_{\omega,\sigma}^k \cup I_{\omega,\sigma'}^k$, so $p \notin I_{\omega,\rho}^k$. Thus $c_{p,\sigma} = c_{p,\rho} = c_{p,\sigma'}$.

For the general case, consider the Newton polyhedron $\Delta_{\varphi_\omega} \subseteq M_{\mathbb{R}}$ (see Definition 1.8) of the function φ_ω . Then Δ_{φ_ω} has vertices $\check{\sigma}, \check{\sigma}'$ with $\check{\sigma} = -m_\sigma$ and $\check{\sigma}' = -m_{\sigma'}$. Suppose $p = (n, r)$. So the fact that $p \notin I_{\omega,\sigma}^k \cup I_{\omega,\sigma'}^k$ says that

$$r + \langle -m_\sigma, n \rangle < k + 1, \quad r + \langle -m_{\sigma'}, n \rangle < k + 1,$$

i.e., both $\check{\sigma}$ and $\check{\sigma}'$ lie on the same side of the hyperplane in $M_{\mathbb{R}}$ defined by

$$\langle \cdot, n \rangle = k + 1 - r.$$

Since Δ_{φ_ω} is convex, one can find a sequence of vertices $\check{\sigma} = \check{\sigma}_1, \check{\sigma}_2, \dots, \check{\sigma}_n = \check{\sigma}'$ all on the same side of this hyperplane, such that $\check{\sigma}_{i-1}$ and $\check{\sigma}_i$ are connected by an edge $\check{\rho}_i$ of Δ_{φ_ω} , corresponding to a codimension one cell ρ_i containing ω . Then $p \notin I_{\omega,\sigma_i}^k$ for $1 \leq i \leq n$, while $\sigma_{i-1} \cap \sigma_i = \rho_i$. Thus we are in the first case considered, with $c_{p,\sigma} = c_{p,\sigma_2} = \dots = c_{p,\sigma'}$, as desired. \square

Since

$$U_\omega^k = \text{Spec } R_\omega / (t^{k+1}),$$

we see that the open sets U_ω^k can be reconstructed from the rings $R_{\omega,\tau}^k$, which are coordinate rings of thickenings of affine subsets of strata of $\pi^{-1}(0)$, for $\pi : \mathbb{P}_{\Delta}^1 \rightarrow \mathbb{A}_{\mathbb{k}}^1$ the usual projection. Of course, X_k can then be reconstructed from the open sets U_ω^k , since if $\omega_1 \subseteq \omega_2$, we have a natural inclusion $U_{\omega_2}^k \subseteq U_{\omega_1}^k$. These are induced by the natural inclusions of rings

$$R_{\omega_1} / (t^{k+1}) \hookrightarrow R_{\omega_2} / (t^{k+1})$$

induced by the inclusions $P_{\omega_1} \subseteq P_{\omega_2}$. In terms of the inverse systems, this can be viewed as being induced by the natural inclusion $R_{\omega_1,\tau}^k \subseteq R_{\omega_2,\tau}^k$ whenever $\omega_1 \subseteq \omega_2 \subseteq \tau$. At any rate, gluing together the open sets U_ω^k via these natural inclusions, one obtains the scheme X_k .

The main point is that while this construction of X_k seems rather more complicated, we have now given a much more local description of X_k which can be used in more general settings, as we explore next.

6.2.2. The global case: no singularities. The first step in generalizing the above approach is to replace B with a general integral tropical manifold without singularities (we still allow B to have boundary or be non-compact). So consider such a B , or rather, a triple $(B, \mathcal{P}, \varphi)$. Note that we now allow φ to be multi-valued. This is important, since, first, we only really need local descriptions of φ defined up to a linear function, and second, there are in general no global strictly convex single-valued piecewise linear functions.

EXAMPLE 6.10. Let $\Gamma \subseteq N$ be a sublattice, and let $B = N_{\mathbb{R}}/\Gamma$. We choose a polyhedral decomposition \mathcal{P} and a multi-valued strictly convex piecewise linear function with integral slopes. This data can, in this case, be viewed on the universal cover $N_{\mathbb{R}}$ of B , in which case \mathcal{P} should be viewed as a Γ -periodic polyhedral

decomposition of $N_{\mathbb{R}}$ and $\varphi : N_{\mathbb{R}} \rightarrow \mathbb{R}$ should be strictly convex, and satisfy for any $\gamma \in \Gamma$

$$\varphi(x + \gamma) = \varphi(x) + \alpha_{\gamma}(x)$$

for some affine linear function $\alpha_{\gamma} : N_{\mathbb{R}} \rightarrow \mathbb{R}$.

In fact, what we are about to say will apply even if B has singularities, so for the moment, we consider the most general case of $(B, \mathcal{P}, \varphi)$ with B an integral tropical manifold *with* singularities.

DEFINITION 6.11. Given $(B, \mathcal{P}, \varphi)$ as above, let $\{(U_i, \varphi_i)\}$ be a choice of representatives for φ on an open cover $\{U_i\}$ of B . Define the sheaf \mathcal{P}_{φ} on B_0 as an extension

$$0 \rightarrow \mathbb{Z} \rightarrow \mathcal{P}_{\varphi} \rightarrow \Lambda \rightarrow 0$$

as follows. First,

$$\mathcal{P}_{\varphi}|_{U_i} \cong \mathbb{Z} \oplus \Lambda|_{U_i}.$$

At $x \in U_i \cap U_j \cap B_0$, we identify

$$(r, m) \in (\mathbb{Z} \oplus \Lambda|_{U_i})_x \text{ with } (r + d(\varphi_j - \varphi_i)(m), m) \in (\mathbb{Z} \oplus \Lambda|_{U_j})_x,$$

giving the gluing of $\mathbb{Z} \oplus \Lambda|_{U_i}$ and $\mathbb{Z} \oplus \Lambda|_{U_j}$ over $U_i \cap U_j \cap B_0$. Note that this makes sense: $\varphi_j - \varphi_i$ is an affine linear function with integral slope, so $d(\varphi_j - \varphi_i)$ is naturally a section of $\check{\Lambda}$ over $U_i \cap U_j \cap B_0$, which can then be evaluated on m a tangent vector.

For a section m of \mathcal{P}_{φ} , we denote its image in Λ under the projection $\mathcal{P}_{\varphi} \rightarrow \Lambda$ by \bar{m} .

REMARK 6.12. In fact, the sheaf \mathcal{P}_{φ} can be described as $\mathcal{A}ff(\check{B}, \mathbb{Z})$, the sheaf of affine linear functions on \check{B} , the discrete Legendre transform of B , after making a canonical identification of B and \check{B} . This description was used in [49]; here we will use the above description to avoid too much use of the discrete Legendre transform.

We next generalize the notion of the monoids P_{τ} to this situation. Previously, these lived in $N \oplus \mathbb{Z}$. Viewing N as the space of integral tangent vectors to $N_{\mathbb{R}}$, it is then natural to imagine that the correct global version of $N \oplus \mathbb{Z}$ is some extension of Λ by \mathbb{Z} . In fact, \mathcal{P}_{φ} is the correct extension.

DEFINITION 6.13. An *exponent* at a point $x \in B_0$ is an element of the stalk of \mathcal{P}_{φ} at x . We define, for any point $x \in B_0$, a monoid

$$P_{\varphi, x} \subseteq \mathcal{P}_{\varphi, x}$$

with $P_{\varphi, x}^{\text{gp}} = \mathcal{P}_{\varphi, x}$, as follows. Suppose $x \in U_i$. For each $\sigma \in \mathcal{P}_{\max}$ such that $x \in \sigma$, let $\varphi_{i, \sigma} \in \check{\Lambda}_{\sigma}$ be the differential of $\varphi_i|_{\sigma}$. Let Σ_x be the fan (of not necessarily strictly convex cones) in $\Lambda_x \otimes_{\mathbb{Z}} \mathbb{R}$ given by the tangent wedges

$$\Sigma_x := \{T_x \tau \mid x \in \tau \in \mathcal{P}\}.$$

Let $\varphi_{i, x} : |\Sigma_x| \rightarrow \mathbb{R}$ be the function which is given by $\varphi_{i, \sigma}$ on the cone $T_x \sigma$ for $\sigma \in \mathcal{P}_{\max}$. Define

$$P_{\varphi, x} := \{(r, m) \mid m \in |\Sigma_x|, r \geq \varphi_{i, x}(m)\} \subseteq \mathcal{P}_{\varphi, x},$$

using the splitting $\mathcal{P}_{\varphi, x} = \mathbb{Z} \oplus \Lambda|_{U_i}$.

The point of the sheaf \mathcal{P}_φ is that the monoid $P_{\varphi,x}$ is independent of the choice of open set U_i used to define it. Indeed, note that

$$\varphi_{j,x} = \varphi_{i,x} + d(\varphi_j - \varphi_i).$$

So if $m \in |\Sigma_x|$, we have

$$\begin{aligned} (r, m) &\in P_{\varphi,x} \text{ as defined in } U_i \\ \Leftrightarrow r &\geq \varphi_{i,x}(m) \\ \Leftrightarrow r &\geq \varphi_{j,x}(m) - d(\varphi_j - \varphi_i)(m) \\ \Leftrightarrow r + d(\varphi_j - \varphi_i)(m) &\geq \varphi_{j,x}(m) \\ \Leftrightarrow (r + d(\varphi_j - \varphi_i)(m), m) &\in P_{\varphi,x} \text{ as defined in } U_j. \end{aligned}$$

This shows that $P_{\varphi,x}$ is well-defined. Furthermore, the same argument shows that the notion of *order* as defined in Definition 6.7 makes sense. If we use the splitting defined using φ_i as above, and write $p \in P_{\varphi,x}$ as (r, m) under this splitting, then for $\sigma \in \mathcal{P}_{\max}$ with $x \in \sigma$, we define

$$\text{ord}_\sigma(p) = r - \varphi_{i,\sigma}(m),$$

and one checks easily as above that this is well-defined independently of the choice of splitting.

Similarly, for $\sigma \in \mathcal{P}_{\max}^\partial$, the set of cells of \mathcal{P} contained in ∂B of dimension $\dim B - 1$, we define, if $x \in \sigma$,

$$\text{ord}_\sigma(p) = \langle n_\sigma, m \rangle,$$

where $n_\sigma \in \check{\Lambda}_x$ is primitive, vanishes on $T_x\sigma$, and is non-negative on $|\Sigma_x|$. This is clearly independent of the splitting, depending only on $m = \bar{p}$.

REMARK 6.14. Although in this subsection we will focus on the case where B has no singularities, let's describe the monodromy of \mathcal{P} locally near a singular point if $\dim B = 2$ and B has positive singularities. Let $p \in B$ be such a singular point, with $p \in \omega$ a one-dimensional face. Then using a representative φ_{U_ω} on U_ω , the open star of ω , we get a splitting

$$\mathcal{P}_\varphi|_{U_\omega} = \mathbb{Z} \oplus \Lambda|_{U_\omega}.$$

So of course the monodromy of \mathcal{P}_φ around the loop depicted in Figure 1 then splits as

$$\begin{aligned} (r, \bar{m}) &\mapsto (r, \bar{m} + \kappa_\omega \langle \check{d}_\omega, \bar{m} \rangle d_\omega) \\ &= (r, \bar{m}) + \kappa_\omega \langle \check{d}_\omega, \bar{m} \rangle (0, d_\omega). \end{aligned}$$

Now recall that the representative φ_{U_ω} for a multi-valued piecewise linear function is given by $\varphi_{U_\omega} = \lambda + \varphi_\omega \circ S_\omega$ (Definition 1.30) where λ is affine linear on U_ω . However, the differential $d\lambda$ is a section of $\check{\Lambda}$. One sees easily from the description of monodromy of Λ that the action of the transpose monodromy on $\check{\Lambda}$ is

$$n \mapsto n + \kappa_\omega \langle n, d_\omega \rangle \check{d}_\omega.$$

Thus, in particular, if λ is well-defined on U_ω , $d\lambda$ must be monodromy invariant. So if $\kappa_\omega \neq 0$, then we must have $\langle d\lambda, d_\omega \rangle = 0$. Thus in particular φ_{U_ω} is constant on ω and for $x \in \text{Int}(\omega) \setminus \{p\}$, $\varphi_{U_\omega,x}$ is zero along the tangent space to ω , generated by d_ω . Thus $(0, d_\omega) \in \mathbb{Z} \oplus \Lambda_{v_+} = \mathcal{P}_{\varphi,v_+}$ can in fact be interpreted as the unique element




FIGURE 5

m_- of $\mathcal{P}_{\varphi, v_+}$ such that $\bar{m}_- = d_\omega$ and $\text{ord}_{\sigma_\pm}(m_-) = 0$. Then the monodromy on $\mathcal{P}_{\varphi, v_+}$ is

$$m \mapsto m + \kappa_\omega \langle \check{d}_\omega, \bar{m} \rangle m_-,$$

and this description is now independent of the choice of splitting. □

Now for $\omega \subseteq \tau \subseteq \sigma$ with $\omega, \tau \in \mathcal{P}$, $\sigma \in \mathcal{P}_{\max}$, we will define a ring $R_{\omega, \tau, \sigma}^k$ as follows. Choose any point $x \in \text{Int}(\omega) \cap B_0$; this gives us the monoid $P_{\varphi, x} \subseteq \mathcal{P}_{\varphi, x}$. We choose a short path from x that passes immediately into the interior of σ with endpoint y , see Figure 5. Via parallel transport in the local system \mathcal{P}_φ , we identify $\mathcal{P}_{\varphi, x}$ with $\mathcal{P}_{\varphi, y}$. The monoid $P_{\varphi, x} \subseteq \mathcal{P}_{\varphi, x}$ then yields, via this identification, a monoid we call $P_{\varphi, \omega, \sigma}$ contained in $\mathcal{P}_{\varphi, y}$. It is easy to see that this is defined independently of the choice of x , since the sheaf \mathcal{P}_φ has no monodromy on the contractible set $\sigma \cap B_0$.

The basic idea here is that we use σ as a reference rather than x ; this will turn out to be better for bookkeeping later on.

Now we define an ideal

$$I_{\omega, \tau, \sigma}^k \subseteq P_{\varphi, x} = P_{\varphi, \omega, \sigma},$$

by

$$I_{\omega, \tau, \sigma}^k = \{p \in P_{\varphi, x} \mid \exists \sigma' \in \mathcal{P}_{\max} \cup \mathcal{P}_{\max}^\partial \text{ with } \tau \subseteq \sigma' \text{ s.t. } \text{ord}_{\sigma'}(p) > k\}.$$

One checks easily that this is well-defined in $P_{\varphi, \omega, \sigma}$ independently of the choice of x . We then define

$$R_{\omega, \tau, \sigma}^k := \mathbb{k}[P_{\varphi, \omega, \sigma}] / I_{\omega, \tau, \sigma}^k.$$

This is completely analogous to the Mumford degeneration of the previous subsection.

So far, as we said, this works even when B has singularities. For the remainder of this subsection, however, we need to assume B has no singularities.

In this case, given $\omega \subseteq \tau \subseteq \sigma, \sigma' \in \mathcal{P}_{\max}$, there is a canonical isomorphism

$$R_{\omega, \tau, \sigma}^k \cong R_{\omega, \tau, \sigma'}^k.$$

This is obtained by parallel transport in \mathcal{P}_φ along a path joining $y \in \sigma$ to $y' \in \sigma'$ as in Figure 6. Thus these rings are defined independently of the reference cells, so we just write them as $R_{\omega, \tau}^k$, as in the Mumford case.




FIGURE 6. Comparing rings with different reference cells.

We then build open sets

$$U_\omega^k := \operatorname{Spec} \varprojlim U_{\omega,\tau}^k$$

using, for $\omega \subseteq \tau_1 \subseteq \tau_2$, the canonical surjections

$$\psi_{\tau_1,\tau_2} : R_{\omega,\tau_2}^k \rightarrow R_{\omega,\tau_1}^k.$$

Furthermore, for $\omega_1 \subseteq \omega_2$, we have inclusions

$$U_{\omega_2}^k \subseteq U_{\omega_1}^k$$

induced by the inclusions $R_{\omega_1,\tau}^k \subseteq R_{\omega_2,\tau}^k$, and so we can glue the various open sets U_ω^k together using these identifications.

This gives a scheme X_k over O_k . If we take $k = 0$, it is not difficult to check that $X_0 = \check{X}_0(B, \mathcal{P})$, i.e., X_0 is obtained by gluing together the \mathbb{P}_σ for $\sigma \in \mathcal{P}_{\max}$. Indeed, consider $\sigma \in \mathcal{P}_{\max}$. For $\omega \subseteq \sigma$ a face, let $x \in \operatorname{Int}(\omega)$ be a point, and suppose $x \in U_i$, where φ is represented by φ_i , so

$$P_{\varphi,x} = \{(r, m) \mid m \in |\Sigma_x|, r \geq \varphi_{i,x}(m)\}.$$

Then, identifying this with $P_{\varphi,\omega,\sigma}$, we see that

$$P_{\varphi,\omega,\sigma} \setminus I_{\omega,\sigma,\sigma}^0 = \{(r, m) \mid m \in |\Sigma_x|, r \geq \varphi_{i,x}(m), r = \varphi_{i,\sigma}(m)\}.$$

Thus $(r, m) \in P_{\varphi,\omega,\sigma} \setminus I_{\omega,\sigma,\sigma}^0$ if and only if m points into σ and $r = \varphi_{i,\sigma}(m)$, by strict convexity of $\varphi_{i,x}$. So we can identify $\mathbb{k}[P_{\varphi,\omega,\sigma}] / I_{\omega,\sigma,\sigma}^0$ with $\mathbb{k}[(T_\omega\sigma) \cap \Lambda_x]$. Identifying Λ_x with Λ_y for $y \in \operatorname{Int}(\sigma)$ a fixed point, we see that $\operatorname{Spec} \mathbb{k}[(T_\omega\sigma) \cap \Lambda_y]$ is the affine open subset of \mathbb{P}_σ specified by the cone $N_\sigma\omega \in \check{\Sigma}_\sigma$, the normal fan to σ . Furthermore, the inclusions $R_{\omega_1,\sigma}^0 \subseteq R_{\omega_2,\sigma}^0$ whenever $\omega_1 \subseteq \omega_2 \subseteq \sigma$, correspond to the natural inclusions

$$\mathbb{k}[(T_{\omega_1}\sigma) \cap \Lambda_y] \subseteq \mathbb{k}[(T_{\omega_2}\sigma) \cap \Lambda_y],$$

which tells us that when we glue together the sets $\operatorname{Spec} R_{\omega,\sigma}^0$ via these inclusions, we obtain \mathbb{P}_σ . Since the sets $\operatorname{Spec} R_{\omega,\sigma}^0$ are the irreducible components of U_ω^0 , we see that X_0 is indeed the gluing of the toric varieties \mathbb{P}_σ .

So this shows that the construction so far is local enough to deal with arbitrary integral tropical manifolds without singularities.




FIGURE 7. The polyhedron $\tilde{\Delta}$.

EXAMPLE 6.15. Let $B = \mathbb{R}/d\mathbb{Z}$ for d a positive integer,

$$\mathcal{P} = \{[i, i + 1] \mid 1 \leq i \leq d\} \cup \{\{i\} \mid 1 \leq i \leq d\}.$$

We can define φ globally on the universal cover \mathbb{R} of B by

$$\varphi(x) = ix - \frac{i(i+1)}{2} \text{ for } x \in [i, i + 1]$$

and then φ satisfies a periodicity condition

$$\varphi(x + d) = \varphi(x) + d \cdot x + \frac{d(d - 1)}{2}.$$

In this case, X_k is a k -th order smoothing of a cycle of d rational curves. In fact, we did not need to use this local description. There is a global construction of X_k given by working on the universal cover of B , as follows.

Let

$$\tilde{\Delta} := \{(n, r) \in \mathbb{R} \oplus \mathbb{R} \mid r \geq \varphi(n)\},$$

as depicted in Figure 7. Let $\check{\Sigma}_{\tilde{\Delta}}$ be the normal fan to $\tilde{\Delta}$. This is an infinite fan, with rays generated by $(i, 1)$, $i \in \mathbb{Z}$, and two-dimensional cones generated by $\{(i, 1), (i + 1, 1)\}$ for $i \in \mathbb{Z}$, see Figure 8.




FIGURE 8. The normal fan of $\tilde{\Delta}$.

This defines a toric variety $X := X_{\check{\Sigma}_{\tilde{\Delta}}}$ which is not of finite type, along with a regular function $\pi : X \rightarrow \mathbb{A}_{\mathbb{k}}^1$ given by projection onto the second coordinate as usual. This is just an infinite version of the Mumford construction, with $\pi^{-1}(0)$ an infinite chain of \mathbb{P}^1 's and the general fibre of π being a copy of \mathbb{G}_m . Note that the group \mathbb{Z} acts linearly on the fan $\check{\Sigma}_{\tilde{\Delta}}$ by

$$\mathbb{Z} \ni 1 \mapsto \begin{pmatrix} 1 & d \\ 0 & 1 \end{pmatrix} \in \text{GL}_2(\mathbb{Z}).$$

Thus \mathbb{Z} acts on X , and this action, restricted to $\pi^{-1}(0)$, is given by shifting the chain of \mathbb{P}^1 's d places.

Now ideally, one would like to divide out by this action, defining X/\mathbb{Z} . However, this quotient is not well-behaved. Indeed, one can see this by noting that if we take t to be the coordinate on \mathbb{A}^1 , then with coordinates z_1, z_2 on the big torus orbit of X corresponding to the standard basis of $\mathbb{R} \oplus \mathbb{R}$, $\pi^{-1}(t)$ is identified with a \mathbb{G}_m with coordinate z_1 and $z_2 = t$. The action of $1 \in \mathbb{Z}$ in these coordinates is $z_1 \mapsto z_1 z_2^d$, $z_2 \mapsto z_2$, and hence the action on $\mathbb{G}_m = \pi^{-1}(t)$ is $z_1 \mapsto z_1 t^d$. This action has fixed points when $t = 1$. As a consequence, this quotient is not well-behaved and does not exist even in the category of complex manifolds.

There are two ways to fix this. The easier way is to work complex analytically, and note that $\pi^{-1}(D)/\mathbb{Z}$ does make sense, where $D = \{t \mid |t| < 1\}$ is the unit disk. This gives us a fibration $\pi^{-1}(D)/\mathbb{Z}$ whose general fibre is an elliptic curve isomorphic to $\mathbb{C}/\langle 1, \frac{1}{2\pi i} \log t \rangle$, and whose fibre over 0 is a cycle of d rational curves. Note that this fibration does not extend to one over \mathbb{A}^1 .

Alternatively, we can set

$$\tilde{X}_k := X \times_{\mathbb{A}^1} O_k,$$

and then $X_k := \tilde{X}_k/\mathbb{Z}$ does make sense. In fact, X_k is precisely the k -th order thickening of $\tilde{X}_0(B, \mathcal{P})$ that we have constructed above. By taking the limit as $k \rightarrow \infty$, one obtains a formal scheme defined over the formal spectrum of $\mathbb{k}[[t]]$. With suitable care, one can then apply the Grothendieck existence theorem ([52], 5.4.5) to show that this arises via the completion along the central fibre of a scheme defined over $\text{Spec } \mathbb{k}[[t]]$. This is the best that we can, in general, achieve in the algebro-geometric setting.

More generally, if $B = \mathbb{R}^g/\Gamma$ for a lattice $\Gamma \subseteq \mathbb{Z}^g \subseteq \mathbb{R}^g$, the same procedure works. This is Mumford's construction of degenerations of abelian varieties, see [84] and [2]; the one-dimensional case just discussed is usually called the *Tate curve*. So in fact we don't really get anything new with this construction until we introduce singularities. \square

6.2.3. Introducing singularities: The strategy. We now consider the case when B is allowed to have singularities. At this point, we shall restrict to the case $\dim B = 2$, as this is the case we will cover here. Much of the complexity of [49] is due to difficulties in higher dimensions. In dimension two, the method of [49] can be described in much the same fashion as the argument of Kontsevich and Soibelman in [70].

So let us consider a local situation where we have an edge $\omega \in \mathcal{P}$ containing a singularity, which we shall assume is positive and simple. With the labelling in Figure 9, the monodromy of a loop γ based at the point x passing *clockwise* around the singularity takes the form

$$T_\gamma(m) = m + \langle \tilde{d}_\omega, m \rangle d_\omega,$$

where d_ω is a primitive tangent vector pointing from v_+ to v_- and \tilde{d}_ω is primitive, orthogonal to d_ω , and is positive on σ_+ . Thus, if we take a basis for Λ_x given by $e_1 = d_\omega$ and e_2 a primitive tangent vector pointing into σ_+ , we find that

$$\begin{aligned} T_\gamma(e_1) &= e_1 \\ T_\gamma(e_2) &= e_1 + e_2. \end{aligned}$$




FIGURE 9

Furthermore, we can take a PL function φ to be 0 on σ_- and to have slope e_2^* on σ_+ , with e_1^*, e_2^* the dual basis of $\check{\Lambda}_x$.

We now have rings $R_{\omega, \sigma_+, \sigma_+}^k$ and $R_{\omega, \omega, \sigma_+}^k$. Identifying $\Lambda_{y_{\pm}}$ with Λ_x via the paths from y_{\pm} to x indicated in Figure 9, we can write, with $\mathbb{Z}^3 = \mathbb{Z} \oplus \Lambda_x = \mathcal{P}_{\varphi, x}$, the generators of $\mathcal{P}_{\varphi, x}$ to be $(1, 0, 1)$, $(0, 0, -1)$, and $(0, \pm 1, 0)$. Set

$$u = z^{(1,0,1)}, \quad v = z^{(0,0,-1)}, \quad w = z^{(0,1,0)},$$

so that

$$\begin{aligned} R_{\omega, \sigma_+, \sigma_+}^k &\cong \mathbb{k}[u, v, w^{\pm 1}]/(v^{k+1}), \\ R_{\omega, \sigma_-, \sigma_-}^k &\cong \mathbb{k}[u, v, w^{\pm 1}]/(u^{k+1}), \\ R_{\omega, \omega, \sigma_{\pm}}^k &\cong \mathbb{k}[u, v, w^{\pm 1}]/(u^{k+1}, v^{k+1}). \end{aligned}$$

Now we have canonical surjections

$$\begin{aligned} R_{\omega, \sigma_+, \sigma_+}^k &\twoheadrightarrow R_{\omega, \omega, \sigma_+}^k, \\ R_{\omega, \sigma_-, \sigma_-}^k &\twoheadrightarrow R_{\omega, \omega, \sigma_-}^k, \end{aligned}$$

and to obtain an inverse system, we need to identify $R_{\omega, \omega, \sigma_-}^k$ and $R_{\omega, \omega, \sigma_+}^k$. However, because of the singularity, we don't have a canonical identification. There are in fact two perfectly natural identifications: one, Ψ , using parallel transport in \mathcal{P}_{φ} along the path from y_+ to y_- via x , as depicted in Figure 9, and the other, Ψ' , coming from a similar path from y_+ to y_- via x' .

Because we identified Λ_{y_+} and Λ_{y_-} by passing through x already,

$$\Psi : R_{\omega, \omega, \sigma_+}^k \rightarrow R_{\omega, \omega, \sigma_-}^k$$

is given by

$$u \mapsto u, \quad v \mapsto v, \quad w \mapsto w.$$

On the other hand, the map from Λ_{y_+} to Λ_{y_-} given by parallel transport through x' is just the inverse to T_{γ} , i.e., $e_1 \mapsto e_1, e_2 \mapsto e_2 - e_1$, so

$$\Psi' : R_{\omega, \omega, \sigma_+}^k \rightarrow R_{\omega, \omega, \sigma_-}^k$$

is given by

$$u \mapsto uw^{-1}, \quad v \mapsto vw, \quad w \mapsto w.$$

Now in these two cases, we can form the inverse limit. In this case, the inverse limit is just the fibre product of rings,

$$R_\omega^k := R_{\omega, \sigma_+, \sigma_+}^k \times_{R_{\omega, \omega, \sigma_-}^k} R_{\omega, \sigma_-, \sigma_-}^k$$

where the map $R_{\omega, \sigma_+, \sigma_+}^k \rightarrow R_{\omega, \omega, \sigma_-}^k$ is the composition of the natural surjection $R_{\omega, \sigma_+, \sigma_+}^k \rightarrow R_{\omega, \omega, \sigma_+}^k$ with either Ψ or Ψ' , and the map $R_{\omega, \sigma_-, \sigma_-}^k \rightarrow R_{\omega, \omega, \sigma_-}^k$ is just the natural surjection. For the fibre product using Ψ , we get

$$R_\omega^k = \mathbb{k}[U, V, W^{\pm 1}, t]/(UV - t, t^{k+1})$$

with

$$U = (u, u), \quad V = (v, v), \quad W = (w, w), \quad t = (uv, uv),$$

while using Ψ' , we get

$$R_\omega^k = \mathbb{k}[U, V, W^{\pm 1}, t]/(UV - tW, t^{k+1})$$

with

$$U = (uw, u), \quad V = (v, vw), \quad W = (w, w), \quad t = (uv, uv).$$

There is no natural way of identifying these two rings, and this is a fundamental problem which needs to be overcome.

The solution is to modify the rings $R_{\omega, \omega, \sigma_\pm}^k$ by localizing them at $1 + w^{-1}$ (or equivalently, $1 + w$), and modifying the maps Ψ and Ψ' by composing each of these with an automorphism of $(R_{\omega, \omega, \sigma_-}^k)_{1+w^{-1}}$. We take the new Ψ to be the old Ψ composed with the automorphism

$$\begin{aligned} u &\mapsto u(1 + w^{-1}), \\ v &\mapsto v(1 + w^{-1})^{-1} \\ w &\mapsto w. \end{aligned}$$

and the new Ψ' to be the old Ψ' composed with the automorphism

$$\begin{aligned} u &\mapsto u(1 + w), \\ v &\mapsto v(1 + w)^{-1} \\ w &\mapsto w. \end{aligned}$$

These new maps Ψ, Ψ' are given by

$$\begin{aligned} \Psi : u &\mapsto u(1 + w^{-1}), \quad v \mapsto v(1 + w^{-1})^{-1}, \quad w \mapsto w \\ \Psi' : u &\mapsto uw^{-1}(1 + w), \quad v \mapsto vw(1 + w)^{-1}, \quad w \mapsto w \end{aligned}$$

Note that these two maps now actually coincide. So the fibre product is now well-defined, and one can see (as will be shown in §6.2.6), using the above choice of Ψ or Ψ' , that

$$R_{\omega, \sigma_+, \sigma_+}^k \times_{(R_{\omega, \omega, \sigma_-}^k)_{1+w^{-1}}} R_{\omega, \sigma_-, \sigma_-}^k$$

is now the ring

$$R_\omega^k := \mathbb{k}[U, V, W^{\pm 1}, t]/(UV + (1 + W^{-1})t, t^{k+1})$$

with

$$U = (u, u(1 + w^{-1})), \quad V = (v(1 + w^{-1}), v), \quad W = (w, w), \quad t = (uv, uv).$$

Note that $\text{Spec } \mathbb{k}[U, V, W^{\pm 1}, t]/(UV + (1 + W^{-1})t)$ now defines a smoothing of $\text{Spec } \mathbb{k}[U, V, W^{\pm 1}]/(UV)$, but the total space of the smoothing is singular at $U = V = t = 0, W = -1$. This singularity reflects the singularity of the affine structure.

This fix may seem rather ad hoc initially, but in fact it is fundamental. The way to view this is as follows. Imagine two rays emanating from the singular point, in opposite directions, along the edge ω . Label the ray going to the right in Figure 9 with $1 + w^{-1}$ and the ray moving to the left with $1 + w$. This instructs us that if we need to cross one of these rays to compare two rings, we twist the map between these two rings using an automorphism defined using the attached function.

In order to continue to glue consistently even away from ω , we need to imagine these rays extending indefinitely in B . However, these various rays intersect at points, destroying compatibility of the gluing because the associated twistings of the gluings don't commute. The question of how to fix this was the crucial problem solved by Kontsevich and Soibelman in [70]. The solution is to add new rays with new attached functions emanating from the intersection points of the old rays in such a way so that compatibility is restored. These in turn generate new collisions, so more rays are added. In the end one obtains for each k a finite set of rays describing the k -th order deformation. As $k \rightarrow \infty$, we normally expect the number of rays to grow without bound.

In the next few sections, we shall make this whole procedure precise.

6.2.4. Structures. In this section we will formalize the data needed to specify the modifications of the gluing sketched in the previous subsection. The data which specifies a way of modifying the gluing is called a *structure*, and consists of a set of *rays*. A ray is a geometric object, essentially a ray in B , which we call a *naked ray*, along with a function, which will specify the gluing when we cross this ray. We first define naked rays.

We fix in this section $(B, \mathcal{P}, \varphi)$ with B an integral tropical manifold with singularities. Furthermore, we assume that B is two-dimensional, positive, and simple. We allow B to have a boundary or be non-compact. Let

$$i : B_0 \hookrightarrow B$$

be the inclusion, and let

$$\Delta := B \setminus B_0.$$

DEFINITION 6.16. A *naked ray* \mathfrak{d} on B is a map

$$\mathfrak{d} : I_{\mathfrak{d}} \rightarrow B,$$

where the interval $I_{\mathfrak{d}}$ is either $[0, +\infty)$ or $[0, L_{\mathfrak{d}}]$ for some $L_{\mathfrak{d}} > 0$, satisfying the following properties:

- (1) \mathfrak{d} is a continuous immersion, i.e., locally on $I_{\mathfrak{d}}$, it is injective. Furthermore, it is differentiable on $\mathfrak{d}^{-1}(B_0)$.
- (2) If $\mathfrak{d}(0) \in B_0$, then $\mathfrak{d}(0)$ has rational coordinates with respect to an integral affine coordinate chart.
- (3) There exists a non-zero global section

$$\bar{m}_{\mathfrak{d}} \in \Gamma(I_{\mathfrak{d}}, \mathfrak{d}^{-1}i_*\Lambda).$$

In particular, if $x \in \mathfrak{d}^{-1}(B_0)$, then the stalk of $\mathfrak{d}^{-1}i_*\Lambda$ at x is canonically identified with the stalk of Λ at $\mathfrak{d}(x)$, so this specifies a tangent vector $\bar{m}_{\mathfrak{d},x}$ at $\mathfrak{d}(x)$. We then require that if t is the coordinate on $I_{\mathfrak{d}}$, then for any point $x \in \mathfrak{d}^{-1}(B_0)$, we have

$$\mathfrak{d}_*(\partial/\partial t) = -\bar{m}_{\mathfrak{d},x}$$

at x . So we can think of $-\bar{m}_{\mathfrak{d},x}$ as specifying the velocity.

- (4) If $I_{\mathfrak{d}} = [0, L_{\mathfrak{d}}]$, then $\mathfrak{d}(L_{\mathfrak{d}}) \in \partial B$ and $\mathfrak{d}([0, L_{\mathfrak{d}}]) \subseteq B \setminus \partial B$. If $I_{\mathfrak{d}} = [0, +\infty)$, then $\mathfrak{d}(I_{\mathfrak{d}}) \subseteq B \setminus \partial B$. So in particular, if B has no boundary, then $I_{\mathfrak{d}} = [0, +\infty)$.

Let us explore what this means. First, the condition that $\mathfrak{d}_*(\partial/\partial t) = -\bar{m}_{\mathfrak{d},x}$ tells us that for any connected open set $U \subseteq I_{\mathfrak{d}}$ with $\mathfrak{d}(U) \subseteq B_0$, the image of U is locally a line segment with rational slope, since it has a locally constant and integral tangent vector.

Second, the behaviour at points $x \in I_{\mathfrak{d}}$ with $\mathfrak{d}(x) \in \Delta$ is restricted. Recall that in a loop based at a point y around a point p of Δ , the local system Λ has monodromy T given in some basis e_1, e_2 of the stalk Λ_y by $T(e_1) = e_1$, $T(e_2) = e_1 + e_2$. This means that if V is a small open neighbourhood of p , then $\Gamma(V \setminus \Delta, \Lambda)$ can be identified with the monodromy invariant elements of Λ_y , so

$$\Gamma(V \setminus \Delta, \Lambda) \cong \mathbb{Z}e_1.$$

Thus $\Gamma(V, i_*\Lambda) = \mathbb{Z}e_1$ by definition of the push-forward, and in particular

$$\Gamma(U, \mathfrak{d}^{-1}i_*\Lambda) = \mathbb{Z}e_1$$

for U a small open neighbourhood of $x \in I_{\mathfrak{d}}$. In particular, $\bar{m}_{\mathfrak{d},x'}$ must, for any $x' \in U \setminus \{x\}$, lie in the subgroup of $\Lambda_{\mathfrak{d}(x')}$ invariant under the monodromy given by parallel transport around a loop about p .

Thus, when the image of \mathfrak{d} passes through a singular point of B , it must pass through in the unique invariant tangent direction. But we know exactly what these invariant tangent directions are: they are given by the edges of \mathcal{P} containing the singular points. Indeed, we saw this explicitly in Example 1.28, (4). Summarizing, the image of \mathfrak{d} is locally a line of rational slope, and when it passes through a singular point, it coincides with the edge of \mathcal{P} containing this singular point.

Note that it is also easy to construct naked rays. Choose a point $y \in B_0 \setminus \partial B$ with rational coordinates in any (or equivalently all) integral affine coordinate charts around y . Choose $\bar{m} \in \Lambda_y$, and define $\mathfrak{d} : I_{\mathfrak{d}} \rightarrow B$ as a solution to the differential equation

$$\mathfrak{d}'(x) = -\bar{m}_x,$$

where \bar{m}_x denotes the parallel transport of \bar{m} along the path determined by \mathfrak{d} . In other words, we just extend \mathfrak{d} locally as a line of rational slope, completely determined once the initial slope $-\bar{m}$ is given. We stop if we reach the boundary of B .

There is one problem: what happens if we hit a singular point p of B ? We are fine if near p , \mathfrak{d} coincides with the edge of \mathcal{P} passing through p ; then we just extend \mathfrak{d} through the singular point. However, we are not allowed to pass through p in any other way. To prevent this, we make the following assumption:

ASSUMPTION 6.17. *For $p \in \Delta$, let $\omega \in \mathcal{P}$ be the edge containing p . Note that ω is identified with a lattice line segment in \mathbb{R} , so it makes sense to demand that p be an irrational point in ω . We assume that all such p are indeed irrational. Since the choice of p inside ω was arbitrary in Construction 1.25, we can always assume this.*

As a consequence of this assumption, we never have to worry about \mathfrak{d} passing through singular points in an illegal way: since $\mathfrak{d}(0)$ is rational and \mathfrak{d} has rational slope, the image of \mathfrak{d} always intersects an edge ω in a rational point, provided \mathfrak{d} is transversal to ω .

If we want to define a naked ray \mathfrak{d} with $\mathfrak{d}(0) \in \Delta$, then we only have two choices for the initial direction of \mathfrak{d} : it must be tangent to the edge containing $\mathfrak{d}(0)$. We can then build this ray as before. Even though $\mathfrak{d}(0)$ is not a rational point, it will still only pass through singular points in the invariant directions, again by the assumption.

We can now define a ray:

DEFINITION 6.18. A ray $(\mathfrak{d}, f_{\mathfrak{d}})$ in B consists of two pieces of data:

- $\mathfrak{d} : I_{\mathfrak{d}} \rightarrow B$ a naked ray,
- $f_{\mathfrak{d}} = 1 + c_m z^m$ where $c_m \in \mathbb{k}$ and m is a section of $\Gamma(I_{\mathfrak{d}}, \mathfrak{d}^{-1} i_* \mathcal{P}_{\varphi})$ such that $\bar{m} \in \Gamma(I_{\mathfrak{d}}, \mathfrak{d}^{-1} i_* \Lambda)$ (see Definition 6.11) is proportional to $\bar{m}_{\mathfrak{d}}$ with a positive constant of proportionality. Furthermore, for each $x \in I_{\mathfrak{d}}$ such that $\mathfrak{d}(x) \in B_0$, the germ m_x of m in $\mathcal{P}_{\varphi, \mathfrak{d}(x)}$ lies in the monoid $P_{\varphi, \mathfrak{d}(x)}$.

Again, let's examine this definition. We are attaching to a naked ray a polynomial $1 + c_m z^m$, with m a section of $\Gamma(I_{\mathfrak{d}}, \mathfrak{d}^{-1} i_* \mathcal{P}_{\varphi})$. Now if $x \in I_{\mathfrak{d}}$ with $\mathfrak{d}(x) \in \Delta$, let V be a small neighbourhood of $\mathfrak{d}(x)$. Then, by the definition of multi-valued piecewise linear function, φ can be represented by a single-valued function on V . Thus \mathcal{P}_{φ} , locally on $V \cap B_0$, splits as $\mathbb{Z} \oplus \Lambda$, so a section of $i_* \mathcal{P}_{\varphi}$ on V can be written as (r, \bar{m}) , where $\bar{m} \in \Gamma(V \cap B_0, \Lambda)$ is tangent to the edge of \mathcal{P} passing through the singular point $\mathfrak{d}(x)$. Thus, similarly we can write $m = (r, \bar{m})$ in a neighbourhood $U \subseteq I_{\mathfrak{d}}$ of x . In particular, there are always liftings of \bar{m} , at least locally, to a section of $\mathfrak{d}^{-1} i_* \mathcal{P}_{\varphi}$.

As a consequence, the section m can be specified simply by giving $m \in \mathcal{P}_{\varphi, \mathfrak{d}(x)}$ for some point $x \in I_{\mathfrak{d}}$ with $\mathfrak{d}(x) \in B_0$. We require that \bar{m} should be positively proportional to $\bar{m}_{\mathfrak{d}}$ at $\mathfrak{d}(x)$. Then by parallel transport, this choice of germ extends to a section of $\mathfrak{d}^{-1} i_* \mathcal{P}_{\varphi}$.

We do have the additional requirement, however, that $m \in P_{\varphi, \mathfrak{d}(x)}$ for all $x \in \mathfrak{d}^{-1}(B_0)$. Since \mathfrak{d} maps $I_{\mathfrak{d}}$ to a ray or line segment which may wind itself around B in some very complicated way, this condition may seem extremely strong. But here is a very important point: As $x \in I_{\mathfrak{d}}$ increases, the order of m with respect to the various maximal cells $\mathfrak{d}(x)$ is contained in *increases*. This is a consequence of the fact that $\mathfrak{d}(x)$ moves in the direction $-\bar{m}$, and is made precise by the following lemma:

- LEMMA 6.19. (1) Let $x \in B_0$, $m \in \mathcal{P}_{\varphi, x}$ such that \bar{m} points into a cell $\sigma \in \mathcal{P}_{\max}$. If $\text{ord}_{\sigma} m \geq 0$ then $m \in P_{\varphi, x}$.
- (2) Let $x \in B_0 \setminus \partial B_0$, $m \in \mathcal{P}_{\varphi, x}$. Let $\sigma_{\pm} \in \mathcal{P}_{\max}$ be two cells containing x with the property that the tangent vector $-\bar{m} \in \Lambda_x$ points into σ_- and \bar{m} points into σ_+ . Then for any $\sigma \in \mathcal{P}_{\max}$ with $x \in \sigma$, we have

$$\text{ord}_{\sigma}(m) - \text{ord}_{\sigma_+}(m) \geq 0,$$

$$\text{ord}_{\sigma_-}(m) - \text{ord}_{\sigma}(m) \geq 0.$$

Furthermore, if \bar{m} is not tangent to $\sigma \cap \sigma_+$, then

$$\text{ord}_{\sigma}(m) - \text{ord}_{\sigma_+}(m) > 0.$$

PROOF. Choose a neighbourhood U of x and a representative φ_U of φ on U , and for $\sigma \in \mathcal{P}_{\max}$ with $x \in \sigma$, let $\varphi_{U, \sigma} \in \check{\Lambda}_x$ be the slope of $\varphi_U|_{\sigma}$. Let $\varphi_{U, x} : |\Sigma_x| \rightarrow \mathbb{R}$ be defined as usual by $\varphi_{U, x}(\bar{m}) = \varphi_{U, \sigma}(\bar{m})$ if $\bar{m} \in T_x \sigma$. For (1), note that \bar{m} pointing into σ implies $\bar{m} \in |\Sigma_x|$, and also $\varphi_{U, x}(\bar{m}) = \varphi_{U, \sigma}(\bar{m})$. Thus $\text{ord}_{\sigma}(m) \geq 0$ implies,

if $m = (r, \bar{m})$ using the splitting of \mathcal{P}_φ given by φ_U , that $r \geq \varphi_{U,\sigma}(\bar{m}) = \varphi_{U,x}(\bar{m})$. So $m \in P_{\varphi,x}$.

For (2), note that since $x \notin \partial B$, Σ_x is a complete fan. Furthermore,

$$\text{ord}_\sigma(r, \bar{m}) = r - \varphi_{U,\sigma}(\bar{m}).$$

Then since \bar{m} points into σ_+ , by strict convexity of $\varphi_{U,x}$, we have $\varphi_{U,\sigma}(\bar{m}) \leq \varphi_{U,\sigma_+}(\bar{m})$ for all $\sigma \in \mathcal{P}_{\max}$ containing x . Thus

$$\text{ord}_\sigma(m) \geq \text{ord}_{\sigma_+}(m).$$

Applying the same argument to $-\bar{m}$ gives $\varphi_{U,\sigma}(-\bar{m}) \leq \varphi_{U,\sigma_-}(-\bar{m})$, or

$$\text{ord}_\sigma(m) \leq \text{ord}_{\sigma_-}(m).$$

Finally, if \bar{m} is not tangent to $\sigma \cap \sigma_+$, then by strict convexity of $\varphi_{U,x}$, $\varphi_{U,\sigma}(\bar{m}) \neq \varphi_{U,\sigma_+}(\bar{m})$, hence the last statement. \square

DEFINITION 6.20. Let $\mathfrak{d} : I_{\mathfrak{d}} \rightarrow B$ be a naked ray, and let $m \in \Gamma(I_{\mathfrak{d}}, \mathfrak{d}^{-1}i_*\mathcal{P}_\varphi)$ be a section such that \bar{m} is a positive multiple of $\bar{m}_{\mathfrak{d}}$. Then define $\text{ord}_m : \mathfrak{d}^{-1}(B_0) \rightarrow \mathbb{Z}$ by

$$\text{ord}_m(x) = \sup\{\text{ord}_\sigma(m_x) \mid \mathfrak{d}(x) \in \sigma \in \mathcal{P}_{\max}\}.$$

Lemma 6.19 implies that this is a step function which increases every time \mathfrak{d} enters a new maximal cell.

We now see it is in fact easy to specify rays:

PROPOSITION 6.21. *Let $\mathfrak{d} : I_{\mathfrak{d}} \rightarrow B$ be a naked ray. If $\mathfrak{d}(0) \in B_0$, let $m \in P_{\varphi,\mathfrak{d}(0)}$ be such that \bar{m} is a negative multiple of $\mathfrak{d}'(0) \in \Lambda_{\mathfrak{d}(0)}$. If $\mathfrak{d}(0) \in \Delta$, let $\epsilon > 0$ be sufficiently small so that $\mathfrak{d}([0, \epsilon])$ is contained in the interior of one edge of \mathcal{P} , and let $m \in P_{\varphi,\mathfrak{d}(\epsilon)}$ be such that \bar{m} is a negative multiple of $\mathfrak{d}'(\epsilon) \in \Lambda_{\mathfrak{d}(\epsilon)}$. Then:*

- (1) *m extends to a section m of $\Gamma(I_{\mathfrak{d}}, \mathfrak{d}^{-1}i_*\mathcal{P}_\varphi)$ such that for all $x \in \mathfrak{d}^{-1}(B_0)$, m defines a germ $m_x \in P_{\varphi,\mathfrak{d}(x)} \subseteq \mathcal{P}_{\varphi,\mathfrak{d}(x)}$.*
- (2) *For every non-negative integer k , there exists an $N \geq 0$ such that if $x \in [N, +\infty) \cap \mathfrak{d}^{-1}(B_0)$, $\tau \in \mathcal{P}$ the unique cell with $\mathfrak{d}(x) \in \text{Int}(\tau)$, $\sigma \in \mathcal{P}_{\max}$ a cell containing τ , then $m_x \in I_{\tau,\sigma}^k$.*

PROOF. The germ m extends to a section of $\Gamma(I_{\mathfrak{d}}, \mathfrak{d}^{-1}i_*\mathcal{P}_\varphi)$ via parallel transport, the discussion after Definition 6.18 showing that this section extends across points of $\mathfrak{d}^{-1}(\Delta)$ precisely because \mathfrak{d} itself is already a naked ray.

The statement that $m_0 \in P_{\varphi,\mathfrak{d}(0)}$ (or $m_\epsilon \in P_{\varphi,\mathfrak{d}(\epsilon)}$ if $\mathfrak{d}(0) \in \Delta$) implies $\text{ord}_\sigma m_0 \geq 0$ for all $\sigma \in \mathcal{P}_{\max}$ containing $\mathfrak{d}(0)$ (or $\text{ord}_\sigma m_\epsilon \geq 0$ for all $\sigma \in \mathcal{P}_{\max}$ containing $\mathfrak{d}(\epsilon)$). But Lemma 6.19, (2) then tells us that this continues to hold, i.e., for all $x \in \mathfrak{d}^{-1}(B_0)$, $\text{ord}_\sigma(m_x) \geq 0$ with $\mathfrak{d}(x) \in \sigma \in \mathcal{P}_{\max}$. Thus, by part (1) of the same lemma, $m_x \in P_{\varphi,\mathfrak{d}(x)}$ for all $x \in \mathfrak{d}^{-1}(B_0)$.

For (2), we just choose an N such that $\text{ord}_m(x) > k$ for all $x \geq N$. \square

We are at last ready to define the notion of a structure.

DEFINITION 6.22. A *structure* \mathcal{S} is a collection of rays satisfying the following properties:

- (1) For every integer $k \geq 0$, let $\mathcal{S}[k]$ be the set of rays $(\mathfrak{d}, 1 + c_m z^m)$ in \mathcal{S} such that there exists an $x \in \mathfrak{d}^{-1}(B_0)$ with $\text{ord}_m(x) \leq k$. Then we require that $\mathcal{S}[k]$ be finite for each k .




FIGURE 10

- (2) If $\omega \in \mathcal{S}$ is an edge containing a singular point p , then \mathcal{S} contains two rays $\mathfrak{d}_{p,\pm}$ such that $\mathfrak{d}_{p,\pm}(0) = p$, $\mathfrak{d}'_{p,\pm}(\epsilon) = \pm \bar{m}$ for ϵ close to zero and \bar{m} a primitive tangent vector to ω . Finally,

$$f_{\mathfrak{d}_{p,\pm}} = 1 + z^{m_{\mp}},$$

where $m_{\mp} \in P_{\varphi,x}$ for any $x \in \text{Int}(\omega) \setminus \{p\}$ is uniquely determined by the requirement that $\bar{m}_{\pm} = \pm \bar{m}$ and $\text{ord}_{\sigma} m_{\pm} = 0$ for σ any maximal cell containing ω . Such a ray is called an *initial ray*. We write the set of initial rays of \mathcal{S} as $\text{Initial}(\mathcal{S})$. For $\mathfrak{d} \in \text{Initial}(\mathcal{S})$ with $\mathfrak{d}(0) \in \omega$ an edge, we define $I_{\mathfrak{d}}^{\text{init}}$ to be the largest subinterval $[0, L] \subseteq I_{\mathfrak{d}}$ such that $\mathfrak{d}([0, L]) \subseteq \omega$.

- (3) If $(\mathfrak{d}, f_{\mathfrak{d}}) \in \mathcal{S}$ is not an initial ray, with $f_{\mathfrak{d}} = 1 + c_{\mathfrak{d}} z^{m_{\mathfrak{d}}}$, then for each $x \in \mathfrak{d}^{-1}(B_0)$, $\text{ord}_{m_{\mathfrak{d}}}(x) > 0$.

EXAMPLE 6.23. Consider the integral tropical manifold depicted in Figure 10. There, the figure on the top depicts the embedding of B minus the depicted cuts into \mathbb{R}^2 , describing the affine structure away from the cuts. In addition, fan structures at the vertices $(-1, 0)$ and $(0, -1)$ are as depicted, specifying the affine structure across these cuts. One can check easily that B is positive and simple. Take the function φ to have slopes $(2, 1)$ on σ_1 , $(1, 2)$ on σ_2 , $(0, 2)$ on σ_3 , $(0, 0)$ on σ_4 , and $(2, 0)$ on σ_5 , and take value zero at the origin; one checks easily that this is strictly convex, well-defined and continuous across the cuts. Figure 11 now depicts a structure \mathcal{S} , consisting of five rays of finite length. Writing elements of Λ as elements of \mathbb{Z}^2 , using the main chart depicted in Figure 10, and using the splitting of $\mathcal{P}_{\varphi} = \mathbb{Z} \oplus \Lambda$




FIGURE 11

given by the above choice of φ , we have two rays emanating from each singular point, i.e., the initial rays, and one additional ray, starting from the origin and heading diagonally northeast.

6.2.5. Compatible structures and gluing. We view a structure as giving a set of data for gluing copies of the various rings $R_{\tau_1, \tau_2, \sigma}^k$. We will now explain how this is done. We will then define the necessary condition to guarantee this gluing makes sense.

For this construction, we fix a structure \mathcal{S} and an integer $k \geq 0$. While \mathcal{S} may contain an infinite number of rays, $\mathcal{S}[k]$ is finite. Furthermore, for each $\mathfrak{d} \in \mathcal{S}[k]$, let $N_{\mathfrak{d}}^k > 0$ be the number promised by Proposition 6.21, (2).

DEFINITION 6.24. Let \mathcal{P}_k be a refinement of the polyhedral decomposition \mathcal{P} which satisfies the following properties:

- (1) The elements of \mathcal{P}_k are convex polyhedra with rational vertices.
- (2) For $\mathfrak{d} \in \mathcal{S}[k]$, $\mathfrak{d}([0, N_{\mathfrak{d}}^k])$ is contained in a union of edges of \mathcal{P}_k and $\mathfrak{d}(N_{\mathfrak{d}}^k)$ is a vertex of \mathcal{P}_k .
- (3) For $\mathfrak{d} \in \mathcal{S}[k] \setminus \text{Initial}(\mathcal{S})$, $\mathfrak{d}(0)$ is a vertex of \mathcal{P}_k .

EXAMPLE 6.25. In Example 6.23, we can take $\mathcal{P}_k = \mathcal{P}$ for all k .

Having chosen \mathcal{P}_k (our construction will not depend on the particular choice), define $\text{Chambers}(\mathcal{S}, k)$ to be the set of maximal cells of \mathcal{P}_k . For each $\mathbf{u} \in \text{Chambers}(\mathcal{S}, k)$, there is a unique $\sigma \in \mathcal{P}_{\max}$ containing \mathbf{u} , which we write as $\sigma_{\mathbf{u}}$.

We now define a category $\underline{\text{Glue}}(\mathcal{S}, k)$ as follows:

- The objects of $\underline{\text{Glue}}(\mathcal{S}, k)$ are triples $(\omega, \tau, \mathbf{u})$ with

$$\omega, \tau \in \mathcal{P}, \mathbf{u} \in \text{Chambers}(\mathcal{S}, k), \text{ and } \omega \subseteq \tau, \omega \cap \mathbf{u} \neq \emptyset, \tau \subseteq \sigma_{\mathbf{u}}.$$

- There is at most one morphism between any two objects, and there is precisely one morphism

$$(\omega, \tau, \mathbf{u}) \rightarrow (\omega', \tau', \mathbf{u}')$$

if and only if $\omega \subseteq \omega'$ and $\tau \supseteq \tau'$.

Note that any morphism in this category decomposes into a composition of morphisms of the following two basic types:

- (I) $\omega \subseteq \omega', \tau \supseteq \tau', \mathbf{u} = \mathbf{u}'$ (*change of strata*).
- (II) $\omega = \omega', \tau = \tau', \dim \mathbf{u} \cap \mathbf{u}' = 1, \omega \cap \mathbf{u} \cap \mathbf{u}' \neq \emptyset$ (*change of chamber*).

We wish to define a functor

$$F_k : \underline{\text{Glue}}(\mathcal{S}, k) \rightarrow \underline{\text{Rings}}$$

from $\underline{\text{Glue}}(\mathcal{S}, k)$ to the category of rings. For an arbitrary choice of structure \mathcal{S} , this functor will not be well-defined, but if \mathcal{S} satisfies a condition that we shall call *compatibility*, the functor becomes well-defined, and we can use it to define the desired modification of the gluing.

First, for any object $(\omega, \tau, \mathbf{u})$ of $\underline{\text{Glue}}(\mathcal{S}, k)$, we associate a ring $R_{\omega, \tau, \mathbf{u}}^k$, as follows:

- If $\omega = \tau$ is an edge of \mathcal{P} containing a singular point $p \in \Delta$, then we set

$$R_{\omega, \tau, \mathbf{u}}^k := (R_{\omega, \tau, \sigma_{\mathbf{u}}})_{f_{\mathfrak{d}_{p, \pm}}}$$

Here the subscript on the right denotes localization at the element $f_{\mathfrak{d}_{p, \pm}} = 1 + z^{m_{\mp}}$. Note that $z^{m_{\pm}}$ is in fact invertible in $R_{\omega, \tau, \sigma_{\mathbf{u}}}^k$, and $f_{\mathfrak{d}_{p, +}} = z^{m_-} f_{\mathfrak{d}_{p, -}}$. Hence it does not matter whether we localize at $f_{\mathfrak{d}_{p, +}}$ or $f_{\mathfrak{d}_{p, -}}$.

- If ω is a vertex and τ is an edge of \mathcal{P} containing a singular point $p \in \Delta$, and $\mathfrak{d} \in \{\mathfrak{d}_{p, +}, \mathfrak{d}_{p, -}\}$ is chosen to be the ray such that ω is the first vertex of \mathcal{P} encountered by that ray, then we set

$$R_{\omega, \tau, \mathbf{u}}^k := (R_{\omega, \tau, \sigma_{\mathbf{u}}})_{f_{\mathfrak{d}}}$$

- Otherwise, we set

$$R_{\omega, \tau, \mathbf{u}}^k := R_{\omega, \tau, \sigma_{\mathbf{u}}}^k$$

We define

$$F_k(\omega, \tau, \mathbf{u}) := R_{\omega, \tau, \mathbf{u}}^k$$

Note that $R_{\omega, \tau, \mathbf{u}}^k$ is an $R_k = \mathbb{k}[t]/(t^{k+1})$ -algebra.

We now wish to associate to every morphism in $\underline{\text{Glue}}(\mathcal{S}, k)$ a homomorphism of the corresponding rings. We shall do this first for the two basic cases: the change of strata and change of chambers.

- I) *The change of strata maps.* Suppose we have

$$(\omega, \tau, \mathbf{u}) \rightarrow (\omega', \tau', \mathbf{u}')$$

a morphism in $\underline{\text{Glue}}(\mathcal{S}, k)$. We then have a composition

$$R_{\omega, \tau, \sigma_{\mathbf{u}}}^k \twoheadrightarrow R_{\omega, \tau', \sigma_{\mathbf{u}}}^k \hookrightarrow R_{\omega', \tau', \sigma_{\mathbf{u}'}}^k$$

of the canonical surjection induced by $I_{\omega, \tau, \sigma_{\mathbf{u}}}^k \subseteq I_{\omega, \tau', \sigma_{\mathbf{u}}}^k$ and the canonical injection induced by $P_{\varphi, \omega, \sigma_{\mathbf{u}}} \subseteq P_{\varphi, \omega', \sigma_{\mathbf{u}'}}$. This induces a map after the appropriate localizations, as we shall now check.

Indeed, consider the relevant localizations of the first map. Here, $R_{\omega, \tau, \mathbf{u}}^k$ is a localization at some element $f = 1 + z^{m_{\pm}}$ of $R_{\omega, \tau, \sigma_{\mathbf{u}}}^k$ if and only if τ is an edge

containing a singularity. If $\tau = \tau'$, then the ring $R_{\omega, \tau', \mathbf{u}}^k$ can be written as a localization of $R_{\omega, \tau', \sigma_{\mathbf{u}}}^k$ at the same element f . If τ' is instead a vertex, then f is easily checked to be invertible in $R_{\omega, \tau', \sigma_{\mathbf{u}}}^k$. Indeed, in this case ω must be a vertex. If $\omega \not\subseteq \partial B$, then by Lemma 6.19, (2), there is a $\sigma \in \mathcal{P}_{\max}$ containing $\tau' = \omega$ such that $\text{ord}_{\sigma}(m_{\pm}) > 0$. On the other hand, if $\omega \subseteq \partial B$, there exists a $\sigma \in \mathcal{P}_{\max}^{\partial}$ such that $\text{ord}_{\sigma}(m_{\pm}) > 0$, since \bar{m}_{\pm} can't be tangent to the boundary. Thus, in any event, some power of $z^{m_{\pm}}$ lies in $I_{\omega, \tau', \sigma_{\mathbf{u}}}^k$, so $1 + z^{m_{\pm}}$ is invertible. Thus the map $R_{\omega, \tau, \mathbf{u}}^k \rightarrow R_{\omega, \tau', \mathbf{u}}^k$ is defined.

Next, consider the relevant localizations of the second map. If $R_{\omega, \tau', \mathbf{u}}^k$ is a localization of $R_{\omega, \tau', \sigma_{\mathbf{u}}}^k$, then $R_{\omega', \tau', \mathbf{u}}^k$ is a localization at the same element. Thus these maps make sense after localizing, and this gives the R_k -algebra homomorphism

$$\psi_{(\omega, \tau), (\omega', \tau')} : R_{\omega, \tau, \mathbf{u}}^k \rightarrow R_{\omega', \tau', \mathbf{u}}^k.$$

This is the change of strata map. □

II) *Change of chambers.* Suppose that $\mathbf{u}, \mathbf{u}' \in \text{Chambers}(\mathcal{S}, k)$ such that

$$\dim \mathbf{u} \cap \mathbf{u}' = 1 \text{ and } \omega \cap \mathbf{u} \cap \mathbf{u}' \neq \emptyset$$

and we have a morphism

$$(\omega, \tau, \mathbf{u}) \rightarrow (\omega, \tau, \mathbf{u}')$$

in $\underline{\text{Glue}}(\mathcal{S}, k)$. Choose a point $y \in \text{Int}(\mathbf{u} \cap \mathbf{u}')$ which lies in a connected component of $(\mathbf{u} \cap \mathbf{u}') \setminus \Delta$ intersecting ω . In particular, y is not a singular point of B . There is a unique $n \in \check{\Lambda}_y$ which is primitive, annihilates the tangent space to $\mathbf{u} \cap \mathbf{u}'$, and is negative on tangent vectors pointing into \mathbf{u}' .

For each pair (\mathfrak{d}, x) such that $\mathfrak{d} \in \mathcal{S}[k]$ and $x \in [0, N_{\mathfrak{d}}^k]$ such that $\mathfrak{d}(x) = y$, we obtain a polynomial

$$f_{(\mathfrak{d}, x)} = 1 + c_{\mathfrak{d}} z^{m_{\mathfrak{d}, x}}$$

where $m_{\mathfrak{d}, x} \in P_{\varphi, \mathfrak{d}(x)}$ is the stalk of $m_{\mathfrak{d}}$ at x . In fact, $m_{\mathfrak{d}, x} \in P_{\varphi, \omega, \sigma_{\mathbf{u}}}$. Indeed, by the choice of y , there is an $x' \in I_{\mathfrak{d}}$ near x such that $\mathfrak{d}(x') \in \omega$, and $m_{\mathfrak{d}, x}$ and $m_{\mathfrak{d}, x'}$ agree under parallel transport. Since $m_{\mathfrak{d}, x'} \in P_{\varphi, \mathfrak{d}(x')}$ by definition of a ray and $P_{\varphi, \mathfrak{d}(x')} \subseteq P_{\varphi, \omega, \sigma_{\mathbf{u}}}$ via parallel transport along \mathfrak{d} , we see that $m_{\mathfrak{d}, x} \in P_{\varphi, \omega, \sigma_{\mathbf{u}}}$. Thus $f_{(\mathfrak{d}, x)} \in R_{\omega, \tau, \mathbf{u}}^k$.

We now consider two cases, defining R_k -algebra homomorphisms $\theta_{\mathbf{u}, \mathbf{u}', y}$.

(1) $\sigma_{\mathbf{u}} = \sigma_{\mathbf{u}'}$. We then define a homomorphism

$$\theta_{\mathbf{u}, \mathbf{u}', y} : R_{\omega, \tau, \mathbf{u}}^k \rightarrow R_{\omega, \tau, \mathbf{u}'}^k$$

by

$$\theta_{\mathbf{u}, \mathbf{u}', y}(z^m) = z^m \prod_{(\mathfrak{d}, x)} f_{(\mathfrak{d}, x)}^{\langle n, \bar{m} \rangle}.$$

Note that this makes sense even if $\langle n, \bar{m} \rangle < 0$. Indeed, since $\sigma_{\mathbf{u}} = \sigma_{\mathbf{u}'}$, we must have $\mathbf{u} \cap \mathbf{u}' \cap \text{Int}(\sigma_{\mathbf{u}}) \neq \emptyset$. Thus no (\mathfrak{d}, x) occurring in this product satisfies $\mathfrak{d} \in \text{Initial}(\mathcal{S})$ and $x \in I_{\mathfrak{d}}^{\text{init}}$. Thus $\text{ord}_{\sigma_{\mathbf{u}}}(m_{\mathfrak{d}, x}) > 0$ by Lemma 6.19, (2), if \mathfrak{d} is an initial ray and by Definition 6.22, (3) if \mathfrak{d} is not an initial ray. From this it follows that each $f_{(\mathfrak{d}, x)}$ appearing in this product is invertible.




FIGURE 12. The X marks the singular point p .

- (2) $\sigma_u \neq \sigma_{u'}$, so that $\sigma_u \cap \sigma_{u'} = \rho$ an edge. In this case, necessarily $\tau \subseteq \rho$. We define

$$\theta_{u,u',y} : R_{\omega,\tau,u}^k \rightarrow R_{\omega,\tau,u'}^k$$

as a composition of the map $R_{\omega,\tau,u}^k \rightarrow R_{\omega,\tau,u'}^k$ obtained by the identification of $P_{\varphi,\omega,\sigma_u}$ and $P_{\varphi,\omega,\sigma_{u'}}$ via parallel transport through y and the map $R_{\omega,\tau,u'}^k \rightarrow R_{\omega,\tau,u'}^k$ defined by

$$z^m \mapsto z^m \prod_{(\mathfrak{d},x)} f_{(\mathfrak{d},x)}^{\langle n, \bar{m} \rangle}.$$

Again this makes sense, as $f_{(\mathfrak{d},x)}$ is always invertible. Indeed, if $\mathfrak{d} \in \text{Initial}(\mathcal{S})$ and $x \in I_{\mathfrak{d}}^{\text{init}}$, then $f_{(\mathfrak{d},x)}$ is invertible, either by definition of the ring $R_{\omega,\tau,u'}^k$ if $\tau = \rho$ and by the argument given in the discussion of the change of strata maps if $\tau \subsetneq \rho$. Otherwise, if $\mathfrak{d} \notin \text{Initial}(\mathcal{S})$, as y is contained in precisely two maximal cells, σ_{\pm} , and $\bar{m}_{\mathfrak{d},x}$ is parallel to ρ , we have $\text{ord}_{\sigma_+} m_{\mathfrak{d},x} = \text{ord}_{\sigma_-} m_{\mathfrak{d},x}$. It then follows as in the previous case that $\text{ord}_{\sigma_{\pm}} m_{\mathfrak{d},x} > 0$ so $f_{(\mathfrak{d},x)}$ is invertible in $R_{\omega,\tau,u'}^k$.

Note that a priori the map $\theta_{u,u',y}$ depends on the choice of point y . However, in the first case, the product $\prod_{(\mathfrak{d},x)} f_{(\mathfrak{d},x)} \in \mathbb{k}[P_{\varphi,\omega,\sigma_u}]$ is independent of the choice of y , by Definition 6.24, (3). Thus the definition of $\theta_{u,u',y}$ is independent of the choice of y , and we can drop the subscript y .

In the second case, the same product depends on y if there is a singular point $p \in \rho$. In particular, $\theta_{u,u',y}$ would appear to depend on whether $y \in \mathfrak{d}_{p,+}(I_{\mathfrak{d}_{p,+}}^{\text{init}})$ or $y \in \mathfrak{d}_{p,-}(I_{\mathfrak{d}_{p,-}}^{\text{init}})$. However, in fact $\theta_{u,u',y}$ is independent of y , as we shall now verify, in essentially the same calculation as carried out in §6.2.3.

Suppose that $p \in u \cap u'$, and choose $y_{\pm} \in u \cap u' \cap \mathfrak{d}_{p,\pm}(\text{Int}(I_{\mathfrak{d}_{p,\pm}}^{\text{init}}))$; see Figure 12. We use parallel transport from σ_u to $\sigma_{u'}$ via y_+ to identify $P_{\varphi,\omega,\sigma_u}$ and $P_{\varphi,\tau,\sigma_{u'}}$. With this identification, the identification of these two monoids via parallel transport from σ_u to $\sigma_{u'}$ via y_- is the monodromy in \mathcal{P}_{φ} given by a counterclockwise

loop about the singular point p . This monodromy on \mathcal{P}_φ is given by

$$m \mapsto m + \langle \check{d}_\omega, \bar{m} \rangle m_-$$

by Remark 6.14, where \check{d}_ω is primitive, annihilates the tangent space to ω , and is positive on σ_u . Thus the map θ_{u,u',y_+} is given by

$$\begin{aligned} z^m &\mapsto z^m f_{\mathfrak{d}_{p,+}}^{\langle \check{d}_\omega, \bar{m} \rangle} \prod_{(\mathfrak{d},x)} f_{(\mathfrak{d},x)}^{\langle \check{d}_\omega, \bar{m} \rangle} \\ &= z^m (1 + z^{m_-})^{\langle \check{d}_\omega, \bar{m} \rangle} \prod_{(\mathfrak{d},x)} f_{(\mathfrak{d},x)}^{\langle \check{d}_\omega, \bar{m} \rangle} \end{aligned}$$

where the product is over all pairs (\mathfrak{d}, x) as usual *excluding* the pair $(\mathfrak{d}_{p,+}, x)$ with $x \in I_{\mathfrak{d}_{p,+}}^{\text{init}}$ such that $\mathfrak{d}_{p,+}(x) = y_+$. On the other hand the map θ_{u,u',y_-} is given as a composition

$$\begin{aligned} z^m &\mapsto z^{m+\langle \check{d}_\omega, \bar{m} \rangle} m_- \\ &\mapsto z^{m+\langle \check{d}_\omega, \bar{m} \rangle} m_- (1 + z^{m_+})^{\langle \check{d}_\omega, \bar{m} \rangle} \prod_{(\mathfrak{d},x)} f_{(\mathfrak{d},x)}^{\langle \check{d}_\omega, \bar{m} \rangle} \\ &= z^m (z^{m_-} + 1)^{\langle \check{d}_\omega, \bar{m} \rangle} \prod_{(\mathfrak{d},x)} f_{(\mathfrak{d},x)}^{\langle \check{d}_\omega, \bar{m} \rangle} \end{aligned}$$

keeping in mind $m_+ = -m_-$. Thus we see that in this case, once again $\theta_{u,u',y}$ is independent of the choice of y , and we drop the y . \square

We have not yet constructed a functor from the category $\underline{\text{Glue}}(\mathcal{S}, k)$ to the category of rings. Rather, we have associated a ring to each object of $\underline{\text{Glue}}(\mathcal{S}, k)$, and a ring homomorphism to each morphism which is either a change of strata or change of chamber homomorphism. More generally, given the unique morphism

$$\mathfrak{e} : (\omega, \tau, \mathbf{u}) \rightarrow (\omega', \tau', \mathbf{u}'),$$

we would like to define

$$F_k(\mathfrak{e}) : R_{\omega, \tau, \mathbf{u}}^k \rightarrow R_{\omega', \tau', \mathbf{u}'}^k.$$

We do this by choosing a sequence

$$\mathbf{u} = \mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_n = \mathbf{u}'$$

such that

$$\mathbf{u}_i \cap \omega' \neq \emptyset, \tau' \subseteq \sigma_{\mathbf{u}_i} \text{ for } 1 \leq i \leq n$$

and

$$\dim \mathbf{u}_{i-1} \cap \mathbf{u}_i = 1 \text{ and } \mathbf{u}_{i-1} \cap \mathbf{u}_i \cap \omega' \neq \emptyset \text{ for } 2 \leq i \leq n.$$

We then get a sequence of change of chamber maps

$$\theta_{\mathbf{u}_{i-1}, \mathbf{u}_i} : R_{\omega', \tau', \mathbf{u}_{i-1}}^k \rightarrow R_{\omega', \tau', \mathbf{u}_i}^k$$

for $2 \leq i \leq n$, and a change of strata map

$$\psi_{(\omega, \tau), (\omega', \tau')} : R_{\omega, \tau, \mathbf{u}_1}^k \rightarrow R_{\omega', \tau', \mathbf{u}_1}^k.$$

We take

$$F_k(\mathfrak{e}) := \theta_{\mathbf{u}_{n-1}, \mathbf{u}_n} \circ \dots \circ \theta_{\mathbf{u}_1, \mathbf{u}_2} \circ \psi_{(\omega, \tau), (\omega', \tau')}$$

There is still a problem: this composition depends on the particular choice of the sequence $\{\mathbf{u}_i\}$. In order for the gluing construction we shall give to work, we need

$F_k(\epsilon)$ to be defined independently of any choices. For this, we need to impose another condition on the structures we consider, which we call *compatibility*.

DEFINITION 6.26. We denote by $\text{Joints}(\mathcal{S}, k)$ the set of vertices of \mathcal{P}_k not contained in ∂B , calling the elements of this set *joints*. For j a joint, σ_j denotes the smallest cell of \mathcal{P} containing j .

DEFINITION 6.27. Let $j \in \text{Joints}(\mathcal{S}, k)$ be a joint. The structure \mathcal{S} is *consistent at j to order k* if for any choice of cyclic ordering u_1, \dots, u_n of the chambers containing j , the composition

$$\theta_{u_n, u_1} \circ \theta_{u_{n-1}, u_n} \circ \dots \circ \theta_{u_1, u_2} : R_{\sigma_j, \sigma_j, u_1}^k \rightarrow R_{\sigma_j, \sigma_j, u_1}^k$$

is the identity.

We say \mathcal{S} is *compatible to order k* if it is consistent to order k at each joint $j \in \text{Joints}(\mathcal{S}, k)$.

Compatibility is sufficient for proving the well-definedness of the morphisms $F_k(\epsilon)$.

THEOREM 6.28. *Suppose \mathcal{S} is compatible to order k .*

- (1) *Given $\epsilon : (\omega, \tau, u) \rightarrow (\omega', \tau', u')$, the homomorphism $F_k(\epsilon)$ is well-defined independently of the choice of $u = u_1, \dots, u_n = u'$.*
- (2) *Given $\epsilon : (\omega, \tau, u) \rightarrow (\omega', \tau', u')$ and $\epsilon' : (\omega', \tau', u') \rightarrow (\omega'', \tau'', u'')$,*

$$F_k(\epsilon' \circ \epsilon) = F_k(\epsilon') \circ F_k(\epsilon).$$

Thus F_k is a functor.

PROOF. Let

$$A := \{u \in \text{Chambers}(\mathcal{S}, k) \mid \omega' \cap u \neq \emptyset, \tau' \subseteq \sigma_u\}.$$

Define Σ to be the abstract two-dimensional cell complex with A as set of vertices, edges connecting $u, u' \in A$ if $\dim u \cap u' = 1$, and a disk glued in for any cycle of chambers u_1, \dots, u_n with a common joint. Note that if there is such a joint, then $\omega' = \tau' = \sigma_j$. An edge with vertices u, u' defines a change of chamber map $\theta_{u, u'} : R_{\omega', \tau', u}^k \rightarrow R_{\omega', \tau', u'}^k$. Consistency says that the composition of these maps following the boundary of a two-cell is the identity. Thus we obtain the desired independence in (1) as long as Σ is simply connected.

To see that Σ is simply connected, we proceed as follows. Let

$$V := \bigcup_{\substack{\sigma \in \mathcal{P} \\ \tau' \subseteq \sigma}} \text{Int}(\sigma)$$

denote the open star of τ' with respect to \mathcal{P} . Then for $u \in \text{Chambers}(\mathcal{S}, k)$, the condition $\tau' \subseteq \sigma_u$ is equivalent to $\text{Int}(u) \subseteq V$, and such chambers define a decomposition of V into polyhedra. Call this decomposition \mathcal{P}_V . Since $\omega' = \tau'$ is topologically a ball the cells of \mathcal{P}_V intersecting ω' form a polyhedral decomposition \mathcal{P}' of a two-ball. Now Σ is the dual polyhedral decomposition of \mathcal{P}' , which is thus simply connected.

For (2), let u_1, \dots, u_n be a sequence of chambers used to compute $F_k(\epsilon)$, and u'_1, \dots, u'_m a sequence of chambers used to compute $F_k(\epsilon')$, so that

$$\begin{aligned} F_k(\epsilon') \circ F_k(\epsilon) &= \theta_{u'_{m-1}, u'_m} \circ \dots \circ \theta_{u'_1, u'_2} \circ \psi_{(\omega', \tau'), (\omega'', \tau'')} \\ &\quad \circ \theta_{u_{n-1}, u_n} \circ \dots \circ \theta_{u_1, u_2} \circ \psi_{(\omega, \tau), (\omega', \tau')}. \end{aligned}$$

One checks easily that $\psi^{(\omega',\tau'),(\omega'',\tau')}$ commutes with each of the $\theta_{\mathbf{u}_{i-1},\mathbf{u}_i}$, $2 \leq i \leq n$, noting that $\theta_{\mathbf{u}_{i-1},\mathbf{u}_i}$ can be interpreted also as a map

$$R_{\omega'',\tau'',\mathbf{u}_{i-1}}^k \rightarrow R_{\omega'',\tau'',\mathbf{u}_i}^k$$

since

$$\begin{aligned} \mathbf{u}_{i-1} \cap \mathbf{u}_i \cap \omega' \neq \emptyset &\Rightarrow \mathbf{u}_{i-1} \cap \mathbf{u}_i \cap \omega'' \neq \emptyset \\ \tau' \subseteq \sigma_{\mathbf{u}_i} &\Rightarrow \tau'' \subseteq \sigma_{\mathbf{u}_i}. \end{aligned}$$

Thus

$$\begin{aligned} F_k(\boldsymbol{\epsilon}') \circ F_k(\boldsymbol{\epsilon}) &= \theta_{\mathbf{u}'_{m-1},\mathbf{u}'_m} \circ \cdots \circ \theta_{\mathbf{u}_1,\mathbf{u}_2} \circ \psi^{(\omega',\tau'),(\omega'',\tau')} \circ \psi^{(\omega,\tau),(\omega',\tau')} \\ &= \theta_{\mathbf{u}'_{m-1},\mathbf{u}'_m} \circ \cdots \circ \theta_{\mathbf{u}_1,\mathbf{u}_2} \circ \psi^{(\omega,\tau),(\omega'',\tau')} \\ &= F_k(\boldsymbol{\epsilon}' \circ \boldsymbol{\epsilon}), \end{aligned}$$

as desired, using the sequence $\mathbf{u}_1, \dots, \mathbf{u}_n = \mathbf{u}'_1, \dots, \mathbf{u}'_m$ to define $F_k(\boldsymbol{\epsilon}' \circ \boldsymbol{\epsilon})$. \square

To summarize, given a structure \mathcal{S} which is compatible to order k , we have constructed a functor

$$F_k : \underline{\text{Glue}}(\mathcal{S}, k) \rightarrow \underline{\text{Rings}},$$

with

$$F_k(\omega, \tau, \mathbf{u}) = R_{\omega,\tau,\mathbf{u}}^k.$$

In fact, F_k maps to the category of R_k -algebras.

EXAMPLE 6.29. Returning to Example 6.23, the structure \mathcal{S} given there is compatible to all orders. The only joint is at the origin, and taking $\mathbf{u}_i = \sigma_i$, one computes

$$\begin{aligned} &\theta_{\mathbf{u}_5,\mathbf{u}_1} \circ \cdots \circ \theta_{\mathbf{u}_1,\mathbf{u}_2}(z^m) \\ &= \theta_{\mathbf{u}_5,\mathbf{u}_1} \circ \cdots \circ \theta_{\mathbf{u}_2,\mathbf{u}_3}(z^m(1 + z^{(0,-1,-1)})^{\langle(1,-1),\bar{m}\rangle}) \\ &= \theta_{\mathbf{u}_5,\mathbf{u}_1} \circ \cdots \circ \theta_{\mathbf{u}_3,\mathbf{u}_4}(z^m(1 + z^{(0,0,-1)})^{\langle(1,0),\bar{m}\rangle}) \\ &\quad \cdot \left(1 + \frac{z^{(0,-1,-1)}}{1 + z^{(0,0,-1)}}\right)^{\langle(1,-1),\bar{m}\rangle} \\ &= \theta_{\mathbf{u}_5,\mathbf{u}_1} \circ \theta_{\mathbf{u}_4,\mathbf{u}_5}(z^m(1 + z^{(0,-1,0)})^{\langle(0,1),\bar{m}\rangle} \left(1 + \frac{z^{(0,0,-1)}}{1 + z^{(0,-1,0)}}\right)^{\langle(1,0),\bar{m}\rangle}) \\ &\quad \cdot \left(1 + \frac{z^{(0,-1,-1)}}{(1 + z^{(0,0,-1)})/(1 + z^{(0,-1,0)})} \right)^{\langle(1,-1),\bar{m}\rangle} \end{aligned}$$

$$\begin{aligned}
&= \theta_{\mathbf{u}_5, \mathbf{u}_1} \circ \theta_{\mathbf{u}_4, \mathbf{u}_5} \left(z^m (1 + z^{(0, -1, 0)})^{\langle (0, 1), \bar{m} \rangle} \left(1 + \frac{z^{(0, 0, -1)}}{1 + z^{(0, -1, 0)}} \right)^{\langle (1, 0), \bar{m} \rangle} \right. \\
&\quad \cdot \left. \left(1 + \frac{z^{(0, -1, -1)}}{1 + z^{(0, -1, 0)} + z^{(0, 0, -1)}} \right)^{\langle (1, -1), \bar{m} \rangle} \right) \\
&= \theta_{\mathbf{u}_5, \mathbf{u}_1} \left(z^m (1 + z^{(0, 0, -1)})^{\langle (-1, 0), \bar{m} \rangle} (1 + z^{(0, -1, 0)} (1 + z^{(0, 0, -1)}))^{\langle (0, 1), \bar{m} \rangle} \right. \\
&\quad \cdot \left(1 + \frac{z^{(0, 0, -1)}}{1 + z^{(0, -1, 0)} (1 + z^{(0, 0, -1)})} \right)^{\langle (1, 0), \bar{m} \rangle} \\
&\quad \cdot \left. \left(1 + \frac{z^{(0, -1, -1)} (1 + z^{(0, 0, -1)})}{1 + z^{(0, -1, 0)} (1 + z^{(0, 0, -1)}) + z^{(0, 0, -1)}} \right)^{\langle (1, -1), \bar{m} \rangle} \right) \\
&= \theta_{\mathbf{u}_5, \mathbf{u}_1} \left(z^m (1 + z^{(0, 0, -1)})^{\langle (-1, 0), \bar{m} \rangle} (1 + z^{(0, -1, 0)} + z^{(0, -1, -1)})^{\langle (0, 1), \bar{m} \rangle} \right. \\
&\quad \cdot \left(\frac{(1 + z^{(0, -1, 0)})(1 + z^{(0, 0, -1)})}{1 + z^{(0, -1, 0)} + z^{(0, -1, -1)}} \right)^{\langle (1, 0), \bar{m} \rangle} \\
&\quad \cdot \left. \left(1 + \frac{z^{(0, -1, -1)}}{1 + z^{(0, -1, 0)}} \right)^{\langle (1, -1), \bar{m} \rangle} \right) \\
&= \theta_{\mathbf{u}_5, \mathbf{u}_1} \left(z^m (1 + z^{(0, -1, 0)} + z^{(0, -1, -1)})^{\langle (0, 1), \bar{m} \rangle} \right. \\
&\quad \cdot \left(\frac{1 + z^{(0, -1, 0)}}{1 + z^{(0, -1, 0)} + z^{(0, -1, -1)}} \right)^{\langle (1, 0), \bar{m} \rangle} \\
&\quad \cdot \left. \left(\frac{1 + z^{(0, -1, 0)} + z^{(0, -1, -1)}}{1 + z^{(0, -1, 0)}} \right)^{\langle (1, -1), \bar{m} \rangle} \right) \\
&= \theta_{\mathbf{u}_5, \mathbf{u}_1} \left(z^m (1 + z^{(0, -1, 0)})^{\langle (0, 1), \bar{m} \rangle} \right) \\
&= z^m.
\end{aligned}$$

Here, the extra diagonal ray in the scattering diagram over and above the initial rays is crucial to guarantee consistency, as the automorphisms associated to the two initial rays passing through this joint do not commute.

6.2.6. k -th order deformations from compatible structures. We fix a structure \mathcal{S} on B which is compatible to order k . This gives a functor

$$F_k : \underline{\text{Glue}}(\mathcal{S}, k) \rightarrow \underline{\text{Rings}}.$$

We now want to define a k -th order deformation of $\check{X}_0(B, \mathcal{P})$. We will do this pretty much as we did in §6.2.2, but with maps given by F_k .

First, fix $\omega \in \mathcal{P}$. We will define a scheme U_ω^k analogous to that defined in §6.2.2. For each $\tau \in \mathcal{P}$ containing ω , choose $\mathbf{u}_\tau \in \text{Chambers}(\mathcal{S}, k)$ such that $\mathbf{u}_\tau \cap \omega \neq \emptyset$, $\tau \subseteq \sigma_{\mathbf{u}_\tau}$. Then given $\omega \subseteq \tau_1 \subseteq \tau_2$, we obtain a morphism

$$\mathbf{e}_{\tau_1, \tau_2} : (\omega, \tau_2, \mathbf{u}_{\tau_2}) \rightarrow (\omega, \tau_1, \mathbf{u}_{\tau_1})$$

in $\underline{\text{Glue}}(\mathcal{S}, k)$, and hence a homomorphism

$$F_k(\mathbf{e}_{\tau_1, \tau_2}) : R_{\omega, \tau_2, \mathbf{u}_{\tau_2}}^k \rightarrow R_{\omega, \tau_1, \mathbf{u}_{\tau_1}}^k.$$

This gives an inverse system of rings: since F_k is a functor, $F_k(\mathbf{e}_{\tau_2, \tau_3}) \circ F_k(\mathbf{e}_{\tau_1, \tau_2}) = F_k(\mathbf{e}_{\tau_1, \tau_3})$. We set

$$R_\omega^k := \varprojlim_{\tau \supseteq \omega} R_{\omega, \tau, \mathbf{u}_\tau}^k$$

and

$$U_\omega^k := \text{Spec } R_\omega^k.$$

These definitions are in fact independent of the choice of the \mathbf{u}_τ 's. Indeed, given a different set of choices \mathbf{u}'_τ , for $\tau \supseteq \omega$, giving maps $\mathbf{e}'_{\tau_1, \tau_2} : (\omega, \tau, \mathbf{u}'_{\tau_2}) \rightarrow (\omega, \tau, \mathbf{u}'_{\tau_1})$, there are unique maps

$$\mathbf{e}_\tau : (\omega, \tau, \mathbf{u}_\tau) \rightarrow (\omega, \tau, \mathbf{u}'_\tau)$$

with

$$F_k(\mathbf{e}'_{\tau_1, \tau_2}) \circ F_k(\mathbf{e}_{\tau_2}) = F_k(\mathbf{e}_{\tau_1}) \circ F_k(\mathbf{e}_{\tau_1, \tau_2}).$$

Furthermore, as $F_k(\mathbf{e}_\tau)$ is given by a sequence of change of chamber maps, $F_k(\mathbf{e}_\tau)$ is an isomorphism, and hence gives an isomorphism between the inverse systems defined by the choice of the \mathbf{u}_τ and the \mathbf{u}'_τ .

LEMMA 6.30. U_ω^0 is isomorphic to $\text{Spec } \mathbb{k}[P_{\varphi, x}]/(t)$ for $x \in \text{Int}(\omega) \cap B_0$.

PROOF. Note that we can take $\mathcal{P}_0 = \mathcal{P}$, as $\mathcal{S}[0]$ can be taken just to consist of the initial rays of \mathcal{S} , and for an initial ray \mathfrak{d} , we can take $N_\mathfrak{d}^0$ so that $[0, N_\mathfrak{d}^0] = I_\mathfrak{d}^{\text{init}}$. Thus we can take $\text{Chambers}(\mathcal{S}, 0) = \mathcal{P}_{\text{max}}$. Now fix ω and $x \in \text{Int}(\omega) \cap B_0$. Note that for each $(\omega, \tau, \mathbf{u}) \in \underline{\text{Glue}}(\mathcal{S}, 0)$, $R_{\omega, \tau, \sigma_\mathbf{u}}^0$ can be identified, via parallel transport from the interior of $\sigma_\mathbf{u}$ to x , with $\mathbb{k}[P_{\varphi, x}]/I_{\omega, \tau, x}^0$, where $I_{\omega, \tau, x}^0$ is the complement of the face of $P_{\varphi, x}$ sitting over the tangent cone T_x in the fan Σ_x . In particular, $R_{\omega, \tau, \sigma_\mathbf{u}}^0$ is an integral domain, and $R_{\omega, \tau, \mathbf{u}}^0$ is a localization of $R_{\omega, \tau, \sigma_\mathbf{u}}^0$, so $R_{\omega, \tau, \sigma_\mathbf{u}}^0 \subseteq R_{\omega, \tau, \mathbf{u}}^0$ (with equality unless $\dim \tau = 1$). A change of stratum map

$$\mathbf{e} : (\omega, \tau_2, \mathbf{u}) \rightarrow (\omega, \tau_1, \mathbf{u})$$

yields

$$F_0(\mathbf{e}) : R_{\omega, \tau_2, \mathbf{u}}^0 \rightarrow R_{\omega, \tau_1, \mathbf{u}}^0,$$

which is just the localization of the canonical surjection

$$\mathbb{k}[P_{\varphi, x}]/I_{\omega, \tau_2, x}^0 \rightarrow \mathbb{k}[P_{\varphi, x}]/I_{\omega, \tau_1, x}^0.$$

On the other hand, the only possibly non-trivial change of chamber map is of the form $\mathbf{e} : (\omega, \tau, \mathbf{u}) \rightarrow (\omega, \tau, \mathbf{u}')$ with $\dim \tau = 1$, and \mathbf{u}, \mathbf{u}' the two different maximal cells containing τ . But if $z^m \in \mathbb{k}[P_{\varphi, x}]/I_{\omega, \tau, x}^0$ is non-zero, then necessarily \bar{m} is tangent to τ and hence by the definition of $\theta_{\mathbf{u}, \mathbf{u}'}$, $\theta_{\mathbf{u}, \mathbf{u}'}(z^m) = z^m$. Hence this change of chamber map is trivial.

From this, we conclude there is an inclusion of inverse systems

$$(\mathbb{k}[P_{\varphi, x}]/I_{\omega, \tau, x}^0)_\tau \rightarrow (R_{\omega, \tau, \mathbf{u}_\tau}^0)_\tau$$

with the homomorphisms in the left-hand inverse system being the canonical ones and the homomorphisms in the right-hand one being the ones defining R_ω^0 . Furthermore, $\mathbb{k}[P_{\varphi, x}]/I_{\omega, \tau, x}^0 \cong R_{\omega, \tau, \mathbf{u}_\tau}^0$ unless $\dim \tau = 1$. Since an element (f_τ) of the inverse limit $\varprojlim \mathbb{k}[P_{\varphi, x}]/I_{\omega, \tau, x}^0$ is completely determined by those f_τ with $\tau \in \mathcal{P}_{\text{max}}$, and the same is true for $\varprojlim R_{\omega, \tau, \mathbf{u}_\tau}^0$, we see in fact the two inverse limits are isomorphic. So

$$R_\omega^0 \cong \varprojlim \mathbb{k}[P_{\varphi, x}]/I_{\omega, \tau, x}^0 \cong \mathbb{k}[P_{\varphi, x}]/(t),$$

the last isomorphism as in Lemma 6.9. \square

We will now show that U_ω^k is a flat deformation of U_ω^0 over $O_k = \text{Spec } \mathbb{k}[t]/(t^{k+1})$. Before we get started, we will need the following observation. As all the rings $R_{\omega,\tau,u}^k$ are $R_k = \mathbb{k}[t]/(t^{k+1})$ -algebras, R_ω^k is also a R_k -algebra. We then have a complex of R_k -modules for each $\ell \leq k$,

$$(6.3) \quad 0 \rightarrow R_\omega^0 \xrightarrow{\cdot t^\ell} R_\omega^\ell \rightarrow R_\omega^{\ell-1} \rightarrow 0.$$

Here the map $R_\omega^0 \rightarrow R_\omega^\ell$ is given by

$$\varprojlim R_{\omega,\tau,u_\tau}^0 \ni (f_\tau)_\tau \mapsto (t^\ell f_\tau)_\tau.$$

Since t^ℓ annihilates any element $z^m \in R_{\omega,\tau,u_\tau}^\ell$ with $z^m \in I_{\omega,\tau,\sigma_{u_\tau}}^0$, one sees that $(t^\ell f_\tau)_\tau \in \varprojlim R_{\omega,\tau,u_\tau}^\ell$. The map $R_\omega^\ell \rightarrow R_\omega^{\ell-1}$ is given by $(f_\tau)_\tau \mapsto (f_\tau \bmod I_{\omega,\tau,\sigma_{u_\tau}}^{\ell-1})$. A priori, (6.3) is not an exact sequence, only a complex.

LEMMA 6.31. *Suppose that (6.3) is exact for each $\ell \leq k$. Then U_ω^k is a flat deformation of U_ω^0 over O_k .*

PROOF. If (6.3) is exact, then inductively the map $R_\omega^\ell \rightarrow R_\omega^0$ given by $(f_\tau)_\tau \mapsto (f_\tau \bmod I_{\omega,\tau,\sigma_{u_\tau}}^0)_\tau$ is surjective. So in particular, the map $R_\omega^0 \rightarrow R_\omega^\ell$ given by multiplication by t^ℓ maps into $t^\ell R_\omega^\ell$. However, it is clear in any event that the image of this map contains $t^\ell R_\omega^\ell$. Hence the image of R_ω^0 in R_ω^ℓ under the map given by multiplication by t^ℓ is $t^\ell R_\omega^\ell$, so exactness of (6.3) implies $R_\omega^{\ell-1} \cong R_\omega^\ell / t^\ell R_\omega^\ell$. Inductively, this shows that $R_\omega^k / t R_\omega^k \cong R_\omega^0$. Furthermore, by [78], Theorem 22.3 (applying condition (4) of that theorem), R_ω^k is a flat R_k -algebra. Thus U_ω^k is a flat deformation of U_ω^0 over O_k . \square

In what follows, we will make use of the toric strata of U_ω^0 for $\omega \in \mathcal{P}$. Given $\tau \in \mathcal{P}$ with $\tau \supseteq \omega$, there is a toric stratum V_τ of U_ω^0 corresponding to τ . In fact, $V_\tau \cong \text{Spec } R_{\omega,\tau,\sigma_{u_\tau}}^0$. Now define $Z_\omega \subseteq U_\omega^0$ to be a subset of the one-dimensional strata of U_ω^0 as follows. If $\tau \supseteq \omega$ is a one-dimensional cell, we take $Z_\omega \cap V_\tau$ to be empty if $\tau \cap \Delta = \emptyset$. If $\tau \cap \Delta = \{p\}$, then one of $f_{\partial_{p,\pm}}$ can be viewed as an element of $R_{\omega,\tau,\sigma_{u_\tau}}^0$, and hence generates an ideal. This ideal is the ideal of $Z_\omega \cap V_\tau$. Of course, given the explicit form of $f_{\partial_{p,\pm}}$, $Z_\omega \cap V_\tau$ consists of just one point. We then take Z_ω to be the union of these sets $Z_\omega \cap V_\tau$ over all one-dimensional τ containing ω .

We can now state the main result of this subsection.

THEOREM 6.32. *U_ω^k is a flat deformation of U_ω^0 over O_k . As a consequence, the underlying topological spaces of U_ω^k and U_ω^0 are the same. In particular, suppose $y \in U_\omega^k$ is, as a point of U_ω^0 , contained in the stratum V_τ for some τ containing ω but not in any $V_{\tau'} \subsetneq V_\tau$. Then if $y \notin Z_\omega$, U_ω^k is étale locally isomorphic in a neighbourhood of y to $\text{Spec } \mathbb{k}[P_{\varphi,x}]/(t^{k+1})$, where $x \in \text{Int}(\tau) \cap B_0$.*

PROOF. We have $0 \leq \dim \omega \leq 2$, so we have three cases for $\dim \omega$, and we will deal with each case separately, going from the easiest to the hardest case.

$\dim \omega = 2$: Then $\omega \in \mathcal{P}_{\max}$, $\sigma_{u_\omega} = \omega$, so $P_{\varphi,\omega,\sigma_{u_\omega}}$ is isomorphic to $\Lambda_x \times \mathbb{N}$ for $x \in \text{Int}(\omega)$. Thus

$$R_{\omega,\omega,u_\omega}^k \cong \mathbb{k}[\Lambda_x] \otimes_{\mathbb{k}} R_k,$$

and so

$$U_\omega^k \cong U_\omega^0 \times_{\mathbb{k}} O_k.$$

This is clearly a flat deformation of U_ω^0 , isomorphic to $\text{Spec } \mathbb{k}[P_{\varphi,x}]/(t^{k+1})$.

$\dim \omega = 1$. In this case, we have $\omega = \sigma_+ \cap \sigma_-$ for two maximal cells σ_\pm , and the inverse limit defining R_ω^k is in fact the fibre product

$$R_\omega^k = R_{\omega, \sigma_-, \mathbf{u}_{\sigma_-}}^k \times_{R_{\omega, \omega, \mathbf{u}_\omega}^k} R_{\omega, \sigma_+, \mathbf{u}_{\sigma_+}}^k.$$

We shall compute this explicitly in the case that \mathbf{u}_{σ_+} and \mathbf{u}_{σ_-} are chosen to be adjacent, i.e., with $\dim \mathbf{u}_{\sigma_+} \cap \mathbf{u}_{\sigma_-} = 1$, with \mathbf{u}_{σ_+} and \mathbf{u}_{σ_-} of course lying on opposite sides of ω . We can also take $\mathbf{u}_\omega = \mathbf{u}_{\sigma_+}$. In this case, we can take a representative for φ which has slope 0 on σ_- and has slope $e\check{d}_\omega$ on σ_+ , where e is a positive integer. As usual $\check{d}_\omega \in \check{\Lambda}_y$ for some $y \in \mathbf{u}_{\sigma_+} \cap \mathbf{u}_{\sigma_-}$ used to define the change of chamber map. This element is primitive, zero on the tangent space of ω , and is positive on tangent vectors pointing into σ_+ . We can identify $P_{\varphi, \omega, \sigma_+}$ with $P_{\varphi, \omega, \sigma_-}$ via parallel transport through y . One sees easily that

$$P_{\varphi, \omega, \sigma_+} \cong \Lambda_\omega \oplus S_e,$$

where $\Lambda_\omega \subseteq \Lambda_y$ is the space of integral tangent vectors to ω at y and S_e is the monoid defined in Example 3.28. We think of $S_e \subseteq \mathbb{Z}^2$ as the monoid generated by $(-1, 0)$, $(0, 1)$ and $(1, e)$, yielding elements of $\mathbb{k}[S_e]$ given by

$$u = z^{(-1, 0)}, \quad v = z^{(1, e)}, \quad t = z^{(0, 1)}.$$

with

$$\begin{aligned} \text{ord}_{\sigma_-} u &= 0, & \text{ord}_{\sigma_-} v &= e, & \text{ord}_{\sigma_-} t &= 1 \\ \text{ord}_{\sigma_+} u &= e, & \text{ord}_{\sigma_+} v &= 0, & \text{ord}_{\sigma_+} t &= 1 \end{aligned}$$

Writing

$$R_\pm := R_{\omega, \sigma_\pm, \mathbf{u}_{\sigma_\pm}}^k, \quad R_\cap := R_{\omega, \omega, \mathbf{u}_\omega}^k,$$

we have

$$\begin{aligned} R_- &= \mathbb{k}[\Lambda_\omega][u, v, t] / \langle uv - t^e, v^\beta t^\gamma \mid \beta e + \gamma \geq k + 1 \rangle, \\ R_+ &= \mathbb{k}[\Lambda_\omega][u, v, t] / \langle uv - t^e, u^\alpha t^\gamma \mid \alpha e + \gamma \geq k + 1 \rangle, \\ R_\cap &= (\mathbb{k}[\Lambda_\omega][u, v, t] / \langle uv - t^e, u^\alpha v^\beta t^\gamma \mid \max\{\alpha, \beta\}e + \gamma \geq k + 1 \rangle)_{f_{\mathfrak{d}, \pm}}. \end{aligned}$$

Because we are taking $\mathbf{u}_\omega = \mathbf{u}_{\sigma_+}$, the map $R_+ \rightarrow R_\cap$ is the canonical surjection followed by the localization. On the other hand, the map $R_- \rightarrow R_\cap$ is the canonical surjection $R_{\omega, \sigma_-, \sigma_-}^k \rightarrow R_{\omega, \omega, \sigma_-}^k$ followed by localization and the change of chamber map, which is given in these coordinates as

$$u \mapsto uf_\omega, \quad v \mapsto vf_\omega^{-1}, \quad t \mapsto t,$$

where

$$f_\omega = \prod_{(\mathfrak{d}, x)} f_{(\mathfrak{d}, x)},$$

the product being over all (\mathfrak{d}, x) with $\mathfrak{d}(x) = y$. We can think of f_ω as living in the ring $\mathbb{k}[\Lambda_\omega][t]$.

LEMMA 6.33. *Let*

$$R_\cup := \mathbb{k}[\Lambda_\omega][U, V, t] / (UV - f_\omega t^e, t^{k+1}).$$

Then

$$\begin{aligned} R_{\cup} &\rightarrow R_{-} \times_{R_{\cap}} R_{+} \\ U &\mapsto (u, f_{\omega}u) \\ V &\mapsto (f_{\omega}v, v) \\ t &\mapsto (t, t) \end{aligned}$$

is an isomorphism of $\mathbb{k}[\Lambda_{\omega}][t]$ -algebras.

PROOF. Note that the rings R_{-} and R_{+} are generated as $\mathbb{k}[\Lambda_{\omega}][t]/(t^{k+1})$ -modules by $1, u^i, v^j$ for $i, j > 0$. Also, the submodules of R_{-} (R_{+}) generated by $u^i, i \geq 0$ ($v^j, j \geq 0$) are free direct summands. So for $g_{\pm} \in R_{\pm}$, we can write uniquely

$$\begin{aligned} g_{-} &= \sum_{i \geq 0} a_i u^i + h_{-}(v, t) \\ g_{+} &= \sum_{j \geq 0} b_j v^j + h_{+}(u, t) \end{aligned}$$

with $a_i, b_j \in \mathbb{k}[\Lambda_{\omega}][t]/(t^{k+1})$ and $h_{\pm}(0, t) = 0$. So $(g_{-}, g_{+}) \in R_{-} \times_{R_{\cap}} R_{+}$ if and only if

$$a_0 = b_0, \quad h_{-}(v, t) = \sum_{j > 0} b_j f_{\omega}^j v^j, \quad h_{+}(u, t) = \sum_{i > 0} a_i f_{\omega}^i u^i$$

as elements of R_{\cap} . But if this is the case then (g_{-}, g_{+}) is the image of $\sum_{i \geq 0} a_i U^i + \sum_{j > 0} b_j V^j \in R_{\cup}$. This shows surjectivity. Injectivity follows easily after noting that R_{\cup} is a free $\mathbb{k}[\Lambda_{\omega}][t]/(t^{k+1})$ -module with basis $U^i, V^j, i \geq 0, j > 0$. \square

As a consequence of this lemma, we see that $R_{\omega}^k = R_{\cup}$, and it then easily follows that (6.3) is exact. This shows U_{ω}^k is a flat deformation of U_{ω}^0 . Furthermore, one checks easily that away from Z_{ω} , which is the locus $x = y = f_{\omega} = t = 0$ in $\text{Spec } R_{\cup}$, U_{ω}^k takes the desired local form.

$\dim \omega = 0$. This is the most difficult case, and we will deal with this in two steps. Fix \mathfrak{u}_{ω} , hence the monoid

$$P := P_{\varphi, \omega, \sigma_{\mathfrak{u}_{\omega}}}.$$

Then

$$U_{\omega}^0 = \text{Spec } \mathbb{k}[P]/(t)$$

by Lemma 6.30.

It is helpful to work locally. To this end, we first define a rather strange object.

First, note that the reduced scheme associated to $\text{Spec } R_{\omega, \tau, \mathfrak{u}}^k$ is $\text{Spec } R_{\omega, \tau, \mathfrak{u}}^0$, which is an open subscheme of $\text{Spec } R_{\omega, \tau, \sigma_{\mathfrak{u}}}^0$. This, in turn, is a closed stratum of U_{ω}^0 . For a scheme X , denote by $|X|$ the underlying topological space. We thus have inclusions

$$i_{\omega, \tau} : |\text{Spec } R_{\omega, \tau, \mathfrak{u}}^k| \rightarrow |U_{\omega}^0|.$$

Let

$$\mathcal{O}_{\omega, \tau}^k := i_{\omega, \tau*} \mathcal{O}_{\text{Spec } R_{\omega, \tau, \mathfrak{u}}^k}.$$

This is a sheaf of rings on $|U_{\omega}^0|$. Note that the maps $F_k(\mathbf{e}_{\tau_1, \tau_2}) : R_{\omega, \tau_2, \mathfrak{u}_{\tau_2}}^k \rightarrow R_{\omega, \tau_1, \mathfrak{u}_{\tau_1}}^k$ induce maps $\mathcal{O}_{\omega, \tau_2}^k \rightarrow \mathcal{O}_{\omega, \tau_1}^k$ defining an inverse system of sheaves $(\mathcal{O}_{\omega, \tau}^k)_{\tau}$. Set

$$\mathcal{O}_{\omega}^k := \varprojlim_{\tau} \mathcal{O}_{\omega, \tau}^k,$$

and denote by V_ω^k the ringed space

$$V_\omega^k := (|U_\omega^0|, \mathcal{O}_\omega^k).$$

Completely analogously to the complex (6.3), one has the complex of sheaves on $|U_\omega^0|$ for $\ell \leq k$,

$$(6.4) \quad 0 \longrightarrow \mathcal{O}_\omega^0 \xrightarrow{\cdot t^\ell} \mathcal{O}_\omega^\ell \longrightarrow \mathcal{O}_\omega^{\ell-1} \longrightarrow 0.$$

Since the inverse limit of sheaves $\varprojlim \mathcal{F}_i$ is defined by $U \mapsto \varprojlim \mathcal{F}_i(U)$, in fact taking global sections of (6.4) yields the complex (6.3): here, we use $\mathcal{O}_\omega^0 \cong \mathcal{O}_{U_\omega^0}$ as follows from the argument of Lemma 6.30. Since $H^1(|U_\omega^0|, \mathcal{O}_{U_\omega^0}) = 0$ as U_ω^0 is an affine scheme, we get (6.3) exact if (6.4) is exact.

Now look at the open set $|U_\tau^0| \subseteq |U_\omega^0|$ where $\tau \in \mathcal{P}$ contains ω and $\dim \tau = 1$. It is easily seen using properties of the inverse limit that the ringed space $(|U_\tau^0|, \mathcal{O}_\omega^k|_{|U_\tau^0|})$ is isomorphic to the affine scheme U_τ^k . We have already shown Theorem 6.32 for U_τ^k , so this shows that (6.4) is exact on $\bigcup_{\substack{\tau \supseteq \omega \\ \dim \tau = 1}} |U_\tau^0|$.

This set is almost all of $|U_\omega^0|$: it is just missing one point, $z \in |U_\omega^0|$, corresponding to the maximal ideal \mathfrak{m}_0 in $\mathbb{k}[P]/(t)$ given by the monoid ideal $P \setminus \{0\} \subseteq P$. Let \mathfrak{m}_k denote the maximal ideal in $\mathbb{k}[P]/(t^{k+1})$ given by the same monoid ideal $P \setminus \{0\}$.

LEMMA 6.34. *There is an isomorphism*

$$\psi : (\mathbb{k}[P]/(t^{k+1}))_{\mathfrak{m}_k} \rightarrow \mathcal{O}_{\omega,z}^k.$$

PROOF. Taking a finite inverse limit and taking stalks commutes, so

$$\mathcal{O}_{\omega,z}^k = \varprojlim (\mathcal{O}_{\omega,\tau}^k)_z.$$

Since $\mathcal{O}_{\omega,\tau}^k$ is supported on a closed subset in a neighbourhood of z , one sees that

$$(\mathcal{O}_{\omega,\tau}^k)_z = (R_{\omega,\tau,u_\tau}^k)_{\mathfrak{m}_\tau} = (R_{\omega,\tau,\sigma_{u_\tau}}^k)_{\mathfrak{m}_\tau},$$

where \mathfrak{m}_τ is the maximal ideal again generated by the monoid ideal $P \setminus \{0\}$. Here, the last equality holds since R_{ω,τ,u_τ}^k is a localization of $R_{\omega,\tau,\sigma_{u_\tau}}^k$ at an element not in \mathfrak{m}_τ . To prevent the notation from becoming too dense, we introduce shorthand

$$R_\tau := (R_{\omega,\tau,\sigma_{u_\tau}}^k)_{\mathfrak{m}_\tau}.$$

We now need to show there is an isomorphism

$$\varprojlim R_\tau \xrightarrow{\cong} (\mathbb{k}[P]/(t^{k+1}))_{\mathfrak{m}_k}.$$

The inverse system on the left is given by maps

$$\varphi_{\tau_1,\tau_2} : R_{\tau_2} \rightarrow R_{\tau_1}$$

which are the localizations of $F_k(\mathbf{e}_{\tau_1,\tau_2})$. This is the composition of a change of stratum map and a number of change of chamber maps. We can identify the monoids $P_{\varphi,\omega,\sigma}$ for various σ with $P_{\varphi,v}$, where $\omega = \{v\}$, by parallel transport in σ to v . This then yields canonical change of stratum maps

$$\psi_{\tau_1,\tau_2} : R_{\tau_2} \rightarrow R_{\tau_1},$$

so that φ_{τ_1,τ_2} is a composition of ψ_{τ_1,τ_2} with a sequence of change of chamber maps

$$\theta_{u_{n-1},u_n} \circ \cdots \circ \theta_{u_1,u_2} : R_{\tau_1} \rightarrow R_{\tau_1}$$

defined using points y near ω (so that the singularities of B play no role here). Note that the collection of rings $\{R_\tau\}$ with the maps φ_{τ_1,τ_2} is an inverse system as usual,

and the same collection of rings with the maps ψ_{τ_1, τ_2} is also an inverse system. The latter inverse system has inverse limit given by Lemma 6.9, namely $(\mathbb{k}[P]/(t^{k+1}))_{\mathfrak{m}_k}$. Thus we just need to give an isomorphism between these two inverse systems, i.e., automorphisms $\psi_\tau : R_\tau \rightarrow R_\tau$ such that whenever $\omega \subseteq \tau_1 \subseteq \tau_2$,

$$(6.5) \quad \begin{array}{ccc} R_{\tau_2} & \xrightarrow{\psi_{\tau_2}} & R_{\tau_2} \\ \varphi_{\tau_1, \tau_2} \downarrow & & \downarrow \psi_{\tau_1, \tau_2} \\ R_{\tau_1} & \xrightarrow{\psi_{\tau_1}} & R_{\tau_1} \end{array}$$

is commutative.

To do so, first note that if u, u' are adjacent chambers, we can write

$$\theta_{u, u'}(z^m) = z^m f_{u, u'}^{(n, \bar{m})}$$

for some n , and $f_{u, u'}$ can be viewed as the image of an element of $\mathbb{k}[P]$ completely specified by the structure \mathcal{S} . Furthermore, $f_{u, u'} \notin \mathfrak{m}_\tau$, so $f_{u, u'}$ can be viewed as an element of R_τ^\times for each τ . Hence we can define

$$\theta_{u, u'}^\tau : R_\tau \rightarrow R_\tau$$

by

$$\theta_{u, u'}^\tau(z^m) = z^m f_{u, u'}^{(n, \bar{m})}.$$

Note that $\psi_{\tau_1, \tau_2} \circ \theta_{u, u'}^{\tau_2} = \theta_{u, u'}^{\tau_1} \circ \psi_{\tau_1, \tau_2}$.

Now for any $\tau \supseteq \omega$, let $u_\tau = u_1, u_2, \dots, u_n = u_\omega$ be the sequence of chambers obtained by passing from u_τ to u_ω . If $\omega \not\subseteq \partial B$, we insist that we do this in a clockwise manner, so that this chain is well-defined. Then define

$$\psi_\tau = \theta_{u_{n-1}, u_n}^\tau \circ \dots \circ \theta_{u_1, u_2}^\tau.$$

One then checks easily that with this choice, (6.5) is commutative □

This lemma now shows that (6.4) is exact at z , and hence, as described above, (6.3) is exact. Hence U_ω^k is a flat deformation of U_ω^0 .

We only need to show it is étale isomorphic to $\mathbb{k}[P]/(t^{k+1})$ at the point z . This follows from the above lemma. The isomorphism given there induces a map

$$\psi : \mathbb{k}[P]/(t^{k+1}) \rightarrow (R_\omega^k)_{\mathfrak{m}_z},$$

where \mathfrak{m}_z is the maximal ideal of the point z in R_ω^k , and since P is finitely generated, one can find an $f \in R_\omega^k \setminus \mathfrak{m}_z$ such that $\psi(z^m) \in (R_\omega^k)_f$ for each $m \in P$. But then we get a map

$$\psi : \mathbb{k}[P]/(t^{k+1}) \rightarrow (R_\omega^k)_f$$

inducing an isomorphism after localizing at \mathfrak{m}_k and \mathfrak{m}_z , hence a map of schemes $V \rightarrow \text{Spec } \mathbb{k}[P]/(t^{k+1})$ which is étale at z . By openness of the étale condition (see, e.g., [83], Proposition 3.8), this gives the desired result. □

To construct X_k the deformation of $X_0 = \check{X}_0(B, \mathcal{P})$, we now glue along open subsets. Indeed, if $\omega_1 \subseteq \omega_2$, then we obtain $U_{\omega_2}^k \subseteq U_{\omega_1}^k$ canonically. This canonical inclusion is obtained as follows. In defining $U_{\omega_2}^k$, we need to choose reference chambers u_τ^2 for each $\tau \supseteq \omega_2$. Similarly, we choose reference chambers u_τ^1 for each $\tau \supseteq \omega_1$. Thus if $\tau \supseteq \omega_2$, we obtain a well-defined map

$$R_{\omega_2, \tau, u_\tau^2}^k \rightarrow R_{\omega_1, \tau, u_\tau^1}^k$$

which is a composition of change of chamber homomorphisms and a change of strata homomorphism. This induces a map on inverse limits, and hence a map $U_{\omega_2}^k \rightarrow U_{\omega_1}^k$ which is seen to be an inclusion of open sets. We leave the details of this to the reader.

Having constructed these canonical inclusions, we can glue together the schemes U_{ω}^k along these common open sets, constructing X_k . Again we leave these straightforward details to the reader. We have thus proved

THEOREM 6.35. *If \mathcal{S} is a structure which is compatible to order k , then there is a scheme $\check{X}_k(B, \mathcal{P})$ flat over O_k which is a deformation of $\check{X}_0(B, \mathcal{P})$.*

EXAMPLE 6.36. Continuing with Example 6.23, we wish to describe the scheme U_{ω}^k for ω the vertex at the origin. The reader may verify the following description of R_{ω}^k , using variables x_1, \dots, x_5 which should be viewed as lifts under the map $R_{\omega}^k \rightarrow R_{\omega}^0$ of $z^{(2,1,0)}$, $z^{(3,1,1)}$, $z^{(2,0,1)}$, $z^{(0,-1,0)}$ and $z^{(0,0,-1)}$ respectively. Here the last two coordinates live in the stalk of Λ at ω , and the first coordinate gives the value of φ_{ω} . Then

$$R_{\omega}^k \cong \mathbb{k}[x_1, \dots, x_5, t]/(t^{k+1}, f_1, \dots, f_5),$$

where

$$\begin{aligned} f_1 &= x_5x_2 - t(x_1 + t^2) \\ f_2 &= x_1x_3 - t(x_2 + t^3) \\ f_3 &= x_2x_4 - t(x_3 + t^2) \\ f_4 &= x_3x_5 - t^2(x_4 + 1) \\ f_5 &= x_4x_1 - t^2(x_5 + 1). \end{aligned}$$

To compactify this, rather than finding charts in neighbourhoods corresponding to the five other vertices of \mathcal{P} , one can just projectize the above chart, adding a variable and obtaining, in $\mathbb{A}^1 \times \mathbb{P}^5$, a degenerating family of del Pezzo surfaces given by the equations

$$\begin{aligned} x_5x_2 - t(x_1x_0 + t^2x_0^2) &= 0 \\ x_1x_3 - t(x_2x_0 + t^3x_0^2) &= 0 \\ x_2x_4 - t(x_3x_0 + t^2x_0^2) &= 0 \\ x_3x_5 - t^2(x_4x_0 + x_0^2) &= 0 \\ x_4x_1 - t^2(x_5x_0 + x_0^2) &= 0. \end{aligned}$$

□

6.3. Achieving compatibility: The tropical vertex group

All that remains to be done is to construct a compatible structure \mathcal{S} to order k , and then let $k \rightarrow \infty$. To do this, we need to understand what happens when rays collide.

6.3.1. The argument of Kontsevich and Soibelman. We will first describe a fairly general setup and general result. This will provide the basic idea of how to produce compatible structures. The definition here is only a slight generalization of the definitions seen in §§5.4.2 and 5.4.3, so should look familiar.

Fix a lattice $M \cong \mathbb{Z}^2$, $N = \text{Hom}_{\mathbb{Z}}(M, \mathbb{Z})$ as usual, and suppose we are given a toric monoid P equipped with a map

$$r : P \rightarrow M.$$

For this chapter, the typical example is $P = P_{\varphi, x}$ for $x \in B_0$, and $r : P_{\varphi, x} \rightarrow M$ is the projection $P_{\varphi, x} \rightarrow \Lambda_x$ given by $m \mapsto \bar{m}$. We then define the *module of (relative) log derivations* of $\mathbb{k}[P]$ to be the module

$$\Theta(\mathbb{k}[P]) := \text{Hom}_{\mathbb{Z}}(M, \mathbb{k}[P]) = \mathbb{k}[P] \otimes_{\mathbb{Z}} N.$$

An element $f \otimes n$ is written as $f\partial_n$, and acts as a derivation on $\mathbb{k}[P]$ by

$$f\partial_n(z^m) = f\langle n, r(m) \rangle z^m.$$

Now denote by \mathfrak{m} both the ideal $P \setminus P^\times$ (where P^\times denotes the group of units in P) and the monomial ideal generated by this monoid ideal in $\mathbb{k}[P]$. For example, if $P = M \oplus \mathbb{N}$, with r the projection onto M , then $P^\times = M \oplus 0$. Let I be a monomial ideal in $\mathbb{k}[P]$ whose radical is \mathfrak{m} ; we shall call such an ideal an \mathfrak{m} -primary ideal. Then given $\xi \in \mathfrak{m}\Theta(\mathbb{k}[P])$, we obtain an element

$$\exp(\xi) \in \text{Aut}(\mathbb{k}[P]/I)$$

via exponentiation of the derivation ξ , i.e.,

$$\exp(\xi)(a) = a + \sum_{i=1}^{\infty} \frac{\xi^i(a)}{i!}.$$

This is a finite sum, since there is some integer n such that $\mathfrak{m}^n \subseteq I$.

One can check, for example, that if $f \in \mathfrak{m}$ can be written as $f = \sum_{m \in S} c_m z^m$ with $\{r(m) \mid m \in S\}$ lying in a rank one sublattice of M , and $n \in N$ annihilates this sublattice, then

$$\exp(f\partial_n)(z^m) = \exp(f)\langle n, r(m) \rangle z^m.$$

In general, however, it is difficult to compute the exponential of a vector field.

There is a natural Lie bracket defined on the module of log derivations:

$$\begin{aligned} [z^m \partial_n, z^{m'} \partial_{n'}] &:= z^{m+m'} (\langle n, r(m') \rangle \partial_{n'} - \langle n', r(m) \rangle \partial_n) \\ &= z^{m+m'} \partial_{\langle n, r(m') \rangle n' - \langle n', r(m) \rangle n}. \end{aligned}$$

Now let

$$\mathfrak{v} := \bigoplus_{\substack{m \in \mathfrak{m} \\ r(m) \neq 0}} z^m \mathbb{k} \otimes r(m)^\perp \subseteq \Theta(\mathbb{k}[P]).$$

This is the \mathbb{k} -vector space of linear combinations of vector fields $z^m \partial_n$ with $m \in \mathfrak{m}$, $r(m) \neq 0$, and $\langle n, r(m) \rangle = 0$. Via exactly the same argument as in §5.4.2, \mathfrak{v} is closed under Lie bracket. For each \mathfrak{m} -primary ideal I , we then get a subgroup of $\text{Aut}(\mathbb{k}[P]/I)$ defined by

$$\mathbb{V}_I := \{\exp(\xi) \mid \xi \in \mathfrak{v}\}.$$

Just as in §5.4.2, we can express elements of this group as Hamiltonian symplectomorphisms. However, this will not be important for us in this chapter.

We will sometimes want to work with power series in a complete situation. Let

$$\widehat{\mathbb{k}[P]} := \varprojlim \mathbb{k}[P]/\mathfrak{m}^k$$

be the completion of $\mathbb{k}[P]$ at the ideal \mathfrak{m} . We can then also define

$$\widehat{\mathbb{V}} := \varprojlim \mathbb{V}_{\mathfrak{m}^k},$$

a pro-nilpotent subgroup of $\text{Aut}(\widehat{\mathbb{k}[P]})$.

Just as in §5.4.3, we can define the notion of a scattering diagram, with essentially the identical definition.

DEFINITION 6.37. Let $r : P \rightarrow M$ be given, and I be an \mathfrak{m} -primary ideal.

(1) A *ray* or *line* is a pair $(\mathfrak{d}, f_{\mathfrak{d}})$ such that

- $\mathfrak{d} \subseteq M_{\mathbb{R}}$ is given by

$$\mathfrak{d} = m'_0 - \mathbb{R}_{\geq 0} m_0$$

if \mathfrak{d} is a ray and

$$\mathfrak{d} = m'_0 - \mathbb{R} m_0$$

if \mathfrak{d} is a line, for some $m'_0 \in M_{\mathbb{R}}$ and $m_0 \in M \setminus \{0\}$. The set \mathfrak{d} is called the *support* of the line or ray. If \mathfrak{d} is a ray, m'_0 is called the *initial point* of the ray, written as $\text{Init}(\mathfrak{d})$.

- For $m_0 \in M \setminus \{0\}$, let

$$P_{m_0} := \{m \in P \mid r(m) = C m_0 \text{ for some rational } C > 0\}.$$

Then $f_{\mathfrak{d}} \in \widehat{\mathbb{k}[P]}$ satisfies

$$f_{\mathfrak{d}} = 1 + \sum_{m \in P_{m_0}} c_m z^m.$$

- $f_{\mathfrak{d}} \equiv 1 \pmod{\mathfrak{m}}$.

(2) A *scattering diagram* \mathfrak{D} over $\mathbb{k}[P]/I$ is a finite collection of lines and rays such that for each $(\mathfrak{d}, f_{\mathfrak{d}}) \in \mathfrak{D}$, $f_{\mathfrak{d}} \in \mathbb{k}[P]$. A *scattering diagram* \mathfrak{D} over $\widehat{\mathbb{k}[P]}$ is a countable collection of lines and rays such that for each $k \geq 1$, there are only a finite number of lines or rays $(\mathfrak{d}, f_{\mathfrak{d}}) \in \mathfrak{D}$ with $f_{\mathfrak{d}} \not\equiv 1 \pmod{\mathfrak{m}^k}$.

We then have the same notions of the sets $\text{Supp}(\mathfrak{D})$, $\text{Sing}(\mathfrak{D})$ as defined in §5.4.3, and given a smooth immersion $\gamma : [0, 1] \rightarrow M_{\mathbb{R}} \setminus \text{Sing}(\mathfrak{D})$, if one fixes an \mathfrak{m} -primary ideal I , one obtains exactly as in §5.4.3 an automorphism

$$\theta_{\gamma, \mathfrak{D}} \in \mathbb{V}_I.$$

Indeed, one needs to check that each automorphism $\theta_{\gamma, \mathfrak{d}_i}$ arising when crossing a ray \mathfrak{d}_i makes sense; this is defined by $\theta_{\gamma, \mathfrak{d}_i}(z^m) = z^m f_{\mathfrak{d}_i}^{\langle n, r(m) \rangle}$ as usual for some $n \in N$. This makes sense, as $f_{\mathfrak{d}_i} \equiv 1 \pmod{\mathfrak{m}}$, so that $f_{\mathfrak{d}_i}$ is invertible in $\mathbb{k}[P]/I$. Thus we obtain, as before, the automorphism $\theta_{\gamma, \mathfrak{D}}$ which only depends on the homotopy type of the path γ in $M_{\mathbb{R}} \setminus \text{Sing}(\mathfrak{D})$.

Furthermore, if \mathfrak{D} is a scattering diagram over $\widehat{\mathbb{k}[P]}$, then $\theta_{\gamma, \mathfrak{D}}$ similarly makes sense: computing $\theta_{\gamma, \mathfrak{D}}$ modulo \mathfrak{m}^k makes sense for any $k > 0$, since this involves only a finite number of compositions. One then takes the limit to get

$$\theta_{\gamma, \mathfrak{D}} \in \widehat{\mathbb{V}}.$$

In §5.4.3, we constructed an explicit scattering diagram \mathfrak{D} derived from the set of all Maslov index zero disks. We found that the automorphisms $\theta_{\gamma, \mathfrak{D}}$ for loops around certain singular points of \mathfrak{D} were the identity. In this section, we

will proceed in quite the opposite fashion, first showing how, using an argument of Kontsevich and Soibelman [70], we can always achieve this condition by adding rays.

THEOREM 6.38. *Fix an \mathfrak{m} -primary ideal I . Let \mathfrak{D} be a scattering diagram over $\mathbb{k}[P]/I$. Then there exists a scattering diagram $\mathfrak{S}_I(\mathfrak{D})$ containing \mathfrak{D} such that $\mathfrak{S}_I(\mathfrak{D}) \setminus \mathfrak{D}$ consists only of rays, and such that $\theta_{\gamma, \mathfrak{S}_I(\mathfrak{D})} = \text{Id} \in \mathbb{V}_I$ for any closed loop γ for which $\theta_{\gamma, \mathfrak{S}_I(\mathfrak{D})}$ is defined.*

Similarly, if \mathfrak{D} is a scattering diagram over $\widehat{\mathbb{k}[P]}$, then there is a scattering diagram $\mathfrak{S}(\mathfrak{D})$ containing \mathfrak{D} such that $\mathfrak{S}(\mathfrak{D}) \setminus \mathfrak{D}$ consists only of rays, and such that $\theta_{\gamma, \mathfrak{S}(\mathfrak{D})} = \text{Id} \in \widehat{\mathbb{V}}$ for any closed loop γ for which $\theta_{\gamma, \mathfrak{S}(\mathfrak{D})}$ is defined.

PROOF. We proceed inductively on k , showing that there exists a \mathfrak{D}_k such that

$$\theta_{\gamma, \mathfrak{D}_k} \equiv \text{Id} \pmod{\mathfrak{m}^{k+1}}$$

for all closed loops γ for which $\theta_{\gamma, \mathfrak{D}_k}$ is defined. In the case where we are working over $\mathbb{k}[P]/I$, since I is \mathfrak{m} -primary, there is some k such that $\mathfrak{m}^{k+1} \subseteq I$, and then we are done, taking $\mathfrak{S}_I(\mathfrak{D}) = \mathfrak{D}_k$. Otherwise we continue indefinitely, taking $\mathfrak{S}(\mathfrak{D})$ to be the (non-disjoint) union of the \mathfrak{D}_k .

We take $\mathfrak{D}_0 = \mathfrak{D}$. To obtain \mathfrak{D}_k from \mathfrak{D}_{k-1} , we proceed as follows. Let \mathfrak{D}'_{k-1} consist of those rays and lines \mathfrak{d} in \mathfrak{D}_{k-1} with $f_{\mathfrak{d}} \not\equiv 1 \pmod{\mathfrak{m}^{k+1}}$. Note that $\text{Sing}(\mathfrak{D}'_{k-1})$ is a finite set. Let $p \in \text{Sing}(\mathfrak{D}'_{k-1})$. Let γ_p be a closed simple loop around p , small enough so it contains no other points of $\text{Sing}(\mathfrak{D}'_{k-1})$. Certainly,

$$\theta_{\gamma_p, \mathfrak{D}_{k-1}} \equiv \theta_{\gamma_p, \mathfrak{D}'_{k-1}} \pmod{\mathfrak{m}^{k+1}}.$$

By the inductive assumption we can write uniquely

$$\theta_{\gamma_p, \mathfrak{D}'_{k-1}} \equiv \exp \left(\sum_{i=1}^s c_i z^{m_i} \partial_{n_i} \right) \pmod{\mathfrak{m}^{k+1}}$$

with $m_i \in \mathfrak{m}^k$, $r(m_i) \neq 0$, $n_i \in r(m_i)^\perp$ primitive and $c_i \in \mathbb{k}$. Let

$$\mathfrak{D}[p] = \{(p - \mathbb{R}_{\geq 0} r(m_i), 1 \pm c_i z^{m_i}) \mid i = 1, \dots, s\}.$$

The sign is chosen in each ray so that its contribution to $\theta_{\gamma_p, \mathfrak{D}[p]}$ is $\exp(-c_i z^{m_i} \partial_{n_i})$ modulo \mathfrak{m}^{k+1} . Since $[c_i z^{m_i} \partial_{n_i}, \xi] \equiv 0 \pmod{\mathfrak{m}^{k+1}}$ for any log derivation $\xi \in \mathfrak{v}$, an automorphism associated to any ray in $\mathfrak{D}[p]$ commutes modulo \mathfrak{m}^{k+1} with any automorphism associated to a ray in \mathfrak{D}_{k-1} . Thus we have

$$\theta_{\gamma_p, \mathfrak{D}_{k-1} \cup \mathfrak{D}[p]} = \text{Id} \pmod{\mathfrak{m}^{k+1}}.$$

Furthermore

$$\mathfrak{D}_k = \mathfrak{D}_{k-1} \cup \bigcup_P \mathfrak{D}[p]$$

now has the desired properties. Indeed, let $p \in \text{Sing}(\mathfrak{D}_k)$. The only rays which contribute to $\theta_{\gamma_p, \mathfrak{D}_k} \pmod{\mathfrak{m}^{k+1}}$ which were not in $\mathfrak{D}_{k-1} \cup \mathfrak{D}[p]$ are in $\bigcup_{p' \neq p} \mathfrak{D}[p']$, and these rays contribute twice, but with inverse automorphisms which cancel. \square

REMARK 6.39. The process described in the proof of the above lemma is very simple, and can be implemented easily on a computer. It is also easy to see that the result is essentially unique. The only non-uniqueness arises because there may be a number of distinct rays or lines with the same support. Given rays or lines $(\mathfrak{d}_1, f_{\mathfrak{d}_1}), \dots, (\mathfrak{d}_n, f_{\mathfrak{d}_n})$ with $\text{Supp}(\mathfrak{d}_1) = \dots = \text{Supp}(\mathfrak{d}_n)$, we can replace these rays or

lines with $(\mathfrak{d}, \prod_{i=1}^n f_{\mathfrak{d}_i})$, with support $\mathfrak{d} = \text{Supp}(\mathfrak{d}_i)$. If one applies this procedure to every set of such rays or lines in a scattering diagram \mathfrak{D} , one obtains a new diagram \mathfrak{D}' such that no two elements have the same support, yet nevertheless the automorphisms $\theta_{\gamma, \mathfrak{D}}$ coincide with $\theta_{\gamma, \mathfrak{D}'}$ for any path γ for which $\theta_{\gamma, \mathfrak{D}}$ is defined. This then leads to the following definition:

DEFINITION 6.40. Two scattering diagrams \mathfrak{D} and \mathfrak{D}' are *equivalent over $\mathbb{k}[P]/I$* (or $\widehat{\mathbb{k}[P]}$) if $\theta_{\gamma, \mathfrak{D}} = \theta_{\gamma, \mathfrak{D}'}$ in \mathbb{V}_I (or $\widehat{\mathbb{V}}$) for every path γ for which both $\theta_{\gamma, \mathfrak{D}}$ and $\theta_{\gamma, \mathfrak{D}'}$ are defined.

The correct uniqueness result then says that $S_I(\mathfrak{D})$ (or $S(\mathfrak{D})$) is unique up to equivalence.

EXAMPLE 6.41. The structure of $S(\mathfrak{D})$ can be very complicated, even for very simple choices of \mathfrak{D} . We will discuss three examples here, and the reader can consult [44] for more details.

We will take in these examples $P = M \oplus \mathbb{N}$, writing $x = z^{(1,0,0)}$, $y = z^{(0,1,0)}$ and $t = z^{(0,0,1)}$, so that $\mathbb{k}[P] = \mathbb{k}[x^{\pm 1}, y^{\pm 1}, t]$ and $\mathfrak{m} = (t)$. Then $\widehat{\mathbb{k}[P]} = \mathbb{k}[x^{\pm 1}, y^{\pm 1}][[t]]$.

Consider the scattering diagram

$$\mathfrak{D} = \{(\mathbb{R}(1, 0), (1 + tx^{-1})^{\ell_1}), (\mathbb{R}(0, 1), (1 + ty^{-1})^{\ell_2})\}$$

for various $\ell_1, \ell_2 > 0$. If $\ell_1 = \ell_2 = 1$, then it can be easily checked by hand that

$$S(\mathfrak{D}) \setminus \mathfrak{D} = \{(\mathbb{R}(1, 1), 1 + t^2 x^{-1} y^{-1})\}.$$

Indeed, this is the same calculation carried out in Example 6.29. One only needs one extra ray. But if $\ell_1 = \ell_2 = 2$, one already needs an infinite number of rays:

$$\begin{aligned} S(\mathfrak{D}) \setminus \mathfrak{D} &= \{(\mathbb{R}(n+1, n), (1 + t^{2n+1} x^{-(n+1)} y^{-n})^2) | n \in \mathbb{Z}, n \geq 1\} \\ &\cup \{(\mathbb{R}(n, n+1), (1 + t^{2n+1} x^{-n} y^{-(n+1)})^2) | n \in \mathbb{Z}, n \geq 1\} \\ &\cup \{(\mathbb{R}(1, 1), (1 - t^2 x^{-1} y^{-1})^{-4})\}. \end{aligned}$$

For $\ell_1 = \ell_2 = 3$, the situation is much more complicated. One finds a certain symmetry, first noticed by Kontsevich:

$$(\mathbb{R}_{\geq 0}(m_1, m_2), f(t^{m_1+m_2} x^{-m_1} y^{-m_2})) \in S(\mathfrak{D})$$

if and only if

$$(\mathbb{R}_{\geq 0}(3m_1 - m_2, m_1), f(t^{4m_1 - m_2} x^{m_2 - 3m_1} y^{-m_1})) \in S(\mathfrak{D}),$$

provided that m_1, m_2 and $3m_1 - m_2$ are all positive. In addition, $S(\mathfrak{D})$ contains rays

$$(\mathbb{R}_{\geq 0}(3, 1), (1 + t^4 x^{-3} y^{-1})^3) \text{ and } (\mathbb{R}_{\geq 0}(1, 3), (1 + t^4 x^{-1} y^{-3})^3)$$

and hence by the above symmetry, there are also rays with support

$$\mathbb{R}_{\geq 0}(8, 3), \mathbb{R}_{\geq 0}(21, 8), \dots \text{ and } \mathbb{R}_{\geq 0}(3, 8), \mathbb{R}_{\geq 0}(8, 21), \dots$$

which converge to the rays of slope $(3 \pm \sqrt{5})/2$, corresponding to the two distinct eigenspaces of the linear transformation $\begin{pmatrix} 3 & -1 \\ 1 & 0 \end{pmatrix}$. Note that by the symmetry, we know what the functions attached to these rays are. On the other hand, inside the cone generated by the rays of slope $(3 \pm \sqrt{5})/2$, every rational slope occurs, and the functions attached to these rays appear to be very complicated. The only known

case is the function attached to the line of slope 1 (and the slopes obtained from this one by symmetry), which was proved by Reineke in [95] to be

$$\left(\sum_{n=0}^{\infty} \frac{1}{3n+1} \binom{4n}{n} t^{2n} x^{-n} y^{-n} \right)^9.$$

While these scattering diagrams are very complicated, note that modulo t^k for any fixed k , there are only a finite number of $f_{\mathfrak{d}}$ occurring with $f_{\mathfrak{d}} \neq 1$.

6.3.2. The enumerative interpretation. In Chapter 5, we saw a relationship between the condition that $\theta_{\gamma, \mathfrak{D}} = \text{Id}$ for a loop γ around a singular point and gluing of Maslov index zero disks to create new Maslov index zero disks. The scattering diagrams we consider here are far more complicated, but in fact there is still a similar interpretation, due to myself, Pandharipande and Siebert, given in [45]. While I will give the precise statement here, as we shall need one aspect of the result, I will not give any details of the proof. Thus this section is essentially the only part of the argument of this chapter which is not entirely self-contained.

We consider the following situation. Let $\ell_1, \dots, \ell_p > 0$ be p positive integers, and let $\ell = \sum_{i=1}^p \ell_i$. Set $P = M \oplus \mathbb{N}^{\ell}$. We denote the monomials in $\mathbb{k}[P]$ corresponding to the generators of \mathbb{N}^{ℓ} as t_{ij} , $1 \leq i \leq p$, $1 \leq j \leq \ell_i$. Let $m_1, \dots, m_p \in M$ be distinct primitive elements. Let \mathfrak{D} be the scattering diagram

$$(6.6) \quad \mathfrak{D} = \left\{ (\mathbb{R}m_i, \prod_{j=1}^{\ell_i} (1 + t_{ij}z^{-m_i})) \mid 1 \leq i \leq p \right\}.$$

The ideal \mathfrak{m} is generated by the variables t_{ij} . We can study the infinite scattering diagram $S(\mathfrak{D})$ over $\widehat{\mathbb{k}[P]}$, i.e., we can work over the ring $\mathbb{k}[M][[t_{ij}]]$. This will, in particular, after reducing modulo any \mathfrak{m} -primary ideal I , yield $S_I(\mathfrak{D})$. Indeed, given $S(\mathfrak{D})$, we can obtain a finite scattering diagram over $\mathbb{k}[P]/I$ by throwing out those rays $(\mathfrak{d}, f_{\mathfrak{d}}) \in S(\mathfrak{D})$ such that $f_{\mathfrak{d}} \equiv 1 \pmod{I}$, and for each $(\mathfrak{d}, f_{\mathfrak{d}}) \in S(\mathfrak{D})$ with $f_{\mathfrak{d}} \not\equiv 1 \pmod{I}$, we truncate $f_{\mathfrak{d}}$ by removing any monomials in I , to ensure that $f_{\mathfrak{d}}$ is a polynomial rather than a power series.

We would like to interpret elements of $S(\mathfrak{D}) \setminus \mathfrak{D}$. Let us assume that $S(\mathfrak{D})$ is chosen so that no two rays have the same support, and let $(\mathfrak{d}, f_{\mathfrak{d}}) \in S(\mathfrak{D}) \setminus \mathfrak{D}$. More specifically, we will interpret $f_{\mathfrak{d}}$. So assume that \mathfrak{d} is now fixed in the discussion.

Let $\Sigma_{\mathfrak{d}}$ be a complete fan in $M_{\mathbb{R}}$ defining a non-singular toric surface $X_{\mathfrak{d}}$, with the property that $\Sigma_{\mathfrak{d}}$ contains, amongst its set of one-dimensional rays,

$$\mathbb{R}_{\geq 0}m_1, \dots, \mathbb{R}_{\geq 0}m_p \text{ and } \mathfrak{d}.$$

(Note that \mathfrak{d} may coincide with one of the other rays). The precise choice of the fan $\Sigma_{\mathfrak{d}}$ will turn out to be irrelevant.

Let D_1, \dots, D_p be the toric divisors of $X_{\mathfrak{d}}$ corresponding to the rays

$$\mathbb{R}_{\geq 0}m_1, \dots, \mathbb{R}_{\geq 0}m_p,$$

and let D_{out} correspond to the ray \mathfrak{d} (possibly $D_i = D_{\text{out}}$ for some i). Choose general points $\{x_{i1}, \dots, x_{i\ell_i}\} \subseteq D_i$. Let

$$\nu : \tilde{X}_{\mathfrak{d}} \rightarrow X_{\mathfrak{d}}$$

be the blow-up of $X_{\mathfrak{d}}$ at the set of points $\{x_{ij}\}$. Let $\tilde{D}_1, \dots, \tilde{D}_p, \tilde{D}_{\text{out}}$ be the proper transforms of the corresponding toric divisors, and let E_{ij} be the exceptional divisor over the point x_{ij} .

Now introduce the additional data of $\mathbf{P} = (\mathbf{P}_1, \dots, \mathbf{P}_p)$, where \mathbf{P}_i denotes a sequence $(p_{i1}, \dots, p_{i\ell_i})$ of ℓ_i non-negative numbers. We will use the notation $\mathbf{P}_i = p_{i1} + \dots + p_{i\ell_i}$, and call \mathbf{P}_i an *ordered partition*. We call $p_{i1}, \dots, p_{i\ell_i}$ the *parts* of \mathbf{P}_i . Define

$$|\mathbf{P}_i| = \sum_{j=1}^{\ell_i} p_{ij},$$

and restrict to those \mathbf{P} such that

$$(6.7) \quad - \sum_{i=1}^p |\mathbf{P}_i| m_i = k_{\mathbf{P}} m_{\mathfrak{d}},$$

where $m_{\mathfrak{d}} \in M$ is a primitive generator of \mathfrak{d} and $k_{\mathbf{P}}$ is a positive integer.

Given this data, consider the class $\beta \in H_2(X_{\mathfrak{d}}, \mathbb{Z})$ specified by the requirement that, if D is a toric divisor of $X_{\mathfrak{d}}$ with $D \notin \{D_1, \dots, D_p, D_{\text{out}}\}$, then $D \cdot \beta = 0$; if $D_{\text{out}} \notin \{D_1, \dots, D_p\}$,

$$D_i \cdot \beta = |\mathbf{P}_i|, \quad D_{\text{out}} \cdot \beta = k_{\mathbf{P}}$$

while if $D_{\text{out}} = D_j$ for some j , then

$$D_i \cdot \beta = \begin{cases} |\mathbf{P}_i| & i \neq j \\ |\mathbf{P}_i| + k_{\mathbf{P}} & i = j \end{cases}$$

That such a class exists follows from (6.7) and Proposition 4.2. We can then define

$$\beta_{\mathbf{P}} = \nu^*(\beta) - \sum_{i=1}^p \sum_{j=1}^{\ell_i} p_{ij} [E_{ij}] \in H_2(\tilde{X}_{\mathfrak{d}}, \mathbb{Z}).$$

This is, roughly speaking, the class of the proper transform of a curve in $\tilde{X}_{\mathfrak{d}}$ which passes through the point x_{ij} transversally to D_i p_{ij} times and intersects D_{out} at $k_{\mathbf{P}}$ points, counted with multiplicity. Finally, let $X_{\mathfrak{d}}^{\circ}$ be obtained from $X_{\mathfrak{d}}$ by removing all zero-dimensional torus orbits, and let $\tilde{X}_{\mathfrak{d}}^{\circ} = \nu^{-1}(X_{\mathfrak{d}}^{\circ})$. Let $\tilde{D}_{\text{out}}^{\circ} = \tilde{D}_{\text{out}} \cap \tilde{X}_{\mathfrak{d}}^{\circ}$.

We now consider, somewhat informally, the moduli space of *relative stable maps* of genus zero representing the class $\beta_{\mathbf{P}}$,

$$\overline{\mathfrak{M}}(\tilde{X}_{\mathfrak{d}}^{\circ}/\tilde{D}_{\text{out}}^{\circ}) \subseteq \overline{\mathfrak{M}}(\tilde{X}_{\mathfrak{d}}/\tilde{D}_{\text{out}}),$$

in which we impose maximal tangency with \tilde{D}_{out} . Roughly, $\overline{\mathfrak{M}}(\tilde{X}_{\mathfrak{d}}/\tilde{D}_{\text{out}})$ is the moduli space of stable maps $f : C \rightarrow \tilde{X}_{\mathfrak{d}}$ with the following properties:

- (1) C is genus zero.
- (2) f represents the class $\beta_{\mathbf{P}}$.
- (3) There is a unique point $p \in C$ such that $f(p) \in \tilde{D}_{\text{out}}$, and if x is a local equation for \tilde{D}_{out} near $f(p)$, then $f^*(x)$ has a zero of order $k_{\mathbf{P}}$.

In general, this isn't quite accurate, because the space of such maps is not compact. To define this moduli space correctly, one must allow bubbling-off phenomena to occur not just to the curve C but also to the space $\tilde{X}_{\mathfrak{d}}$. This was done in the algebro-geometric context by Jun Li in [72] and [73]. See [35] for a summary of this construction. The reader unfamiliar with this concept, however, need not worry too much about it, as we shall not need the details of the construction here.

We then define $\overline{\mathfrak{M}}(\tilde{X}_{\mathfrak{d}}^{\circ}/\tilde{D}_{\text{out}}^{\circ})$ to be the open subspace consisting of those relative stable maps with image contained in $\tilde{X}_{\mathfrak{d}}^{\circ}$ rather than $\tilde{X}_{\mathfrak{d}}$. While $\overline{\mathfrak{M}}(\tilde{X}_{\mathfrak{d}}/\tilde{D}_{\text{out}})$

is compact because the target space $\tilde{X}_{\mathfrak{d}}$ is compact, it turns out that the open subset $\overline{\mathfrak{M}}(\tilde{X}_{\mathfrak{d}}^{\circ}/\tilde{D}_{\text{out}}^{\circ})$ is also compact: such curves cannot converge to curves passing through one of the deleted points. This is proved in [45], Proposition 5.1. Furthermore, the space $\overline{\mathfrak{M}}(\tilde{X}_{\mathfrak{d}}^{\circ}/\tilde{D}_{\text{out}}^{\circ})$ (generally a Deligne-Mumford stack) has expected dimension zero, and a virtual fundamental class. We then define the number

$$N_{\mathbf{P}} := \int_{[\overline{\mathfrak{M}}(\tilde{X}_{\mathfrak{d}}^{\circ}/D_{\text{out}}^{\circ})]^{\text{vir}}} 1.$$

Informally, this is the number of rational curves in the class $\beta_{\mathbf{P}}$ which are maximally tangent to \tilde{D}_{out} . Of course, these numbers can be more subtle, involving the usual sorts of virtual counts which appear in Gromov-Witten theory: multiple covers and stacky phenomena can produce fractional numbers.

With these definitions, we can now state the main theorem of [45]:

THEOREM 6.42.

$$\log f_{\mathfrak{d}} = \sum_{\mathbf{P}} k_{\mathbf{P}} N_{\mathbf{P}} t^{\mathbf{P}} z^{-k_{\mathbf{P}} m_{\mathfrak{d}}},$$

where the sum is over all \mathbf{P} satisfying (6.7) and $t^{\mathbf{P}}$ denotes the monomial $\prod_{ij} t_{ij}^{p_{ij}}$.

EXAMPLES 6.43. (1) Let's consider the case $p = 2$, $m_1 = (1, 0)$, $m_2 = (0, 1)$; this is essentially the case considered in Example 6.41, but with the variable t replaced with variables $t_{11}, \dots, t_{1\ell_1}, t_{21}, \dots, t_{2\ell_2}$. First, with $\ell_1 = \ell_2 = 1$, we have one output ray, $(\mathbb{R}_{\geq 0}(1, 1), 1 + t_{11}t_{21}x^{-1}y^{-1})$. (To get the original $1 + t^2x^{-1}y^{-1}$ given in Example 6.41, just substitute $t = t_{11} = t_{12}$.) The rays $\mathbb{R}_{\geq 0}(-1, 0)$, $\mathbb{R}_{\geq 0}(0, -1)$ and $\mathbb{R}_{\geq 0}(1, 1)$ are the one-dimensional cones in a fan defining \mathbb{P}^2 , and the three toric divisors, D_1, D_2 and D_{out} , are the three coordinate axes. Fixing one point on D_1 and one point on D_2 , there is precisely one line going through these two points, and taking $\mathbf{P} = (1, 1)$, we thus see $N_{\mathbf{P}} = 1$. Note, on the other hand, that

$$\log(1 + t_{11}t_{21}x^{-1}y^{-1}) = \sum_{k=1}^{\infty} \frac{(-1)^{k+1}}{k^2} k t_{11}^k t_{21}^k x^{-k} y^{-k}.$$

Thus we see that $N_{(k,k)} = (-1)^{k+1}/k^2$. For $k = 1$, this is just the single line counted above; all other contributions come from multiple covers of this line totally ramified at the point of intersection between the line and D_{out} .

(2) Next consider the case that $\ell_1 = \ell_2 = 2$, and consider the output ray $(\mathbb{R}_{\geq 0}(1, 1), (1 - t^2x^{-1}y^{-1})^{-4})$ in $\mathcal{S}(\mathfrak{D})$ as described in Example 6.41. Note

$$\log(1 - t^2x^{-1}y^{-1})^{-4} = 4 \sum_{k=1}^{\infty} k \cdot \frac{1}{k^2} t^{2k} x^{-k} y^{-k}.$$

Since we have substituted t for the four variables t_{ij} , $1 \leq i, j \leq 2$, each coefficient represents a sum over various \mathbf{P} . In particular, $N_{(1+0,1+0)} + N_{(0+1,1+0)} + N_{(1+0,0+1)} + N_{(0+1,0+1)} = 4$ is the coefficient of $t^2x^{-1}y^{-1}$. This is the count of lines passing through two points, one of which is chosen to be in $\{x_{11}, x_{12}\}$, and the other chosen to be in $\{x_{21}, x_{22}\}$. The coefficient of $t^4x^{-2}y^{-2}$ is 1, which is $N_{(1+1,1+1)} + N_{(2+0,2+0)} + \dots + N_{(0+2,0+2)}$. Here invariants of the form $N_{(2+0,1+1)}$ are easily seen to be zero, and the ones of the form $N_{(2+0,2+0)}$ just yield the multiple cover contribution from lines, which is $-1/4$ for each of the four lines. Thus $N_{(1+1,1+1)} = 2$, which accounts for the number of conics passing through the four points $\{x_{ij}\}$ and tangent to D_{out} .

(3) Finally, let's look at the case $\ell_1 = \ell_2 = 3$, again looking at the ray of slope 1, with attached function $f_{\mathfrak{d}}$. Then from the explicit formula given in Example 6.41, one calculates

$$\log f_{\mathfrak{d}} = 9t^2x^{-1}y^{-1} + 2 \cdot \frac{63}{4}t^4x^{-2}y^{-2} + 3 \cdot 55t^6x^{-3}y^{-3} + \dots$$

The accounting is as follows. The first coefficient, 9, counts the number of lines passing through one point of $\{x_{11}, x_{12}, x_{13}\}$ and one point of $\{x_{21}, x_{22}, x_{32}\}$. The second coefficient, $63/4$, accounts for the double covers of these lines, for a total of $-9/4$, and the number of conics passing through two points of each of these two sets, and tangent to D_{out} , for a total of $3 \times 3 \times 2 = 18$. Note $18 - 9/4 = 63/4$. The accounting for cubics, as can be checked by computing the scattering diagram not using the single variable t but using the variables t_{ij} , is as follows. One has $N_{(1+1+1, 1+1+1)} = 18$. This is the count of nodal cubics passing through all six points $\{x_{ij}\}$, and triply tangent to D_{out} . The number $N_{(2+1+0, 1+1+1)} = 3$: this is the number of nodal cubics triply tangent to D_{out} whose node coincides with x_{11} and which pass through all of the six points except for x_{13} . Note there are $2 \times 3 \times 2 = 12$ choices of \mathbf{P} of this type, so that the total count from genuine cubics is $18 + 36 = 54$. On the other hand, we also have triple covers of the 9 lines, each contributing $1/9$, hence a total of $54 + 1 = 55$. \square

In fact [45] gives formulas for more general scattering diagrams. Suppose we are given for each $1 \leq i \leq p$ a positive integer p_i and non-negative integers ℓ_{ij} for each $1 \leq j \leq p_i$. Consider

$$(6.8) \quad \mathfrak{D} = \left\{ \left(\mathbb{R}m_i, \prod_{j=1}^{p_i} \prod_{k=1}^{\ell_{ij}} (1 + t_{ijk}z^{-jm_i}) \right) \mid 1 \leq i \leq p \right\}.$$

As before, $m_1, \dots, m_p \in M$ are distinct primitive elements. Again, we consider $S(\mathfrak{D})$ over $\mathbb{k}[M][[t_{ijk}]]$, and interpret rays $(\mathfrak{d}, f_{\mathfrak{d}}) \in S(\mathfrak{D})$.

We define a *graded partition* to be a finite sequence $\mathbf{G} = (\mathbf{P}_1, \dots, \mathbf{P}_d)$ of ordered partitions with each part of \mathbf{P}_i being divisible by i . We write $|\mathbf{G}| = \sum_{i=1}^d |\mathbf{P}_i|$. Let $\mathbf{G} = (\mathbf{G}_1, \dots, \mathbf{G}_p)$ be a p -tuple of graded partitions, with each \mathbf{P}_{ij} , the j -th piece of \mathbf{G}_i , being of the form $p_{ij1} + \dots + p_{ij\ell_{ij}}$. Restrict to those \mathbf{G} such that

$$(6.9) \quad - \sum_{i=1}^p |\mathbf{G}_i| m_i = k_{\mathbf{G}} m_{\mathfrak{d}}$$

for some positive integer $k_{\mathbf{G}}$. Given this data, consider the class $\beta \in H_2(X_{\mathfrak{d}}, \mathbb{Z})$ specified by the requirement that, if $D \notin \{D_1, \dots, D_p, D_{\text{out}}\}$, then $D \cdot \beta = 0$; if $D_{\text{out}} \notin \{D_1, \dots, D_p\}$,

$$D_i \cdot \beta = |\mathbf{G}_i|, \quad D_{\text{out}} \cdot \beta = k_{\mathbf{G}};$$

while if $D_{\text{out}} = D_j$ for some j , then

$$D_i \cdot \beta = \begin{cases} |\mathbf{G}_i| & i \neq j \\ |\mathbf{G}_i| + k_{\mathbf{G}} & i = j \end{cases}$$

Pick general points $x_{ijk} \in D_i$, for $1 \leq i \leq p$, $1 \leq j \leq p_i$, $1 \leq k \leq \ell_{ij}$. We will now be slightly less precise than we were before. Define $N_{\mathbf{G}}$ to be the virtual number of rational curves which have tangency to order $k_{\mathbf{G}}$ with D_{out} at precisely one point, and have p_{ijk}/j branches of the curve tangent to order j with D_i at

each point x_{ijk} . This latter incidence condition is obtained as before by blowing up each point x_{ijk} , but performing an orbifold blow-up instead of an ordinary blow-up of the point x_{ijk} whenever $j > 1$, see [45], §5.5 for details. Then we have the analogous theorem from [45], §5.7:

THEOREM 6.44.

$$\log \tilde{f}_\mathfrak{d} = \sum_{\mathbf{G}} k_{\mathbf{G}} N_{\mathbf{G}} t^{\mathbf{G}} z^{-k_{\mathbf{G}} m_0}$$

where the sum is over all \mathbf{G} satisfying (6.9) and $t^{\mathbf{G}}$ denotes the monomial $\prod t_{ijk}^{p_{ijk}/j}$.

REMARK 6.45. Relative stable maps in $\overline{\mathfrak{M}}(\tilde{X}_\mathfrak{d}^\circ/\tilde{D}_{\text{out}}^\circ)$ can be fairly complicated objects, involving maps $f : C \rightarrow \tilde{X}_\mathfrak{d}^\circ$ where $\tilde{X}_\mathfrak{d}^\circ$ is a reducible scheme. However, this reducible scheme comes with a map $\hat{X}_\mathfrak{d}^\circ \rightarrow \tilde{X}_\mathfrak{d}^\circ$, and composing f with this map, we get a map $\tilde{f} : C \rightarrow \tilde{X}_\mathfrak{d}$. It follows from the arguments of Lemmas 4.1 and 4.2 of [44] that the image of \tilde{f} only intersects the proper transform of the toric boundary under the blow-up map $\nu : \tilde{X}_\mathfrak{d} \rightarrow X_\mathfrak{d}$ at $\tilde{D}_{\text{out}}^\circ$. Here ν is the ordinary blow-up of $X_\mathfrak{d}$ at the points x_{ij} in the case considered in Theorem 6.42, and is the weighted blow-up of points x_{ijk} in the case considered in Theorem 6.44.

We can also compose \tilde{f} with the blow-down ν , in which case we obtain a map $\bar{f} : C \rightarrow X_\mathfrak{d}$. Note that neither \tilde{f} nor \bar{f} needs to be stable. The domain C may have many different components, on some of which \bar{f} may be constant. However, from the above observation that the image of \tilde{f} is disjoint from \tilde{D}_i , and the fact that $\tilde{f}(C)$ is a divisor in the class $\beta_{\mathbf{P}}$ (or a similarly defined class $\beta_{\mathbf{G}}$ in the more general case), it follows that the intersection multiplicity of the divisor $\bar{f}_*(C)$ on $X_\mathfrak{d}$ with D_i at the point x_{ij} is precisely p_{ij} (or at the point x_{ijk} is precisely p_{ijk}). Furthermore, while $\bar{f}_*(C)$ may have a number of irreducible components, there is a point $q \in D_{\text{out}}$ such that

$$\bar{f}_*(C) \cap \partial X_\mathfrak{d} = \{q\} \cup \{x_{ij} \mid 1 \leq i \leq p, 1 \leq j \leq \ell_i\}.$$

(or

$$\bar{f}_*(C) \cap \partial X_\mathfrak{d} = \{q\} \cup \{x_{ijk} \mid 1 \leq i \leq p, 1 \leq j \leq p_i, 1 \leq k \leq \ell_{ij}\}.$$

There is also a constraint on this point q , which we describe in the more general case. Indeed, $\bar{f}_*(C)$ is a divisor representing $\beta \in H_2(X_\mathfrak{d}, \mathbb{Z}) \cong H^2(X_\mathfrak{d}, \mathbb{Z}) \cong \text{Pic } X_\mathfrak{d}$, as $X_\mathfrak{d}$ is two-dimensional and rational, and restricting this element to $\partial X_\mathfrak{d}$ gives a well-defined element of $\text{Pic}(\partial X_\mathfrak{d})$. Since $\bar{f}_*(C)$ is in the linear system defined by β , $\bar{f}_*(C)|_{\partial X_\mathfrak{d}}$ is in the linear equivalence class given by $\beta|_{\partial X_\mathfrak{d}}$. Note that

$$\bar{f}_*(C)|_{\partial X_\mathfrak{d}} = \sum_{i,j,k} p_{ijk} x_{ijk} + k_{\mathbf{G}} q.$$

As $\partial X_\mathfrak{d}$ is a (singular) elliptic curve, for fixed choices of x_{ijk} there are only a finite number of choices of q for which this divisor is in the linear equivalence class given by $\beta|_{\partial X_\mathfrak{d}}$.

REMARK 6.46. While the scattering diagram (6.8) may still look rather special, it in fact can be used to give the results of Theorem 6.38 for essentially any scattering diagram. Let us explain this process here in the case that we have some given P , $r : P \rightarrow M$ as usual, and a scattering diagram

$$\mathfrak{D} = \{(\mathbb{R}r(m_i), f_i)\}$$

with $f_i = 1 + c_i z^{-m_i}$, $c_i \in \mathbb{k}$ and $m_i \in P$ with $r(m_i) \neq 0$. We wish to describe $S_I(\mathfrak{D})$ for an \mathfrak{m} -primary ideal I .

First, \mathfrak{D} may have a number of overlapping lines, so replace \mathfrak{D} with an equivalent scattering diagram which we describe as follows, after renaming the m_i . There are $\bar{m}_1, \dots, \bar{m}_p \in M$ distinct primitive vectors, positive integers p_i for each $1 \leq i \leq p$, and positive integers ℓ_{ij} for each $1 \leq i \leq p$, $1 \leq j \leq p_i$, such that

$$\mathfrak{D} = \left\{ \left(\mathbb{R}\bar{m}_i, \prod_{j=1}^{p_i} \prod_{k=1}^{\ell_{ij}} (1 + c_{ijk} z^{-m_{ijk}}) \right) \mid 1 \leq i \leq p \right\}$$

where $r(m_{ijk})$ is positively proportional to \bar{m}_i and the index of $r(m_{ijk})$ is j . Let $\ell = \sum_{ij} \ell_{ij}$, and set

$$\mathfrak{D}' = \left\{ \left(\mathbb{R}\bar{m}_i, \prod_{j=1}^{p_i} \prod_{k=1}^{\ell_{ij}} (1 + t_{ijk} z^{-r(m_{ijk})}) \right) \mid 1 \leq i \leq p \right\}$$

to be the scattering diagram over the ring $\mathbb{k}[M][[t_{ijk}]]$; this is the scattering diagram considered in (6.8). Thus the rays of $S(\mathfrak{D}')$ are given by Theorem 6.44.

Let $Q \subseteq M \oplus \mathbb{N}^\ell$ be the submonoid freely generated by

$$\{(-r(m_{ijk}), e_{ijk})\}$$

where $e_{ijk} \in \mathbb{N}^\ell$ corresponds to t_{ijk} . Then \mathfrak{D}' is in fact defined over the subring $\mathbb{k}[Q]$ of $\mathbb{k}[M \oplus \mathbb{N}^\ell]$. Define a ring homomorphism

$$\begin{aligned} \varphi : \mathbb{k}[Q] &\rightarrow \mathbb{k}[P] \\ t_{ijk} z^{-r(m_{ijk})} &\mapsto c_{ijk} z^{-m_{ijk}} \end{aligned}$$

Let $I' = \varphi^{-1}(I)$. We obtain a scattering diagram $S_{I'}(\mathfrak{D}')$, over $\mathbb{k}[Q]/I'$, which by uniqueness can be taken to be obtained from $S(\mathfrak{D}')$ by reduction modulo I' .

Applying φ to the function attached to each element of $S_{I'}(\mathfrak{D}')$, we get a scattering diagram $\varphi(S_{I'}(\mathfrak{D}'))$ defined over $\mathbb{k}[P]$. Furthermore, $\theta_{\gamma, \varphi(S_{I'}(\mathfrak{D}'))} \equiv \text{Id} \pmod{I}$ for γ a loop around the origin, as $\theta_{\gamma, S_{I'}(\mathfrak{D}')}$ is the identity in $\mathbb{k}[Q]/I'$. Thus, by uniqueness, in fact $\varphi(S_{I'}(\mathfrak{D}'))$ must be equivalent to $S_I(\mathfrak{D})$. In this sense, $S(\mathfrak{D}')$ is a universal scattering diagram.

Note that the particular process described in the proof of Theorem 6.38 produces new rays whose attached functions are binomials $1 + c_\mathfrak{d} z^{m_\mathfrak{d}}$. We can achieve this here: given a ray $(\mathfrak{d}, f_\mathfrak{d}) \in \varphi(S_{I'}(\mathfrak{D}'))$, we can factor $f_\mathfrak{d}$ in $\mathbb{k}[P]/I$ uniquely into a product of such terms using the following mechanism. Suppose that we have factored

$$f_\mathfrak{d} \equiv \prod_{i=1}^d (1 + c_i z^{p_i}) \pmod{\mathfrak{m}^k}.$$

Then modulo \mathfrak{m}^{k+1} ,

$$f_\mathfrak{d} - \prod_{i=1}^d (1 + c_i z^{p_i}) \equiv \sum_j d_j z^{q_j},$$

with $z^{q_j} \in \mathfrak{m}^k \setminus \mathfrak{m}^{k+1}$. Then modulo \mathfrak{m}^{k+1} , we can write

$$f_\mathfrak{d} \equiv \left(\prod_{i=1}^d (1 + c_i z^{p_i}) \right) \left(\prod_j (1 + d_j z^{q_j}) \right).$$

We continue like this until we reach a k such that $\mathfrak{m}^k \subseteq I$. Thus we can replace the single ray $(\mathfrak{d}, f_{\mathfrak{d}})$ with the collection of rays $\{(\mathfrak{d}, 1 + c_i z^{p_i})\}$. This new scattering diagram satisfies the conclusions of Theorem 6.38, and in fact is the scattering diagram produced by the algorithm given there.

We need one technical observation that comes easily from Theorem 6.42.

PROPOSITION 6.47. *Let \mathfrak{D} be the scattering diagram given in (6.8). Let J be a monomial ideal in $R = \mathbb{k}[\{t_{ijk} \mid (i, j, k) \neq (1, 1, 1)\}]$ such that R/J is Artinian. Also denoting by J the ideal generated by J in $\mathbb{k}[M] \otimes_{\mathbb{k}} \mathbb{k}[\{t_{ijk}\}]$, let $I_e = (t_{111}^e) + J$ be the monomial ideal in this latter ring, so that I_e is an \mathfrak{m} -primary ideal. Let \mathfrak{D}_e be the scattering diagram $\mathbf{S}_{I_e}(\mathfrak{D})$, obtained via the algorithm of the proof of Theorem 6.38, such that $\mathbf{S}_{I_e}(\mathfrak{D})$ contains no elements of the form $(\mathfrak{d}, f_{\mathfrak{d}})$ with $\mathfrak{d} \equiv 1 \pmod{I_e}$. Note that by construction, we then have $\mathfrak{D}_e \subseteq \mathfrak{D}_{e+1}$. Then the sequence $\mathfrak{D}_1, \mathfrak{D}_2, \dots$ stabilizes.*

PROOF. Consider the set Γ consisting of collections of graded partitions $\mathbf{G} = (\mathbf{G}_1, \dots, \mathbf{G}_p)$ with

$$\prod_{\substack{i,j,k \\ (i,j,k) \neq 1}} t_{ijk}^{p_{ijk}/j} \notin J$$

and $p_{ijk} > 0$ for some $(i, j, k) \neq (1, 1, 1)$. Since R/J is Artinian, for any given choice of p_{111} , there are a finite number of $\mathbf{G} \in \Gamma$ with this p_{111} , but Γ itself is infinite as there is no bound on p_{111} . In particular, if $\mathbf{G} \in \Gamma$, then all p_{ijk} are bounded by some number N except for p_{111} . We will first show that there are only a finite number of $\mathbf{G} \in \Gamma$ such that $N_{\mathbf{G}} \neq 0$.

Suppose $\mathbf{G} \in \Gamma$ with $N_{\mathbf{G}} \neq 0$. There is a primitive $m_{\mathfrak{d}} \in M$ such that (6.9) holds, and hence a ray $(\mathfrak{d}, f_{\mathfrak{d}})$ in the scattering diagram \mathfrak{D}_e for some e with support $\mathfrak{d} = \mathbb{R}_{\geq 0}m_{\mathfrak{d}}$. Let $\Sigma_{\mathfrak{d}}$ be as usual. We can always refine $\Sigma_{\mathfrak{d}}$, and can thus assume that $\Sigma_{\mathfrak{d}}$ contains both the ray $\mathbb{R}_{\geq 0}(-m_1)$ and the ray $\mathbb{R}_{\geq 0}(m_1)$, so that the projection $M_{\mathbb{R}} \rightarrow M_{\mathbb{R}}/\mathbb{R}m_1$ gives a toric morphism $\pi : X_{\mathfrak{d}} \rightarrow \mathbb{P}^1$. The morphism π has a section given by the divisor D_1 corresponding to $\mathbb{R}_{\geq 0}m_1$.

By Remark 6.45, the fact that $N_{\mathbf{G}} \neq 0$ means that there is a map $\bar{f} : C \rightarrow X_{\mathfrak{d}}$ such that the divisor $\bar{f}_*(C)$ has intersection multiplicity p_{111} with D_1 at x_{111} . Also, we can assume that the fibre $\pi^{-1}(\pi(x_{111}))$ is not contained in the image of \bar{f} . Indeed, this fibre intersects $\partial X_{\mathfrak{d}}$ at another point other than x_{111} , and by assuming the points x_{ijk} are generally chosen, we can assume this point does not coincide with any of the x_{ijk} 's. Hence, in the notation of Remark 6.45, we can call this point q . Now by that remark, the divisor class of the divisor on $\partial X_{\mathfrak{d}}$ given by $\sum p_{ijk}x_{ijk} + k_{\mathbf{G}}q$ is determined completely by \mathbf{G} . However, there is some $(i, j, k) \neq (1, 1, 1)$ with $p_{ijk} \neq 0$, so by varying that x_{ijk} and keeping all other points fixed, we obtain a contradiction. Thus for general choice of this x_{ijk} , $\bar{f}_*(C)$ cannot contain $\pi^{-1}(\pi(x_{111}))$.

Now $\bar{f}_*(C)$ represents the class β , and $\bar{f}_*(C)$ represents a class $\beta_{\mathbf{G}} = \nu^*(\beta) - p_{111}E_{111} - \dots$. Since the total transform of the fibre $\pi^{-1}(\pi(x_{111}))$ consists of two irreducible components, one being E_{111} , and $E_{111} \cdot \beta_{\mathbf{G}} = p_{111}$, we must have the intersection multiplicity of $\pi^{-1}(\pi(x_{111}))$ with $\bar{f}_*(C)$ being at least p_{111} . From this, we conclude that the intersection number $\beta \cdot F$, where F is the class of a fibre of π , is at least p_{111} .

Next observe that at least one of $\pi^{-1}(0)$ and $\pi^{-1}(\infty)$ does not contain the divisor D_{out} ; assume it is $\pi^{-1}(0)$. Then the proper transform of $\pi^{-1}(0)$ in $\tilde{X}_{\mathfrak{d}}$ is disjoint from $\tilde{f}(C)$. Thus $\beta \cdot \pi^{-1}(0)$ is completely determined by $\mathbf{G}_2, \dots, \mathbf{G}_p$, and hence is bounded. Since $\pi^{-1}(0)$ is linearly equivalent to F , we see that p_{111} is bounded. This argument did not depend on \mathfrak{d} , so we see that p_{111} is bounded independently of \mathfrak{d} . Hence the set $\{\mathbf{G} \in \Gamma \mid N_{\mathbf{G}} \neq 0\}$ is finite.

In particular, by Theorem 6.44 and (6.9), one sees that at most a finite number of distinct directions for rays \mathfrak{d} can occur in $\bigcup_e \mathfrak{D}_e$.

Suppose now we have taken \mathfrak{D}_e not to have more than one ray with any given support, so that for a ray $(\mathfrak{d}, f_{\mathfrak{d}}) \in \mathfrak{D}_e$, $\log f_{\mathfrak{d}}$ is given by Theorem 6.44. Consider a ray $(\mathfrak{d}, f_{\mathfrak{d}})$ which appears in \mathfrak{D}_e for some e . Since this is a ray, any coefficient must involve some t_{ijk} for $(i, j, k) \neq 1$; otherwise, it would have to arise when computing $S(\{\{\mathbb{R}m_1, 1 + t_{111}z^{-m_1}\}\})$, but there are no rays in this diagram. Thus in particular, the above argument shows that there are only a finite number of terms in $\log f_{\mathfrak{d}}$ modulo J which can appear, and since every coefficient in $\log f_{\mathfrak{d}}$ involves some t_{ijk} with $t_{ijk}^N \in J$, $f_{\mathfrak{d}} = \exp(\log f_{\mathfrak{d}})$ only has a finite number of terms modulo J . Thus there is some bound on the number of terms in $f_{\mathfrak{d}}$ independent of e . Finally, applying the factorization process given in Remark 6.46, we can replace each ray $(\mathfrak{d}, f_{\mathfrak{d}})$ with a finite number of rays $(\mathfrak{d}, f_{\mathfrak{d}})$ with $f_{\mathfrak{d}}$ having the form $1 + cz^m$. This shows the desired result. \square

REMARK 6.48. The above proposition can be interpreted as follows. The scattering diagram

$$S(\mathfrak{D}) = \bigcup_{e=1}^{\infty} S_{I_e}(\mathfrak{D})$$

can be viewed as a scattering diagram for the ring

$$\hat{A} := \mathbb{k}[M] \otimes_{\mathbb{k}} (\mathbb{k}[\{t_{ijk} \mid (i, j, k) \neq (1, 1, 1)\}]/J)[[t_{111}]],$$

and in this ring $\theta_{\gamma, S(\mathfrak{D})} = \text{Id}$. However, \hat{A} contains as a subring the ring

$$A := (\mathbb{k}[M] \otimes_{\mathbb{k}} \mathbb{k}[\{t_{ijk}\}]/J)_{1+t_{111}z^{-m_1}},$$

and $\theta_{\gamma, \mathfrak{D}}$ makes sense as an automorphism of A . The fact that $S(\mathfrak{D})$ is finite means that $\theta_{\gamma, S(\mathfrak{D})}$ also makes sense as an automorphism of A , and is of course still the identity. This tells us that the Kontsevich-Soibelman lemma (Theorem 6.38) makes sense over the localized ring A in this particular case.

6.3.3. Making structures compatible. We can now construct a compatible structure. Start with $(B, \mathcal{P}, \varphi)$ as usual: B is an integral tropical manifold of dimension two with positive and simple singularities. We will construct a sequence of structures $\mathcal{S}_0, \mathcal{S}_1, \dots$ with \mathcal{S}_k compatible to order k .

We start with \mathcal{S}_0 : we have no choice here but to set

$$\mathcal{S}_0 = \{\mathfrak{d}_{p, \pm} \mid p \in \Delta\},$$

where $\mathfrak{d}_{p, \pm}$ denotes the two rays associated to the singular point p required by the definition of structure. It is easy to see that \mathcal{S}_0 is compatible to order 0. Let us check this. First, for $\mathfrak{d}_{p, \pm}$, we can take $N_{\mathfrak{d}_{p, \pm}}^0$, the number promised by Proposition 6.21, (2), to coincide with the non-zero endpoint of the interval $I_{\mathfrak{d}_{p, \pm}}$ defined in Definition 6.22, (2). As a consequence, we can take $\mathcal{P}_0 = \mathcal{P}$. Thus the only joints

are the vertices of \mathcal{P} . If v is a vertex, then $f_{\mathfrak{d}} \equiv 1 \pmod{I_{v,v,\sigma}^0}$ for any $\mathfrak{d} \in \mathcal{S}_0$, so compatibility is trivial.

We now wish to construct \mathcal{S}_k inductively, so assume that we have constructed \mathcal{S}_{k-1} compatible to order $k-1$. We will add new rays to obtain \mathcal{S}_k compatible to order k . For this purpose, we use the results of the previous two subsections. In particular, the argument is purely local, in that at each joint we add a collection of rays dictated by the argument of Kontsevich and Soibelman.

For a given joint \mathfrak{j} , take $M = \Lambda_{\mathfrak{j}}$, the stalk of Λ at the point \mathfrak{j} , and take $P = P_{\varphi,\sigma_{\mathfrak{j}},\sigma}$ for some $\sigma \in \mathcal{P}_{\max}$ containing \mathfrak{j} . Let $I = I_{\sigma_{\mathfrak{j}},\sigma}^k$. Note that I is always \mathfrak{m} -primary, where $\mathfrak{m} = P \setminus P^\times$.

Construct a scattering diagram $\mathfrak{D}_{\mathfrak{j}}$ as follows. The elements of $\mathfrak{D}_{\mathfrak{j}}$ are in one-to-one correspondence with pairs (\mathfrak{d}, x) with $\mathfrak{d} \in \mathcal{S}_{k-1}$ such that $x \in [0, N_{\delta}^k]$ and $\mathfrak{d}(x) = \mathfrak{j}$. Given such a pair (\mathfrak{d}, x) , we associate to it an element of $\mathfrak{D}_{\mathfrak{j}}$ of the form

$$\begin{cases} (\mathbb{R}_{\geq 0}\mathfrak{d}'(x), 1 + c_{\mathfrak{d}}z^{m_{\mathfrak{d},x}}) & x = 0 \\ (\mathbb{R}\mathfrak{d}'(x), 1 + c_{\mathfrak{d}}z^{m_{\mathfrak{d},x}}) & x \neq 0 \end{cases}$$

Here $\mathfrak{d}'(x) \in \Lambda_{\mathfrak{d}(x)} = \Lambda_{\mathfrak{j}}$ denotes the tangent vector to \mathfrak{d} at x . We now need to consider two cases.

$\dim \sigma_{\mathfrak{j}} = 2$ and $\dim \sigma_{\mathfrak{j}} = 0$: In this case, $\mathfrak{D}_{\mathfrak{j}}$ is actually a scattering diagram over $\mathbb{k}[P]/I$. All the conditions are obvious from the definition of a ray in a structure except for the requirement that $f_{\mathfrak{d}} \equiv 1 \pmod{\mathfrak{m}}$ for $\mathfrak{d} \in \mathfrak{D}_{\mathfrak{j}}$. However, if \mathfrak{d} corresponds to a pair $(\bar{\mathfrak{d}}, x)$ and $\dim \sigma_{\mathfrak{j}} = 2$, then $\bar{\mathfrak{d}}$ cannot be an initial ray, and hence $\text{ord}_{\sigma_{\mathfrak{j}}} m_{\bar{\mathfrak{d}},x} > 0$ by Definition 6.22, (3). Thus $1 + c_{\bar{\mathfrak{d}}}z^{m_{\bar{\mathfrak{d}},x}} \equiv 1 \pmod{\mathfrak{m}}$. On the other hand, if $\dim \sigma_{\mathfrak{j}} = 0$, and \mathfrak{d} corresponds to $(\bar{\mathfrak{d}}, x)$, then there is some $\sigma \in \mathcal{P}_{\max}$ containing \mathfrak{j} with $\text{ord}_{\sigma} m_{\bar{\mathfrak{d}},x} > 0$, either by Definition 6.22, (3) if $\bar{\mathfrak{d}}$ is not an initial ray, and by Lemma 6.19, (2), if $\bar{\mathfrak{d}}$ is an initial ray (noting that a joint never lies in ∂B). So again $1 + c_{\bar{\mathfrak{d}}}z^{m_{\bar{\mathfrak{d}},x}} \equiv 1 \pmod{\mathfrak{m}}$.

Thus we get a finite scattering diagram $S_I(\mathfrak{D}_{\mathfrak{j}})$. Note that the particular construction of $S_I(\mathfrak{D}_{\mathfrak{j}})$ given in Theorem 6.38 always added rays $(\mathfrak{d}, f_{\mathfrak{d}})$ with $f_{\mathfrak{d}}$ of the form $1 + c_{\mathfrak{d}}z^{m_{\mathfrak{d}}}$, and we can assume that $m_{\mathfrak{d}} \in I_{\sigma_{\mathfrak{j}},\sigma}^{k-1}$ for each added ray. Indeed, the fact that \mathcal{S}_{k-1} is compatible implies

$$\theta_{\gamma,\mathfrak{D}_{\mathfrak{j}}} \equiv \text{Id} \pmod{I_{\sigma_{\mathfrak{j}},\sigma}^{k-1}}$$

for a loop γ around the origin.

Now for each $(\mathfrak{d}, 1 + c_{\mathfrak{d}}z^{m_{\mathfrak{d}}}) \in S_I(\mathfrak{D}_{\mathfrak{j}}) \setminus \mathfrak{D}_{\mathfrak{j}}$, let $(\bar{\mathfrak{d}}, f_{\bar{\mathfrak{d}}})$ be the ray on B obtained by taking $\bar{\mathfrak{d}}(0) = \mathfrak{j}$, $\bar{\mathfrak{d}}'(0) = -\bar{m}_{\mathfrak{d}}$, and $f_{\bar{\mathfrak{d}}} = 1 + c_{\mathfrak{d}}z^{m_{\bar{\mathfrak{d}}}}$, where $m_{\bar{\mathfrak{d}}}$ is obtained by parallel transport of $m_{\mathfrak{d}}$. It follows from Proposition 6.21 that this is a ray. Set

$$(6.10) \quad \mathcal{S}_{\mathfrak{j}} = \{(\bar{\mathfrak{d}}, f_{\bar{\mathfrak{d}}}) \mid \mathfrak{d} \in S_I(\mathfrak{D}_{\mathfrak{j}}) \setminus \mathfrak{D}_{\mathfrak{j}}\}.$$

$\dim \sigma_{\mathfrak{j}} = 1$: If the edge $\sigma_{\mathfrak{j}}$ contains a singular point, then, without loss of generality, we can assume $\mathfrak{j} \in \mathfrak{d}_{p,+}(I_{\mathfrak{d}_{p,+}}^{\text{init}})$ (as opposed to $\mathfrak{d}_{p,-}(I_{\mathfrak{d}_{p,-}}^{\text{init}})$). Then for any σ containing $\sigma_{\mathfrak{j}}$, $\text{ord}_{\sigma} m_{-} = 0$, so $f_{\mathfrak{d}_{p,+}} \not\equiv 1 \pmod{\mathfrak{m}}$. However, this is the only ray which causes any problem. In particular, given a path γ in $M_{\mathbb{R}}$, $\theta_{\gamma,\mathfrak{D}_{\mathfrak{j}}}$ does not make sense in the ring $\mathbb{k}[P]/I$ because $f_{\mathfrak{d}_{p,+}}$ is not invertible. It only makes sense in the ring $(\mathbb{k}[P]/I)_{f_{\mathfrak{d}_{p,+}}}$, but we haven't proved a version of Theorem 6.38 over this ring. Nevertheless, we can interpret $\mathfrak{D}_{\mathfrak{j}}$ in a different way to make sense of this, using Remark 6.46 and Proposition 6.47.

In particular, apply the construction of Remark 6.46 to the scattering diagram

$$\overline{\mathfrak{D}}_j = \{\mathfrak{d} \in \mathfrak{D}_j \mid \mathfrak{d} \text{ is a line}\}.$$

We replace it as in the remark with an equivalent scattering diagram

$$\overline{\mathfrak{D}}_j = \left\{ \left(\mathbb{R}\overline{m}_i, \prod_{j=1}^{p_i} \prod_{k=1}^{\ell_{ij}} (1 + c_{ijk} z^{-m_{ijk}}) \right) \right\},$$

where we can assume that $1 + c_{111} z^{-m_{111}} = 1 + z^{m-}$ is the term causing trouble for us. We then follow the procedure of Remark 6.46, obtaining $S(\mathfrak{D}')$. The only subtlety is that $t_{111} z^{-r(m_{111})}$ now maps under φ to z^{m-} , which does not lie in \mathfrak{m} . Nevertheless, it then follows from Proposition 6.47 that if $I \subseteq P$ is an \mathfrak{m} -primary ideal, then $\varphi(S(\mathfrak{D}'))$, modulo I , only contains a finite number of rays. However, it is still true that $\theta_{\gamma, \varphi(S(\mathfrak{D}'))} \equiv \text{Id} \pmod{I}$, but now in the localized ring $(\mathbb{k}[P]/I)_{1+z^{m-}}$, where this automorphism makes sense, as in Remark 6.48.

Now define $S_I(\mathfrak{D}_j)$ to be obtained by factorizing all the rays in $\varphi(S(\mathfrak{D}'))$ as in Remark 6.46. This is a unique procedure, and all rays in this scattering diagram which are not trivial modulo $I_{\sigma_j, \sigma_j, \sigma}^{k-1}$ will inductively already appear in \mathfrak{D}_j . So as before use (6.10) to define \mathcal{S}_j . \square

Let

$$\mathcal{S}_k = \mathcal{S}_{k-1} \cup \bigcup_{j \in \text{Joints}(\mathcal{S}_{k-1}, k)} \mathcal{S}_j.$$

THEOREM 6.49. \mathcal{S}_k is compatible to order k .

PROOF. Let $j \in \text{Joints}(\mathcal{S}_k, k)$, and let

$$\theta_j = \theta_{u_n, u_1} \circ \cdots \circ \theta_{u_1, u_2} : R_{\sigma_j, \sigma_j, u_1}^k \rightarrow R_{\sigma_j, \sigma_j, u_1}^k$$

be the automorphism which must be the identity modulo $I = I_{\sigma_j, \sigma_j, \sigma_{u_1}}^k$ in order to achieve compatibility. By assumption, $\theta_j \equiv \text{Id} \pmod{I_{\sigma_j, \sigma_j, \sigma_{u_1}}^{k-1}}$, since all rays $\mathfrak{d} \in S_I(\mathfrak{D}_j) \setminus \mathfrak{D}_j$ satisfy $f_{\mathfrak{d}} \equiv 1 \pmod{I_{\sigma_j, \sigma_j, \sigma_{u_1}}^{k-1}}$. (If $j \notin \text{Joints}(\mathcal{S}_k, k-1)$, then we take $S_I(\mathfrak{D}_j) \setminus \mathfrak{D}_j$ to be empty.) Furthermore, by construction of $S_I(\mathfrak{D}_j)$, θ_j can only fail to be the identity modulo I if there is some ray $\mathfrak{d} \in \mathcal{S}_k$ such that $\mathfrak{d}(x) = j$ for some $x \neq 0$ and the pair (\mathfrak{d}, x) did not appear in the construction of \mathfrak{D}_j .

There are two reasons why the pair (\mathfrak{d}, x) would fail to appear in the construction. First, it could be that $x \notin [0, N_{\mathfrak{d}}^k]$, in which case $z^{m_{\mathfrak{d}, x}} \in I_{\sigma_j, \sigma_j, \sigma_{u_1}}^k$ and thus (\mathfrak{d}, x) does not contribute to θ_j . Second, it could be that $\mathfrak{d} \in \mathcal{S}_k \setminus \mathcal{S}_{k-1}$, in which case $z^{m_{\mathfrak{d}, x}} \in I_{\sigma_j, \sigma_j, \sigma_{u_1}}^{k-1}$. We now analyze the situation based on the possible dimensions of σ_j .

$\dim \sigma_j = 2$: In this case, $I_{\sigma_j, \sigma_j, \sigma_{u_1}}^0 \cdot I_{\sigma_j, \sigma_j, \sigma_{u_1}}^{k-1} \subseteq I_{\sigma_j, \sigma_j, \sigma_{u_1}}^k$. Thus the automorphism induced by any such ray commutes with the automorphisms associated with all the other rays passing through j , modulo $I_{\sigma_j, \sigma_j, \sigma_{u_1}}^k$. Since \mathfrak{d} contributes twice to θ_j , but with inverse automorphisms, in fact \mathfrak{d} does not affect $\theta_j \equiv \text{Id} \pmod{I_{\sigma_j, \sigma_j, \sigma_{u_1}}^k}$.

$\dim \sigma_j = 1$: In this case, either \mathfrak{d} is parallel to the edge σ_j , or it is transversal to this edge. In the first case the automorphism associated to \mathfrak{d} commutes with the automorphism associated to $\mathfrak{d}_{p, \pm}$, where $p \in \sigma_j$ is the singular point, if there is one. This is because \mathfrak{d} is parallel to $\mathfrak{d}_{p, \pm}$. On the other hand, let σ_{\pm} be the maximal cells containing σ_j . Writing $f_{(\mathfrak{d}, x)} = 1 + cz^m$, one sees that we must have $\text{ord}_{\sigma_{\pm}} m = k$.

Thus the automorphism associated to \mathfrak{d} commutes with any other automorphism contributing to θ_j modulo $I_{\sigma_j, \sigma_j, \sigma_{u_1}}^k$.

In the second case, suppose without loss of generality that \mathfrak{d} passes say, from the maximal cell σ_+ into the maximal cell σ_- . By construction, we must have $\text{ord}_{\sigma_+}(m_{\mathfrak{d},x}) \geq k$ and thus $\text{ord}_{\sigma_-}(m_{\mathfrak{d},x}) > k$, by Lemma 6.19. So the automorphism associated to this ray is trivial modulo $I_{\sigma_j, \sigma_j, \sigma_{u_1}}^k$ and has no impact on θ_j .

$\dim \sigma_j = 0$: In this case, \mathfrak{d} passes from a maximal cell σ_+ containing j to a cell σ_- containing j . By the same argument, $\text{ord}_{\sigma_+}(m_{\mathfrak{d},x}) \geq k$ and $\text{ord}_{\sigma_-}(m_{\mathfrak{d},x}) > k$, and again \mathfrak{d} has no impact on θ_j . \square

Proof of Theorem 6.6. The above construction produces, by Theorem 6.49, a structure \mathcal{S}_k compatible to order k . By Theorem 6.35, this gives a scheme $\tilde{X}_k(B, \mathcal{P})$ flat over O_k . Furthermore, by construction, $\tilde{X}_k(B, \mathcal{P}) \times_{R_k} R_{k-1} \cong \tilde{X}_{k-1}(B, \mathcal{P})$. Hence we can take the limit of schemes $\tilde{X}_k(B, \mathcal{P})$ over all k , getting a formal scheme $\tilde{\mathfrak{X}}(B, \mathcal{P})$ flat over $\text{Spf } \mathbb{k}[[t]]$, as desired. \square

6.4. Remarks and generalizations

The argument given in the previous two sections has a lot in common with the original argument given by Kontsevich and Soibelman. They did not work with a polyhedral decomposition, however, and instead of constructing a degenerating family of varieties, they constructed a rigid analytic space via a similar gluing construction. The data controlling the two gluing constructions are very similar. A key difference, though, is that Kontsevich and Soibelman work in what we would call the fan picture. As a result, rays are not straight, but have to be taken to be gradient flow lines of certain functions with respect to a special metric. Thus their construction looks less tropical than the one given here, where the union of all rays looks like a union of tropical trees. Ignoring the discrete data of a polyhedral decomposition also makes it easier to describe the construction, since there are no case-by-case analyses to conduct. Despite this, there is one unpleasant aspect of their argument: they need to construct a choice of metric so that none of the gradient flow lines returns to a small neighbourhood of the singular points. This is needed to guarantee convergence of the construction.

This issue causes pain here too, but is all subsumed in Proposition 6.47. While the argument here is simple, it is of course using the full strength of [45]. In [49], we give a different, longer but more elementary argument to deal with this case. In any event, it seems to be impossible to avoid some technical issues about convergence caused by singularities.

The paper [49] gives results in all dimensions. The approach taken in this chapter can be viewed as a hybrid of Kontsevich and Soibelman's version and the approach of [49]. There, in general, instead of considering rays, one considers codimension one walls emanating from the singular locus, with attached functions determining gluings. Such objects can be topologically very complicated, so it proved much easier to build walls out of building blocks, order by order, replacing objects of possibly infinite extent with walls sitting inside cells of \mathcal{P} . This already makes [49] somewhat more difficult to read. Now, joints are codimension two polyhedra, and one adds new walls at each joint.

Far more subtle are the issues of convergence that arise from order 0 terms on walls coming directly out of the singular locus. In particular, the arguments that

guarantee suitable versions of Theorem 6.38 for joints contained in codimension one and codimension two cells of \mathcal{P} are very difficult, and run to about 40 pages. It is hopeful that, with a better enumerative understanding of the higher dimensional case, these arguments might become more conceptual, if not easier. On the other hand, I know of no way of generalizing Kontsevich and Soibelman's approach to higher dimensions; the convergence issues seem even harder. Another issue which arises is that codimension three polyhedra where different joints intersect play a role. Such polyhedra are called interstices. Fortunately, the algorithm takes care of the interplay between various joints meeting at an interstice relatively easily; the argument of the proof of Theorem 5.39 uses a very similar idea.

That being said, what can we hope to accomplish with this construction? The main point is that we have a description of degenerations of varieties determined by $(B, \mathcal{P}, \varphi)$ in terms of what are essentially Maslov index zero tropical disks. Rational tropical curves can then be obtained by gluing together such tropical disks, very much as was carried out in Chapter 5. One needs to perform a more general version of the period calculations done in Chapter 5 to see how these period calculations extract all ways of gluing together these disks. Once this is done, one should have a precise tropical interpretation for the B-model period calculations in general. Coupling this with a calculation of Gromov-Witten invariants using the theory of log Gromov-Witten invariants currently under development, in the spirit of Chapter 4, one can then hope to prove mirror symmetry. At the time of writing, there remains much to be done in this direction.

6.5. References and further reading

The material in this section is drawn from joint work of myself and Siebert, except for the material of §6.3.2, which recounts work carried out in [45]; see also [44] for a survey. To learn more about the details of the Gross-Siebert program, the papers [47], [41], and [50] are intended as expository papers.

The technical details of the program itself have been developed in [48], [51], and [49]. The latter paper, giving the smoothing result, should be accessible after reading this chapter without reading [48].

As mentioned earlier, the argument given in this chapter can be viewed as a variation of the original argument given in the rigid analytic situation by Kontsevich and Soibelman in [70]. Since the latter paper was written, the authors have found applications of the group of symplectomorphisms which arose originally in [70] (here called the tropical vertex group) to wall-crossing formulas: see [71]. Furthermore, scattering diagrams as discussed here also have an interpretation in terms of Euler characteristics of moduli spaces of quiver representations: see the work of Reineke in [94] and [95]. See also related work of Gaiotto, Moore and Neitzke, [29], [30]. The ideas described in this chapter have also been used recently in joint forthcoming work of myself with Paul Hacking and Sean Keel to prove a conjecture of Looijenga on the smoothability of cusp singularities and to construct canonical theta functions for K3 surfaces.

Bibliography

1. *Revêtements étales et groupe fondamental*, Springer-Verlag, Berlin, 1971, Séminaire de Géométrie Algébrique du Bois Marie 1960–1961 (SGA 1), Dirigé par Alexandre Grothendieck. Augmenté de deux exposés de M. Raynaud, Lecture Notes in Mathematics, Vol. 224. MR MR0354651 (50 #7129)
2. Valery Alexeev and Iku Nakamura, *On Mumford's construction of degenerating abelian varieties*, *Tohoku Math. J. (2)* **51** (1999), no. 3, 399–420. MR MR1707764 (2001g:14013)
3. P. Aspinwall, T. Bridgeland, A. Craw, M. Douglas, M. Gross, A. Kapustin, G. Moore, G. Segal, B. Szendrői, and P.M.H. Wilson, *Dirichlet branes and mirror symmetry*, Clay Mathematics Monographs, vol. 4, American Mathematical Society, Providence, RI, 2009.
4. Serguei Barannikov, *Semi-infinite Hodge structures and mirror symmetry for projective spaces*, arXiv:math/0010157, 2000.
5. ———, *Quantum periods. I. Semi-infinite variations of Hodge structures*, *Internat. Math. Res. Notices* (2001), no. 23, 1243–1264. MR MR1866443 (2002k:32017)
6. Victor V. Batyrev, *Dual polyhedra and mirror symmetry for Calabi-Yau hypersurfaces in toric varieties*, *J. Algebraic Geom.* **3** (1994), no. 3, 493–535. MR MR1269718 (95c:14046)
7. Victor V. Batyrev and Lev A. Borisov, *On Calabi-Yau complete intersections in toric varieties*, *Higher-dimensional complex varieties (Trento, 1994)*, de Gruyter, Berlin, 1996, pp. 39–65. MR MR1463173 (98j:14052)
8. Kai Behrend and Barbara Fantechi, *The intrinsic normal cone*, *Invent. Math.* **128** (1997), no. 1, 45–88. MR MR1437495 (98e:14022)
9. Kai Behrend and Yuri Manin, *Stacks of stable maps and Gromov-Witten invariants*, *Duke Math. J.* **85** (1996), no. 1, 1–60. MR MR1412436 (98i:14014)
10. Philip Candelas, Xenia C. de la Ossa, Paul S. Green, and Linda Parkes, *A pair of Calabi-Yau manifolds as an exactly soluble superconformal theory*, *Nuclear Phys. B* **359** (1991), no. 1, 21–74. MR MR1115626 (93b:32029)
11. Philip Candelas, Gary T. Horowitz, Andrew Strominger, and Edward Witten, *Vacuum configurations for superstrings*, *Nuclear Phys. B* **258** (1985), no. 1, 46–74. MR MR800347 (87k:83091a)
12. Philip Candelas, Monika Lynker, and Rolf Schimmrigk, *Calabi-Yau manifolds in weighted \mathbf{P}_4* , *Nuclear Phys. B* **341** (1990), no. 2, 383–402. MR MR1067295 (91m:14062)
13. Michael Carl, Max Pumperla, and Bernd Siebert, *A tropical view of Landau-Ginzburg models*, In preparation.
14. Kwokwai Chan and Naichung Conan Leung, *Mirror symmetry for toric Fano manifolds via SYZ transformations*, *Adv. Math.* **223** (2010), no. 3, 797–839. MR MR2565550
15. Shiu Yuen Cheng and Shing-Tung Yau, *The real Monge-Ampère equation and affine flat structures*, *Proceedings of the 1980 Beijing Symposium on Differential Geometry and Differential Equations, Vol. 1, 2, 3 (Beijing, 1980)* (Beijing), Science Press, 1982, pp. 339–370. MR MR714338 (85c:53103)
16. Cheol-Hyun Cho and Yong-Geun Oh, *Floer cohomology and disc instantons of Lagrangian torus fibers in Fano toric manifolds*, *Asian J. Math.* **10** (2006), no. 4, 773–814. MR MR2282365 (2007k:53150)
17. Tom Coates, Hiroshi Iritani, and Hsian-Hua Tseng, *Wall-crossings in toric Gromov-Witten theory. I. Crepant examples*, *Geom. Topol.* **13** (2009), no. 5, 2675–2744. MR MR2529944
18. David A. Cox and Sheldon Katz, *Mirror symmetry and algebraic geometry*, *Mathematical Surveys and Monographs*, vol. 68, American Mathematical Society, Providence, RI, 1999. MR MR1677117 (2000d:14048)

19. Pierre Deligne and David Mumford, *The irreducibility of the space of curves of given genus*, Inst. Hautes Études Sci. Publ. Math. (1969), no. 36, 75–109. MR MR0262240 (41 #6850)
20. Antoine Douai and Claude Sabbah, *Gauss-Manin systems, Brieskorn lattices and Frobenius structures. I*, Proceedings of the International Conference in Honor of Frédéric Pham (Nice, 2002), vol. 53, 2003, pp. 1055–1116. MR MR2033510 (2005h:32073)
21. ———, *Gauss-Manin systems, Brieskorn lattices and Frobenius structures. II*, Frobenius manifolds, Aspects Math., E36, Vieweg, Wiesbaden, 2004, pp. 1–18. MR MR2115764 (2006e:32037)
22. Geir Ellingsrud and Stein Arild Strømme, *The number of twisted cubic curves on the general quintic threefold*, Math. Scand. **76** (1995), no. 1, 5–34. MR MR1345086 (96g:14045)
23. Kenji Fukaya, Yong-Geun Oh, Hiroshi Ohta, and Kaoru Ono, *Lagrangian Floer theory on compact toric manifolds II: Bulk deformations*, 2008.
24. ———, *Lagrangian intersection Floer theory: anomaly and obstruction. Part I*, AMS/IP Studies in Advanced Mathematics, vol. 46, American Mathematical Society, Providence, RI, 2009. MR MR2553465
25. ———, *Lagrangian intersection Floer theory: anomaly and obstruction. Part II*, AMS/IP Studies in Advanced Mathematics, vol. 46, American Mathematical Society, Providence, RI, 2009. MR MR2548482
26. ———, *Lagrangian Floer theory on compact toric manifolds. I*, Duke Math. J. **151** (2010), no. 1, 23–174. MR MR2573826
27. William Fulton, *Introduction to toric varieties*, Annals of Mathematics Studies, vol. 131, Princeton University Press, Princeton, NJ, 1993, The William H. Roever Lectures in Geometry. MR MR1234037 (94g:14028)
28. William Fulton and Rahul Pandharipande, *Notes on stable maps and quantum cohomology*, Algebraic geometry—Santa Cruz 1995, Proc. Sympos. Pure Math., vol. 62, Amer. Math. Soc., Providence, RI, 1997, pp. 45–96. MR MR1492534 (98m:14025)
29. Davide Gaiotto, Gregory Moore, and Andrew Neitzke, *Four-dimensional wall-crossing via three-dimensional field theory*, 2008, arXiv:0807.4723.
30. ———, *Wall-crossing, Hitchin systems, and the WKB approximation*, 2009, arXiv:0907.3987.
31. Andreas Gathmann and Hannah Markwig, *Kontsevich’s formula and the WDVV equations in tropical geometry*, Adv. Math. **217** (2008), no. 2, 537–560. MR MR2370275
32. I. M. Gelfand, M. M. Kapranov, and A. V. Zelevinsky, *Discriminants, resultants and multidimensional determinants*, Modern Birkhäuser Classics, Birkhäuser Boston Inc., Boston, MA, 2008, Reprint of the 1994 edition. MR MR2394437 (2009a:14065)
33. Alexander B. Givental, *Homological geometry and mirror symmetry*, Proceedings of the International Congress of Mathematicians, Vol. 1, 2 (Zürich, 1994) (Basel), Birkhäuser, 1995, pp. 472–480. MR MR1403947 (97j:58013)
34. ———, *Equivariant Gromov-Witten invariants*, Internat. Math. Res. Notices (1996), no. 13, 613–663. MR MR1408320 (97e:14015)
35. Tom Graber and Ravi Vakil, *Relative virtual localization and vanishing of tautological classes on moduli spaces of curves*, Duke Math. J. **130** (2005), no. 1, 1–37. MR MR2176546 (2006j:14035)
36. Brian Greene and Ronen Plesser, *Duality in Calabi-Yau moduli space*, Nuclear Phys. B **338** (1990), no. 1, 15–37. MR MR1059831 (91h:32018)
37. Mark Gross, *Special Lagrangian fibrations. I. Topology*, Winter School on Mirror Symmetry, Vector Bundles and Lagrangian Submanifolds (Cambridge, MA, 1999), AMS/IP Stud. Adv. Math., vol. 23, Amer. Math. Soc., Providence, RI, 2001, pp. 65–93. MR MR1876066
38. ———, *Special Lagrangian fibrations. II. Geometry. A survey of techniques in the study of special Lagrangian fibrations*, Winter School on Mirror Symmetry, Vector Bundles and Lagrangian Submanifolds (Cambridge, MA, 1999), AMS/IP Stud. Adv. Math., vol. 23, Amer. Math. Soc., Providence, RI, 2001, pp. 95–150. MR MR1876067
39. ———, *Topological mirror symmetry*, Invent. Math. **144** (2001), no. 1, 75–137. MR MR1821145 (2002c:14062)
40. ———, *Toric degenerations and Batyrev-Borisov duality*, Math. Ann. **333** (2005), no. 3, 645–688. MR MR2198802

41. ———, *The Strominger-Yau-Zaslow conjecture: from torus fibrations to degenerations*, Algebraic geometry—Seattle 2005. Part 1, Proc. Sympos. Pure Math., vol. 80, Amer. Math. Soc., Providence, RI, 2009, pp. 149–192. MR MR2483935
42. Mark Gross, *Mirror symmetry for \mathbb{P}^2 and tropical geometry*, Adv. Math. **224** (2010), 169–245.
43. Mark Gross, Daniel Huybrechts, and Dominic Joyce, *Calabi-Yau manifolds and related geometries*, Universitext, Springer-Verlag, Berlin, 2003, Lectures from the Summer School held in Nordfjordeid, June 2001. MR MR1963559 (2004c:14075)
44. Mark Gross and Rahul Pandharipande, *Quivers, curves, and the tropical vertex*, Port. Math. **67** (2010), 211–259.
45. Mark Gross, Rahul Pandharipande, and Bernd Siebert, *The tropical vertex*, Duke Math. J. **153** (2010), 297–362.
46. Mark Gross and Bernd Siebert, *Logarithmic Gromov-Witten invariants*, To appear.
47. ———, *Affine manifolds, log structures, and mirror symmetry*, Turkish J. Math. **27** (2003), no. 1, 33–60. MR MR1975331 (2004g:14041)
48. ———, *Mirror symmetry via logarithmic degeneration data. I*, J. Differential Geom. **72** (2006), no. 2, 169–338. MR MR2213573
49. ———, *From real affine geometry to complex geometry*, arXiv:math/0703822, 2007.
50. ———, *An invitation to toric degenerations*, 2008, arXiv:0808.2749.
51. ———, *Mirror symmetry via logarithmic degeneration data II*, J. Algebraic Geom. **19** (2010), 679–780.
52. Alexander Grothendieck, *Éléments de géométrie algébrique. III. Étude cohomologique des faisceaux cohérents. II*, Inst. Hautes Études Sci. Publ. Math. (1963), no. 17, 91. MR MR0163911 (29 #1210)
53. ———, *Éléments de géométrie algébrique. IV. Étude locale des schémas et des morphismes de schémas IV*, Inst. Hautes Études Sci. Publ. Math. (1967), no. 32, 361. MR MR0238860 (39 #220)
54. Christian Haase and Ilia Zharkov, *Integral affine structures on spheres and torus fibrations of Calabi-Yau toric hypersurfaces I*, arXiv:math/0205321, 2002.
55. ———, *Integral affine structures on spheres: complete intersections*, Int. Math. Res. Not. (2005), no. 51, 3153–3167. MR MR2187503 (2006g:14063)
56. Joe Harris and Ian Morrison, *Moduli of curves*, Graduate Texts in Mathematics, vol. 187, Springer-Verlag, New York, 1998. MR MR1631825 (99g:14031)
57. Robin Hartshorne, *Algebraic geometry*, Springer-Verlag, New York, 1977, Graduate Texts in Mathematics, No. 52. MR MR0463157 (57 #3116)
58. Kentaro Hori and Cumrun Vafa, *Mirror symmetry*, arXiv:hep-th/0002222, 2000.
59. Luc Illusie, *Logarithmic spaces (according to K. Kato)*, Barsotti Symposium in Algebraic Geometry (Abano Terme, 1991), Perspect. Math., vol. 15, Academic Press, San Diego, CA, 1994, pp. 183–203. MR MR1307397 (95j:14023)
60. Kenneth Intriligator, *Fusion residues*, Modern Phys. Lett. A **6** (1991), no. 38, 3543–3556. MR MR1138873 (92k:81180)
61. Hiroshi Iritani, *An integral structure in quantum cohomology and mirror symmetry for toric orbifolds*, Adv. Math. **222** (2009), no. 3, 1016–1079. MR MR2553377
62. Ilia Itenberg, Grigory Mikhalkin, and Eugenii Shustin, *Tropical algebraic geometry*, second ed., Oberwolfach Seminars, vol. 35, Birkhäuser Verlag, Basel, 2009. MR 2508011 (2010d:14086)
63. Fumiharu Kato, *Log smooth deformation theory*, Tohoku Math. J. (2) **48** (1996), no. 3, 317–354. MR MR1404507 (99a:14012)
64. ———, *Log smooth deformation and moduli of log smooth curves*, Internat. J. Math. **11** (2000), no. 2, 215–232. MR MR1754621 (2001d:14016)
65. Kazuya Kato, *Logarithmic structures of Fontaine-Illusie*, Algebraic analysis, geometry, and number theory (Baltimore, MD, 1988), Johns Hopkins Univ. Press, Baltimore, MD, 1989, pp. 191–224. MR MR1463703 (99b:14020)
66. Sheldon Katz, *On the finiteness of rational curves on quintic threefolds*, Compositio Math. **60** (1986), no. 2, 151–162. MR MR868135 (88a:14047)
67. L. Katzarkov, M. Kontsevich, and T. Pantev, *Hodge theoretic aspects of mirror symmetry*, From Hodge theory to integrability and TQFT tt*-geometry, Proc. Sympos. Pure Math., vol. 78, Amer. Math. Soc., Providence, RI, 2008, pp. 87–174. MR 2483750 (2009j:14052)

68. Maxim Kontsevich, *Homological algebra of mirror symmetry*, Proceedings of the International Congress of Mathematicians, Vol. 1, 2 (Zürich, 1994) (Basel), Birkhäuser, 1995, pp. 120–139. MR MR1403918 (97f:32040)
69. Maxim Kontsevich and Yan Soibelman, *Homological mirror symmetry and torus fibrations*, Symplectic geometry and mirror symmetry (Seoul, 2000), World Sci. Publishing, River Edge, NJ, 2001, pp. 203–263. MR MR1882331 (2003c:32025)
70. ———, *Affine structures and non-Archimedean analytic spaces*, The unity of mathematics, Progr. Math., vol. 244, Birkhäuser Boston, Boston, MA, 2006, pp. 321–385. MR MR2181810 (2006j:14054)
71. ———, *Stability structures, motivic Donaldson–Thomas invariants and cluster transformations*, (2008), arXiv:0811.2435.
72. Jun Li, *Stable morphisms to singular schemes and relative stable morphisms*, J. Differential Geom. **57** (2001), no. 3, 509–578. MR MR1882667 (2003d:14066)
73. ———, *A degeneration formula of GW-invariants*, J. Differential Geom. **60** (2002), no. 2, 199–293. MR MR1938113 (2004k:14096)
74. Jun Li and Gang Tian, *Virtual moduli cycles and Gromov–Witten invariants of algebraic varieties*, J. Amer. Math. Soc. **11** (1998), no. 1, 119–174. MR MR1467172 (99d:14011)
75. Bong H. Lian, Kefeng Liu, and Shing-Tung Yau, *Mirror principle. I*, Asian J. Math. **1** (1997), no. 4, 729–763. MR MR1621573 (99e:14062)
76. Yuri I. Manin, *Frobenius manifolds, quantum cohomology, and moduli spaces*, American Mathematical Society Colloquium Publications, vol. 47, American Mathematical Society, Providence, RI, 1999. MR MR1702284 (2001g:53156)
77. Hannah Markwig and Johannes Rau, *Tropical descendant Gromov–Witten invariants*, Manuscripta Math. **129** (2009), no. 3, 293–335. MR MR2515486
78. Hideyuki Matsumura, *Commutative ring theory*, second ed., Cambridge Studies in Advanced Mathematics, vol. 8, Cambridge University Press, Cambridge, 1989, Translated from the Japanese by M. Reid. MR MR1011461 (90i:13001)
79. Grigory Mikhalkin, *Counting curves via lattice paths in polygons*, C. R. Math. Acad. Sci. Paris **336** (2003), no. 8, 629–634. MR MR1988122 (2004d:14077)
80. ———, *Enumerative tropical algebraic geometry in \mathbb{R}^2* , J. Amer. Math. Soc. **18** (2005), no. 2, 313–377 (electronic). MR MR2137980 (2006b:14097)
81. ———, *Introduction to tropical geometry, notes from the IMPA lectures, summer 2007*, arXiv:0709.1049, 2007.
82. Grigory Mikhalkin and Ilia Zharkov, *Tropical curves, their Jacobians and theta functions*, **465** (2008), 203–230. MR MR2457739
83. James S. Milne, *Étale cohomology*, Princeton Mathematical Series, vol. 33, Princeton University Press, Princeton, N.J., 1980. MR MR559531 (81j:14002)
84. David Mumford, *An analytic construction of degenerating abelian varieties over complete rings*, Compositio Math. **24** (1972), 239–272. MR MR0352106 (50 #4593)
85. Takeo Nishinou, *Disc counting on toric varieties via tropical curves*, 2006.
86. Takeo Nishinou and Bernd Siebert, *Toric degenerations of toric varieties and tropical curves*, Duke Math. J. **135** (2006), no. 1, 1–51. MR MR2259922 (2007h:14083)
87. Tadao Oda, *Convex bodies and algebraic geometry*, Ergebnisse der Mathematik und ihrer Grenzgebiete (3) [Results in Mathematics and Related Areas (3)], vol. 15, Springer-Verlag, Berlin, 1988, An introduction to the theory of toric varieties, Translated from the Japanese. MR 922894 (88m:14038)
88. Arthur Ogus, *Lectures on logarithmic algebraic geometry*, In preparation.
89. Arthur Ogus and Vadim Vologodsky, *Nonabelian Hodge theory in characteristic p* , Publ. Math. Inst. Hautes Études Sci. (2007), no. 106, 1–138. MR MR2373230 (2009e:14033)
90. Rahul Pandharipande, *Rational curves on hypersurfaces (after A. Givental)*, Astérisque (1998), no. 252, Exp. No. 848, 5, 307–340, Séminaire Bourbaki. Vol. 1997/98. MR MR1685628 (2000e:14094)
91. Brett Parker, *Exploded fibrations*, Proceedings of Gökova Geometry-Topology Conference 2006, Gökova Geometry/Topology Conference (GGT), Gökova, 2007, pp. 52–90. MR MR2404950 (2009h:53202)
92. ———, *Exploded manifolds*, arXiv:0910.4201, 2009.
93. ———, *Holomorphic curves in exploded manifolds: compactness*, arXiv:0911.2241, 2009.
94. Markus Reineke, *Poisson automorphisms and quiver moduli*, 2008, arXiv:0804.3214.

95. ———, *Cohomology of quiver moduli, functional equations, and integrality of Donaldson-Thomas type invariants*, 2009, arXiv:0903.0261.
96. Jürgen Richter-Gebert, Bernd Sturmfels, and Thorsten Theobald, *First steps in tropical geometry*, Idempotent mathematics and mathematical physics, Contemp. Math., vol. 377, Amer. Math. Soc., Providence, RI, 2005, pp. 289–317. MR MR2149011 (2006d:14073)
97. Wei-Dong Ruan, *Lagrangian torus fibration of quintic hypersurfaces. I. Fermat quintic case*, Winter School on Mirror Symmetry, Vector Bundles and Lagrangian Submanifolds (Cambridge, MA, 1999), AMS/IP Stud. Adv. Math., vol. 23, Amer. Math. Soc., Providence, RI, 2001, pp. 297–332. MR MR1876075 (2002m:32041)
98. ———, *Lagrangian torus fibration of quintic Calabi-Yau hypersurfaces. II. Technical results on gradient flow construction*, J. Symplectic Geom. **1** (2002), no. 3, 435–521. MR MR1959057 (2004b:32040)
99. ———, *Lagrangian torus fibration of quintic Calabi-Yau hypersurfaces. III. Symplectic topological SYZ mirror construction for general quintics*, J. Differential Geom. **63** (2003), no. 2, 171–229. MR MR2015547 (2004k:32043)
100. Yongbin Ruan, *Topological sigma model and Donaldson-type invariants in Gromov theory*, Duke Math. J. **83** (1996), no. 2, 461–500. MR MR1390655 (97d:58042)
101. Yongbin Ruan and Gang Tian, *A mathematical theory of quantum cohomology*, J. Differential Geom. **42** (1995), no. 2, 259–367. MR MR1366548 (96m:58033)
102. Claude Sabbah, *On a twisted de Rham complex*, Tohoku Math. J. (2) **51** (1999), no. 1, 125–140. MR MR1671743 (2000c:14014)
103. ———, *Isomonodromic deformations and Frobenius manifolds*, Universitext, Springer-Verlag London Ltd., London, 2007, An introduction. MR MR2368364 (2008i:32015)
104. Kyoji Saito, *Period mapping associated to a primitive form*, Publ. Res. Inst. Math. Sci. **19** (1983), no. 3, 1231–1264. MR MR723468 (85h:32034)
105. Eugenii Shustin, *Patchworking singular algebraic curves, non-archimedean amoebas and enumerative geometry*, 2002, arXiv:math/0209043.
106. David Speyer, *Uniformizing tropical curves I: genus zero and one*, arXiv:0711.2677, 2007.
107. Joseph Steenbrink, *Limits of Hodge structures*, Invent. Math. **31** (1975/76), no. 3, 229–257. MR MR0429885 (55 #2894)
108. Andrew Strominger, Shing-Tung Yau, and Eric Zaslow, *Mirror symmetry is T-duality*, Nuclear Phys. B **479** (1996), no. 1-2, 243–259. MR MR1429831 (97j:32022)
109. Bernard Teissier, *The hunting of invariants in the geometry of discriminants*, Real and complex singularities (Proc. Ninth Nordic Summer School/NAVF Sympos. Math., Oslo, 1976), Sijthoff and Noordhoff, Alphen aan den Rijn, 1977, pp. 565–678. MR MR0568901 (58 #27964)
110. Cumrun Vafa, *Topological mirrors and quantum rings*, Essays on mirror manifolds, Int. Press, Hong Kong, 1992, pp. 96–119. MR MR1191421 (94c:81193)
111. Magnus Dehli Vigeland, *Tropical lines on smooth tropical surfaces*, arXiv:07083847, 2007.
112. Edward Witten, *Two-dimensional gravity and intersection theory on moduli space*, Surveys in differential geometry (Cambridge, MA, 1990), Lehigh Univ., Bethlehem, PA, 1991, pp. 243–310. MR MR1144529 (93e:32028)
113. Shing Tung Yau, *On the Ricci curvature of a compact Kähler manifold and the complex Monge-Ampère equation. I*, Comm. Pure Appl. Math. **31** (1978), no. 3, 339–411. MR MR480350 (81d:53045)

Index of Symbols

- $(\cdot, \cdot)_{\mathcal{E}}$, 55
 \circ , 43
 $*$, 40

 A , 60
 $\mathcal{A}ff(B, \mathbb{R})$, 28
 $\mathcal{A}ff(B, \mathbb{Z})$, 28
 $\text{Aff}(M)$, 20
 $\text{Aff}(M_{\mathbb{R}})$, 20
 $\text{Aff}(\mathbb{R}^n)$, xiii
 $\text{Aff}(\mathbb{Z}^n)$, xiii
 $\text{Asym}(\Delta)$, 96
 A_X , 60

 $B_i(d)$, 222
 (B, \mathcal{P}) , 24
 $(\tilde{B}, \tilde{\mathcal{P}}, \tilde{\varphi})$, 28

 $C(\Delta)$, 96
 $\text{Chambers}(\mathcal{S}, k)$, 276
 $\mathbb{C}\{\hbar\}$, 51
 $\mathbb{C}\{\hbar, \hbar^{-1}\}$, 51

 \mathfrak{D} , 147
 \mathfrak{d} , 271
 $\partial^- E, \partial^+ E$, 147
 (D, Dlog) , 112
 Δ_{φ} , 12
 $D_i(d, m)$, 225
 $D_i(d, n_0, n_1, n_2)$, 223
 $\mathfrak{D}(\Sigma, P_1, \dots, P_k)$, 201
 $\partial\sigma$, 4
 D_{τ} , 136
 ∂_X , 106

 F_k , 277
 $f^* \mathcal{M}_Y$, 100

 \mathbf{G} , 298
 Γ , 13
 $\bar{\Gamma}$, 13
 Γ_f , 142
 $\tilde{\Gamma}_f$, 141
 $\hat{\Gamma}$, 147
 $\mathbb{G}(L)$, 95
 $\text{Glue}(\mathcal{S}, k)$, 276
 Gr_E , 46

 \mathcal{H} , 57
 $\mathcal{H}_{\mathcal{M}}$, 60
 \mathcal{H}_- , 59

 $I_{\mathfrak{d}}$, 271
 $I_{\mathfrak{d}}^{\text{init}}$, 275
 $I_{\omega, \tau}^k$, 258
 $\text{Init}(\mathfrak{d})$, 198
 $\text{Initial}(\mathcal{S})$, 275
 $\text{Int}(\sigma)$, 4

 \mathbb{J} , 58
 $\text{Jac}(W)$, 69
 $\text{Joints}(\mathcal{S}, k)$, 281
 $J_{\mathbb{P}^2}^{\text{trop}}$, 185
 J_X , 66

 K , 227
 \mathbb{k} , xvi
 K_d , 228
 $\mathbb{k}[P]$, 292

 Λ , xiii
 $\hat{\Lambda}$, xiii
 $L_{\mathfrak{d}}$, 271
 $L_{i, \gamma, \omega \rightarrow \tau}^d$, 231
 $L_{i, \sigma}^d$, 228
 L_i^d , 226
 $\mathcal{L}_{(\mathbf{u}, \mathbf{w})}$, 144

 M , 4
 \tilde{m} , 263
 $\tilde{m}_{\mathfrak{d}}$, 271
 $\overline{\mathcal{M}}_{g, n}$, 33
 $\overline{\mathcal{M}}_{g, n}(X, \beta)$, 36
 $[\overline{\mathcal{M}}_{g, n}(X, \beta)]^{\text{vir}}$, 37
 $MI(h)$, 176
 $\text{Mono}(h)$, 176
 $M_{\mathbb{R}}$, 4
 \mathfrak{m}_{R_k} , 196
 \widetilde{M} , 135
 \widehat{M} , 51, 75
 $\mathcal{M}_{\Delta, n}^{\text{trop}}(X_{\Sigma}, P_1, \dots, P_k, \psi^{\nu} S)$, 180
 $\text{Mult}_V(h)$, 19
 \mathcal{M}_X , 98
 m_x , 273

- N , 4
 \mathbb{N} , xvi
 $N_{\Delta, \Sigma}^{0, \text{hol}}$, 134
 $\widehat{\nabla}_X(Y)$, 46
 $N_{\Delta}(\sigma)$, 12
 $N_{\mathfrak{o}}^k$, 276
 $N_{\mathbf{G}}$, 298
 $N_{\mathbf{P}}$, 297
 $N_{\mathbb{R}}$, 4
 \widetilde{N} , 135

 Ω , 177
 $\Omega_{X^\dagger/S^\dagger}^1$, 114
 $\overline{\Omega}$, 57
 $\mathcal{O}_{\widetilde{\mathcal{M}}}\{\hbar\}$, $\mathcal{O}_{\widetilde{\mathcal{M}}}\{\hbar, \hbar^{-1}\}$, 51
 $\mathcal{O}(\mathbb{P}^1 \setminus \{0\})$, 51
 ord_m , 274
 $\text{ord}_\sigma(p)$, 258

 \mathcal{P} , 6
 \underline{P} , 106
 \mathbf{P} , 296
 $\langle P_1, \dots, P_{3d-2-\nu}, \psi^\nu S \rangle_{0,d}^{\text{trop}}$, 182
 $\langle P_1, \dots, P_{3d-\nu-(2-\dim S)}, \psi^\nu S \rangle_{d,\sigma}^{\text{trop}}$, 184
 $\mathcal{P}_{\max}^\partial$, 258
 Φ , 147
 φ_Δ , 12
 $\varphi(\tau)$, 12
 $\mathcal{P}^{[k]}$, 7
 \mathcal{P}_k , 276
 $\mathcal{P}\mathcal{L}_{\mathcal{P}}(B, \mathbb{R})$, 28
 $\mathcal{P}\mathcal{L}_{\mathcal{P}}(B, \mathbb{Z})$, 28
 \mathcal{P}_{\max} , 7
 \mathcal{P}_φ , 263
 $\psi_i(n_0, n_1, n_2)$, 223
 $\psi_i(m)$, 225
 P_τ , 257

 R_k , 175
 $R_{\omega, \tau}^k$, 258
 R_ω , 258
 \mathbb{R}^{trop} , 3

 \mathcal{S} , 274
 $\mathbf{S}(\mathfrak{D})$, 293
 S_e , 110
 $\mathbf{S}_I(\mathfrak{D})$, 293
 $\tilde{\sigma}$, 227
 $\tilde{\sigma}_d$, 228
 $\tilde{\Sigma}_\Delta$, 12
 σ_j , 281
 $\Sigma(\tau)$, 12
 Σ_τ , 22
 σ_u , 276
 σ^\vee , 29
 $\text{Sing}(\mathfrak{D})$, 199
 S_k , 225
 $\mathcal{S}[k]$, 274

 $\text{Spec } k^\dagger$, 100
 $\text{Supp}(\mathfrak{D})$, 198

 $\langle T_2^m, \psi^\nu T_i \rangle_{0,d}^{\text{trop}}$, 185
 τ , 60
 $\theta_{\gamma, \mathfrak{D}}$, 199
 $\Theta(k[P])$, 291
 $\Theta_{X^\dagger/S^\dagger}$, 112
 $\text{Trees}(\Sigma, P_1, \dots, P_k)$, 201
 $T_\tau\sigma$, 12

 $u_{(\partial^- E, E)}$, 147
 U_ω^k , 283

 $\text{Val}(x)$, 180
 $V(f)$, 3
 \widehat{V} , 292
 V_I , 291
 $\mathbb{V}_{\Sigma, k}$, 196
 $\mathfrak{v}_{\Sigma, k}$, 196
 V_τ , 285

 $W_k(Q)$, 177

 $X(B)$, xiii
 $\tilde{X}(B)$, xiii
 X_f , 197
 X_σ , 93
 X_Σ , 93

General Index

- affine length, 8, 16
- affine linear, 28
- affine linear automorphisms, 20
- affine manifold, xiii
 - tropical, with singularities, 20
 - integral, 20
 - integral, with singularities, 20
 - tropical, 20
- balancing condition, 9, 13
- Batyrev-Borisov construction, 26
- big torus, 95
- boundary, 4
- broken line, 204, 215
 - deformation of, 207
 - degenerate, 207
 - family of, 207
 - local, 213
- Calabi-Yau manifold, xi, 26, 37, 247
- cell, 6
- change of chamber map, 278
- change of strata map, 277
- chart, 106
- combinatorial type, 16
- compatibility, 277
- compatible
 - fan structure, 24
 - structure, 281
- complete, 11
- cone picture, 252
- consistent, 281
- convex
 - multi-valued PL function, 28
- cubic surface, 32
- degree
 - of a tropical disk, 175
- Dilaton Axiom, 39
- discrete Legendre transform, 7, 28, 253, 254
- discriminant locus, 20
- Divisor Axiom, 39
- double point, 109
- Dubrovin connection, 46
- equivalent charts, 106
- equivalent scattering diagrams, 294
- étale topology, 99
- Euler vector field, 45
- expected dimension, 36
- face, 4
- fan, 11
- fan picture, 248
- fan structure, 22
- first structure connection, 46
- flag, 13
- flat coordinates, 64, 87
- Frobenius manifold, 43
 - identity on, 44
- Fundamental Class Axiom, 39
- geometric point, 99
- ghost sheaf, 101
- Givental J -function, 66
- good decomposition, 134
- graded partition, 298
- grading, 55
- graph, 13
 - marked, 13
- gravitational descendent invariants, 38
- Gromov-Witten invariants, xii, xiv, 37
 - descendent, 38
- Gromov-Witten potential, 40
- homogeneous, 45
- Homological Mirror Symmetry, xii
- homomorphism
 - integral, 109
- index, 4
- indistinguishable, 141
- initial point, 198, 292
- initial ray, 275
- integral, 12
 - homomorphism of monoids, 109
 - monoid, 92
 - morphism of log schemes, 109
- integral affine linear, 28
- integral tropical manifold, 24
- interior, 4
- interstice, 215
- Jacobian ideal, 69

- joint, 215, 281
- Landau-Ginzburg
 - model, 67
 - potential, 67
 - potential, k -pointed, 177
- Lefschetz thimble, 72
- line, 144, 198, 292
 - bivalent, 144
 - trivalent, 144
- log deformation theory, 117
- log derivation, 112, 291
- log differentials, 114
- log geometry, xiv
- log Gromov-Witten invariants, 250
- log marked point, 109
- log scheme, 98
 - morphism of, 99
- log smooth curve, 109
- log smooth morphism, 107
- log structure, 98
 - associated to a pre-log structure, 100
 - chart for, 106
 - divisorial, 100
 - fine, 106
 - fine saturated, 106
 - pull-back, 100
 - trivial, 100
- logarithmic differentials, 91
- logarithmic normal sheaf, 121
- map of fans, 94
- Maslov index
 - of a tropical disk, 176
 - of a tropical tree, 201
- Milnor ring, 69
- miniversal, 62
- mirror map, 87
- mirror of a toric Fano variety, 174
- mirror symmetry, xi
 - for \mathbb{P}^n , 87
- monoid, 92
 - fibre co-product of, 93
 - finitely generated, 92
 - integral, 92
 - saturated, 106
- multiplicity, 19
 - of a tropical disk, 175
- Mumford degeneration, 97
- naked ray, 271
- Newton polyhedron, 12
- normal cone, 12
- normal crossings, 91
- normal fan, 12
- open star, 22
- opposite subspace, 59
- order, 258
- ordered partition, 296
- oscillatory integral, 70
- overvalence, 16
- penguins
 - complete lack of, 185
- PL function, 28
 - multi-valued, 28
- Point Mapping Axiom, 39
- polyhedral decomposition, 6, 21
- polyhedron, 4
- polytope, 4
- positive tropical manifold, 255
- pre-Frobenius structure, 43
- pre-log structure, 98
- Pressley-Segal Grassmannian, 57
- primitive, 4
- quantum cohomology, xii, 40
 - of \mathbb{P}^2 , 42
- quantum differential equation, 47
- quintic threefold, xi
- quotient fan, 12
- ray, 198, 271, 292
- reflexive polytope, 24
- regular decomposition, 7
- regular singular point, 54
- relative stable maps, 296
- scattering diagram, 198, 292
- scheme-theoretically trivial, 151
- semi-infinite variation of Hodge structure, 55
- sheaf of log derivations, 112
- simple curve, 17
- simple tropical manifold, 255
- singular locus, 20
- smooth point, 109
- stable curve
 - n -pointed, 33
- stable manifold, 73
- stable map
 - n -pointed, 35
- stable reduction, 36
- stack
 - Deligne-Mumford, 33
- standard cone, 93
- standard log point, 100
- stationary phase approximation, 73
- strict, 107
- strictly convex, 11
- strictly convex rational polyhedral cone, 11
- string theory, xi
- structure, 271, 274
- superabundant, 17
- support, 11, 198, 292
- SYZ conjecture, xii, 247

- tangent wedge, 12
- Tate curve, 268
- Topological Recursion Relation, 47
- toric boundary, 106
- toric stratum, 95
- toric variety, 93
 - affine, 93
- torically transverse, 133
- torically transverse log curve, 138
- torically transverse pre-log curve, 147
- tropical
 - J -function, 185
 - affine manifold, 20
 - Bézout theorem, 10
 - curve in a tropical manifold, 31
 - descendent invariants, 182
 - disk, 174
 - disk in X_Σ , 175
 - hypersurface, 3
 - Jacobians, 20
 - manifold, 24
 - marked parametrized curve, 13
 - regular curve, 17
 - semi-ring, 3
 - simple curve, 17
 - tree, 200
- tropical vertex group, 290
- tropicalization, 248
- twisted de Rham complex, 68

- universal unfolding, 74
- unstable manifold, 73

- Vafa-Intriligator formula, 43
- virtual dimension, 37
- virtual fundamental class, 37

- wall-crossing, 192, 208, 231
- WDVV equation, 41
- weight, 4