

Mirror Symmetry and Cluster Varieties

Travis Mandel

August 5, 2017

(In progress.) These mostly expository notes are based on a class taught by the author on mirror symmetry and cluster algebras, particularly on the Gross-Hacking-Keel-Kontsevich construction [GHKK] of canonical bases of theta functions for cluster algebras.

Contents

1	Introduction	3
1.1	Historical context	3
1.2	Outline	4
1.3	Acknowledgements	6
2	Toric Varieties	7
2.1	Constructing an atlas for a toric variety from a fan	7
2.2	Cones correspond to torus orbits	8
2.3	Singularities	8
2.4	Maps of fans	9
2.5	The quotient construction	9
2.6	The Dual Pairing	10
2.7	Piecewise Linear Functions and Polytopes	11
2.8	Projective toric varieties from convex polytopes	13
2.9	Moment maps	14
2.10	Toric degenerations	15
3	Motivation from Mirror Symmetry	16
3.1	Mirror symmetry and T-duality in the complement of an anticanonical divisor	16
3.2	The Gross-Siebert Program	19
4	Mirror Symmetry for Log Calabi-Yau Surfaces	23
4.1	Log Calabi-Yau Surfaces with Maximal Boundary	23
4.2	Tropicalizing U	24

4.3	The Mumford Degeneration	25
4.4	The consistent scattering diagram	27
4.5	Broken Lines and Theta Functions	30
4.6	Compactifications	31
4.7	More scattering diagrams	32
5	Cluster Algebras and Cluster Varieties	36
5.1	Some Motivation	36
5.2	Basic Definitions	38
5.3	Other Structures on Cluster Varieties	42
5.4	Cluster varieties with principal coefficients	46
6	Canonical Bases and Cluster Varieties	49
6.1	The Fock-Goncharov Conjecture	49
6.2	The Naive Initial Scattering Diagrams	50
6.3	The relation to the log Calabi-Yau surface constructions	52
6.4	Using Principal Coefficients	56
7	Examples and Applications of Cluster Algebras	57
7.1	Cluster structures on double Bruhat cells of semisimple groups	57
7.2	Cluster varieties and surfaces	57
A	Some Algebraic Geometry Background	62
A.1	Divisors	62
A.2	Formal schemes	63

Chapter 1

Introduction

These notes are based on a course taught by the author at the Center for Quantum Geometry of Moduli Spaces in Fall, 2014. The constructions are algebro-geometric, but students in the class are not assumed to have a strong algebraic geometry background. As such, these notes will introduce most of the relevant algebraic geometry background as needed or in the appendix. The main goal is to explain the [GHKK] construction of dual canonical bases for cluster varieties, including the motivation from mirror symmetry.

1.1 Historical context

Motivated by ideas from string theory, mirror symmetry roughly relates the symplectic geometry of one Calabi-Yau manifold Y (the A-model) to the algebraic geometry of another Calabi-Yau manifold X (the B-model), and vice versa. In 1991, [CdLOGP92] motivated mathematical interest in mirror symmetry by using B-model data (i.e., period integrals on X) to predict certain Gromov-Witten invariants (i.e., A-model data) of Y . The first solid mathematical description of mirror symmetry was given in 1994 by Kontsevich in [Kon95], where he defined mirror symmetry as a derived equivalence between the Fukaya category $\mathrm{Fuk}(Y)$ and the bounded derived category of coherent sheaves $D^b\mathrm{Coh}(X)$ of X . We will not discuss this viewpoint, but it is worth mentioning that the existence of certain canonical bases of global regular functions (called theta functions) on X can, under certain conditions, be realized as a consequence of this—intersection points of Lagrangians form canonical bases for the Hom-spaces in $\mathrm{Fuk}(Y)$, and the theta functions are supposed to be mirror to such a basis.

In 1996, [SYZ96] gave another, more constructive interpretation of mirror symmetry, in which Y and X are seen to roughly admit dual special Lagrangian torus fibrations over a common affine manifold \mathcal{B} . Our interest is in the Gross-Siebert program,¹ which uses a tropicalized version of the SYZ-viewpoint to create a very hands-on relationship between Y and X . In particular, in [GHK11], Gross, Hacking, and Keel use the techniques of Gross and Siebert to define canonical “theta functions” on X in terms of tropical disks in the tropicalization of Y . This is the piece of mirror symmetry in which we will be interested.

The other side of our story is that of cluster algebras. These combinatorially defined algebras were introduced in 2001 by Fomin and Zelevinsky [FZ02] with the hope of better understanding

¹The Gross-Siebert program has been developed in [GS03], [GS06], and [GS11a]. A good overview is given in [Gro13].

Lusztig’s dual canonical bases [Lus90] and total positivity [Lus94]. Since then, many objects of interest have been found to admit cluster structures, including (but not limited to) double Bruhat cells of semisimple Lie groups [BFZ05], Grassmanians [Sco06] and more general (partial) flag varieties [GLS08], and generalizations of Teichmüller space [FG06]. In addition to giving special coordinate systems that help with understanding positivity and canonical bases, these cluster structures are useful for studying Poisson geometry and generalizations of Weil-Petersson forms ([GSV05],[FG06],[GSV10]), tropicalization [FG09], and quantum deformations ([BZ05],[FG09]).

In [FG09], Fock and Goncharov formalized the framework for a geometric approach to cluster algebras (i.e., cluster varieties). They also conjectured the existence of canonical bases of global regular functions on cluster varieties, parameterized by certain tropical points of their Langland’s duals. Gross, Hacking, and Keel reinterpreted cluster varieties from the viewpoint of birational geometry, realizing that cluster varieties are an ideal candidate for beginning to generalize their log Calabi-Yau surface constructions [GHK11] to higher dimensions. Together with Kontsevich in [GHKK], they used their techniques to construct the bases conjectured by Fock and Goncharov (to the extent that they actually exist).² In the process, they prove other significant conjectures from cluster theory, including the positive Laurent phenomenon³. Understanding this construction of [GHKK] is the driving goal of these notes.

1.2 Outline

In **Chapter 2**, we give an overview of the theory of toric varieties. This is the main tool for the Gross-Siebert program, and the chapters that follow will frequently rely on this background. As mentioned above, the reader is not assumed to have much algebraic geometry background, so this background is introduced as needed or in Appendix A. Toric varieties may be viewed as trivial cluster varieties (i.e., cluster varieties with no non-trivial mutations), and the canonical bases here are simply the usual bases of monomials.

Chapter 3 introduces the ideas from mirror symmetry that motivate the Gross-Hacking-Keel construction. This chapter is not strictly necessary, but it serves as the main motivation for the constructions of Chapters 4 and 6. We first follow [Aur] to explain the SYZ picture of mirror symmetry. Very briefly, let Y be a smooth compact Kähler manifold with Kähler form ω and non-vanishing holomorphic n -form Ω . Choose an anti-canonical divisor D (i.e. Ω has simple poles along D and is a holomorphic volume form on $Y \setminus D$), and a special Lagrangian torus $L_0 \subset Y$. The mirror X to the pair (Y, D) can be viewed as the (quantum corrected) moduli space of special Lagrangian deformations of L_0 , decorated with the class of a trivial flat $U(1)$ -connection ∇ . X has local holomorphic coordinates given by

$$z^\Gamma(L, \nabla) := \exp\left(-\int_\Gamma \omega\right) \text{hol}_\nabla(\partial\Gamma),$$

²[GHK13a] showed that the bases conjectured by Fock and Goncharov cannot always exist. From the mirror symmetry viewpoint, the bases may only be defined in a formal neighborhood of a large complex structure limit. [GHKK] does construct formal versions of these bases in general and gives nearly optimal conditions under which these extend to algebraic bases.

³The positive Laurent phenomenon was previously proved by very different means in [LS14] for all skew-symmetric cluster algebras. [GHKK] does not assume skew-symmetry, but does assume geometric type.

where $\Gamma \in H_2(X, L, \mathbb{Z})$ with $\partial\Gamma \neq 0 \in H_1(L, \mathbb{Z})$, and ω is the Kähler form on X . One can then hope to define the “superpotential” (and, by a slight modification, the theta functions) W on X by

$$W(L, \nabla) := \sum_{\beta, \beta \cdot D=1} n_\beta(L) z_\beta.$$

Here, n_β is a Gromov-Witten count of the holomorphic disks in Y of class β whose boundary passes through a generically specified point of L . The problem with this definition is that defining z^Γ requires keeping track of how Γ deforms as we vary L , and this cannot usually be done globally. Instead, one can define the mirror with these coordinates and this W locally, and then glue via “wall-crossing” automorphisms which take into account so-called instanton/quantum corrections in order to get a global mirror with a well-defined W (and well-defined theta functions).

Next, in §3.2 we sketch the Gross-Siebert approach to SYZ fibrations and the construction the mirror. We mostly follow the much more detailed survey [Gro13] of the Gross-Siebert program. Basically, since finding explicit special Lagrangian fibrations and counting holomorphic disks is often too difficult to actually do in practice, Gross and Siebert describe combinatorial versions of this data. The base of the SYZ fibration is taken to be a tropicalization of the manifold, and counts of (virtual) holomorphic disks are replaced by counts of tropical curves (which [GHK11] then condenses into broken lines).

Chapter 4 is where we actually describe the Gross-Hacking-Keel construction for log Calabi-Yau surfaces, following [GHK11]. This construction is essentially an explicit application of the Gross-Siebert program, although the affine manifolds used in [GHK11] typically have worse singularities than Gross and Siebert usually allow. This change is not too significant though since the singularities can be factored into the type (focus-focus) which Gross and Siebert do consider. If we sacrifice some boundary divisors of the mirror family, we can even remove the singularities completely, as we do in §4.7.3, and as [GHK11] does in their §3. This is the picture we will use when generalizing to cluster varieties (although it should be possible to generalize the singular version too).

In **Chapter 5** we introduce the basic definitions and properties of cluster algebras and cluster varieties, following [FG09]. We also explain the birational geometry viewpoint of [GHKK]. Basically, [GHKK] show that the operation of “mutation” for cluster varieties can be understood as a blowup of a certain “hypertorus” on the boundary of a toric variety, followed by a certain blowdown to another toric variety. Cluster varieties are thus seen to roughly be certain blowups of toric varieties.

Cluster varieties come in pairs \mathcal{A} and \mathcal{X} , where \mathcal{X} is a Poisson manifold, and \mathcal{A} is a manifold with a (possibly degenerate) skew-symmetric form. There is a natural map from \mathcal{A} to a certain symplectic leaf of \mathcal{X} with the skew form being the pull-back of the symplectic form. We will see that all the log Calabi-Yau surfaces from Chapter 4 appear (up to codimension 2) as symplectic leaves of some \mathcal{X} -space, so we will view these symplectic leaves as the thing generalizing the log Calabi-Yau surfaces.

Taking this perspective in **Chapter 6**, we will explain a slightly modified version of the [GHKK] construction as a direct generalization of the [GHK11] construction. This interpretation seems to be new. From this perspective, \mathcal{A} (or at least its tropicalization) plays the role of the symplectic data for a symplectic leaf U of \mathcal{X} , while \mathcal{X} is essentially the moduli space of complex structures. One views \mathcal{A} as mirror to the “Langland’s dual” \mathcal{X} -space \mathcal{X}^\vee , and \mathcal{X} as mirror to \mathcal{A}^\vee . This exchange of symplectic and complex structures is what one expects in mirror symmetry.

Finally, in Chapter 7, we sketch some of the main applications of cluster algebras, including the Fock-Goncharov constructions of higher Teichmüller theory. [I MAY ADD MORE]. This Chapter

could easily precede Chapter 6, and can in fact be read independently of the rest of the notes, except for the basic definitions of cluster algebras from Chapter 5.

1.3 Acknowledgements

Chapter 2

Toric Varieties

My main reference is [Ful93]. Chapter 7 of [HKK⁺03] also has a good introduction from a different, more physics oriented point of view. Another good book that is more recent and covers quite a bit more is [CJS11]. [Gro11] also includes a good brief introduction to toric varieties which is well-tailored to our needs in these notes. I will use this for the toric degenerations section, and I will mention some other references as we go. For the reader unfamiliar with the language of divisors and linear systems in algebraic geometry, see §A.1 for an overview.

2.1 Constructing an atlas for a toric variety from a fan

Definition 2.1.1. A **toric variety** is a complex¹ algebraic variety X containing $T := (\mathbb{C}^*)^r$ as a dense open subset such that the torus action extends to all of X . We'll always assume that X is normal (these are the cases which can be constructed from fans). The complement of the big torus orbit is called the **(toric) boundary**.

Let $N := \mathbb{Z}^n$ be a finite rank lattice (the **cocharacter lattice**) and $M := N^* := \text{Hom}(N, \mathbb{Z}) \cong \mathbb{Z}^n$ the dual lattice (the **character lattice**). We denote the dual pairing by $\langle \cdot, \cdot \rangle : N \oplus M \rightarrow \mathbb{Z}$. For any lattice L and abelian group A , let $L_A := L \otimes A$. We may also denote $T_L := L_{\mathbb{C}^*}$.

For any monoid σ , define

$$\mathbb{C}[\sigma] := \mathbb{C}[z^u \mid u \in \sigma] / \langle z^u \cdot z^v = z^{u+v} \rangle,$$

where the addition in the exponent is the monoid addition.

Given a strictly convex (i.e., not containing any line through the origin) rational polyhedral cone (a cone with apex at the origin generated by a finite number of vectors in N) $\sigma \subset N_{\mathbb{R}}$, let

$$\sigma^\vee := \{m \in M \mid \langle n, m \rangle \geq 0 \ \forall n \in \sigma\}.$$

The monoid σ then determines an affine toric variety $\text{TV}(\sigma) := \text{Spec } \mathbb{C}[\sigma^\vee]$. Alternatively, we may define (the geometric points of) $\text{TV}(\sigma)$ as the space of semigroup homomorphisms $\text{Hom}_{\text{sg}}(\sigma^\vee, \mathbb{C})$

¹The base field can be generalized quite a bit depending on the situation. Most of what we will do works over any algebraically closed characteristic 0 field \mathbb{k} , and sometimes even more generally than this. In fact, many ideas in cluster theory, including the definition of the theta functions, work over \mathbb{Z} , and even over semi-fields.

(recall that a semigroup is a set with an associative binary operation i.e., a monoid without identity, so \mathbb{C} is a semigroup under multiplication).

Example 2.1.2. Let $N = \mathbb{Z}^2$, $\sigma = \langle e_1, e_2 \rangle$, so $\sigma^\vee = \langle e_1^*, e_2^* \rangle$. Then $\mathbb{C}[\sigma^\vee] = \mathbb{C}[x, y]$ where $x := z^{e_1^*}$, $y := z^{e_2^*}$, and $\text{Spec } \mathbb{C}[\sigma^\vee] = \mathbb{A}^2$. We will continue to use this notation in future examples.

Definition 2.1.3. A fan Σ is a set of rational, strictly convex, polyhedral cones in $N_{\mathbb{R}}$ such that each face of a cone in Σ is also a cone in Σ , and the intersection of any two cones is a face in each.

If $\rho \subset \sigma$, then $\sigma^\vee \subset \rho^\vee$, and $\text{Spec } \mathbb{C}[\rho^\vee] \hookrightarrow \text{Spec } \mathbb{C}[\sigma^\vee]$. Thus, if $\rho = \sigma_1 \cap \sigma_2$, then we can glue $\text{Spec } \mathbb{C}[\sigma_1^\vee]$ to $\text{Spec } \mathbb{C}[\sigma_2^\vee]$ along $\text{Spec } \mathbb{C}[\rho^\vee]$. In this way, we construct a toric variety from Σ , denoted $\text{TV}(\Sigma)$.

Example 2.1.4. Every fan Σ contains the origin, and $\{0\}^\vee = M$, so every toric variety contains the algebraic torus $\text{Spec } \mathbb{C}[M] \cong T_N$. This is the torus T required in Definition 2.1.1. Note that we are following the common practice of conflating a toric variety with its geometric points (i.e., saying Spec when we should perhaps be saying $m \text{Spec}$).

Example 2.1.5. Let $N = \mathbb{Z}^n$, and let Σ be the fan consisting of the $n+1$ rays generated by e_1, \dots, e_n and $-(\sum_{i=1}^n e_i)$, along with $\{0\}$ and the cones that these rays bound. Then $\text{TV}(\Sigma) = \mathbb{P}^n$.

2.2 Cones correspond to torus orbits

There is an order-reversing correspondence between cones $\sigma \in \Sigma$ and orbits of the action of T_N on $\text{TV}(\Sigma)$: σ corresponds to the orbit

$$O_\sigma := \text{TV}(\sigma) \setminus \bigcup_{\tau \subsetneq \sigma} \text{TV}(\tau) \cong \text{Spec } \mathbb{C}[\sigma^\perp \cap M].$$

More importantly are the orbit closures, $C_\sigma := \overline{O_\sigma}$. Note that if $\dim(\sigma) = r$, then $\dim(O_\sigma) = \dim(C_\sigma) = n - r$. These orbit closures are called the **toric strata** of $\text{TV}(\Sigma)$. In particular, if $\rho \in \Sigma^{[1]}$ (the set of rays of Σ), then C_ρ is a divisor, called a **boundary divisor** of $\text{TV}(\Sigma)$. For ρ a ray generated by $u \in N$, we may write D_ρ or D_u for the divisor C_ρ . The union $D := \bigcup_{\rho \in \Sigma^{[1]}} D_\rho = \text{TV}(\Sigma)$ is called the **toric boundary**.

2.3 Singularities

Proposition 2.3.1. *$\text{TV}(\Sigma)$ is nonsingular if and only if each $\sigma \in \Sigma$ is generated by part of a basis for N . If Σ is simplicial, i.e., each cone is generated by independent vectors, then $\text{TV}(\Sigma)$ is an orbifold, i.e., it has only quotient singularities.*

Example 2.3.2. $N = \mathbb{Z}^2$, σ generated by $(1, 0)$ and $(1, 2)$. Then σ^\vee is generated by $(0, 1)$ and $(2, -1)$ (using the standard inner product to identify N with M), and $\mathbb{C}[\sigma^\vee] \cong \mathbb{C}[x^2 y^{-1}, x, y] \subset \mathbb{C}[x^{\pm 1}, y^{\pm 1}]$. Letting $U^2 = x^2 y^{-1}$, $V^2 = y$, we get $\mathbb{C}[\sigma^\vee] \cong \mathbb{C}[U^2, UV, V^2] = \mathbb{C}[U, V]^{\mathbb{Z}/2\mathbb{Z}}$, that is, the ring of invariants under the $\mathbb{Z}/2\mathbb{Z}$ -action that negates U and V . Thus, $\text{Spec } \mathbb{C}[\sigma^\vee]$ is the quotient of \mathbb{C}^2 by $\mathbb{Z}/2\mathbb{Z}$ (acting by negation). This is the A_1 singularity.

Similarly, replacing $(1, 2)$ replaced by $(1, n)$ and the $\mathbb{Z}/2\mathbb{Z}$ action by a $\mathbb{Z}/n\mathbb{Z}$ action (multiplying by n^{th} roots of unity) we get an A_{n-1} singularity.

2.4 Maps of fans

A map of fans $\varphi : \Sigma_1 \rightarrow \Sigma_2$ is a homomorphism $\varphi : N_1 \rightarrow N_2$, Σ_i a fan in N_i , such that for all $\sigma_1 \in \Sigma_1$, there exists some $\sigma_2 \in \Sigma_2$ containing $\varphi(\sigma_1)$. Such a φ induces a morphism $\varphi : \text{TV}(\Sigma_1) \rightarrow \text{TV}(\Sigma_2)$.

Let $|\Sigma| \subset N_{\mathbb{R}}$ denote the support of Σ (i.e., the union of all the cones).

Proposition 2.4.1. *φ is proper (i.e., fibers are compact) if and only if $\varphi^{-1}(|\Sigma_2|) = |\Sigma_1|$. In particular, $\text{TV}(\Sigma)$ is complete (i.e., compact, or proper over $\text{Spec } \mathbb{C}$) if and only if $|\Sigma| = N_{\mathbb{R}}$.*

Examples 2.4.2.

- We have seen that $\rho \subset \sigma$ induces $\text{TV}(\rho) \hookrightarrow \text{TV}(\sigma)$.
- **Blowups:** Adding rays corresponds to taking blowups. For example, we can resolve the A_n singularities of Example 2.3.2 by adding the rays generated by $(1, 1), \dots, (1, n-1)$.
- **Very important for our geometric picture of cluster varieties:** Let $u \in N$, and suppose $\mathbb{R}_{\geq 0}(\pm u)$ are rays in Σ . Then $N \rightarrow N/\langle u \rangle$ induces a map of fans $\Sigma \rightarrow \Sigma/\langle u \rangle$, and the corresponding map of toric varieties gives a \mathbb{P}^1 -fibration of $\text{TV}(\Sigma)$ with D_u and D_{-u} as sections. If $\mathbb{R}_{\geq 0}(\pm u)$ are the only rays in Σ , then $\text{TV}(\Sigma) \cong \mathbb{P}^1 \times (\mathbb{C}^*)^{n-1}$, with $D_{\pm u}$ being the sections corresponding to 0 and ∞ in \mathbb{P}^1 . With cluster varieties, “mutation” will correspond to blowing up some locus in D_u and then contracting some fibers of this fibration.
- **Products:** We can take a product of two fans Σ_i in N_i , $i = 1, 2$. The cones in $\Sigma_1 \times \Sigma_2$ are the cones in $N_1 \times N_2$ of the form $\{(x_1, x_2), x_i \in N_i, |x_i \in \sigma_i \text{ for some cone } \sigma_i \text{ in } \Sigma_i\}$. Then $\text{TV}(\Sigma_1 \times \Sigma_2) = \text{TV}(\Sigma_1) \times \text{TV}(\Sigma_2)$.
- **More General Fiber Bundles:** See the exercise at the bottom of [Ful93, pg 41]. The fiber bundles showing up in the previous two examples and the quotient construction below are enough for our purposes.

2.5 The quotient construction

We now construct toric varieties as quotients, generalizing the standard quotient construction of \mathbb{P}^n . This is the primary approach in [HKK⁺03], and some constructions with cluster varieties (the map from \mathcal{A} to its image in \mathcal{X}) can be viewed as a generalization of this.

Let $\Sigma(1)$ denote the rays of Σ , generated by v_1, \dots, v_m . Let $\mathbb{Z}^{\Sigma(1)}$ denote the lattice freely generated by these rays, with generators denoted $\tilde{v}_1, \dots, \tilde{v}_m$. We define a fan $\tilde{\Sigma}$ in $\mathbb{Z}^{\Sigma(1)}$ as follows: If $\sigma = \langle v_1, \dots, v_s \rangle$ is a cone in Σ generated by v_{i_1}, \dots, v_{i_s} , then $\tilde{\Sigma}$ contains a cone generated by the corresponding $\tilde{v}_{i_1}, \dots, \tilde{v}_{i_s}$. Alternatively, let $Z \subset \text{Spec } \mathbb{C}[\tilde{v}_1^*, \dots, \tilde{v}_m^*]$ be the union of all sets of the form

$$V(\{z^{\tilde{v}_{i_j}^*} \mid j = 1, \dots, k-s \text{ such that the } v_{i_j} \text{'s do not span a cone in } \Sigma\}).$$

Then $\text{TV}(\tilde{\Sigma})$ is $\text{Spec } \mathbb{C}[\tilde{v}_1^*, \dots, \tilde{v}_m^*] \setminus Z = \mathbb{C}^m \setminus Z$.

We have an obvious map $\pi : \mathbb{Z}^{\Sigma(1)} \rightarrow N$ induced by $\tilde{v}_i \mapsto v_i$. By construction, this is a map of fans, so we get a map $\text{TV}(\tilde{\Sigma}) \rightarrow \text{TV}(\Sigma)$. We would like to view this as a quotient of $\text{TV}(\tilde{\Sigma})$ by the action of some group G —this way we get a construction of $\text{TV}(\Sigma)$ that does not depend on our prior construction involving dual cones and gluing. We treat the nonsingular case first.

The kernel K of π is the trivial fan in the lattice of relations among the v_i 's. The short exact sequence $0 \rightarrow K \rightarrow \mathbb{Z}^{\Sigma(1)} \rightarrow N \rightarrow 0$ induces (in the nonsingular cases) a short exact sequence of

the corresponding toric varieties which realizes $\mathrm{TV}(\Sigma)$ as the quotient of $\mathrm{TV}(\tilde{\Sigma})$ by the action of the torus $G := T_K$. Explicitly, the element $(\sum a_i v_i) \otimes \lambda \in K \otimes \mathbb{C}^*$ acts on $\mathrm{TV}(\tilde{\Sigma})$ by $(x_1, \dots, x_n) \mapsto (\lambda^{a_1} x_1, \dots, \lambda^{a_n} x_n)$.

Now suppose we are in a non-singular case. We will see in §2.6 that $\mathbb{Z}^{\Sigma(1)}$ is the lattice of T -invariant Weil-divisors, and M can be identified with the principal T -invariant divisors. The quotient is the group A_{n-1} of divisors up to linear equivalence (see [Ful93], §3.4. Note that we don't need to say T -invariant here). That is, we have a short exact sequence

$$0 \rightarrow M \rightarrow \mathbb{Z}^{\Sigma(1)} \rightarrow A_{n-1} \rightarrow 0$$

Let $G := \mathrm{Hom}(A_{n-1}, \mathbb{C}^*)$. In the nonsingular cases, this is just T_K . However, in the singular cases A_{n-1} may have torsion elements, and G is then isomorphic to $T_K \times \mathrm{Hom}((A_{n-1})_{\mathrm{tor}}, \mathbb{Q}/\mathbb{Z})$. Taking $\mathrm{Hom}(\cdot, \mathbb{C}^*)$ of the above short exact sequence, we get

$$0 \rightarrow G \rightarrow (\mathbb{C}^*)^{\Sigma(1)} \rightarrow T_N \rightarrow 0.$$

$G \subset (\mathbb{C}^*)^{\Sigma(1)}$ acts on $\mathbb{C}^{\Sigma(1)}$ by $g \cdot (v_1, \dots, v_n) = (g([D_{v_1}]v_1), \dots, g([D_{v_n}]v_n))$, where $[D_{v_i}]$ denotes the linear equivalence class of the Weil divisor corresponding to v_i .

Cox explains this quite well (I used his explanation when writing up the singular cases above), and he gives some very interesting additional structure, in [Cox95]. Here he shows that this quotient is a “categorical quotient,” and it is a “geometric quotient” iff Σ is simplicial.

2.6 The Dual Pairing

Recall that a ray $\rho \subset N_{\mathbb{R}}$ corresponds to a (T -invariant) Weil divisors $D_{\rho} \subset \mathrm{TV}(\Sigma)$, $\Sigma \ni \rho$. Alternatively, let v be the primitive² generator of ρ , and define $D_v := D_{\rho}$. More generally, for $k \in \mathbb{Z}_{\geq 0}$, define $D_{kv} := kD_v = kD_{\rho}$.

Proposition 2.6.1. *Let $n \in N$, $m \in M$. Then $\mathrm{val}_{D_n}(z^m) = \langle n, m \rangle$.*

Checking this in a simple example should suffice to see why it is true.

Remark 2.6.2. Note that even if $n \in N$ is not in a ray of Σ , we can still think of it as corresponding to a divisor in some toric blowup (i.e., a blowup corresponding to refining the fan) of $\mathrm{TV}(\Sigma)$, and thus to a discrete valuation on the function field determined by the formula in the above proposition. As a generalization, Gross, Hacking, and Keel define the integer points in the tropicalization of a log Calabi-Yau variety Y as the set of “divisorial” discrete valuations on the function field of Y along which the holomorphic volume form $(\sum d \log z_i$ for toric varieties) has a pole.

We see that $\mathbb{Z}^{\Sigma(1)}$ corresponds to T -invariant Weil divisors. Also, M corresponds to T -invariant principal divisors via: $\mathrm{div}(m) := -\sum_{\rho} \langle m, n_{\rho} \rangle D_{\rho}$, which is $-(z^m)$ the notation of §A.1.

Now, let $W = \sum_{n_i \in \Sigma(1)} a_i D_{n_i}$ be a Cartier divisor. We have a line bundle $\mathcal{O}(W)$ whose global sections may be identified with rational functions f on $\mathrm{TV}(\Sigma)$ such that $\mathrm{val}_{D_{n_i}}(f) \geq -a_i$ for all i . That is,

$$H^0(\mathrm{TV}(\Sigma), \mathcal{O}(W)) = \langle z^m \mid m \in M, \langle n_i, m \rangle \geq -a_i \forall i \rangle.$$

²Primitive meaning not a positive integer multiple of any other element in the lattice.

These are the values of $m \in M$ lying in the **polytope** $P_W := \bigcap_i \{ \langle n_i, \cdot \rangle \geq -a_i \subset M_{\mathbb{R}} \}$. P_W is called the *Newton polytope* of W , or of the line bundle L_W corresponding to W , or of a generic f section of L_W . Note that if f is expressed as a sum $\sum a_i z^{m_i}$, $a_i \neq 0$, then the Newton polytope of f is the convex hull of the m_i 's appearing in this sum.

Let $\Delta \subset N_{\mathbb{R}}$ be a polytope. In the next section, we will need the **polar polytope**

$$\Delta^\circ := \{ u \in M_{\mathbb{R}} \mid \langle u, v \rangle \geq -1 \ \forall v \in \Delta \}.$$

Polar polytopes are convex. Taking polar polytopes is order reversing on the dimensions of faces. The polar polytope of a rational polytope (one with vertices in a lattice Λ) is also rational (with vertices in the dual lattice Λ^*).

2.7 Piecewise Linear Functions and Polytopes

Let $f \in \mathbb{C}(M)$ be a rational function on T_N . Define $f^{\text{trop}} : N_{\mathbb{R}} \rightarrow \mathbb{Z}$ by $f^{\text{trop}}(n) = \text{val}_{D_n}(f)$ for $n \in N$, and extend piecewise-linearly to all of $N_{\mathbb{R}}$. Consider a nonsingular fan Σ such that f is Σ -piecewise linear (i.e., linear on the complement of Σ). Let σ_1, σ_2 be two maximal dimensional cones whose intersection $\sigma_1 \cap \sigma_2$ is a codimension 1 cone ρ . Let $\rho = \langle v_1, \dots, v_{n-1} \rangle$, and $\sigma_i = \langle \rho, u_i \rangle$, $i = 1, 2$.

Let $W_f := \sum_{v \in \Sigma(1)} f^{\text{trop}}(v) D_v$ (in the notation of §A.1, these are the boundary components of (f))³. What is $W_f \cdot C_\rho$? We easily see that D_{u_i} , $i = 1, 2$ as above satisfy $D_{u_i} \cdot C_\rho = 1$ (they intersect only at the point C_{σ_i}). For $v \notin \sigma_1 \cup \sigma_2$, we easily see $D_v \cdot C_\rho = 0$. What about for $v_i \in \rho$? Since f^{trop} is linear on σ_i , we can find a monomial z^m such that $(f \cdot z^m)^{\text{trop}}|_{\sigma_1} = 0$. Since Σ is nonsingular, we have $u_1 + u_2 \in \pm \rho \cap N$, so $(z^m)^{\text{trop}}(u_2) = -[f^{\text{trop}}|_{\rho}(u_1 + u_2) - f^{\text{trop}}(u_1)]$. Hence,

$$W_f \cdot C_\rho = (f \cdot z^m)^{\text{trop}}(u_2) = f^{\text{trop}}(u_2) - f^{\text{trop}}|_{\rho}(u_1 + u_2) + f^{\text{trop}}(u_1).$$

Now, let m_ρ be a primitive element of M which is 0 on ρ and positive on, say, σ_2 . The **bending parameter** of f^{trop} along ρ is the integer b_ρ such that

$$f^{\text{trop}}|_{\sigma_2} = f^{\text{trop}}|_{\sigma_1} + b_\rho m_\rho. \tag{2.1}$$

Here, by the restriction to a cone, I really mean the linear extension of the restriction. With that interpretation,

$$b_\rho = f^{\text{trop}}|_{\sigma_2}(u_2) - f^{\text{trop}}|_{\sigma_1}(u_2) = f^{\text{trop}}(u_2) - [f^{\text{trop}}|_{\rho}(u_1 + u_2) - f^{\text{trop}}(u_1)] = W_f \cdot C_\rho.$$

This proves:

Proposition 2.7.1. *The bending parameter of f^{trop} along ρ is given by $W_f \cdot D_\rho$.*

Corollary 2.7.2. *f^{trop} is strictly convex (i.e., all bending parameters are negative) if and only if $-W_f$ is ample.*

³There are different sign conventions here, and people sometimes take W_f to be negative of our definition. I believe I've taken Fulton's convention. I've tried to be careful, but it's very possible that I've mixed things up somewhere and gotten sign errors, so watch out for that.

Proof. This proof might be using bigger machinery than should be required, but it introduces some things we will want later on.

The Kleiman condition for ampleness says that a divisor W in a variety Y is ample if and only if $D \cdot C > 0$ for all $C \in \overline{\text{NE}}(Y)$ (see also the Nakai-Moishezon criterion). $\text{NE}(Y)$ here is the Mori cone, which is the cone generated by numerical equivalence classes of effective curves—i.e., the dual to the cone of numerically effective divisors. Perhaps some intuition to help this seem somewhat reasonable is to think of having positive intersection as meaning that $|D|$ sees every curve, so none are getting contracted when applying the map to $|D|^*$.

For toric varieties, it turns out that $\text{NE}(Y)$ is generated by the codimension 1 cells of Σ . The claim now follows immediately. \square

Let $\Delta_f := \{f^{\text{trop}} \geq -1\}$. Recall the polytope P_{W_f} corresponding to W_f .

Proposition 2.7.3. *If f^{trop} is strictly convex, then $P_{W_f} = \Delta_f^\circ$. Thus, strictly convex integral polytopes P in $M_{\mathbb{R}}$ (those with points in M as vertices) correspond to ample divisors—translate P so that it contains the origin (this corresponds to multiplying the sections of the line bundle by some z^m , m in the interior of P), then define a strictly convex integral piecewise linear function φ by saying that it equals -1 along the boundary of $P^\circ \subset N_{\mathbb{R}}$.*

We also note the following corollary of Proposition 2.7.1:

Corollary 2.7.4. *Let v_1, v_2, v_3 be three consecutive (under a counterclockwise ordering) primitive generators of rays for the fan of a nonsingular complete toric surface. We have $v_1 + D_{v_2}^2 v_2 + v_3 = 0$.*

More generally,

Proposition 2.7.5. *For any curve class $[C]$ in a nonsingular complete toric variety $\text{TV}(\Sigma)$, $\sum_{v \in \Sigma(1)} ([C] \cdot D_v) v = 0$.*

2.7.1 Reflexive Polytopes and Fano Varieties

If P is rational and contains the origin in its interior, and P° is also a rational polytope (equivalently, P contains no interior lattice points other than the origin), then we say that P (and P°) is **reflexive**. Such polytopes correspond to Fano toric varieties. In this case, φ can be taken to be the tropicalization of a function ψ (the sum of the monomials corresponding to the vertices of P) called the superpotential, and the “Landau-Ginzburg model” $((\mathbb{C}^*)^n, \psi)$ is sometimes viewed as “mirror” to the toric variety corresponding to the polytope P° .

Alternatively, people study mirror symmetry of certain Calabi-Yau hypersurfaces of Fano toric varieties (Batyrev’s construction). Let D denote the toric boundary of a toric variety Y . For any complete toric variety, D is an anti-canonical divisor, meaning that it can be realized as the poles of some meromorphic volume form—in this case, $\sum d \log z_i$. For Fano toric varieties, D is ample, so we can deform it to a linearly equivalent divisor M which does not contain any boundary components. By the adjunction formula, M is Calabi-Yau: $K_M = (K_Y + D)|_M$, and $K_Y + D = 0$.

Dually, if Δ_D is the reflexive polytope corresponding to D , we can do the same thing for the toric variety corresponding to the polar polytope Δ_D° to get a Calabi-Yau M^\vee . Note that we might get something singular this way, but we can sub-divide (blow up) to make it nonsingular. The claim is that the family of M ’s in Y is mirror dual to the family of M^\vee ’s.

2.8 Projective toric varieties from convex polytopes

Given two faces $\tau \subset \sigma$ of Δ (face can mean any codimension for us), define $K_\tau(\sigma) := \langle u-v \mid u \in \sigma, v \in \tau \rangle$ (i.e., vectors pointing from τ into σ with their tail translated to the origin).

Note that there is an order-reversing relationship between fans and polytopes: In particular, maximal cones of the fan correspond to vertices of the polytope.

Lemma 2.8.1. *Let Δ be a strictly convex integral polytope, so it corresponds to an ample divisor. The corresponding divisor is in fact very ample if and only if the integral points of $K_{p(\sigma)}(\Delta)$ generate $\sigma^\vee \cap M$ for all maximal cones $\sigma \in \Sigma$, where $p(\sigma)$ is the vertex of Δ corresponding to σ .*

In particular (and more usefully):

Proposition 2.8.2. *On complete toric varieties which are either two-dimensional or nonsingular, T -invariant ample divisors are very ample. Thus, they give an embedding to projective space.*

From now on, we might as well assume we're in these cases.

The above section shows that ample divisors correspond to strictly convex integral polytopes in $M_{\mathbb{R}}$. Given such a polytope Δ , this means that we can construct the associated toric variety $\text{TV}(\Delta)$ as follows: Consider \mathbb{P}^n with homogeneous coordinates denoted by X_0, X_1, \dots, X_m . Let $\{v_0, \dots, v_m\} := \Delta \cap M$. Then we have a map $f : \text{TV}(\Delta) \hookrightarrow \mathbb{P}^n$ defined by $f^*(X_i) = z^{v_i}$, $i = 0, 1, \dots, m$. So $\text{TV}(\Delta)$ is the subvariety of \mathbb{P}^n given by requiring the X_i 's to satisfy the same *homogeneous* relations as the corresponding z^{v_i} 's. For example, if $v_1 + v_2 + v_3 = 3v_0$, then we have the relation $X_1 X_2 X_3 = X_0^3$.

Example 2.8.3. Let Δ be the convex hull of $(0, 0)$, $(1, 0)$, and $(0, 1)$. There are no homogeneous relations amongst these vertices, so this just corresponds to \mathbb{P}^2 . The corresponding ample line bundle is $\mathcal{O}(1)$.

Here is another version of the above construction which makes sense even for unbounded polytopes (cf. [Gro11], §3.1.2). Let $\Delta \subseteq M_{\mathbb{R}}$ be a not necessarily bounded convex integral polytope. Define the *cone* over Δ by:

$$C(\Delta) = \overline{\{(rn, r) \in M_{\mathbb{R}} \oplus \mathbb{R} \mid n \in \Delta, r \geq 0\}}.$$

Taking the closure above (denote by the overline) adds the asymptotic (i.e., unbounded) directions:

$$\text{Asym}(\Delta) := C(\Delta) \cap (M_{\mathbb{R}}, 0)$$

$\mathbb{C}[C(\Delta) \cap (M \oplus \mathbb{Z})]$ is graded by $\deg z^{m,d} := d \in \mathbb{Z}_{\geq 0}$. Define $\mathbb{P}_{\Delta} := \text{Proj } \mathbb{C}[C(\Delta) \cap (M \oplus \mathbb{Z})]$ (with respect to the above grading). This is projective over (i.e., fibers are projective) Spec of the degree 0 part, $\text{Spec } \mathbb{C}[\text{Asym}(\Delta) \cap M]$. If Δ is bounded, this is the same as the above construction. In any case, \mathbb{P}_{Δ} is a toric variety, and the fan can be obtained by any of the below methods.

2.8.1 Recovering the fan from a polytope

Lemma 2.8.1 implicitly suggests another way to recover $\text{TV}(\Delta)$ (without the embedding into projective space) from Δ : $\text{TV}(\Delta) = \text{TV}(\Sigma)$, where Σ is the fan whose maximal cones are dual to those generated by the $K_p(\Delta)$'s, p a vertex of Δ (for other cones in Σ , one can use higher dimensional faces of Δ).

Example 2.8.4. Check this for the \mathbb{P}^2 case above.

This Σ can also be constructed as (the negation of) the **normal fan** of Δ —choose a point in the interior of Δ and take rays normal to faces of Δ . This is certainly the fastest way to get Σ from Δ .

Alternatively, if Δ contains the origin in its interior, one can take Δ° and then define Σ to be the fan consisting of cones over the faces of Δ° .

Example 2.8.5. Check these for the convex hull of $(2, -1)$, $(-1, 2)$, and $(-1, -1)$, which corresponds to $\mathcal{O}(3)$ on \mathbb{P}^2 .

2.9 Moment maps

2.9.1 General theory of moment maps

We shouldn't actually need this, and I don't have much background in it, but it shows up enough to be worth mentioning.

Let M be a manifold with symplectic form ω . Suppose the Lie group G acts on M by symplectomorphisms. $\zeta \in \mathfrak{g}$ induces a left invariant vector field $L(\zeta)$ on M . The moment map $\mu : M \rightarrow \mathfrak{g}^*$ is defined by

$$d\langle \mu(x), \zeta \rangle = \iota_{L(\zeta)_x} \omega$$

for all $x \in M$, $\zeta \in \mathfrak{g}$.

If the action is Hamiltonian, meaning that there is a Lie algebra homomorphism $\zeta \mapsto f_\zeta$ from \mathfrak{g} to smooth functions on M (with the Poisson bracket), then we can define μ by:

$$\langle \mu(x), \zeta \rangle = f_\zeta(x).$$

Theorem 2.9.1 (Atiyah, Guillemin, Sternberg). *For Hamiltonian (compact) k -torus actions on closed symplectic manifolds, the image of the moment map is a convex polytope in \mathbb{R}^k .*

Theorem 2.9.2. *If M is $2k$ -dimensional and G is a compact k -torus, then the regular fibers of μ are Lagrangian tori (i.e., they are k -dimensional and ω restricted to them is the 0 form).*

2.9.2 Moment maps for toric varieties

Recall that rather than defining the affine toric variety associated to σ to be $\text{Spec } \mathbb{C}[\sigma^\vee]$, we could have said that it is $\text{Hom}_{sg}(\sigma^\vee, \mathbb{C})$. We could easily replace \mathbb{C} here by any other monoid. For example, define $\text{TV}_\geq(\Sigma)$ to be the toric variety constructed using $\text{Hom}(\sigma^\vee, \mathbb{R}_{\geq 0})$ for each cone σ . This forms a closed subspace of $\text{TV}(\Sigma)$. Consider the compact torus S_N in $T_N := N \otimes \mathbb{C}^* = (\mathbb{C}^*)^n$ corresponding to values of norm 1.

Proposition 2.9.3. *The quotient map for the action of S_N on $\text{TV}(\Sigma)$ gives a retraction to $\text{TV}_\geq(\Sigma)$.*

Various choices of ω can be used to identify this with a moment map as before. I am taking these examples from [Rud] but I believe this is not the original source of the examples.

If $\text{TV}(\Sigma)$ is simply $(\mathbb{C}^*)^n$, we can take $\omega = -\frac{1}{(2\pi)^2} \sum d \log r_i \wedge d\theta_i$, and then $\mu : (\mathbb{C}^*)^n \rightarrow \mathbb{R}^n$ is given by $-\frac{1}{2\pi}(\log |z_1|, \dots, \log |z_n|)$. The fibers are Lagrangian, and if we take $\Omega = \frac{1}{(2\pi i)^n} d \log z_1 \wedge \dots \wedge d \log z_n$,

then the fibers are actually **special Lagrangian** (i.e., in addition to $\omega|_L = 0$ for each fiber L , we also have $\text{Im}(\Omega)|_L = 0$).

Alternatively, suppose we take a rational convex polytope $P \subset M_{\mathbb{R}}$ with $m + 1$ lattice points, and consider the corresponding projective embedding φ of the toric variety M into \mathbb{P}^m . So lattice points $v_i \in P$ are associated with homogeneous coordinates X_i for \mathbb{P}^m , and $X_i|_M = z^{v_i}$. Then the moment map is

$$\mu(w) = -\frac{\sum |z^{v_i}(w)|^2 v_i}{\sum |z^{v_i}(w)|^2}.$$

(Fulton does not have the minus sign or the squares in either the numerator or the denominator. I don't know how much that matters.) The image of this is $-P$. Removing the minus sign makes it just P , which is more common, so I'll pretend we've done this.

The preimage of a real k -dimensional face of P is a complex k -dimensional orbit in M , with the preimage of a point being a real k -dimensional compact torus. Fulton has a couple nice pictures of this. In particular, the fibers over the interior of P are special Lagrangian tori.

2.10 Toric degenerations

Let $\Delta \subset M_{\mathbb{R}}$ be a lattice polytope. Give Δ a polyhedral decomposition⁴ \mathcal{P} which is the singular locus of some convex piecewise-linear function $\varphi : \Delta \rightarrow \mathbb{R}$ with integral slopes. Define a polyhedron⁵

$$\tilde{\Delta} := \{(m, r) \in M_{\mathbb{R}} \times \mathbb{R} \mid m \in \Delta, r \leq \varphi(m)\},$$

and consider $C(\tilde{\Delta})$. If Δ was bounded, then $\text{Asym}(\tilde{\Delta}) = (0, \mathbb{R}_{\leq 0})$, and we get that $\mathbb{P}_{\tilde{\Delta}} := \text{Proj}(\mathbb{C}[C(\tilde{\Delta}) \cap (M \times \mathbb{Z} \times \mathbb{Z})])$ is projective over \mathbb{A}^1 . In fact, even if Δ was not bounded, we still have a morphism $f : \mathbb{P}_{\tilde{\Delta}} \rightarrow \mathbb{A}^1$ coming from realizing $\mathbb{C}[C(\tilde{\Delta}) \cap (M \times \mathbb{Z} \times \mathbb{Z})]$ as a $\mathbb{C}[t]$ -algebra by setting $t \mapsto z^{(0, -1, 0)}$ (first 0 in the exponent being $0 \in M$ and the final 0 being in the last copy of \mathbb{Z}).

On the complement of $t = 0$, we can divide by t , which corresponds to “lifts the top off” of the polytope $\tilde{\Delta}$. One sees from this (since products of polytopes correspond to products of the corresponding toric varieties) that $f^{-1}(\mathbb{A}^1 \setminus \{0\}) = \mathbb{P}_{\Delta} \times \mathbb{C}^*$, with the restriction of f being projection to the second factor.

On the other hand, one can show that $f^{-1}(0)$ is a union of toric divisors of $\mathbb{P}_{\tilde{\Delta}}$: the irreducible components are the toric varieties corresponding to the maximal chambers of the polyhedral decomposition \mathcal{P} , with intersections coming from \mathcal{P} as one would expect. See [Gro11] (Example 3.5) or [GS11b] (§1.2) for more details.

⁴i.e., a decomposition into a union of polytopes such that the intersection of two polytopes is a face of each and any face of a polytope is also part of the decomposition. We further assume that each of these polytopes is integral, meaning that they have lattice points as vertices.

⁵Warning: People sometimes use $r \geq \varphi(m)$ instead of \leq , but with our convention of convex meaning *negative* bending parameters, we want to look at everything *below* the graph of φ .

Chapter 3

Motivation from Mirror Symmetry

This chapter is not strictly necessary for the constructions of [GHK11] and [GHKK]. However, it gives the motivation for much of what goes into those constructions, as well as context for those constructions.

3.1 Mirror symmetry and T-duality in the complement of an anticanonical divisor

This section will closely follow Denis Auroux’s paper [Aur]. Let (X, ω, J) denote a smooth compact Kähler manifold X of complex dimension n . Let Ω be a non-vanishing meromorphic volume form with D (an anti-canonical divisor) as the divisor of poles of Ω , so $\Omega|_{X \setminus D}$ (which we will also denote by Ω) is a holomorphic volume form on $X \setminus D$.

3.1.1 Setup

Definition 3.1.1. $L \subset X \setminus D$ of real dimension n is *special Lagrangian* (sLag for short) with phase $\phi \in \mathbb{R}/2\pi$ if $\omega|_L = 0$ (Lagrangian) and $\text{Im}(e^{-i\phi}\Omega)|_L = 0$.

By replacing Ω with $e^{-i\phi}\Omega$, we will always assume $\phi = 0$.

The main conjecture is:

Conjecture 3.1.2. *A “mirror” manifold M may be constructed as the moduli space of sLag tori in $X \setminus D$ equipped with a flat $U(1)$ -connection, up to gauge, along with a holomorphic superpotential $W : M \rightarrow \mathbb{C}$ (the m_0 obstruction to Floer homology). Furthermore, fibers¹ of W are mirror to D .*

The possibility of singular sLag torus fibers prevents this from holding as stated—quantum corrections² are needed.

¹The pair (M, W) is called a Landau-Ginzburg model.

²i.e., instanton corrections, i.e., wall-crossing formulas, i.e., generalized mutations.

3.1.2 The tangent space to the moduli space

Let g denote the Kähler metric on X . This induces a volume form vol_g on L . There exists a function $\psi \in C^\infty(L, \mathbb{R}_+)$ such that $\text{Re}(\Omega)|_L = \psi vol_g$ (for the compact Calabi-Yau situation where $D = 0$, $\psi = 1$).

Definition 3.1.3. $\alpha \in \Omega^1(L, \mathbb{R})$ is ψ -harmonic if $d\alpha = d^*(\psi\alpha) = 0$. Let $\mathcal{H}_\psi^1(L)$ denote the space of ψ -harmonic 1-forms on L .

ψ -harmonic 1-forms are a lot like the usual harmonic 1-forms:³

Lemma 3.1.4. *Every cohomology class has a unique ψ -harmonic representative.*

Proposition 3.1.5. *A section of the normal bundle $v \in C^\infty(NL)$ determines a 1-form $\alpha := -\iota_v\omega$ and an $(n-1)$ -form $\beta := \iota_v\Omega$ on L satisfying $\beta = \psi *_g \alpha$ ($*_g$ the Hodge-star operator with respect to g) and the infinitesimal deformation corresponding to v is sLag iff α and β are both closed (i.e., iff $\alpha \in \mathcal{H}_\psi^1(L)$). That is, infinitesimal sLag deformations correspond bijectively to ψ -harmonic 1-forms. Furthermore, the deformations are unobstructed, meaning that infinitesimal deformations really are derivatives of actual deformations.*

Definition 3.1.6. An integral affine structure⁴ on an m -dimensional real manifold M is an atlas with transition functions in $GL_m(\mathbb{Z}) \ltimes \mathbb{R}^m$. Equivalently,⁵ it is a collection of sections of the tangent bundle which form a full rank lattice \mathbb{Z}^m in each fiber.

We see that, at least locally, the moduli space \mathcal{B} of sLag deformations of L is a smooth manifold with two affine structures, corresponding to identifying the tangent space with either $H^1(L, \mathbb{R})$ or $H^{n-1}(L, \mathbb{R})$ and taking the lattices of integral forms, $H^1(L, \mathbb{Z})$ or $H^{n-1}(L, \mathbb{Z})$, respectively.

3.1.3 Constructing M

We are still following [Aur], but in the Calabi-Yau ($D = 0$) case, the following is due to Hithin [Hit97]. Let M denote the space of pairs (L, ∇) where $L \in \mathcal{B}$ as above, and ∇ is a flat⁶ $U(1)$ -connection on the trivial complex line bundle over L , up to gauge. I.e., ∇ corresponds to an element $\text{hol}_\nabla \in \text{Hom}(H_1(L), U(1)) \cong H^1(L, \mathbb{R})/H^1(L, \mathbb{Z})$, the fibers over \mathcal{B} are indeed tori.

We can represent ∇ by $d + iA$ for some ψ -harmonic 1-form A on L . We can identify $T_{(L, \nabla)}M$ with the set of pairs $(v, \alpha) \in C^\infty(NL) \oplus \Omega^1(L, \mathbb{R})$ such that v gives an infinitesimal sLag deformation and α is (represented by) a ψ -harmonic 1-form, viewed as an infinitesimal deformation of ∇ . Mapping $(v, \alpha) \mapsto -\iota_v\omega + i\alpha$ identifies $T_{(L, \nabla)}M$ with $\mathcal{H}_\psi^1 \otimes \mathbb{C}$, giving M a complex structure $J^\vee(v, \alpha) = (a, -\iota_v\omega)$, where a is a normal vector field such that $\iota_a\omega = \alpha$.

³In fact, if $n \neq 2$, ψ -harmonic for g is the same as harmonic for $\psi^{2/(n-2)}g$.

⁴There are different conventions for terminology here. For example, Mark Gross might call this a tropical structure, and say an affine structure has transition maps in $\text{Aff}(\mathbb{R}^n)$, and an integral affine structure has them in $\text{Aff}(\mathbb{Z}^m)$.

⁵ \mathbb{R}^m contains the lattice \mathbb{Z}^m which can be identified with a lattice in any tangent space of \mathbb{R}^m . Pulling back along charts gives lattices for the tangent spaces of M which are preserved by the action of the affine group. Conversely, integrating the 1-forms corresponding to the dual lattice give coordinates for the charts.

⁶More generally, people consider X with “complexified Kähler structures” $\omega_{\mathbb{C}} = B + i\omega$ for some $B \in H^2(X, \mathbb{R})$, called the B -field. Then ∇ is required to have curvature $-iB$.

Proposition 3.1.7. *Let $\Gamma \in H_2(X, L, \mathbb{Z})$ be a relative homology class with boundary $\partial\Gamma \neq 0 \in H_1(L, \mathbb{Z})$. Then*

$$z_\Gamma := \exp\left(-\int_\Gamma \omega\right) \text{hol}_\nabla(\partial\Gamma) : M \rightarrow \mathbb{C} \quad (3.1)$$

is holomorphic.

Proof. $d \log z_\Gamma$ is $(v, \alpha) \mapsto \int_{\partial\Gamma} -\iota_v \omega + i\alpha$, which is \mathbb{C} -linear. \square

Keeping track of how Γ changes as we vary L , we thus obtain locally well-defined coordinates z_Γ . These may be globally multi-valued if \mathcal{B} has non-trivial monodromy. The [GHK11] and [GHKK] constructions essentially describe how to find sums of z_Γ 's which are invariant under the monodromy action, thus giving well-defined global functions.

Remark 3.1.8. Here is another point of view which may be better for us (cf. [Rud]). To get coordinates for the affine structures on the base, let $\{\gamma_i\}$ be a basis for $H_1(L, \mathbb{Z})$, $\{\gamma_i^*\}$ the dual basis, $\Gamma_i \in H_2(X, L, \mathbb{Z})$ cylinders (as above) traced out by the γ_i 's as we move L , and Γ_i^* traced out by the γ_i^* 's. Then we have coordinates for one affine structure given by $y_i := \int_{\Gamma_i} \omega$, and for the other given by $\hat{y}_i := \int_{\Gamma_i^*} \text{Im}(\Omega)$.

Now, let $x_i = dy_i$ and $\hat{x}_i := \partial_{\hat{y}_i}$. We can view M locally as $T\mathcal{B}/\Lambda$ (where Λ is the lattice of integral tangent vectors) with complex structure given by $z_i := x_i + iy_i$ and holomorphic volume form $\Omega = dz_1 \wedge \cdots \wedge dz_n$. On the other hand, we can view X locally as $T^*\mathcal{B}/(\Lambda^*)$ with symplectic form $\omega = d\hat{x}_i \wedge d\hat{y}_i$. The data of a Hessian metric on \mathcal{B} allows one to obtain the complimentary data.

The following proposition summarizes the rest of §2 of [Aur]:

Proposition 3.1.9.

$$\omega^\vee((v_1, \alpha_1), (v_2, \alpha_2)) = \int_L \alpha_2 \wedge \iota_{v_1} \text{Im}(\Omega) - \alpha_1 \wedge \iota_{v_2} \text{Im}(\Omega)$$

defines a Kähler form on M , compatible with J^\vee , with respect to which the fibers of $\pi : (L, \nabla) \mapsto L$ are Lagrangian. If L is a torus, then $\dim M = \dim X = n$, and M has a holomorphic volume form

$$\Omega^\vee((v_1, \alpha_1), \dots, (v_n, \alpha_n)) = \int_L (\iota_{v_1} \omega + i\alpha_1) \wedge \cdots \wedge (\iota_{v_n} \omega + i\alpha_n).$$

Equivalently, $\Omega^\vee = d \log z_{\Gamma_1} \wedge \cdots \wedge d \log z_{\Gamma_n}$ for Γ_i 's constructed from a basis for $H_1(L, \mathbb{Z})$ as in Remark 3.1.8. The fibers of π are now sLag with phase $n\pi/2$.

If ψ -harmonic 1-forms on L have no zeroes (automatic for $n \leq 2$), then in a neighborhood of L , (X, J, ω, Ω) and $(M, J^\vee, \omega^\vee, \Omega^\vee)$ do indeed admit dual sLag torus fibrations over \mathcal{B} .

3.1.4 The Superpotential

Since we are not dealing with a compact Calabi-Yau, we may run into an obstruction when defining its Fukaya category. Namely, ∂^2 might not be 0. One instead uses a “twisted” version of the Fukaya category. The usual Fukaya category is an A_∞ category: that is, rather than just having an associative multiplication/composition of pairs of morphisms, we have maps m_k which “compose” k morphisms. m_1 is the chain map ∂ , and the multiplication m_2 only associative “up to higher homotopy,” i.e., up to higher m_k terms. Roughly, m_k counts (pseudo-)holomorphic disks with their boundary on

Lagrangians L_0, L_1, \dots, L_k . For twisted Fukaya categories, we also have an obstruction term m_0 which counts holomorphic disks with boundary on a single L . This term corresponds to the superpotential $W : M \rightarrow \mathbb{C}$ on the mirror. If there are no Maslov index⁷ zero disks,

$$W(L, \nabla) := \sum_{\beta, \mu(\beta)=2} n_\beta(L) z_\beta.$$

Here, the sum is over all classes $\beta \in H_2(X, L, \mathbb{Z})$ of Maslov index $\mu(\beta) = 2$, and z_β is as in Equation 3.1. $n_\beta(L)$ is a “virtual” count of the number of holomorphic disks of class β whose boundary passes through a generically specified point of L .

The condition of there being no Maslov index 0 disks is satisfied for toric varieties. However, if there are Maslov index 0 disks, then W as above isn’t well-defined globally. The issue is that as we vary L , some “disk bubbling” might occur (roughly, we can glue Maslov index 0 disks to Maslov index 2 disks to get new Maslov index 2 disks). To deal with this, as we cross a “wall” in \mathcal{B} where this gluing happens, we change the z_β ’s according to a “wall-crossing formula.” We’ll see this when we cover [GHK11].

The theta functions in [GHK11] are, at least heuristically, defined in essentially the same way. [GHK11] deals with surfaces obtained by blowing up points on the boundary D of a toric variety \bar{X} . Let N be the cocharacter lattice for \bar{X} , $n_i \in N$ corresponding to a boundary divisor $D_i \subset D$. Then on the mirror we will have theta functions ϑ_{kn_i} which heuristically are given in local coordinates by

$$\vartheta_{kn_i}(L, \nabla) := \sum_{\beta} n_\beta(L) z_\beta,$$

where the sum is over classes $\beta \in H_2(X, L, \mathbb{Z})$ satisfying $\beta \cdot D_i = k$ and $\beta \cdot D_j = 0$ for each component $D_j \subset D$, $j \neq i$, and n_β counts curves which hit D_i at a single point with full multiplicity k . W then will be given by $\sum \vartheta_{n_i}$, where the sum is over the n_i such that D_{n_i} is a component of D (i.e., over primitive generators for the rays of the fan for \bar{X}). The actual definition will use “broken lines” (an abridged version of tropical curves) in place of actual holomorphic disks.

3.2 The Gross-Siebert Program

This section is mainly based on parts of [Gro13], which gives a nice detailed overview of the Gross-Siebert program. This program is an approach to SYZ mirror symmetry which replaces the difficult differential geometry with some more tractable toric and tropical geometry. This section give only a rough sketch of the program.

3.2.1 Toric Degenerations

For toric varieties, we have seen that finding sLag fibrations is rather easy. However, in general, it is quite hard, if not impossible. A key idea behind the Gross-Siebert program is to consider large complex structure limits which are “toric degenerations” of the Calabi-Yau manifold X .⁸ Roughly,

⁷One can show that, for L sLag, $\mu(\beta) = 2\beta \cdot D$. I intend to just use this as our definition of Maslov index.

⁸My understanding is that they more generally use a “log-structure” on X instead of the degeneration data. When we have a nontrivial anti-canonical divisor D as before, this divisor is used to get at least part of the log structure. I unfortunately don’t know any details about how this works, but Mark Gross has given me the impression that he knows how to do things from this perspective in the [GHK11] situations we will see later.

in mirror symmetry, one expects the “large volume limit” on the symplectic side (corresponding to letting $\omega \rightarrow \infty$) to be mirror to the “large complex structure limit” on the complex side.⁹ A toric degeneration is essentially a (proper flat) family $f : \mathcal{X} \rightarrow D$ (D an analytic disk) with fibers over $D \setminus \{0\}$ being nice (normal, not too singular) Calabi-Yau varieties, and special fiber \mathcal{X}_0 over $0 \in D$ being a union of toric varieties, such that any $x \in \mathcal{X}_0 \setminus Z$ (Z some locus in \mathcal{X}_0 of codimension ≥ 2) has an analytic open neighborhood in \mathcal{X} on which the degeneration is isomorphic to an analytic open subset of a toric degeneration as in 2.10.

Example 3.2.1. Consider the space $\mathcal{X}_t := \{z_0 z_1 z_2 z_3 + t f = 0\} \subset \mathbb{P}^3 \times D$, where f is a generic degree 4 polynomial in the homogeneous coordinates z_i $i = 0, 1, 2, 3$, and $t \in D$. At $t = 0$, this becomes just the coordinate hyperplanes in \mathbb{P}^3 , i.e., a union of toric varieties (\mathbb{P}^2 's). \mathcal{X}_t is singular at $\{t = f = 0\} \cap \text{Sing}(\mathcal{X}_0)$, which includes 4 points on each coordinate line, for a total of 24 singular points forming the set Z .

3.2.2 Constructing B and \widehat{B}

The toric degeneration data is then used to construct the affine base B , which we will sometimes refer to as the *tropicalization* of X —It is constructed as the dual-intersection complex of \mathcal{X}_0 . That is, k -dimensional strata¹⁰ of \mathcal{X}_0 correspond to codimension k cells of B , glued in the obvious inclusion-reversing way. Note that this gives B a polyhedral decomposition \mathcal{P} . B has an obvious affine structure on the interiors of the maximal faces of \mathcal{P} , but we still need to extend this over their intersections (there will be singular points Z we cannot extend over, but this should have codimension ≥ 2). We do this by assigning a “fan structure” to each vertex of B . Basically, a vertex $v \in B$ corresponds to an irreducible component of \mathcal{X}_0 , and a neighborhood of v in B is identified with a neighborhood of $0 \in N_{\mathbb{R}}$ in the fan corresponding to this irreducible component.

If \mathcal{X} is equipped with a relatively (i.e., fiberwise) ample line bundle \mathcal{L} , we can also get the dual affine manifold \widehat{B} with polyhedral decomposition $\widehat{\mathcal{P}}$. \widehat{B} is constructed as the intersection complex— k -dimensional strata of \mathcal{X}_0 correspond to k -dimensional cells σ of \widehat{B} which are equal to the Newton polytopes associated to $L|_{\sigma}$. $\widehat{\mathcal{P}}$ is now this decomposition into Newton polytopes. The fan structure at a vertex $v \in \widehat{\mathcal{P}}$ comes from the normal fan to the corresponding cell σ_v in the dual intersection complex.

The relatively ample line bundle also gives us a “multi-valued” strictly convex piecewise-linear function φ on B . A multi-valued piecewise-linear function may be defined as a section of the sheaf $\mathcal{P}\mathcal{L}/\mathcal{L}$, where \mathcal{L} is the sheaf of linear functions and $\mathcal{P}\mathcal{L}$ is the sheaf of piecewise-linear functions. Such functions are uniquely determined by their bends. If D denotes a divisor of \mathcal{X}_0 corresponding to $\mathcal{L}|_{\mathcal{X}_0}$, then the bending parameter of φ along a codimension 1 cell τ of \mathcal{P} is $-D \cdot C_{\tau}$, where C_{τ} is the curve associated with τ . Ampleness implies this bend is negative (yielding strict convexity).

The data $(B, \mathcal{P}, \varphi)$ is actually sufficient to construct $(\widehat{B}, \widehat{\mathcal{P}})$, along with a function $\widehat{\varphi}$ on \widehat{B} which makes this invertible. This is called the discrete Legendre transform. When dealing with actual SYZ

⁹Very roughly, a large complex structure limit is a degeneration such that the monodromies of the middle cohomology around the singular loci are maximally unipotent—see Auroux’s lecture notes for more details. I often just picture them as toric degenerations, although this is not always accurate. We will see below what approaching this limit means for the affine base.

¹⁰By strata of \mathcal{X}^0 , I mean intersections of irreducible components. By strata of a polytope, I mean intersections of collections of faces (e.g., faces, edges, vertices, etc.)

fibrations, one uses a non-discrete version of this involving a “Hessian” metric g on B in place of \mathcal{P} and φ (cf. [Gro13] and [Rud]), but we will not need this.

Remark 3.2.2. Before moving on, we clarify that $(B^\vee, \mathcal{P}^\vee, \varphi^\vee)$ should correspond to the affine structure coming from the symplectic form—i.e., affine coordinates should be given by $\int_{\Gamma_i} \omega$ as in 3.1.8. On the other hand, the affine coordinates on B should be given by $\int_{\Gamma_i^*} \text{Im}(\Omega)$. To remember this, note that we constructed B before specifying a Kähler structure (i.e., an ample line bundle), so clearly the symplectic structure was not important on this side. Newton polytopes, on the other hand, are bases for moment maps and are thus linked to the symplectic structure.

3.2.3 The Reconstruction Problem

The other key step in the Gross-Siebert program is an algorithm for using the data $(B, \widehat{\mathcal{P}}, \varphi)$ or $(\widehat{B}, \widehat{\mathcal{P}}, \widehat{\varphi})$ to actually construct \mathcal{X}^\vee , at least in a formal neighborhood¹¹ of the large complex structure limit. This was carried out in [GS11a], and is also explained in [Gro11]. The construction can be quite complicated, particularly in higher dimensions, but the essential ideas are visible in dimension 2. We will soon carry out this procedure for log Calabi-Yau surfaces, following [GHK11].

The main idea is that, by construction, \mathcal{X} should at least locally look like a toric degeneration coming from the data $(B^\vee, \mathcal{P}^\vee, \varphi^\vee)$ (I have not totally justified this, but at least recall that our toric degeneration construction used the Newton polytope of the generic fiber as its base, with cells of \mathcal{P} corresponding to strata of the singular fiber). The mirror should have $(\widehat{B}, \widehat{\mathcal{P}}, \widehat{\varphi})$ as the dual intersection complex and $(B, \mathcal{P}, \varphi)$ as the intersection complex, so it should locally look like a toric degeneration constructed from the data $(B, \mathcal{P}, \varphi)$. Unfortunately, doing this more than just locally is in general quite difficult, particularly when there are singularities in the affine structure. Dealing with the singularities requires a structure called a scattering diagram.

Basically, holomorphic curves in X are supposed to map to shapes called “amoebas” in B . As we approach the large complex structure limit, these amoebas should contract to so-called “tropical curves.” There are various theorems about the exact correspondence between counts of tropical curves and counts of holomorphic curves—cf. [Mik05] for toric surfaces and [NS06] for higher dimensional toric varieties. As we will see in a little more detail soon, a scattering diagram is some combinatorial data in B which [GPS09] shows in some cases records data about counts of tropical curves, and hence of holomorphic disks. These disks will correspond to the Maslov index 0 disks that showed up when we needed to make quantum corrections of the superpotential in §3.1.4. Thus, the scattering diagram will tell us how to make these quantum corrections.

3.2.4 Tropical Curves

I have mentioned tropical curves a few times without ever saying what they are, so before moving on I’ll take a quick moment to introduce them via some examples.

Recall the affine manifold B from before. Let Λ_B denote the lattice of integer tangent vectors on B . Then X should locally look like TB/Λ_B . Let y_1, \dots, y_n denote local affine coordinates on B such that $x_j := dy_j$ forms a basis for Λ_B^* . Then we can write complex coordinates on X as $z_j := e^{2\pi i(x_j + iy_j)}$. Taking the large complex structure limit corresponds to replacing Λ with $\epsilon\Lambda$ and letting $\epsilon \rightarrow 0$. With

¹¹See §A.2 for background on formal schemes.

x_j and y_j defined with respect to Λ as before, the complex structure for the manifold X_ϵ locally defined by $TB/\epsilon\Lambda$ is locally given by $z_j^{1/\epsilon} := e^{2\pi i(x_j + iy_j)/\epsilon}$.

Now recall from §2.9.2 that $(\mathbb{C}^*)^2$ has a moment map given (up to a factor of $\frac{1}{2\pi}$ which I'll ignore for simplicity) by $\mu : (z_1, z_2) \mapsto -(\log |z_1|, \log |z_2|)$. If we consider the curve $C := \{z_1 = 1\} \subset (\mathbb{C}^*)^2$, we see that $\mu(C) = (0, \mathbb{R})$, i.e., the y -axis. This is a rather trivial example of a tropical curve. More generally, any linear subspace of B with rational slope should correspond to some relatively simple complex submanifold of X (cf. [Gro13], §4, where most of this subsection is coming from). The idea of tropical geometry is that, as we approach the large complex structure limit, we should be able to piece together these linear subspaces in a way that gives us more interesting complex submanifolds.

For example,¹² consider $C_\epsilon = \{z_1^{1/\epsilon} + z_2^{1/\epsilon} = 1\}$. If we keep the same moment map μ as before but express it in terms of the complex coordinates $z_j^{1/\epsilon}$, then we get

$$\mu(z_1^{1/\epsilon}, z_2^{1/\epsilon}) := -(\log |z_1|, \log |z_2|) = -\epsilon(\log |z_1^{1/\epsilon}|, \log |z_2^{1/\epsilon}|).$$

Suppose we fix, say, $-\epsilon \log |z_1^{1/\epsilon}| = a$, so $-\epsilon \log |z_2^{1/\epsilon}| \in [-\epsilon \log |1 - e^{-a/\epsilon}|, -\epsilon \log |1 + e^{-a/\epsilon}|]$. As $\epsilon \rightarrow 0^+$, one checks that this interval approaches a if $a < 0$ and 0 if $a > 0$. Applying similar calculations for a fixed value of $-\epsilon \log |z_2^{1/\epsilon}|$, one finds that $\mu(C_\epsilon)$ is the union of the positive x -axis, the positive y -axis, and $\mathbb{R}_{\geq 0}(-1, -1)$. This is an example of a tropical curve. In the non-limiting cases, the image is called an amoeba, and as $\epsilon \rightarrow 0$, we see that $\mu(C_N)$ converges to this tropical curve.

Here is another way to get this tropical curve, at least up to a shift. In the equation $z_1 + z_2 - 1$, replace addition with min and multiplication (which we don't have in this example) with addition, and replace z_1 and z_2 with the standard real coordinates y_1 and y_2 for \mathbb{R}^2 . Then we get a piecewise-linear function on \mathbb{R}^2 given by $\min(-1, y_1, y_2)$. The singular locus of this is exactly the tropical curve from above shifted by the vector $(-1, -1)$ (The shift surprises me, so hopefully I haven't made a mistake). More generally, tropical varieties are defined to be the intersections (with some multiplicity recording changes of slope) of the singular loci of "tropical polynomials," i.e., those polynomials in the y_i 's over \mathbb{R} with multiplication and addition being replaced by addition and min. Many properties of the varieties can be seen in the corresponding tropical curves

Alternatively, one might define "parametrized tropical curves." Briefly, if Γ is a weighted graph with its univalent vertices removed and weight function $w : \Gamma^1 \rightarrow \mathbb{Z}_{\geq 0}$, then a weighted tropical curve in B is a map $h : \Gamma \rightarrow B$ which contracts weight 0 edges, linearly embeds all other edges, and satisfies the "balancing condition:" if V is a vertex in the edges E_1, \dots, E_s of Γ and v_1, \dots, v_s are primitive generators for the images of the E_i 's, oriented to point away from $h(V)$, then $\sum w(E_i)v_i = 0$.

As mentioned above, [Mik05] and [NS06] have theorems relating certain counts of parametrized tropical curves to actual counts of holomorphic curves. [GPS09] relates the data of scattering diagrams to counts of tropical curves and then uses the results of [NS06] to relate this to counts of holomorphic curves and then to relative Gromov-Witten invariants.

¹²I hope I've avoided errors in these calculations, but I make no guarantees.

Chapter 4

Mirror Symmetry for Log Calabi-Yau Surfaces

This section is a good warm-up for the cluster situations we will attack later. Log Calabi-Yau surfaces are essentially (up to codimension 2 issues) the same as the fibers of rank 2 cluster \mathcal{X} -varieties (those for which the skew-form has rank 2). In fact, we will see that the mirror construction of [GHK11] applied to a fiber U of rank 2 \mathcal{X} is really the same as the mirror construction of [GHKK] applied to the \mathcal{A} -space. In either case, the \mathcal{X} -space turns out to basically be the mirror (under some affineness assumption).

4.1 Log Calabi-Yau Surfaces with Maximal Boundary

Definition 4.1.1. Let (Y, D) denote a smooth projective surface Y over \mathbb{C} (or more generally, over an algebraically closed field \mathbb{k} with characteristic 0), and a singular nodal anti-canonical divisor $D = D_1 + \dots + D_s$ (D_i the irreducible components, cyclically ordered). Denote $U := Y \setminus D$.

We may call (Y, D) a Looijenga pair, and U a Looijenga interior. Alternatively, we may call U a log Calabi-Yau surface,¹ and (Y, D) a minimal model for U . D will be referred to as the boundary. We will call a divisor D -ample if it has positive intersection with every irreducible component of D .

Examples 4.1.2.

- Y a complete nonsingular toric surface, D the toric boundary.
- $Y \cong \mathbb{P}^2$, D a nodal cubic.
- Blowups: Given a Looijenga pair (\bar{Y}, \bar{D}) , we can take:
 - *Toric blowups*, where Y is the blowup of \bar{Y} at a nodal point of \bar{D} and D is the inverse image of \bar{D} . Note that toric blowups do not change U .
 - *Non-toric blowups*, where \tilde{Y} is the blowup of \bar{Y} at a non-nodal point of \bar{D} , and \tilde{D} is the proper transform of \bar{D} .

¹[GHK] calls U a log Calabi-Yau surface *with maximal boundary*, in contrast to the compact Calabi-Yau case where $D = 0$. Briefly, maximal boundary means that D has a 0-stratum in a neighborhood of which D is the zero set of a product of $\dim(U)$ local analytic coordinate functions.

We will allow the term *toric blowup* to refer to a sequence of toric blowups. A *toric model* will mean a sequence of non-toric blowups of a toric variety.

Proposition 4.1.3 ([GHK11], Prop. 1.19). *Every Looijenga pair has a toric blowup which admits a toric model.*

4.2 Tropicalizing U

We want to construct a mirror family \mathcal{X} to (Y, D) . This family will be a “formal” smoothing of $\mathbb{V}^n := \mathbb{A}_{x_1, x_2}^2 \cup \mathbb{A}_{x_2, x_3}^2 \cup \dots \cup \mathbb{A}_{x_s, x_1}^2 \subset \mathbb{A}^s$. If D supports a D -ample divisor, we can actually extend this to an actual algebraic family containing the original U as a fiber (non-canonically). We can even partially compactify the mirror so that (Y, D) can be identified with a fiber. Following the Gross-Siebert perspective, this suggests² that the base B of the SYZ fibration, together with the polyhedral decomposition \mathcal{P} , should look like the fan for (\bar{Y}, \bar{D}) —i.e., cones for each $\mathbb{A}_{x_i, x_{i+1}}^2$, glued together in agreement with the intersections. There will, however, be a singular point at the origin.

We now describe the construction of B , together with its affine structure, as carried out in [GHK11]. We will from now on denote B by U^{trop} . U^{trop} will in fact have more than just an affine structure—it will be an (oriented³) integral linear manifold (with singularity at 0), meaning that transition functions will be in $\text{SL}_n(\mathbb{Z})$ ($n = 2$ for this case).⁴ It may be helpful to keep in mind that U^{trop} should generalize $N_{\mathbb{R}}$ for (Y, D) and $M_{\mathbb{R}}$ for the mirror.

Let $(\tilde{Y}, \tilde{D}) \rightarrow (\bar{Y}, \bar{D})$ be a toric model of a toric blowup of (Y, D) . Let N be the cocharacter lattice corresponding to \bar{Y} , and $\tilde{\Sigma} \subset N_{\mathbb{R}} := N \otimes \mathbb{R}$ the corresponding fan. As a topological space, U^{trop} is canonically identified with $N_{\mathbb{R}}$, but it will have a more interesting linear structure. The fan Σ with rays corresponding to components of D (rather than components of \bar{D} or \tilde{D} , as with $\tilde{\Sigma}$, which is a refinement of Σ) will play the role of \mathcal{P} .⁵ To simplify notation (i.e., to avoid all the tildes) we will from now on assume that $(\tilde{Y}, \tilde{D}) = (Y, D)$ (i.e., assume that no toric blowups are needed to get a toric model).

Let ρ_i denote the ray corresponding to D_i , and let $\sigma_{i, i+1}$ denote the 2-cells corresponding to $D_i \cap D_{i+1}$. Let v_i be primitive generator of ρ_i . Define⁶ $\sigma_i := \sigma_{i-1, i} \cup \sigma_{i, i+1}$.

Now for our charts we can take $\psi_i : \sigma_i \rightarrow \mathbb{R}^2$ defined to be linear on each cell of Σ and satisfying

$$\psi_i(v_{i-1}) = (1, 0) \quad \psi_i(v_i) = (0, 1) \quad \psi_i(v_{i+1}) = (-1, -\tilde{D}_i^2).$$

²As I’ve mentioned before, my understanding is that in the Gross-Siebert perspective, B and \mathcal{P} should actually be constructed using a log structure on U induced by the boundary D . The approach I’m suggesting here is more accurately a description of how to tropicalize the mirror, which is cheating since a priori we don’t know what the mirror is.

³The orientation is determined by the cyclic ordering of the components of D , so we should note that this ordering is part of our data. This corresponds to orienting the fibers of the SYZ fibration, or equivalently, to choosing a sign for the holomorphic volume form Ω .

⁴More generally, for cluster varieties with rank $2k$ skew-form, we will actually have transition functions in $\text{Sp}(2k, \mathbb{Z})$. However, in the cluster situations we will typically avoid using this $\text{Sp}(2k, \mathbb{Z})$ -structure in favor of a choice of vector space structure corresponding to some choice of seed.

⁵Note that the singular point being the vertex of \mathcal{P} makes this somewhat different from the situation of §3.2.2.

⁶We will continue using the notation σ_i later on, but if (Y, D) does not admit a toric model before taking toric blowups, then it is possible that $s \leq 2$, in which case this definition must be modified. If $s = 2$ the necessary modifications are pretty obvious but notationally messy, and for $s = 1$ a more complicated construction is needed (although [GHK11] implies that it is still doable).

In other words, we are modifying the affine structure of $N_{\mathbb{R}}$ so that in σ_i we have

$$v_{i-1} + D_i^2 v_i + v_{i+1} = 0. \quad (4.1)$$

This is motivated by Corollary 2.7.4, which also shows that if (Y, D) is toric then the linear structure really is that of $N_{\mathbb{R}}$.

Alternatively, we can say that U^{trop} and $N_{\mathbb{R}}$ agree as Σ -piecewise-linear manifolds, and that defining the linear structure can be done by identifying which piecewise-linear functions are actually linear functions. Σ -piecewise-linear functions φ correspond to Weil divisors $W_\varphi := \sum_i \varphi(v_i) D_i$, and the bending parameter of φ along ρ_i is defined to be $W_\varphi \cdot D_i$. In particular, φ is linear along ρ_i if and only if this intersection is linear. This definition is very nice because it tells us the bending parameters of piecewise-linear functions, and because it easily generalizes easily for higher dimensions log Calabi-Yau varieties.

U^{trop} has a canonical set of integral points $U^{\text{trop}}(\mathbb{Z})$ defined by $U^{\text{trop}}(\mathbb{Z}) := \bigcup_i \psi_i^{-1}(\mathbb{Z}^2)$. Alternatively, we can simply define $U^{\text{trop}}(\mathbb{Z})$ to be the set N viewed as a subset of U^{trop} . We will denote $U_0^{\text{trop}} := U^{\text{trop}} \setminus \{0\}$.

Remarks 4.2.1. • Note that U^{trop} depends only on the intersection matrix $(D_i \cdot D_j)_{ij}$. Thus, the choice of toric model is unimportant.

- Toric blowups correspond to refinements of the fan, but do not change the linear structure. Thus, U^{trop} really depends only on U .
- Points of $U^{\text{trop}}(\mathbb{Z})$ correspond to non-negative multiples of boundary divisors on (Y, D) and its toric blowups, with divisors on different toric blowups identified if they correspond to the same discrete valuation of the function field.
- Integral linear manifolds come with a flat connection, pulled back from the flat connection on \mathbb{R}^n which is given by identifying \mathbb{R}^n with its tangent spaces, with $0 \in \mathbb{R}^n$ always identified with 0 in the tangent space. The lattice Λ in TU^{trop} coming from the integral linear structure has non-trivial monodromy about 0. We can identify cones $\sigma \ni p$ in U^{trop} with cones in $T_p U^{\text{trop}}$. This identification takes integer points of U^{trop} to points in Λ and is equivariant under parallel transport within σ .

Example 4.2.2. When trying to draw U^{trop} , I like to consider the universal cover \tilde{U}^{trop} of U_0^{trop} , and draw a linear immersion of a couple sheets of \tilde{U}^{trop} into the plane. This is called the developing map of U^{trop} . For example, if I have $D = D_1 + \dots + D_5$ with each $D_i^2 = -1$ (cf. Example 4.3.2 below), I would draw Figure 4.2.1.

4.3 The Mumford Degeneration

Let P^{gp} be a finite-rank free Abelian group, $P_{\mathbb{R}}^{gp} := P^{gp} \otimes \mathbb{R}$, and $P \subset P^{gp}$ a sub-monoid. We say a function $\varphi : U^{\text{trop}} \rightarrow P_{\mathbb{R}}^{gp}$ is integral Σ -piecewise-linear if it is piecewise linear with bends only along rays of Σ , and $\varphi(U^{\text{trop}}(\mathbb{Z})) \subset P^{gp}$. We say it is convex (resp. strictly convex) if the bending parameters are in P (resp. $P \setminus P^\times$, where P^\times denotes the invertible elements of P). Let φ be a multi-valued strictly⁷ convex integral Σ -piecewise-linear function on U^{trop} . Recall that φ is uniquely determined by its bending parameters.

⁷Strictness of the convexity is not really necessary for much of what we will say.




Figure 4.2.1: A developing map of U^{trop} for $U = \mathcal{M}_{0,5}$. ρ_i^j denotes the ray corresponding to D_i on the i^{th} sheet of \tilde{U}^{trop} .

Examples 4.3.1.

- We will often take $P^{gp} = A_1(Y)$, $P = \text{NE}(Y)$, φ_{NE} having bending parameter $[D_i]$ along ρ_i . However, it is important that P is rational polyhedral. A version of the Cone Theorem (Theorem 3.7 of [KM98]) implies that if D supports a D -ample divisor, then $\text{NE}(Y)$ is indeed rational polyhedral. In general, however, it is necessary to take a rational polyhedral cone σ containing $\text{NE}(Y)$, and define $P = \sigma \cap A_1(Y)$. To simplify things, we will assume from now on that D does support a D -ample divisor, unless otherwise noted.
- Let $\eta : A_1(Y) \rightarrow P^{gp}$ be a homomorphism such that $\eta[\text{NE}(Y)] \subset P$. Then take $\varphi := \eta \circ \varphi_{\text{NE}}$. For example, we could have $P^{gp} = \mathbb{Z}$, $P = \mathbb{Z}_{\leq 0}$, $\eta([C]) := -W \cdot [C]$, where W is the class of a D -ample divisor on Y . We denote this function by φ_W . From our definition of the integral linear structure, if $W = \sum a_i D_i$, then $\varphi_W(v_i) = -a_i$. In particular, φ_W can be represented by a single-valued function.

From now on, φ will be given by some $\eta \circ \varphi_{\text{NE}}$ as above. $p : \mathbb{P} \rightarrow U^{\text{trop}}$ will denote a bundle over U^{trop} which φ may be viewed as a single-valued section of. When viewing φ as a section instead of a $P_{\mathbb{R}}^{gp}$ -valued function, we will denote it by $\tilde{\varphi}$ (i.e., $\tilde{\varphi} := (\text{Id}, \varphi) : U^{\text{trop}} \rightarrow \mathbb{P}$). [GHK11] constructs such a bundle, but rather than giving the construction now, I will wait until we move on to the cluster situation, where we will see that \mathbb{P} can be identified with $\mathcal{A}^{\text{trop}}$.

Let $\Gamma_{\mathbb{R}} \subset \mathbb{P}$ be the graph of φ , and $\Gamma := \Gamma_{\mathbb{R}} \cap [U^{\text{trop}}(\mathbb{Z}) \times P^{gp}]$. If (Y, D) is toric (so U^{trop} is non-singular), then $\Gamma + P$ is a monoid, and the desired mirror family is $\mathcal{X} := \text{Spec } \mathbb{k}[\Gamma + P] \rightarrow \text{Spec } \mathbb{k}[P]$, with the morphism coming from the inclusion of $P \mapsto (0, P)$ (note that this is just a minor generalization of the construction from §2.10). The fiber over 0 is just \mathbb{V}^n , and the general fiber is just U (which in this case is $(\mathbb{C}^*)^2$).

Unfortunately, the singularity at 0 complicates this in general. Instead, we do a local version of this construction as follows: For each $\sigma \in \Sigma$, we can canonically identify $\sigma \times P_{\mathbb{R}}^{gp}$ with a cone in $T_v(\mathbb{P})$ for any $v \in p^{-1}(\sigma)$. Let σ be any cone (not necessarily top dimensional) in Σ containing ρ_i . Using this identification, define a cone

$$\Gamma_{\rho_i, \sigma, \mathbb{R}} := \{x - y \in T_{\tilde{\varphi}(v_i)}(\mathbb{P}) \mid x \in \varphi(\rho_i), y \in \varphi(\sigma)\}. \quad (4.2)$$

Let $\Gamma_{\rho_i, \sigma}$ denote the integer points. Define $R_{\rho_i, \sigma} := \mathbb{k}[\Gamma_{\rho_i, \sigma}]$, $U_{\rho_i, \sigma} = \text{Spec } \mathbb{k}[R_{\rho_i, \sigma}]$.

Now, we want to glue each U_{ρ_i, ρ_i} to $U_{\rho_{i+1}, \rho_{i+1}}$ by identifying $U_{\rho_i, \sigma_{i, i+1}}$ with $U_{\rho_{i+1}, \sigma_{i, i+1}}$. Fixing a representative of φ on $U_i \cup U_{i+1}$ and using parallel transport in $\sigma_{i, i+1}$ to identify tangent spaces, the naive approach is to note that $\Gamma_{\rho_i, \sigma_{i, i+1}} = \Gamma_{\rho_{i+1}, \sigma_{i, i+1}}$ and then use this obvious identification.

Unfortunately, this does not give a satisfactory mirror. [GHK11] shows that this approach, modulo some ideal in $\mathbb{k}[P]$, can be used to get a *formal* smoothing of \mathbb{V}^0 , but this does not extend across the origin. The issue is essentially that the non-trivial monodromy of Λ around $0 \in U^{\text{trop}}$ prevents functions from patching correctly.⁸ The solution is to modify the gluing with something called a *scattering diagram*.

Example 4.3.2. We have just introduced a lot of notation, so let me work a specific example to help clarify what is going on. Consider the toric pair (\bar{Y}, \bar{D}) corresponding to the fan with rays generated by $\rho(\pm 1, 0)$, $(0, \pm 1)$, and $(-1, -1)$. Label the rays ρ_i , $i = 1, \dots, 5$, in counterclockwise order, beginning with $\rho_1 := \mathbb{R}_{\geq 0}(1, 0)$. Take one non-toric blowup on D_1 and one on D_2 to obtain a Looijenga pair (Y, D) with $D_i^2 = -1$ for each i . We can define a chart mapping $\sigma_{1,2} \cup \sigma_{2,3}$ to \mathbb{R}^2 which takes v_1 to $(1, 0)$, v_2 to $(0, 1)$, and v_3 to $(-1, 1)$ (and similarly for any shift in the indices).

Take $W = \sum D_i$, so $\varphi_W(v_i) = -1$ for each i . We can identify \mathbb{P} in this case with $U^{\text{trop}} \times \mathbb{R}$. We will write down R_{ρ_2, ρ_2} explicitly. Using the above chart, $\tau := (\sigma_{1,2} \cup \sigma_{2,3}) \times \mathbb{R}$ can be identified with a cone in \mathbb{R}^3 (i.e., the tangent space to some point in τ) as follows: $\tilde{\varphi}(v_1) = (1, 0, 1)$, $\tilde{\varphi}(v_2) = (0, 1, 1)$, and $\varphi(v_3) = (-1, 1, 1)$. Let us denote $x := z^{(1,0,0)}$, $y := z^{(0,1,0)}$, and $z^{(0,0,-1)}$. Then $R_{\rho_2, \rho_2} := \mathbb{k}\langle xz, yz, x^{-1}yz, z, y^{-1}z^{-1} \rangle$, with relations coming from the embedding in $\mathbb{k}(x, y, z)$. The $y^{-1}z^{-1}$ is to take care of the terms we subtract in Equation 4.2. To simplify this, let $u := xz$, $v = x^{-1}yz$, and $t = yz$. Note that $uvt^{-1} = z$, so we can remove this generator, and then we in fact have no relations. Thus, $R_{\rho_2, \rho_2} = \mathbb{k}[u, v, t^{\pm 1}]$, and $U_{\rho_2, \rho_2} := \text{Spec } R_{\rho_2, \rho_2}$ is (using $m\text{Spec}$) just $\mathbb{k}^2 \times \mathbb{k}^*$. The map to $\text{Spec } \mathbb{k}[P] = \text{Spec } \mathbb{k}[z] \cong \mathbb{A}$ comes from the ring homomorphism $z \mapsto uvt^{-1}$.

Now, $R_{\rho_2, \sigma_{1,2}} = (R_{\rho_2, \rho_2})_u$ (i.e., we adjoin u^{-1}), and Spec of this corresponds to taking the complement of $Z(u)$ (the zero set of u). Note that this ring is canonically equal to $R_{\rho_1, \sigma_{1,2}}$. Thus, we have an obvious canonical way to glue U_{ρ_1, ρ_1} to U_{ρ_2, ρ_2} along this common open subset. This is what I said above is the *wrong* gluing.

4.4 The consistent scattering diagram

When we deal with the general cluster variety situation, I will introduce a more general and precise description of what a scattering diagram is. For now, let us just say that a scattering diagram \mathfrak{d} includes the data of a set of rays in U^{trop} with associated functions which satisfy certain conditions. These functions are used to define certain ring automorphisms, and for the “consistent” scattering diagram which we will define, these automorphisms make it possible to construct the scheme we were after in §4.3.

For a ray $\rho \subset U^{\text{trop}}$ with rational slope, let D_ρ be the corresponding boundary divisor in (\tilde{Y}, \tilde{D}) (some toric blowup π of (Y, D)). Let $\beta \in H_2(\tilde{Y}, \mathbb{Z})$ with $k_\beta := \beta \cdot D_\rho \in \mathbb{Z}$, and $\beta \cdot D_{\rho'} = 0$ for $\rho \neq \rho'$ (in particular, the Maslov index is $2k_\beta$). Let $F_\rho := \overline{D} \setminus D_\rho$, $\tilde{Y}_\rho^\circ := \tilde{Y} \setminus F_\rho$, and $D_\rho^0 := D \setminus F_\rho$.

⁸I am surprised that the gluing morphisms are even compatible in the sense necessary to get a well-defined scheme, although [GHK11] claims this is “easily checked.” Perhaps modding out by the ideal is what makes it possible?

Now, define $\overline{\mathcal{M}}(\tilde{Y}_\rho^\circ/D_\rho^\circ, \beta)$ to be the moduli space of stable relative maps⁹ of genus 0 curves to \tilde{Y}_ρ° , representing the class β and intersecting D_ρ° at one unspecified point with multiplicity k_β . This moduli space has a virtual fundamental class with virtual dimension 0. Furthermore, $\overline{\mathcal{M}}(\tilde{Y}_\rho^\circ/D_\rho^\circ, \beta)$ is proper over $\text{Spec } k$ (cf. Theorem 4.2 of [GPS09] or Lemma 3.2 of [GHK11]). Thus, we can define the relative Gromov-Witten invariant N_β as

$$N_\beta := \int_{[\overline{\mathcal{M}}(\tilde{Y}_\rho/D_\rho, \beta)]^{vir}} 1.$$

If $N_\beta \neq 0$, we call β an \mathbb{A}^1 -class.

Remark 4.4.1. We of course have not covered enough to necessarily know what all this means. The point though is that N_β is a “virtual” count of the number of curves in \tilde{Y} of class β which intersect D at precisely one point on D_ρ° . We will later see a more combinatorial construction of scattering diagrams. The purpose of this approach is to relate the scattering diagram to the counts of holomorphic disks that mirror symmetry tells us to consider, and also to make it clear that the construction is canonical (it takes some work to see that the combinatorial approach does not depend on the choice of toric model).

Recall that η denotes a homomorphism from $\text{NE}(Y)$ to P . We now define

$$f_\rho := \exp \left[\sum_{\beta} k_\beta N_\beta z^{\eta(\pi_*(\beta)) - \tilde{\varphi}(k_\beta v_\rho)} \right] \in R_{\rho, \rho}.$$

Here, the sum is over all $\beta \in \text{NE}(\tilde{Y})$ which have 0 intersection with all boundary divisors except for D_ρ .

Example 4.4.2. Let $\bar{Y} = \mathbb{P}^2$ and let $\bar{D} = \bar{D}_1 + \bar{D}_2 + \bar{D}_3$ be a triangle of generic lines in \bar{Y} . Consider the pair (Y, D) obtained by performing a single non-toric blowup at a point on \bar{D}_1 . Let $\beta = E_1$ be the exceptional divisor. Then $N_\beta = 1$. Due to the stacky nature of $\overline{\mathcal{M}}(Y_\rho/D_\rho, \beta)$, N_β might not always be a positive integer. For example, with Y and β as above, we have $N_{k\beta} = \frac{(-1)^{k-1}}{k^2}$ (see [GPS09], Proposition 6.1).

These multiple covers of E_1 are the only \mathbb{A}^1 classes for D_1 , so we can compute f_{ρ_1} . Suppose $P^{gp} := A_1(Y)$ and $\eta := \text{Id}$. We have

$$\begin{aligned} f_{\rho_1} &= \exp \left[\sum_{k \in \mathbb{Z}_{>0}} k \left(\frac{(-1)^{k-1}}{k^2} \right) z^{k[E_1] - \varphi(kv_{\rho_1}) - kv_{\rho_1}} \right] \\ &= 1 + z^{[E_1] - \varphi(v_{\rho_1}) - v_{\rho_1}}. \end{aligned}$$

More generally, if the only \mathbb{A}^1 -classes hitting D_ρ are a set $\{E_1, \dots, E_k\}$ of (-1) -curves, along with their multiple covers, then

$$f_\rho = \prod_{i=1}^k \left(1 + z^{\eta(E_i) - \tilde{\varphi}(v_\rho)} \right) \tag{4.3}$$

⁹For details on relative Gromov-Witten invariants, see [Li02], or see [GPS09] for a treatment of this particular situation.

Example 4.4.3. Now consider the situation of Example 4.3.2 above. Here, for each i , there is exactly one interior (-1) -curve E_i (an exceptional divisor for some toric model) hitting each D_i , and the multiple covers of these are the only \mathbb{A}^1 -classes. So Equation 4.3 applies for each ρ_i and becomes the very simple function $f_{\rho_i} = 1 + z^{(-v_i, 0)}$. I note that the rays of the scattering diagram are exactly those drawn in Figure 4.2.1, although it is rare for the scattering diagram and the fan to coincide like this.

4.4.1 Modifying the gluing

Now for the modified gluing, to glue $U_{\rho_i, \sigma_{i, i+1}}$ to $U_{\rho_{i+1}, \sigma_{i, i+1}}$ we use the *path-ordered product* along a path γ in $\sigma_{i, i+1}$ from v_i to v_{i+1} which transversely crosses rays in the counterclockwise direction. That is, whenever γ crosses a scattering ray ρ , apply

$$z^v \mapsto z^v f_{\rho}^{n_{\rho}(v)}. \quad (4.4)$$

n_{ρ} here is the primitive element of the cotangent space containing ρ in its kernel and taking positive values on $-\gamma'(t)$. If ρ_i is a scattering ray, we have to specify whether or not the scattering automorphism for crossing ρ_i should be applied. I like to use the notation $U_{\rho_i, \sigma}^+$ (resp. $U_{\rho_i, \sigma}^-$) to indicate that we view the endpoint of γ on ρ_i as actually being infinitesimally counterclockwise (resp. clockwise) of ρ_i when the scattering automorphism for crossing ρ_i counterclockwise has been applied.

The problem now is that there are often infinitely many scattering rays between ρ_i and ρ_{i+1} . We can deal with this by working modulo \mathfrak{m}^k . Modulo this ideal, there are only finitely many rays with non-trivial attached function—this is because there are only finitely many points in $P \setminus k\mathfrak{m}_P$, and \mathbb{A}^1 -classes with non-vanishing contributions live in $\text{NE}(Y) \setminus k\mathfrak{m}_{\text{NE}(Y)}$.

Now, gluing all of the $U_{\rho_i, \sigma, k}^{\pm}$'s as above (using the subscript k to denote that we are working modulo \mathfrak{m}^k), we construct an infinitesimal deformation \mathcal{X}_0^k of \mathbb{V}_0^n , over the base $\text{Spec } \mathbb{k}[P]/\mathfrak{m}^k$. Let $R_k := \Gamma(\mathcal{X}_0^k, \mathcal{O}_{\mathcal{X}_0^k})$. The scattering diagram we have used is “consistent,” as we will explain below, and [GHK11] shows this implies that

$$\mathcal{X}^k := \text{Spec } R_k$$

gives an infinitesimal smoothing of \mathbb{V}^n , flat over $\text{Spec } \mathbb{k}[P/\mathfrak{m}^k]$. Let $\widehat{R} := \varprojlim R_k$, and define

$$\widehat{\mathcal{X}} := \text{Spf } \widehat{R}$$

(see Appendix A.2 for a brief explanation of formal schemes). [GHK11] shows that this¹⁰ yields a flat formal smoothing of \mathbb{V}^n over $\text{Spf } \mathbb{k}[[P]]$. Furthermore, when D supports a D -ample divisor, then this extends to an algebraic family

$$\mathcal{X} := \text{Spec } \Gamma(\widehat{\mathcal{X}}, \mathcal{O}_{\widehat{\mathcal{X}}}),$$

giving a flat affine smoothing of \mathbb{V}^n over $\text{Spec } \mathbb{k}[P]$.

All of this of course requires that \mathcal{X}_k actually has global regular functions, which was not the typically case before we introduced the scattering diagram. We now describe the theta functions, which give a canonical basis of such global regular functions.

¹⁰If $\text{NE}(Y)$ is not rational polyhedral (so in particular, D does not support a D -ample divisor), then [GHK11] actually requires completing with respect to a slightly different ideal J .

4.5 Broken Lines and Theta Functions

Before defining the theta functions, we must understand “broken lines,” as these record the monomial terms in local expressions of the theta functions.

Definitions 4.5.1. Let $q \in U^{\text{trop}}(\mathbb{Z})$, and $Q \in U_0^{\text{trop}}$. A *broken line* γ with $\text{Ends}(\gamma) = (q, Q)$ is the data of a continuous map $\gamma : (-\infty, 0] \rightarrow U^{\text{trop}}$, values $-\infty < t_0 < t_1 < \dots < t_s = 0$, and for each $t \neq t_i$, $i = 0, \dots, s$, an associated monomial $c_t z^{m_t} \in R_{\gamma(t)} := \mathbb{k}[\Lambda_{\gamma(t)} \mathcal{P}]$ with $c_t \in \mathbb{k}$ and $r_*(m_t) = -\gamma'(t)$, such that:

- $\gamma(0) = Q$
- $\gamma_0 := \gamma|_{(-\infty, t_0]}$ and $\gamma_i := \gamma|_{[t_{i-1}, t_i]}$ are geodesics (i.e., straight lines with constant velocities).
- For all $t \ll t_0$, $\gamma(t)$ is in some fixed convex cone σ_q containing q , and $m_t = \tilde{\varphi}(q)$ under parallel transport in σ_q .
- For all $a \in (t_{i-1}, t_i)$ (or $(-\infty, t_0)$ for $i = 0$) and $b \in (t_i, t_{i+1})$, and all relevant $R_{\gamma(t)}$ ’s identified using parallel transport along γ , we have that $\gamma(t_i)$ is contained in a scattering ray ρ , and

$$c_b z^{m_b} = (c_a z^{m_a})(c_\rho z^{m_\rho})$$

where $c_\rho z^{m_\rho}$ is any term in the formal power series expansion of $f_\rho^{\langle n_\rho, r_*(m_a) \rangle}$ (so $c_b z^{m_b}$ is a monomial term from the expansion of Equation 4.4).

We are at last ready to define the theta functions. Define $\vartheta_0 := 1$. For $q \neq 0$, define $\vartheta_q|_{U_{\rho, \rho}}$ (maybe with subscript k and superscript \pm) by picking a point Q infinitesimally close¹¹ to ρ and defining

$$\vartheta_q|_{U_{\rho, \rho}} := \sum_{\gamma| \text{Ends}(\gamma)=(q, Q)} c_\gamma z^{m_\gamma} \quad (4.5)$$

where $c_\gamma z^{m_\gamma}$ denotes the monomial attached to the final straight segment of γ .

For this to be well-defined, it must commute with the gluings by path-ordered products. When this happens, we say the scattering diagram is “consistent.” Much of §3 of [GHK11] is devoted to relating the scattering diagram we defined above with the scattering diagrams from [GPS09] and using this to prove consistency. Eventually I’ll give a sketch of how this goes. For now, assuming the consistency of our scattering diagram, we have:

Theorem 4.5.2.

$$\Gamma(\mathcal{X}^k, \mathcal{O}_{\mathcal{X}^k}) = \bigoplus_{q \in U^{\text{trop}}(\mathbb{Z})} (\mathbb{k}[P]/\mathfrak{m}^k) \cdot \vartheta_q.$$

If D supports a D -ample divisor, then

$$\Gamma(\mathcal{X}, \mathcal{O}_{\mathcal{X}}) = \bigoplus_{q \in U^{\text{trop}}(\mathbb{Z})} \mathbb{k}[P] \cdot \vartheta_q.$$

¹¹When working modulo \mathfrak{m}^k , we just need there to be no scattering rays between Q and ρ . When dealing with \mathcal{X} it is more difficult to state precisely what is meant by “infinitesimally close.” [Manb] gives a precise definition, but it may be intuitively obvious what is meant by this, and one could also just work modulo \mathfrak{m}^k for each k and take the limit.

Example 4.5.3. Consider an example with only finitely many scattering rays, each of the form from Equation 4.3. This is the case in the situation from Examples 4.3.2 and 4.4.3. Here, we can avoid modding out by the ideals and just work directly with the $R_{\sigma_{i,i+1}, \sigma_{i,i+1}}$'s—if we use the R_{ρ_i, ρ_i} 's, we actually get a compactification of \mathcal{X} , as is apparent already in Example 4.3.2.

In this situation, a broken line which begins in some cone $\sigma_{i,i+1}$ will never return to $\sigma_{i,i+1}$ after leaving it, so for $q_i \in \sigma_{i,i+1}$, $\vartheta_{q_i}|_{U_{\sigma_{i,i+1}, \sigma_{i,i+1}}} = z^{\tilde{\varphi}(q_i)}$. If we cross into, say, $\sigma_{i+1, i+2}$, then this transforms into $z^{\tilde{\varphi}(q_i)} f_{\rho_{i+1}}^{n_{\rho_{i+1}}(q_i)}$. In the cluster language, $z^{\tilde{\varphi}(q)}$ is a product of cluster monomials, and the transformation above is just a mutation.

The definition easily implies the following multiplication formula:

Theorem 4.5.4. *Given $q_1, q_2, q \in U^{\text{trop}}(\mathbb{Z})$, the ϑ_q -coefficient of $\vartheta_{q_1} \cdot \vartheta_{q_2}$ is given (modulo \mathfrak{m}^k if we are talking about \mathcal{X}^k) by*

$$\sum_{\substack{\gamma_i, i=1,2 \\ \text{Ends}(\gamma_i)=(q_i, Q) \\ m_{\gamma_1} + m_{\gamma_2} = q}} c_{\gamma_1} c_{\gamma_2}, \quad (4.6)$$

where Q is a point infinitesimally close to ρ_q .

The key to the proof of this is to note that ϑ_q is the only theta function with a z^q term along ρ_q . Similar formulas are easy to come up with for products of more than two theta functions. Note that Theorem 4.5.4 together with Theorem 4.5.2 completely determine \mathcal{X}^k and \mathcal{X} .

Exercise 4.5.5. Here is a fun exercise that offers practice with the multiplication formula 4.6 and also proves something interesting. Suppose (Y, D) is the cubic surface obtained from $(\mathbb{P}^2, \overline{D} = \overline{D}_1 + \overline{D}_2 + \overline{D}_3)$ by blowing up two points on each \overline{D}_i . Each $D_i^2 = -1$, and one easily checks that the monodromy of U^{trop} is $\mu = -\text{Id}$ (i.e. for any ray $\rho \subset U^{\text{trop}}$, $U^{\text{trop}} \setminus \rho$ can be identified with a half-plane in \mathbb{R}^2). The mirror family can be identified with the moduli space of flat $\text{SL}_2(\mathbb{C})$ -connections on a sphere S^2 with 4 punctures. According to [FG06] (cf. §7.2.3), primitive points in $U^{\text{trop}}(\mathbb{Z})$ correspond bijectively with isotopy classes of simple loops γ in S which are not contractible or homotopic to a boundary loop. The theta function $\vartheta_{k\gamma}$ corresponding to $k\gamma$, γ primitive and $k \in \mathbb{Z}_{\geq 0}$, should take the value $\text{Tr}(\text{hol}_{k\gamma} \nabla)$ at the connection ∇ in the moduli space. Assuming this is correct for ϑ_γ with γ primitive (this has been proven in [GHK]), prove it for $\vartheta_{k\gamma}$.

Hint: Use Equation 4.6 to write $\vartheta_{k\gamma}$ as a polynomial in ϑ_γ . Only straight lines show up in this computation, so it is not too difficult. On the other hand, if $\text{hol}_\gamma(\nabla)$ has eigenvalues λ_1 and λ_2 , then $\text{Tr}(\text{hol}_{k\gamma} \nabla) = \lambda_1^k + \lambda_2^k$ can be written as a polynomial in $\text{Tr}(\text{hol}_\gamma \nabla) = \lambda_1 + \lambda_2$ (since $\lambda_1 \lambda_2 = 1$). The goal now is to show that these two polynomials are the same. With either approach, one finds that the coefficient of $\vartheta_{c\gamma}$ in ϑ_γ^k should be the number of k -tuples $(\epsilon_1, \dots, \epsilon_k)$ such that each $\epsilon_i = \pm 1$ and $\sum \epsilon_i = c$.

4.6 Compactifications

When D supports a D -ample divisor W , we can in fact compactify the mirror family \mathcal{X} to get a family of deformations of the original (Y, D) . Here is one way to define this compactification: For $W = \sum a_i D_i$, let Δ_W denote the polytope in U^{trop} containing 0 and bounded by lines denoted $L_{\vartheta_i}^{a_i}$.

This is the line which goes to infinity parallel to v_i with the origin on the left and with “lattice distance” a_i from the origin—i.e., if we parallel transport v_i to a point p on the line, then $p \wedge v_i = a_i$. Note that the monodromy of U^{trop} might cause this straight line to actually self-intersect. One can show that D -ampleness of W implies strict convexity of Δ_W .

Now, much like in the toric situation, we have a couple ways to define the compactification. One way is to locally apply the normal fan construction to the polytope $\tilde{\Delta}_W := [\tilde{\varphi}(\Delta_W) + P]$. This is essentially the same as the construction of \mathcal{X} , but now with extra affine open subsets (containing the compactifying divisors) corresponding to the vertical faces of $\tilde{\Delta}_W$ (i.e., the faces living over the faces of Δ_W).

Alternatively, one can define

$$\mathcal{X}_W := \text{Proj} \bigoplus_{\substack{k \in \mathbb{Z}_{\geq 0} \\ q \in k\Delta_W}} \mathbb{k}[P] \cdot (\vartheta_q, z^k)$$

Multiplication of the (ϑ_q, t) 's is the usual multiplication of theta functions in the first variable and is given by the obvious addition of exponents in the second. The grading is of course given by k . Of course, for this definition to make sense, one must check that if $q_i \in k_i\Delta_W$, $i = 1, 2$, then $\vartheta_{q_1}\vartheta_{q_2} \in (k_1 + k_2)\Delta_W$. I proved this in [Manb].

[GHK] shows that (Y, D) can be identified with a fiber of \mathcal{X}_W (after making certain choices which they describe).¹² This means that we get theta functions on the original space log Calabi-Yau surface! Furthermore, the line bundle \mathcal{L}_W corresponding to W has a canonical basis of global sections given by the ϑ_q 's with $q \in \Delta_W$. Thus, Δ_W generalizes the “normal polytope” from the toric situation for \mathcal{L}_W (since U^{trop} has $\text{SL}_2(\mathbb{Z})$ -monodromy, we must in general use the *parallel polytope* rather than the normal polytope).

4.7 More scattering diagrams

Here we look at scattering diagrams in vector spaces and relate them to the consistent scattering diagram we saw in §4.4. As before, N and P^{gp} denote free Abelian groups, $M := \text{Hom}(N, \mathbb{Z})$, and P is a finitely generated submonoid of P^{gp} (that is, $P = \sigma_P \cap P^{gp}$ for some convex rational polyhedral cone $\sigma_P \subset P_{\mathbb{R}}^{gp}$). We do not assume that N has rank 2. Consider the ring

$$R := \varprojlim_k (\mathbb{k}[N] \otimes \mathbb{k}[P]/\mathfrak{m}^k),$$

where \mathfrak{m} is the maximal ideal of $\mathbb{k}[P]$ generated by all z^p with $p \in P \setminus P^\times$.

4.7.1 The Tropical Vertex Group

I think we will only use this subsection for one step in the proof of Theorem 4.7.4, which we only prove in the two-dimensional case anyways. so the reader can get away with skipping this. [I MIGHT REMOVE THIS AND THE 2D PROOF FROM THE FINAL DRAFT]

¹²I believe that [GHK]'s proof involves their Torelli theorem for Looijenga pairs from [GHK13b]. In the cases where Example 4.5.3 applies, this can be checked directly.

The *module of log derivations* is defined to be $\Theta(R) := R \otimes_{\mathbb{Z}} M$, with action on R give by $(f \otimes m)(z^n) = f\langle n, m \rangle z^n$. $\langle \cdot, \cdot \rangle$ here denotes the dual pairing between N and M . One often denotes $f \otimes m$ by $f\partial_m$.

The commutator makes $\Theta(R)$ into a Lie algebra with bracket

$$[z^{n_1}\partial_{m_1}, z^{n_2}\partial_{m_2}] = z^{n_1+n_2}\partial_{\langle n_2, m_1 \rangle m_2 - \langle n_1, m_2 \rangle m_1}.$$

Let $d : P^{gp} \rightarrow \mathbb{Z}$ be a linear function which is positive on $P \setminus P^\times$ and such that $P_{\leq k} := \{n \in P \mid d(n) \leq k\}$ is finite for all k . For example, let \mathfrak{m}_P denote the monoid ideal of P generated by all elements of $P \setminus P^\times$, and let $d(p)$ be the largest value of k such that $p \in \mathfrak{m}_P^k$.

Consider the Lie subalgebra (over \mathbb{k}) of $\mathfrak{m}\Theta(\widehat{\mathbb{k}[P]})$ defined by

$$\mathfrak{v}^{>k} := \bigoplus_{d(n) > k} z^n (\mathbb{k} \otimes n^\perp).$$

Let $\mathfrak{v}^{\leq k} := \mathfrak{v} / \mathfrak{v}^{>k}$, and define $\mathbb{V}^{\leq k} := \exp(\mathfrak{v}^{\leq k})$. The group $\mathbb{V} := \varprojlim \mathbb{V}^{\leq k} \subset \text{Aut}(\widehat{\mathbb{k}[P]})$ is called the *tropical vertex group*.

Now, suppose $M_{\mathbb{R}}$ is equipped with an integral skew-symmetric form $\{\cdot, \cdot\}$. We define \mathbb{V}^s similarly to \mathbb{V} , taking an inverse limit of the exponentials of the Lie algebras

$$\mathfrak{v}^{>k,s} := \bigoplus_{d(n) > k} z^n (\mathbb{k} \otimes \{n, \cdot\}).$$

I might call \mathbb{V}^s the *skew tropical vertex group*.

Example 4.7.1. An element $\log(f)\partial_n \in \varprojlim \mathfrak{v}^{\leq k}$ exponentiates to the automorphism $z^m \mapsto z^m f^{\langle n, m \rangle}$.

4.7.2 Scattering Diagrams

Suppose that N is equipped with a skew-symmetric form $\{\cdot, \cdot\}$.

Definition 4.7.2. A *wall* in $N_{\mathbb{R}}$ is the data $(\mathfrak{d}, f_{\mathfrak{d}})$, where

- $f_{\mathfrak{d}} = \sum_{k \in \mathbb{Z}_{\geq 0}} c_k z^{kn_{\mathfrak{d}}} \in R$, where $c_k \in \mathbb{k}[P]$ and $n_{\mathfrak{d}}$ is a primitive vector in N . $-n_{\mathfrak{d}}$ is called the direction of the wall. Let $m_{\mathfrak{d}} := \{n_{\mathfrak{d}}, \cdot\} \in M$. We assume $m_{\mathfrak{d}}$ is primitive.
- $\mathfrak{d} \subset m_{\mathfrak{d}}^\perp \subset N_{\mathbb{R}}$ is a full-dimensional (in $m_{\mathfrak{d}}^\perp$) convex (but not necessarily strictly convex) rational polyhedral cone spanning a hyperplane $j \subset m_{\mathfrak{d}}^\perp$. We say the wall is incoming if $n_{\mathfrak{d}} \subset \mathfrak{d}$, and outgoing otherwise. \mathfrak{d} is called the support of the wall.¹³
- $f_{\mathfrak{d}} \equiv 1 \pmod{\mathfrak{m}}$.

We might sometimes denote a wall $(\mathfrak{d}, f_{\mathfrak{d}})$ as just \mathfrak{d} .

Definition 4.7.3. A *scattering diagram* \mathfrak{D} is a set of walls such that for each integer $k > 0$, there are only finitely many walls $(\mathfrak{d}, f_{\mathfrak{d}}) \in \mathfrak{D}$ with $f_{\mathfrak{d}} \not\equiv 1 \pmod{\mathfrak{m}^k}$.

¹³We explain some of the terminology. $-n_{\mathfrak{d}}$ is called the direction of the wall because it is the direction in which broken lines can bend when crossing the wall. In dimension 2, walls are supported on rays or lines containing the origin. Being outgoing in this case means that the wall is a ray whose direction points away from the origin (outward), while being incoming means that there is some point at which the direction points towards the origin (inward).

Let $\text{Joints}(\mathfrak{D})$ denote the union of all the boundaries of walls of \mathfrak{D} and all codimension 2 (in $N_{\mathbb{R}}$) intersections of walls of \mathfrak{D} . For a path γ not intersecting $\text{Joints}(\mathfrak{D})$, one may define the path-ordered product $\theta_{\gamma, \mathfrak{D}}$ just like we did in §4.4.1. Note that by Example 4.7.1 the automorphism for crossing a wall \mathfrak{d} is just

$$\exp(\log(f_{\mathfrak{d}})\partial_{\pm\{n_{\mathfrak{d}}, \cdot\}}),$$

where the sign is chosen to make $\pm\{n_{\mathfrak{d}}, -\gamma'(t_0)\} > 0$ (t here denotes the time at which γ crosses \mathfrak{d}). Thus, these automorphisms live in \mathbb{V}^s .

Given a scattering diagram \mathfrak{D} , we can combine several walls with the same support into one by multiplying together their attached functions—the new scattering diagram is said to be *equivalent* to \mathfrak{D} (similarly in the reverse direction with factoring).

Theorem 4.7.4 ([GS11a],[KS06]). *Let \mathfrak{D} be a scattering diagram with only finitely many walls. Then there is a scattering diagram $S(\mathfrak{D})$ containing \mathfrak{D} such that $S(\mathfrak{D}) \setminus \mathfrak{D}$ consists only of outgoing walls, and each path-ordered product depends only on the endpoints of the path. $S(\mathfrak{D})$ is unique up to equivalence.*

Proof. We only give the complete proof in dimension 2, and then comment on the general situation.

We will inductively construct scattering diagrams \mathfrak{D}_k such that the claim holds modulo \mathfrak{m}^{k+1} , and then define $S(\mathfrak{D}) := \bigcup_k \mathfrak{D}_k$ (note that this is not a *disjoint* union). Let $\mathfrak{D}_0 := \mathfrak{D}$. The claim holds for $k = 0$ because all the scattering functions are required to be trivial modulo \mathfrak{m} . We now describe how to obtain \mathfrak{D}_k from \mathfrak{D}_{k-1} .

For any joint $p \in \text{Joints}(\mathfrak{D}_k)$ (in the 2-dimensional situation, $\{0\}$ is the only joint), let γ_p be a simple closed loop around p which does not contain any other points in $\text{Joints}(\mathfrak{D}_k)$. Then by the inductive assumption, we have a unique expansion

$$\theta_{\gamma_p, \mathfrak{D}_k} = \exp\left(\sum_{i=1}^s c_i z^{n_i} \partial_{\{n_i, \cdot\}}\right),$$

with $c_i \in \mathfrak{m}^k$ and $n_i \in N \setminus \{0\}$ (this is where we use §4.7.1). Let

$$\mathfrak{D}_k[p] := \{p + \mathbb{R}_{\geq 0}n_i, 1 \pm c_i z^{n_i}\},$$

where the sign is chosen so that the contribution to $\theta_{\gamma, \mathfrak{D}_k}$ is $\exp(-c_i z^{n_i} \partial_{\{n_i, \cdot\}})$ (this is determined by the direction in which γ crosses the wall). Let

$$\mathfrak{D}_k := \mathfrak{D}_{k-1} \bigcup_{p \in \text{Joints}(\mathfrak{D}_{k-1})} \mathfrak{D}_k[p].$$

Since automorphisms coming from $\mathfrak{D}_k[p]$ are central modulo \mathfrak{m}^{k+1} , we easily see in dimension 2 that \mathfrak{D}_k has the desired properties. In higher dimensions, checking this is significantly more complicated because walls extending from one joint might contain other joints. \square

In this proof, we followed [GPS09], which was itself following [KS06]. [GS11a] followed this approach to prove the higher dimensional cases (much more generally). I believe that [GHKK], in their most recent rewrite, follow the quite different approach of [KS13], with the methods of [GS11a] appearing in their appendix.

4.7.3 Relation to the [GHK11] situation

Recall that each toric model $\pi : (\tilde{Y}, \tilde{D}) \rightarrow (\bar{Y}, \bar{D})$ induces a linear structure on U^{trop} by identifying it with $N_{\mathbb{R}, \bar{Y}} := N_{\bar{Y}} \otimes \mathbb{R}$, where $N_{\bar{Y}}$ is the cocharacter lattice of \bar{Y} . Let $\bar{\varphi} : N_{\bar{Y}} \rightarrow A_1(\bar{Y}, \mathbb{R})$ be an integral Σ -piecewise-linear function with bending parameter $[\bar{D}_i]$ along ρ_i ($\bar{\varphi}$ exists and is unique up to a choice of linear function, cf. [GHK11], Lemma 1.13). Let $P_0^{gp} := \pi^*(A_1(\bar{Y}, \mathbb{Z}))$, and let $E \subset A_1(Y, \mathbb{Z})$ be the lattice generated by the exceptional divisors of π . Now $P^{gp} := A_1(Y, \mathbb{Z}) = P_0^{gp} \oplus E$, and $P := \pi^*(\text{NE}(\bar{Y})) \oplus E$. Define $\varphi_\pi : U^{\text{trop}} \rightarrow P^{gp}$ by $\varphi_\pi(u) := (\pi^*(\bar{\varphi}(u)), 0)$.

Suppose that π includes b_i non-toric blowups on D_{n_i} , with corresponding exceptional divisors E_{ij} , $j = 1, \dots, b_i$. Now, let \mathfrak{D}_0 be the scattering diagram in $N_{\mathbb{R}}$ with walls

$$\left\{ \mathbb{R}n_i, \prod_{j=1}^{b_i} \left(1 + z^{\bar{\varphi}_\pi(n_i) - [E_{ij}]} \right) \mid i = 1, \dots, n \right\}. \quad (4.7)$$

Now consider $\mathfrak{D} := S(\mathfrak{D}_0)$. Since all of the rays of $\mathfrak{D} \setminus \mathfrak{D}_0$ are outgoing, any broken line crossing these scattering rays can only bend away from the origin. Thus, it is only broken lines crossing $\mathbb{R}_{\geq 0}n_i$ that can bend towards the origin.

U^{trop} with its canonical integral linear structure now comes from modifying $N_{\bar{Y}, \mathbb{R}}$ so that lines which take the maximal allowed bend towards the origin are actually straight. Furthermore, if we break our initial scattering rays up into two outgoing rays by negating the exponents of the $\mathbb{R}_{\geq 0}n_i$ parts of the initial rays, then $S(\mathfrak{D}_0)$ becomes our consistent scattering diagram \mathfrak{D} in U^{trop} from before. The interpretation of $S(\mathfrak{D}_0)$ in terms of counts of holomorphic curves is the main result of [GPS09]. Path-ordered products depending only on the endpoints of paths implies consistency of the scattering diagram by a result in [CPS]. §3 of [GHK11] works through this in detail in order to prove the consistency of their scattering diagrams.

Chapter 5

Cluster Algebras and Cluster Varieties

In this chapter we will define cluster algebras and cluster varieties, and we will give some motivating examples and applications. For a more combinatorial approach, along with an introduction to categorification of cluster algebras (which I do not plan to discuss), [Kel12] gives a good survey.

5.1 Some Motivation

Cluster algebras can be somewhat difficult to motivate—the definition appears to be some ugly combinatorial formulas that just come from nowhere. As with anything in math, a good motivation is one that the reader finds interesting or relevant. Since we should by now be familiar with toric models (recall §4.1), we will use Example 5.1.2 below to try and introduce cluster algebras as a natural framework for studying varieties with toric models (this is the viewpoint of [GHK13a]).

First, to hopefully convince the reader that cluster algebras are at least worth studying, we touch on some of the more standard motivation for their study. The original motivation in [FZ02] was to develop an algebraic framework for understanding Lusztig’s dual canonical bases and total positivity (cf. [Lus10] and [Lus94]). We will touch slightly on these applications (avoiding quantum groups), but since we have not covered (and I do not know) the background for this, we will mostly take a more geometric viewpoint. However, we do begin by briefly mentioning some of the standard motivating examples.

The basic idea behind the typical algebraic approach is that for each “seed,” we have cluster variables A_1, \dots, A_s which freely generate a part of the cluster algebra. We can “mutate” the seed in a combinatorial way to obtain a new seed, with all but one of the cluster variables remaining the same. The other is replaced by $A'_i := \frac{1}{A_i}(M_1 + M_2)$ (called an exchange relation) for M_1 and M_2 certain monomials in the A_j ’s, $j \neq i$, specified by the seed data.

Example 5.1.1. $\mathrm{SL}_2(\mathbb{C})$ admits a natural cluster structure. We view the group as matrices with entries $a, b, c, d \in \mathbb{C}$ satisfying $ad - bc = 1$. We can take one seed to have cluster variables a, b , and c , and mutate to another seed with cluster variables d, b , and c . The exchange relation is given by $d = \frac{1}{a}(1 + bc)$. b and c are called frozen variables since we do not mutate with respect to them.

Now view $\mathrm{SL}_2(\mathbb{C})$ as a variety. For each seed, we have an affine open subset $\mathrm{Spec} \mathbb{C}[A^{\pm 1}, b, c] \cong \mathbb{C} \times \mathbb{C}^2$ where $A := a$ or d (we will often invert the frozen variables too so that each seed corresponds to an algebraic torus, but obtaining partial compactifications by not inverting frozen variables is common). The first seed (with $A = a$) misses points where $a = 0$, while the second misses points where $d = 0$. Together, they cover all of $\mathrm{SL}_2(\mathbb{C})$ except for the codimension 2 subset where $a = d = 0$ (and so $bc = -1$). Missing codimension 2 subsets is common, but not important—we are mainly interested in the spaces of global sections of line bundles on cluster varieties, and these are not affected by changes in codimension 2.

We will see that many other well known spaces admit cluster structures like $\mathrm{SL}_2(\mathbb{C})$. For example, [FZ02] mentions Grassmanians $Gr_{2,n}$ (generalized to $Gr(k, n)$ in [Sco06]), where certain Plücker coordinates form the cluster variables. They also mention examples of double Bruhat cells of semisimple complex Lie groups (including 5.1.1), whose cluster structure was described in general in [BFZ05] (we may cover this application later on).

On the other hand, Fock and Goncharov were motivated in [FG06] and [FG09] by coordinates on moduli of local systems on punctured Riemann surfaces. In particular, [Pen87] (well before the invention of cluster algebras) described coordinates called λ -lengths on decorated Teichmüller space, which may be viewed as a moduli space of (decorated) $\mathrm{PSL}_2(\mathbb{R})$ local systems. The λ -lengths depend on a choice of “ideal triangulation” of the punctured Riemann surface, and changing the triangulation by a “Ptolemy transform” (i.e., taking a quadrilateral in the triangulation and flipping the diagonal to obtain a new triangulation) produces a new coordinate system. Fock and Goncharov’s cluster coordinates can be viewed as a complexification (and generalization) of these λ -lengths, with mutation generalizing the Ptolemy transforms. We will see this in more detail later on.

From the perspective Gross, Hacking, and Keel (cf. [GHK13a]), cluster varieties are essentially the log Calabi-Yau varieties which admit toric models like those from §4.1 (and thus include all log Calabi-Yau surfaces by Proposition 4.1.3). This is the viewpoint most useful for us.

Example 5.1.2. Consider a lattice $\bar{N} = \mathbb{Z}^2$, $u \in \bar{N}$ primitive. Let Σ be the fan in \bar{N} with rays corresponding to u and $-u$. Recall from Example 2.4.2 that $\bar{N} \rightarrow \bar{N}/\langle u \rangle$ induces a \mathbb{P}^1 fibration π of $\mathrm{TV}(\Sigma)$ over $\mathrm{TV}(\bar{N}/\langle u \rangle) \cong \mathbb{C}^*$. Suppose we blow up a point H^+ in the boundary divisor D_u of $\mathrm{TV}(\Sigma)$ to obtain a variety \bar{U} (in higher dimensions H^+ will be what we call a hypertorus). Let \tilde{E} be the exceptional divisor, and \tilde{D} the proper transform of the toric boundary. Let F be the fiber of π containing H^+ . F has self-intersection $F^2 = 0$, so the proper transform \tilde{F} has $\tilde{F}^2 = -1$, and can thus be blown down to a point $H^- \subset D_{-u}$. Let E denote the image of \tilde{E} after the blowdown.

This blowup-blowdown procedure results in a new toric variety $\mathrm{TV}(\Sigma')$. Restricting to the complement of the boundary divisors, we view this as a birational map $\mu_u : (\mathbb{C}^*)^2 \dashrightarrow (\mathbb{C}^*)^2$. μ_u tells us how to glue the two tori along the complement of E and F . The result of this gluing is *essentially* just $U := \bar{U} \setminus \tilde{D}$ —I say “essentially” because we are missing the point $p := \tilde{E} \cap \tilde{F}$, but this has codimension 2. We can of course repeat this procedure, possibly for other primitive generators of \bar{N} , to obtain other non-toric blowups of compactifications of $(\mathbb{C}^*)^2$ (up to codimension 2). By Proposition 4.1.3, all log Calabi-Yau surfaces can be obtained this way.

See Figure 5.1 for an illustration of this. The key observation of [GHK13a] is that all mutations are given by blowup-blowdown procedures like this μ_u . Let us use this to motivate the definition of a seed. The important data is a set of vectors in \bar{N} telling us which boundary divisors we perform our blowups on. We may also have other vectors, called *frozen vectors*, which do not indicate blowups




Figure 5.1.1: A mutation μ_u involves blowing up a hypertorus H^+ in D_u (left arrow) and then contracting the proper transform \tilde{F} of the fibers F which hit H^+ (right arrow), down to a hypertorus H^- in D_{-u} . \tilde{E} denotes the exceptional divisor, with E being its image after the contraction of \tilde{F} . The locus $p = \tilde{E} \cap \tilde{F}$ has codimension 2 and does not appear in the cluster variety.

(they might, for example, indicate some boundary divisors with which we compactify U). Together, all these vectors are called *seed vectors*. In Definition 5.2.1 below, N can be thought of as the lattice freely generated by these seed vectors in \overline{N} .

Suppose in this two-dimensional example that we fix an orientation on \overline{N} and take the corresponding standard symplectic form $\cdot \wedge \cdot$. The data of a seed will include the skew form $\langle \cdot, \cdot \rangle$ on N induced by $\cdot \wedge \cdot$. μ_u as above decreases D_u^2 by 1 and increases D_{-u}^2 by 1. By Corollary 2.7.4, and in particular by the formula $v_{i+1} = -v_{i-1} - D_i^2 v_i$ which we used when defining U^{trop} (Equation 4.1), we can see that the tropicalization of μ_u (i.e., the induced map on \overline{N}) can be given by

$$v \mapsto v + \max(0, v \wedge u)u. \quad (5.1)$$

A version of this formula will tell us how mutation changes our seed. Seeds will include one other piece of data—a set of “multipliers” $d_i \in \mathbb{Q}_{>0}$ —which may be viewed as allowing us to deal with non-primitive $u \in \overline{N}$, or blowups along non-reduced loci $H^+ \subset D_u$.

5.2 Basic Definitions

[MUCH OF THIS SECTION HAS BEEN COPY-AND-PASTED FROM ONE OF MY PAPERS AND HAS SO FAR ONLY BEEN PARTIALLY EDITTED FOR OUR PURPOSES]

Definition 5.2.1. A *seed* is a collection of data

$$S = (N, I, E := \{e_i\}_{i \in I}, F, \langle \cdot, \cdot \rangle, \{d_i\}_{i \in I}),$$

where N is a finitely generated free Abelian group, I is a finite index set, E is a basis for N indexed by I , F is a subset of I , $\langle \cdot, \cdot \rangle$ is a skew-symmetric \mathbb{Q} -valued bilinear form, and the d_i ’s are positive

rational numbers called *multipliers*. We call e_i a *frozen* vector if $i \in F$. The *rank* of a seed or of a cluster variety will mean the rank of $\langle \cdot, \cdot \rangle$.

We define another bilinear form on N by

$$(e_i, e_j) := \epsilon_{ij} := d_j \langle e_i, e_j \rangle,$$

and we require that $\epsilon_{ij} \in \mathbb{Z}$ for all¹ $i, j \in I$. Let $M = N^*$. Define

$$p_1^* : N \rightarrow M, \quad v \mapsto (v, \cdot), \quad p_2^* : N \rightarrow M, \quad v \mapsto (\cdot, v).$$

Let $K_i := \ker(p_i^*)$, $\overline{N}_i := \text{im}(p_i^*) \subseteq M$, $\overline{e}_i := p_1^*(e_i)$, and $v_i := p_2^*(e_i)$. For each $i \in I$, we define a “modified multiplier” d'_i by saying that v_i is d'_i times a primitive vector in M .

Remark 5.2.2. Given only the matrix (e_i, e_j) and the set F , we can recover the rest of the data, up to a rescaling of $\langle \cdot, \cdot \rangle$ and a corresponding rescaling of the d_i 's. This rescaling does not affect the constructions below, and it is common to take the scaling out of the picture by assuming that the d_i 's are relatively prime integers—I personally prefer not making this assumption so we do not get caught up in irrelevant arithmetic issues. Also, notice that $\langle \cdot, \cdot \rangle$ and $\{d'_i\}$ together determine $\{d_i\}$, so when describing a seed we may at times give $\{d'_i\}$ instead of $\{d_i\}$.

Remark 5.2.3. By Remark 5.2.2 above, a seed is essentially just the data of a skew-symmetrizable matrix. Alternatively, a seed can be viewed as the data of a decorated **quiver** as follows: Each seed vector e_i corresponds to a vertex V_i of the quiver. The number of arrows from V_i to V_j is equal to $\langle e_i, e_j \rangle$, with a negative sign meaning that the arrows actually go from V_j to V_i (we may have to rescale to make sure these $\langle e_i, e_j \rangle$'s are integers). Each vertex V_i is decorated with the number d_i . Furthermore, the vertices corresponding to frozen vectors are boxed. Observe that all the data of the seed can be recovered from this quiver.

A seed is called *acyclic* if the corresponding quiver contains no directed paths that do not pass through any frozen (boxed) vertices. A cluster variety will be called acyclic if any of the corresponding seeds are acyclic.

For $a \in \mathbb{R}$, define $[a]_+ := \max(a, 0)$ and $[a]_- := \min(a, 0)$. Given a seed S as above and a choice of $j \in I \setminus F$, we can use a *mutation* to define a new seed $\mu_j(S) := (N, I, E' = \{e'_i\}_{i \in I}, F, \langle \cdot, \cdot \rangle, \{d_i\})$, where the (e'_i) 's are defined by

$$e'_i = \mu_j(e_i) := \begin{cases} e_i + [\epsilon_{ij}]_+ e_j & \text{if } i \neq j \\ -e_j & \text{if } i = j \end{cases} \quad (5.2)$$

Note the resemblance to Equation 5.1, which should be viewed as the induced map on \overline{N}_2 (except for the negation of e_j). We also note the dual equation on M :

$$\mu_j(e_i^*) := \begin{cases} -e_i^* + \sum_{k \in I} [\epsilon_{kj}]_+ e_k^* & \text{if } i = j \\ e_j & \text{if } i \neq j \end{cases} \quad (5.3)$$

¹The construction of cluster varieties does not depend on the values of $\langle e_i, e_j \rangle$ or ϵ_{ij} for $i, j \in F$, and so it is common to not include these coefficients in the data. When they are included in the data, as in [FG09] and [GHK13a], they are not typically required to be integers. However, as [GHK13a] points out, if these are not integers, then the image of p_i^* is not contained in M . [GHK13a] takes a slightly different fix to this (in which the ϵ_{ij} with $i, j \in F$ are again irrelevant), but it is essentially equivalent to our fix if we dropped the requirement that $\langle e_i, e_j \rangle = -\langle e_j, e_i \rangle$ when $i, j \in F$.

Corresponding to a seed S , we can define a so-called seed \mathcal{X} -torus $X_S := T_M = \text{Spec } \mathbb{k}[N]$, and a seed \mathcal{A} -torus $A_S := T_N = \text{Spec } \mathbb{k}[M]$. We define *cluster monomials* $X_i := z^{e_i} \in \mathbb{k}[N]$ and $A_i := z^{e_i^*} \in \mathbb{k}[M]$, where $\{e_i^*\}_{i \in I}$ is the dual basis to E . We may also write $X_v := z^v$ for any $v \in N$, and similarly $A_u := z^u$ for any $u \in M$.

Remark 5.2.4. We are departing somewhat from a common convention. In place of M , other authors typically use the superlattice $(M)^\circ \subset M \otimes \mathbb{Q}$ spanned over \mathbb{Z} by vectors $f_i := d_i^{-1} e_i^*$. They then take $A_i := (z^{f_i}) \in \mathbb{k}[M^\circ]$. It seems to me that this significantly complicates the exposition and the formulas that follow, with little benefit, and so we do not follow this convention.

For any $j \in I$, we have a birational morphism $\mu_j^{\mathcal{X}} : \mathcal{X}_S \rightarrow \mathcal{X}_{\mu_j(S)}$ (called a cluster \mathcal{X} -mutation) defined by

$$(\mu_j^{\mathcal{X}})^* X'_i = X_i \left(1 + X_j^{\text{sign}(-\epsilon_{ij})}\right)^{-\epsilon_{ij}} \quad \text{for } i \neq j; \quad (\mu_j^{\mathcal{X}})^* X'_j = X_j^{-1}. \quad (5.4)$$

Similarly, we can define a cluster \mathcal{A} -mutation $\mu_j^{\mathcal{A}} : \mathcal{A}_S \rightarrow \mathcal{A}_{\mu_j(S)}$,

$$A_j(\mu_j^{\mathcal{A}})^* A'_j = \prod_{i: \epsilon_{ji} > 0} A_i^{\epsilon_{ji}} + \prod_{i: \epsilon_{ji} < 0} A_i^{-\epsilon_{ji}}; \quad (\mu_j^{\mathcal{A}})^* A'_i = A_i \quad \text{for } i \neq j. \quad (5.5)$$

Now, the cluster \mathcal{X} -variety \mathcal{X} is defined by using compositions of \mathcal{X} -mutations to glue $\mathcal{X}_{S'}$ to \mathcal{X}_S for every seed S' which is related to S by some sequence of mutations. Similarly for the cluster \mathcal{A} -variety \mathcal{A} , with \mathcal{A} -tori and \mathcal{A} -mutations. The *cluster algebra* is the subalgebra of $\mathbb{k}[M]$ generated by the cluster variables A_i of every seed that we can get to by some sequence of mutations. In this context, the well-known Laurent phenomenon simply says that all the cluster variables are regular functions on \mathcal{A} . The ring of all global regular functions on \mathcal{A} is called the *upper cluster algebra*.

On the other hand, the X_i 's do not always extend to global functions on \mathcal{X} . When a monomial on a seed torus (i.e., a monomial in the X_i 's for a fixed seed) does extend to a global function on \mathcal{X} , we call it a *global monomial*, as in [GHK13a].

5.2.1 The Interpretation in Terms of Toric Models

As in [GHK13a], for a lattice L with dual L^* and with $u \in L$, $\psi \in L^*$, and $\psi(u) = 0$, define

$$m_{u, \psi, L} : T_L \dashrightarrow T_L \quad (5.6)$$

$$m_{u, \psi, L}^*(z^\varphi) = z^\varphi (1 + z^\psi)^{-\varphi(u)} \quad \text{for } \varphi \in L^*. \quad (5.7)$$

One can check that the mutations above satisfy

$$(\mu_j^{\mathcal{X}})^* = m_{(\cdot, e_j), e_j, M}^* : z^v \mapsto z^v (1 + z^{e_j})^{-(v, e_j)} \quad (5.8)$$

$$(\mu_j^{\mathcal{A}})^* = m_{e_j, (e_j, \cdot), N}^* : z^\gamma \mapsto z^\gamma (1 + z^{(e_j, \cdot)})^{-\gamma(e_j)}. \quad (5.9)$$

The following Lemma, compiled from §3 of [GHK13a], is what leads to the nice geometric interpretations of mutations and cluster varieties.

Lemma 5.2.5 ([GHK13a]). *Suppose that u is primitive in a lattice L . Let Σ be a fan in L with rays corresponding to u and $-u$. Recall that the toric variety $TV(\Sigma)$ admits a \mathbb{P}^1 fibration π with D_u and D_{-u} as sections, corresponding to the projection $L \rightarrow L/\mathbb{Z}\langle u \rangle$.*

The mutation $\mu_{u,\psi,L}$ is the birational map on $T_L \subset TV(\Sigma)$ coming from blowing up the set

$$H^+ := \{1 + z^\psi = 0\} \cap D_u$$

and then contracting the proper transforms of the fibers F of π which intersect this hypertorus. Furthermore, $\mu_j^{\mathcal{X}}$ preserve the centers of the blowups corresponding to $\mu_i^{\mathcal{X}}$ for each $i \neq j$. The same is true of $\mu_i^{\mathcal{A}}$ if S is “totally coprime,” which in particular holds if $\langle \cdot, \cdot \rangle$ is non-degenerate.

Definition 5.2.6. We call a set of the form $\{a + z^\psi = 0\} \cap D_u$, $a \in \mathbb{C}^*$, a *hypertorus*. An *irreducible hypertorus* is one for which ψ is primitive.

Thus, a cluster \mathcal{X} -mutation $(\mu_j^{\mathcal{X}})^*$ corresponds to blowing up $\{\mathcal{X}_j = -1\} \cap D_{(\cdot, e_j)}$, followed by blowing down some fibers of a certain \mathbb{P}^1 fibration, and repeating for a total of d'_j times (since (\cdot, e_j) is d'_j times a primitive vector, and $m_{(\cdot, e_j), e_j, M} = [m_{(\cdot, e_j)/d'_j, e_j, M}]^{d'_j}$). The new seed torus is only different from the old one in that it is missing the blown-down fibers of the initial \mathbb{P}^1 fibration, but has gained the exceptional divisor from the final blowup (except for the lower-dimensional set of points where this exceptional divisor intersects a blown-down fiber, represented by p in Figure 5.1.1).

Since the centers of the blowups corresponding to the other mutations have not changed, this shows that the cluster \mathcal{X} -variety can be constructed (up to codimension 2) as follows: For any seed S , take a fan in M with rays generated by $\pm(\cdot, e_i)$ for each i , and consider the corresponding toric variety. For each $i \in I \setminus F$, blow up the hypertorus $\{X_i = -1\} \cap D_{(\cdot, e_i)}$ d'_i times, and then remove the first $(d'_i - 1)$ exceptional divisors. The cluster \mathcal{X} variety is then the complement of the proper transform of the toric boundary.

Remark 5.2.7. In this construction of \mathcal{X} , the centers for the hypertori we blow up may intersect $(\cdot, e_i) = (\cdot, e_j)$ for some $i \neq j$, so some care must be taken regarding the ordering of the blowups. Fortunately, this issue only matters in codimension at least 2 (cf. [GHK13a] for more details). However, when we consider fibers of \mathcal{X} below, it is possible that some special fibers will have discrepancies in codimension 1. We will use the notation \mathcal{X}^{ft} to denote that we are restricting to the variety constructed as above for some fixed ordering of the blowups, and keep in mind that while $\mathcal{X} \setminus \mathcal{X}^{\text{ft}}$ is codimension 2 in \mathcal{X} , there may be special fibers of \mathcal{X} whose intersection with $\mathcal{X} \setminus \mathcal{X}^{\text{ft}}$ is codimension 1 in the fiber. As we will see below, \mathcal{A} is a torsor over what is perhaps the “most special” fiber of \mathcal{X} . The failure of mutations to preserve the centers of blowups for \mathcal{A} may be viewed as a consequence of such codimension 1 discrepancies in the special fiber.

We will similarly write \mathcal{A}^{ft} and U^{ft} (U a symplectic leaf of \mathcal{X} as we explain below) to denote the subvarieties of \mathcal{A} and U obtained by performing only one blowup for each non-frozen seed vector.

Remark 5.2.8. We have seen that codimension 2 issues arise as a result of missing points like p in Figure 5.1.1, and also as a result of reordering the blowups. There are also missing contractible complete subvarieties—the $(d'_j - 1)$ exceptional divisors we remove when applying $(\mu_j^{\mathcal{X}})^*$. These issues are relatively unimportant, since they do not affect the sheaf of regular functions on \mathcal{X} . When we are interested in \mathcal{X} or its fibers up to these issues, we will say “up to irrelevant loci.”

5.2.2 Langland’s dual seeds

Observations 5.2.9.

- K_1 is also equal to $\ker(v \mapsto \langle v, \cdot \rangle)$, so $\langle \cdot, \cdot \rangle$ induces non-degenerate skew-symmetric form on \overline{N}_1 . This also means that we could have equivalently defined the rank to be that of (\cdot, \cdot) .
- Define another skew-symmetric bilinear form on N by $[e_i, e_j] := d_i d_j \langle e_i, e_j \rangle$. Then $K_2 = \ker(v \mapsto [\cdot, v])$, so $[e_i, e_j]$ induces a non-degenerate skew-symmetric form on \overline{N}_2 . We can extend this to $\overline{N}_2^{\text{sat}}$ (the saturation in M of \overline{N}_2), and after possibly rescaling $[\cdot, \cdot]$ (and adjusting the d_i 's accordingly) we can identify this with the standard skew-symmetric form on $\overline{N}_2^{\text{sat}}$ with the induced orientation. We will denote this form and the induced symplectic form on $\overline{N}_{2, \mathbb{R}}$ by $(\cdot \wedge \cdot)$.
- Since $(\cdot, e_i) = -d_i \langle e_i, \cdot \rangle$, we see that $\text{im}(p_2^*)$ and $\text{im}(v \mapsto \langle v, \cdot \rangle)$ span the same subspace of $M_{\mathbb{R}}$. Thus, there is a canonical isomorphism $\overline{N}_{2, \mathbb{R}} \cong \overline{N}_{1, \mathbb{R}}$. We easily see that this is a symplectomorphism with respect to the symplectic forms induced by $[\cdot, \cdot]$ and $\langle \cdot, \cdot \rangle$.

Definition 5.2.10. The seed obtained from S by replacing $\langle \cdot, \cdot \rangle$ with the form $[\cdot, \cdot]$ defined above and d_i with d_i^{-1} produces the *Langland's dual seed* S^\vee described in [FG09]. Note that $\epsilon_{ij}^\vee = -\epsilon_{ji}$. So switching to S^\vee essentially has the effect of switching the roles of (and negating) p_1^* and p_2^* .

Note that although E changes under mutation, the form $\langle \cdot, \cdot \rangle$ does not, so K_1 and $\overline{N}_1^{\text{sat}}$ are invariant under mutation. The same is true for K_2 and $\overline{N}_2^{\text{sat}}$, as can similarly be seen using the Langland's dual seed and $[\cdot, \cdot]$ —one can check that the procedure for obtaining S^\vee from S commutes with mutation.

5.3 Other Structures on Cluster Varieties

5.3.1 Skew-forms and Poisson structures

For this section I believe we do have to make the assumption that the d_i 's are relatively prime integers.

Each seed \mathcal{X} -torus carries a Poisson structure defined by \mathcal{X}_S by

$$\{X_v, X_w\} = \langle v, w \rangle X_v X_w.$$

In fact, this commutes with mutations to give a Poisson structure on \mathcal{X} .

Similarly, \mathcal{A} carries a closed 2-form $\tilde{\Omega}$ given on a seed \mathcal{A} -torus \mathcal{A}_S by

$$\sum_{i, j \in I} [e_i, e_j] d \log A_i \wedge d \log A_j.$$

This can be viewed as an analog of the Weil-Petersson form on Teichmüller space.

5.3.2 An Exact Sequence of Cluster Varieties

Observe that for each seed S , there is a *not necessarily exact* sequence

$$0 \rightarrow K_2 \rightarrow N \xrightarrow{p_2^*} M \rightarrow K_1^* \rightarrow 0.$$

Here, $M \rightarrow K_1^*$ is the map dual to the inclusion $K_1 \hookrightarrow N$. Tensoring with \mathbb{k}^* yields an exact sequence, and one can check (cf. Lemma 2.10 of [FG09]) that this sequence commutes with mutation. Thus, one obtains the exact sequence

$$1 \rightarrow T_{K_2} \rightarrow \mathcal{A} \xrightarrow{p_2^*} \mathcal{X} \xrightarrow{\lambda} T_{K_1^*} \rightarrow 1. \quad (5.10)$$

One can check that the fibers of λ are symplectic leaves of the Poisson structure on \mathcal{X} . Let $\mathcal{U} := p_2(\mathcal{A}) \subset \mathcal{X}$. The symplectic form Ω on \mathcal{U} induced by the Poisson structure of \mathcal{X} is the same as the one induced by the 2-form $\tilde{\Omega}$ on \mathcal{A} (p_2 is the null-foliation of $\tilde{\Omega}$).

The sequence $1 \rightarrow T_{K_2} \rightarrow \mathcal{A} \rightarrow \mathcal{U} \rightarrow 1$ should be viewed as a generalization of the construction of toric varieties as quotients from §2.5, with \mathcal{U} being the generalization of the toric variety.²

Lemma 5.3.1. *The sets of the form $\{1 + X_j = 0\} \cap D_{(\cdot, v_j)} \cap \lambda^{-1}(\phi)$, $\phi \in T_{K_1^*}$, are unions of k irreducible hypertori, where k is the index of $p_1(e_j)$ in \bar{N}_1 .*

Proof. The argument is the same as that of Lemma 5.1 in [GHK13a]. Briefly, we choose a splitting $N \cong \bar{N}_1 \oplus K_1$, and decompose $e_j = (e'_j, e''_j)$ with respect to this splitting. Then $z^{e_j}|_{\lambda^{-1}(\phi)} = z^{e''_j}(\phi)z^{e'_j}$, so the degree of this restriction is indeed the index of $e'_j = p_1(e_j)$ in \bar{N}_1 . \square

Example 5.3.2. Consider the case where Y is a cubic surface, obtained by blowing up 2 points on each boundary divisor of $(\bar{Y} \cong \mathbb{P}^2, \bar{D} = D_1 + D_2 + D_3)$. We can take

$$\bar{E} = \{(1, 0), (1, 0), (0, 1), (0, 1), (-1, -1), (-1, -1)\},$$

with each $d_i = d'_i = 1$ and F empty. Then the fibers of the resulting \mathcal{X} -variety \mathcal{X}_1 correspond to the different possible choices of blowup points on the D_i 's. The fiber \mathcal{U} is very special, having four

(-2) -curves. If we instead take $\bar{E} = \{(1, 0), (0, 1), (-1, -1)\}$ with $\langle \cdot, \cdot \rangle$ given by $\begin{pmatrix} 0 & 1 & -1 \\ -1 & 0 & 1 \\ 1 & -1 & 0 \end{pmatrix}$, and

each $d_i = d'_i = 2$, then the fibers of the resulting \mathcal{X} -variety \mathcal{X}_2 include only the surfaces constructed by blowing up the same point twice on each D_i and then removing the three resulting (-2) -curves. \mathcal{U} is the fiber where the blowup points are colinear and so there is one remaining (-2) -curve.

The deformation type of the fibers of \mathcal{X}^{ft} has only changed by the removal of certain (-2) -curves, i.e., by some irrelevant loci. Note that $\mathcal{X}_2^{\text{ft}} = \mathcal{X}_2$, and that \mathcal{X}_2 can be identified (after filling in the removed (-2) -curves) with a subfamily of $\mathcal{X}_1^{\text{ft}}$ whose fibers do not agree with those of \mathcal{X}_1 in codimension 1.

These examples are well-known: we will see that the former corresponds to a moduli space of flat $\text{PGL}_2(\mathbb{C})$ -connections on a four-punctured sphere, while the latter corresponds to the analogous space for the once-punctured torus (cf. §2.7 of [FG09] or Exercise 7.2.11).

5.3.3 Semifield-valued points and tropicalizations of cluster varieties

Note that the formulas for mutations of cluster variables are always rational functions which can be expressed as quotients of two polynomials with only positive integral coefficients. A space with such an atlas is called a *positive space*. For any positive space, it makes sense to talk about the \mathbb{F} -valued points for any *semifield* \mathbb{F} (see §1.1 of [FG09] for more details). Here, a semifield is a set P with two operations addition and multiplication such that addition is commutative and associative, multiplication makes P into an Abelian group, and multiplication distributes over addition.

For each seed S , $\mathcal{A}_S(\mathbb{F}) := N \otimes_{\mathbb{Z}} \mathbb{F} \cong \mathbb{F}^n$, where $n := \dim N$, and the Abelian group structure on \mathbb{F} is the one given by the semifield multiplication. Similarly, $\mathcal{X}_S(\mathbb{F}) := M \otimes_{\mathbb{Z}} \mathbb{F} \cong \mathbb{F}^n$. Mutation acts

²This sequence really only generalizes the construction for toric varieties without boundary (i.e., just algebraic tori), but I believe it can be modified using frozen vectors to also deal with cases with boundary.

on the \mathbb{F} -valued points by interpreting the addition and multiplication in the definitions of mutation as addition and multiplication in the semifield. In fact, mutations act bijectively on these sets of \mathbb{F} -valued points, so we actually have identifications $\mathcal{A}(\mathbb{F}) \cong \mathbb{F}^n$ and $\mathcal{X} \cong \mathbb{F}^n$. Keep in mind, however, that the identifications depends on the choice of seed—we will use the subscript S to indicate that we have chosen the identification corresponding to S .

Examples 5.3.3.

- The tropical real numbers \mathbb{R}^t with multiplication defined to be the usual addition, and addition defined to be \min , form a semifield. Similarly for \mathbb{Z}^t and \mathbb{Q}^t . When Fock and Goncharov talk about the tropicalizations of \mathcal{A} and \mathcal{X} , they usually view these as $\mathcal{A}(\mathbb{R}^t)$ and $\mathcal{X}(\mathbb{R}^t)$ (with \mathbb{Z}^t and \mathbb{Q}^t valued points giving the usual integral and rational points of the tropicalizations). We will not worry about this point of view since it is conceptually easier to just identify these spaces (as piecewise-linear manifolds) with $N_{\mathbb{R}}$ and $M_{\mathbb{R}}$.
- The positive real numbers $\mathbb{R}_{>0}$ with the usual multiplication and addition form a semifield.

We will often be interested in the tropical points of \mathcal{A} and \mathcal{X} . We will usually fix a base seed so that we can simply view $\mathcal{A}(R^t)$ ($R = \mathbb{R}, \mathbb{Q}$, or \mathbb{Z}) as $N \otimes R$ and $\mathcal{X}(R^t)$ as $M \otimes R$. The identifications $\mu_j^{\mathcal{A}^t} : \mathcal{A}_S(\mathbb{R}^t) \rightarrow \mathcal{A}_{\mu_j(S)}(\mathbb{R}^t)$ and $\mu_j^{\mathcal{X}^t} : \mathcal{X}_S(\mathbb{R}^t) \rightarrow \mathcal{X}_{\mu_j(S)}(\mathbb{R}^t)$ can be described as follows: tropicalizing Equation 5.6 yields $(m_{u,\psi,L}^t)^*(\varphi) = \varphi - \varphi(u) \min(0, \psi)$, where $\varphi, \psi \in L^*$ are viewed as coordinates on $L \otimes \mathbb{R}$. Dualizing yields $m_{u,\psi,L}^t(x) = x - [\psi(x)]_- u$ (i.e., $[(m_{u,\psi,L}^t)^*(\varphi)](x) = \varphi(m_{u,\psi,L}^t(x))$).

In particular, we have:

$$\begin{aligned}\mu_j^{\mathcal{A}^t}(n) &= n - [(e_j, n)]_- e_j \\ \mu_j^{\mathcal{X}^t}(m) &= m - [m(e_j)]_- (\cdot, e_j).\end{aligned}$$

Since $-[a]_- = [-a]_+$, we see that $\mu_j^{\mathcal{A}^t}$ is the same as the seed mutation formula from Equation 5.2 without the negation of e_j if we replace ϵ_{ij} with $-\epsilon_{ji}$. In other words:

Proposition 5.3.4. *Aside from the negation of e_j , the seed mutation μ_j is the tropicalization of the Langland's dual³ cluster mutation $\mu_j^{(\mathcal{A}^\vee)^t}$.*

Similarly, letting $v_i := (\cdot, e_i)$, we have $\mu_j^{\mathcal{X}^t}(v_i) = v_i + [\epsilon_{ij\vee}]_+ v_j$, so the action of mutation on the v_i 's is just the one induced by p_2^* and the action on the e_i 's.

We will also be interested in the fibers U of λ and their tropicalizations. For each seed S , we can canonically identify $U(\mathbb{Z}^t)$ ⁴ with $\overline{N_2}^{\text{sat}} = p_2^*(\mathcal{A}(\mathbb{Z}^t))^{\text{sat}} \subset \mathcal{X}(\mathbb{Z}^t)$. Since $\overline{N_2}$ is the span of the v_i 's, the above shows that as an identification of sets this is independent of the choice of seed.

In particular, if U^{ft} is a log Calabi-Yau surface, then $U(\mathbb{Z}^t)$ as defined here is the same (as a piecewise-linear manifold) as $U^{\text{trop}}(\mathbb{Z})$ as defined in §4.2. In §4.7.3, we saw that a toric model $\pi : (Y, D) \rightarrow (\overline{Y}, \overline{D})$ for a minimal model (Y, D) of U determines an identification of $U^{\text{trop}}(\mathbb{Z})$ with $N_{\overline{Y}}$. If p_i is the toric model corresponding to a seed S , then $U(\mathbb{Z}^t)_S = N_{\overline{Y}}$ as lattices (see Construction 2.11 of [Mana] for more details).

³[GHK13a] does not seem to observe the appearance of Langland's duality here. Using the conventions mentioned in Remark 5.2.4, $\mathcal{A}(\mathbb{Z}^t)$ is actually identified with a sublattice $N^\circ \subset N$, and then μ_j^* is (ignoring e_j) the piecewise-linear map on N induced by $\mu_j^{\mathcal{A}^t}$, with no mention of Langland's duality necessary. However, since vectors in E correspond to cluster monomials in \mathcal{X} , our perspective is consistent with the viewpoint of Fock and Goncharov's Conjecture 6.1.1.

⁴Another perspective which might be worth exploring would be to identify the tropicalizations of different fibers of λ with different fibers of λ^* , with only λ_e corresponding to what we call $U(\mathbb{Z}^t)$ here.

5.3.4 Integral linear structures in general

We note that just as in the two-dimensional situation, we can always view points in $*(\mathbb{Z}^t)$ ($*$ denoting \mathcal{X} , \mathcal{A} , or U) as divisorial valuations on the function field of $*$. The construction of the integral linear structure from §4.2 also generalizes. Namely, let U be any log Calabi-Yau variety⁵ and choose a minimal model (Y, D) for U . (Y, D) induces⁶ an integral linear structure on U^{trop} , which we call the GHK integral linear structure (it has previously been observed by Gross, Hacking, and Keel).

As a piecewise-linear manifold, U^{trop} can be canonically identified with the real cone over the dual complex B of D . Here, B consists of a real k -cell for each $(n - k)$ stratum of D . $U^{\text{trop}}(\mathbb{Z})$ similarly corresponds to the integer points of the cone over B . Note that B naturally induces a complete fan structure Σ on the cone over B . If U is a fiber of some \mathcal{X} , then this can be identified as a topological space with $\mathcal{U}^{\text{trop}} := p_2^*(\mathcal{A}^{\text{trop}}) \subset \mathcal{X}^{\text{trop}}$ (or with any fiber of $\mathcal{X}^{\text{trop}}$).

We now define integral linear charts on U^{trop} as follows: let σ_0, σ_1 be two maximal dimensional cells of Σ intersecting along some nonempty codimension 1 cell ρ . Let v_2, \dots, v_n be primitive generators of the boundary rays of ρ , and let v_0, v_1 be the primitive generators for the additional boundary rays of σ_0 and σ_1 , respectively. Let D_{v_i} denote the boundary component corresponding to v_i , and C_ρ the 1-stratum of D corresponding to ρ . Then we define a chart $\varphi : \sigma_1 \cup \sigma_2 \rightarrow \mathbb{R}^n$ which is linear on the cones of Σ , which takes $v_i, i \neq 0$, to the standard basis vector e_i , and takes v_0 to $-e_1 - \sum_{i=2}^n (C_\rho \cdot D_{v_i})v_i$. This defines a (non-oriented) integral linear structure without singularities on the complement of the codimension 2 cells of Σ . We may be able to extend smoothly over some codimension 2 strata, but in general these will be singular loci. If we fix an orientation of B or U^{trop} (for example, using the skew-pairing coming from the cluster structure) and restrict to charts which agree with this orientation, we obtain an oriented integral linear structure.

This definition in terms of charts is equivalent to saying that an integral piecewise-linear function f taking values a_i along v_i is linear across a codimension 1 cell ρ if and only if the Weil divisor $W_f := \sum a_i D_{v_i}$ satisfies $W_f \cdot C_\rho = 0$, where C_ρ is the 1-stratum of D corresponding to ρ .

5.3.5 The Cluster Modular Group

A *seed isomorphism* $h : S \rightarrow S'$ is an isomorphism of the underlying lattices which respects all the seed data in the obvious way—that is, it takes (frozen) seed vectors to (frozen) seed vectors (thus inducing a bijection $h : I \rightarrow I'$ taking F to F'), such that $d_i = d_{h(i)}$ and $\langle e_i, e_j \rangle = \langle h(e_i), h(e_j) \rangle'$. Note that although we typically write I as $\{1, \dots, n\}$, this ordering need not be preserved. h then induces a *cluster isomorphism* $h : \mathcal{X} \rightarrow \mathcal{X}'$ and $h : \mathcal{A} \rightarrow \mathcal{A}'$ given by $h^* X'_{h(e_i)} = X_i$ and $h^* A'_{h(e_i)} = A_i$. A *seed transformation* is a composition of seed mutations and seed isomorphisms, and a *cluster transformation* is a composition of cluster mutations and cluster isomorphisms (i.e., the corresponding maps on the varieties). [FG09] defines the *cluster modular group* Γ to be the group of cluster automorphisms of a base seed S , modulo trivial cluster automorphisms (those which are the identity on both \mathcal{A} and \mathcal{X}).

Example 5.3.5. For any finite type seed, Γ is a finite group. For example, for the A_2 quiver $\bullet \rightarrow \bullet$, all seeds are isomorphic, and there are 5 distinct seeds, so $\Gamma \cong \mathbb{Z}/5\mathbb{Z}$. We revisit this in Example 5.3.8 below.

⁵We will particularly care about when U is a fiber of \mathcal{X}^{ft} , but we could also take $U = \mathcal{X}^{\text{ft}}$ or $U = \mathcal{A}^{\text{ft}}$.

⁶Unlike the two-dimensional situation, the integral linear structure will typically depend on the choice of minimal model.

5.3.6 The Cluster Complex

Definition 5.3.6. A seed S with seed vectors e_1, \dots, e_n determines a cone $C_S \subset \mathcal{X}_S^{\text{trop}} := M_{S, \mathbb{R}}$ given by

$$C_S := \{m \in M_{S, \mathbb{R}} \mid e_i(m) \geq 0 \text{ for all } i \in I \setminus F\}.$$

The collection of all such cones in \mathcal{X} for every seed mutation equivalent to S is called the **cluster complex** \mathcal{C} .

Recall that $U_S^{\text{trop}} := \overline{N_{2, \mathbb{R}, S}}$ has a natural symplectic structure induced by $[\cdot, \cdot]$.

Proposition 5.3.7. *For a fixed base seed S and two mutation equivalent seeds S_1 and S_2 , let $g_l : M_{S, \mathbb{R}} \rightarrow M_{S_l, \mathbb{R}}$, $l = 1, 2$, be vector space isomorphisms which identify C_S with C_{S_l} and restrict to symplectomorphisms $U_S^{\text{trop}} \rightarrow U_{S_l}^{\text{trop}}$ taking $v_{i, S}$ to $v_{\sigma_l(i), S_l}$ for some permutation σ of I . As automorphisms of $\mathcal{X}^{\text{trop}}$, these g_l 's preserve \mathcal{C} . Consequently, there is a seed S_3 such that $g_1 \circ g_2$ has the above form for S_3 . Composition thus gives a group structure to these automorphisms of $\mathcal{X}^{\text{trop}}$, and this group is canonically identified with the cluster modular group Γ .*

Proof. We first describe the bijection with Γ as sets. Let $W_{i, S}$ denote the wall of C_S contained in e_i^\perp . Then if $g_l : W_{i, S} = W_{j, S_l}$, the element of Γ corresponding to g_l is the seed isomorphism from S to S_l taking $e_i \in E_S$ to $e_j \in E_{S_l}$ (or rather, this composed with the chain of mutation we take to see that S_l is mutation equivalent to S). We note that we could choose σ_l so that $\sigma_l(i) = j$. The condition that $g|_{U^{\text{trop}}}$ gives a symplectomorphism means that it preserves the pairing $[\cdot, \cdot]$. This together with the d_i 's being preserved (which follows from the v_i 's being preserved) implies that $\langle \cdot, \cdot \rangle$ and $\{d_i\}$ are preserved (up to the usual rescaling issues). The condition of taking C_S to C_{S_l} tells us that F is preserved, since the walls corresponding to F are the unbounded ones.

We thus see that g_l does give an element of Γ . Since Γ is the group of automorphisms of \mathcal{A} and \mathcal{X} preserving the cluster structure, it must also preserve the cluster complex. This gives the claim about compositions. It remains to check that trivial automorphisms of $\mathcal{X}^{\text{trop}}$ correspond to trivial seed transformations. This will be an immediate consequence of Lemma 6.2.3. \square

Example 5.3.8. For the A_2 quiver (recall this corresponds to the $\mathcal{M}_{0,5}$ example), the cluster complex consists of 5 chambers in $\mathcal{X}^{\text{trop}}$. These are exactly the chambers of the scattering diagram for this log Calabi-Yau variety. We will see that this is always the case for finite-type cluster varieties. In general, the cluster complex can be viewed as a subset of the corresponding scattering diagram.

5.4 Cluster varieties with principal coefficients

Consider a seed $S = \{N, I, E, F, \langle \cdot, \cdot \rangle, \{d_i\}\}$ for the cluster varieties to \mathcal{A} and \mathcal{X} . In the next chapter, it will be useful to use certain enlargements of \mathcal{A} and \mathcal{X} called $\mathcal{A}^{\text{prin}}$ and $\mathcal{X}^{\text{prin}}$, respectively.

Definition 5.4.1. $\mathcal{A}^{\text{prin}}$ and $\mathcal{X}^{\text{prin}}$ are the \mathcal{A} and \mathcal{X} -spaces corresponding to the seed S^{prin} defined as follows:

- $N_{S^{\text{prin}}} := N \oplus M$.
- $I_{S^{\text{prin}}}$ is the disjoint union of two copies of I . We will call them I_1 and I_2 to distinguish between them.




Figure 5.3.2: The cluster complex in $\mathcal{X}_S^{\text{trop}}$ for the A_2 -quiver. Note that these are exactly the cones of the scattering diagram from Example 4.4.3—This figure only looks different from Figure 4.2.1 because we here we are using the vector space structure on $\mathcal{X}^{\text{trop}}$ corresponding to the seed S , whereas before we used a developing map for $U^{\text{trop}} = \mathcal{X}^{\text{trop}}$.

- $E_{S^{\text{prin}}} := \{(e_i, 0) | i \in I_1\} \cup \{(0, e_i^*) | i \in I_2\}$
- $F_{S^{\text{prin}}} := F_1 \cup I_2$, where F_1 is simply F viewed as a subset of I_1 .
- $\langle (n_1, m_1), (n_2, m_2) \rangle_{S^{\text{prin}}} := \langle n_1, n_2 \rangle_S + m_2(n_1) - m_1(n_2)$.
- The d_i 's are the same as before (viewing i as an element of I).

Remark 5.4.2. As a quiver, this is quite easy to describe. Let Q be the quiver corresponding to S . For each vertex $e_i \in Q$ we add a new frozen vertex e_i' along with an arrow from e_i to e_i' . We give e_i' the same multiplier as e_i . This gives the quiver corresponding to S^{prin} .

Lemma 5.4.3. $\langle \cdot, \cdot \rangle_{S^{\text{prin}}}$ is unimodular. Also, $(S^{\text{prin}})^\vee = (S^\vee)^{\text{prin}}$.

Proof. Let B denote the skew-symmetric matrix corresponding to S . Then the skew-symmetric matrix for S^{prin} is $B^{\text{prin}} := \left(\begin{array}{c|c} B & \text{Id} \\ \hline -\text{Id} & 0 \end{array} \right)$, and this has determinant 1. The second statement is also easy to check using this form for B^{prin} . \square

Recall that \mathcal{A} is a T_{K_2} -torsor over the special fiber $U := p_2(\mathcal{A}) \subset \mathcal{X}$. $\mathcal{A}^{\text{prin}}$ serves to extend this to the whole \mathcal{X} -space, as made precise in the following proposition:

Proposition 5.4.4. *We have the following maps:*

$$\begin{array}{ccccccc}
 T_N & \hookrightarrow & \mathcal{A}^{\text{prin}} & \xrightarrow{\tilde{p}} & \mathcal{X} & & \\
 \text{Id} \downarrow & & \downarrow \pi & & \downarrow \lambda & & \\
 T_{K_2} & \hookrightarrow & T_N & \xrightarrow{\epsilon} & T_M & \xrightarrow{l} & T_{K_1^*}
 \end{array} \tag{5.11}$$

- $\pi : \mathcal{A}^{\text{prin}} \rightarrow T_M$ given on the cocharacter lattice corresponding to a seed by $(n, m) \mapsto m$.
- $\tilde{p} : \mathcal{A}^{\text{prin}} \rightarrow \mathcal{X}$ given on cocharacter lattices by $(n, m) \mapsto m - p_2^*(n)$.
- $\iota : T_N \hookrightarrow \mathcal{A}^{\text{prin}}$ given on cocharacter lattices by $n \mapsto (n, p_2^*(n))$.
- $l : T_M \rightarrow T_{K_1^*}$ defined by the dual to $K_1 \hookrightarrow N$, just as λ was defined in §5.3.2.

- $\epsilon : T_N \rightarrow T_M$ induced by p_2^* .

These maps satisfy the following:

- The rows are exact. In particular, \tilde{p} realizes $\mathcal{A}^{\text{prin}}$ as a T_N -torsor over \mathcal{X} .
- For $t \in T_M$, \tilde{p} restricts to a map $p_t : \pi^{-1}(t) \rightarrow \lambda^{-1}(l(t))$.
- T_N acts equivariantly with respect to π . Since $\epsilon(T_N)$ is the kernel of l , we have $p_{t_1}(\mathcal{A}_{t_1}) = p_{t_2}(\mathcal{A}_{t_2})$ if and only if t_1 and t_2 are in the same orbit of the $\epsilon(T_N)$ -action on T_M .

Proof. See §2 of [GHK13a] or Appendix B of [GHKK] for the proof and many more details. For statements involving l , these sources require that there are no frozen vectors, but this assumption can be easily removed with our setup (because we have defined (\cdot, \cdot) to be \mathbb{Z} -valued for frozen vectors, and we use this pairing when defining p_2^*). \square

As an application, we show how [GHK13a] gives a short proof of the Laurent phenomenon (a special case of their Corollary 3.11):

Theorem 5.4.5 (The Laurent phenomenon). *Any cluster monomial $A_i := z^{e_i^*}$, for any seed S , can be expressed as a Laurent polynomial in the cluster monomials of any other seed. In other words, every A_i (and hence the entire cluster algebra) is contained in the upper cluster algebra $\Gamma(\mathcal{A}, \mathcal{O}_{\mathcal{A}})$.*

Proof. It is immediate from the definition in Equation 5.5 that A_i is a Laurent polynomial in the cluster variables of any mutation adjacent seed. By Lemma 5.2.5, $\mathcal{A}^{\text{prin}}$ is covered, up to codimension 2, by the seed torus \mathcal{A}_S along with the seed tori $\mathcal{A}_{S'}$ for all S' mutation adjacent to S . Thus, the claim holds for $\mathcal{A}^{\text{prin}}$. \mathcal{A} is the subvariety of $\mathcal{A}^{\text{prin}}$ obtained by setting $z^{(0, e_i^*)} = 1$ for each $i \in I_2$, with $z^{e_i^*}$ being the restriction of $z^{(e_i, 0)^*}$, so the claim for \mathcal{A} follows. \square

Chapter 6

Canonical Bases and Cluster Varieties

In this chapter we at last explain how to get canonical bases for the rings of global sections of cluster varieties.

6.1 The Fock-Goncharov Conjecture

Recall that for each seed $S = \{N, I, E, F, \langle \cdot, \cdot \rangle, \{d_i\}_{i \in I}\}$, we have an identification of $\mathcal{A}^{\text{trop}}(\mathbb{Z})$ with N and of $\mathcal{X}^{\text{trop}}(\mathbb{Z})$ with $M = \text{Hom}(N, \mathbb{Z})$ (and for the rational or real tropical points, we tensor with \mathbb{Q} or \mathbb{R} , respectively). Points in N correspond to monomials on the torus \mathcal{X}_S , and points in M correspond to monomials on the torus \mathcal{A}_S . In fact, these monomials form \mathbb{C} -module bases for the rings of regular functions on these tori. However, “most” of these monomials do not extend to global functions. Fock and Goncharov conjecture the following:

Conjecture 6.1.1 (Fock-Goncharov Dual Basis Conjecture). *There is a canonical basis for $\Gamma(\mathcal{X}^\vee, \mathcal{O}_{\mathcal{X}^\vee})$ indexed in a natural way by $\mathcal{A}^{\text{trop}}(\mathbb{Z})$, and similarly, there is a canonical basis for $\Gamma(\mathcal{A}, \mathcal{O}_{\mathcal{A}})$ indexed by $(\mathcal{X}^\vee)^{\text{trop}}(\mathbb{Z})$.¹*

They also conjecture many properties that they expect these dual bases to satisfy (cf. [FG09], Conjecture 4.1), but I will not take the time to list these.

[GHK13a] shows that the Fock-Goncharov conjecture, as stated above, is false. This is already easy for us to see using the ideas of Chapter 5. Consider, for example, a Looijenga pair (Y, D) obtained by taking 4 non-toric blowups on each boundary divisor of the toric pair $(\mathbb{P}^2, \bar{D} = \bar{D}_1 + \bar{D}_2 + \bar{D}_3)$. Then D does not support an ample divisor. In fact, the intersection matrix $(D_i \cdot D_j)$ being negative definite implies that D can be contracted to a point p (a cusp singularity). Since the complement of this point is the same as $U = Y \setminus D$, we see that U is equal to a compact variety up to codimension 2. Thus, Hartog’s theorem tells us that U has no nonconstant global regular functions.

¹In our setup, $\mathcal{A}^{\text{trop}}$ and $(\mathcal{A}^\vee)^{\text{trop}}$ are the same set, as are $\mathcal{X}^{\text{trop}}$ and $(\mathcal{X}^\vee)^{\text{trop}}$, so taking Langland’s duals may appear somewhat meaningless. See Remark 6.2.1 for an explanation of the sense in which the Langland’s dual geometry really is what controls the theta functions.

However, we saw in Chapter 5 that a generic such U will appear (up to codimension 2) as a fiber of some $\lambda: \mathcal{X} \rightarrow T_{K_1^*}$ (e.g., if we take $\epsilon_{ij} = \begin{pmatrix} 0 & 4 & -4 \\ -4 & 0 & 4 \\ 4 & -4 & 0 \end{pmatrix}$, then some deformations of U appear as fibers, up to codimension 2 and the removal of (-2) curves which do not affect the global sections). Let $r = \text{rank}(K_1)$, so $\text{rank}(N) = 2 + r$. Conjecture 6.1.1 then says that $\mathcal{A}^{\text{trop}}(\mathbb{Z})$ parametrizes global sections of \mathcal{X} , but since fibers are deformations of U and thus have no non-constant global sections, the only global sections of \mathcal{X} are those pulled back from $T_{K_1^*}$. These are of course parametrized by K_1 , which is a proper subset of $\mathcal{A}^{\text{trop}}(\mathbb{Z})$, so the conjecture fails here.

In the other direction, [GHK13a] shows that the upper cluster algebra $\Gamma(\mathcal{A}, \mathcal{O}_{\mathcal{A}})$ is often not be finitely generated (in fact, [Spe13] had previously shown this for the case with $\epsilon_{ij} = \begin{pmatrix} 0 & 3 & -3 \\ -3 & 0 & 3 \\ 3 & -3 & 0 \end{pmatrix}$).

They suggest that the failure of the upper cluster algebra to be finitely generated is related (perhaps equivalent) to the failure above of the affineness condition needed for Conjecture 6.1.1 to hold. As far as I know this is not yet completely understood.

Fortunately, as we will see, the conjecture does hold in many cases, and a formal version holds in general.

6.2 The Naive Initial Scattering Diagrams

The main ideas behind the construction of the canonical bases in [GHKK] are essentially the same as those in [GHK11] (i.e., what we saw in Chapter 4), although we will see that some interesting tricks have to be used to get around certain technical difficulties. That is, we will define initial scatterings diagrams $\mathcal{D}_{\mathcal{X}}^0$ in $\mathcal{A}^{\text{trop}}$ and $\mathcal{D}_{\mathcal{A}^\vee}^0$ in $\mathcal{X}^{\text{trop}}$, hope that we can extend these to consistent scattering diagrams $\mathcal{D}_{\mathcal{A}}$ and $\mathcal{D}_{\mathcal{X}^\vee}$ via Theorem 4.7.4, and then define the theta functions using broken lines with respect to these scattering diagrams.

For a given choice of base seed S , we try the initial scattering diagrams:

$$\mathcal{D}_{\mathcal{X}, S}^0 := \{e_i^\perp, 1 + z^{(\cdot, e_i)} \mid i \in I \setminus F\} \subset \mathcal{X}_S^{\text{trop}} \quad (6.1)$$

$$\mathcal{D}_{\mathcal{A}, S}^0 := \{(e_i, \cdot)^\perp, (1 + z^{e_i})^{\text{ind}_M((e_i, \cdot))} \mid i \in I \setminus F\} \subset \mathcal{A}_S^{\text{trop}}. \quad (6.2)$$

Here, $\text{ind}_M((e_i, \cdot))$ is the index of (e_i, \cdot) in M (so if \overline{N}_1 is saturated in M , this agrees with the index used in 5.3.1). Comparing to Equation 5.8 and using that $(\cdot, e_i) = (e_i, \cdot)^\vee$, we see that $\mathcal{D}_{\mathcal{X}, S}^0$ is defined so that the wall-crossing automorphism corresponding to crossing from $e_i > 0$ to $e_i < 0$ corresponds to applying the mutation μ_i to \mathcal{A}_{S^\vee} . Similarly for $\mathcal{D}_{\mathcal{A}, S}^0$, crossing from $v_i > 0$ to $v_i < 0$ and applying μ_i to \mathcal{X}_{S^\vee} .

Remark 6.2.1. We of course could have used $1 + z^{(e_i, \cdot)}$ for $\mathcal{D}_{\mathcal{X}, S}^0$ above and used $\text{ind}((\cdot, e_i))$ in the exponent for the definition of $\mathcal{D}_{\mathcal{A}, S}^0$, so it may seem silly to make these changes that force us to consider the Langland's dual seeds. However, the scattering diagrams 6.1 above really do record geometric information corresponding to S , rather than its Langland's dual. the geometry of the \mathcal{X} and \mathcal{A} -spaces, respectively. For example, recall from §5.2.1 that the mutation $\mu_i^{\mathcal{X}}$ corresponds to blowing up some locus on the boundary divisor $D_{(\cdot, e_i)}$ of \mathcal{X} . In a rough analogy with Equation 4.7, this suggests that we should indeed have $1 + z^{(\cdot, e_i)}$ as the corresponding scattering function in $\mathcal{D}_{\mathcal{X}, S}^0$.

Similarly, for the \mathcal{A} -space, the blowup corresponding to μ_i is along D_{e_j} intersecting with the locus $\{1 + z^{(e_j, \cdot)} = 0\}$, and thus includes $\text{ind}((e_j, \cdot))$ irreducible hypertori. This accounts for the $\text{ind}(e_i, \cdot)$ in the exponents. See §6.3 for another geometric interpretation of $\mathcal{D}_{\mathcal{A}, S}^0$.

6.2.1 The naive \mathcal{A} -scattering diagram

$\mathcal{A}_{\mathcal{X}, S}^0$ is indeed a well-defined scattering diagram in the sense of §4.7. The lattice N in §4.7 is actually just the 0-dimensional lattice here, and P^{gp} from §4.7 is the lattice N here. We can take the monoid P from §4.7 to be the one spanned by the e_i 's, $i \in I \setminus F$. The skew-form $\{\cdot, \cdot\}$ in §4.7 is of course the form $\langle \cdot, \cdot \rangle$ from the seed data. One easily checks now that $\mathcal{D}_{\mathcal{X}, S}^0$ satisfies the conditions needed to apply Theorem 4.7.4. We thus obtain a consistent scattering diagram $\mathcal{D}_{\mathcal{X}, S} \subset N$, as desired. One can now construct the mirror to \mathcal{A} as in §4.4.1, equipped with theta functions constructed as in §4.5.

We expect that there is some degeneration of \mathcal{X}^\vee such that the mirror family can be identified with a formal neighborhood of the degenerate fiber. Under certain conditions (like D supporting a D -ample divisor, where D is the boundary divisor of a minimal model of a generic fiber $U \subset \mathcal{X}$), we expect that this formal family can be extended to include all of \mathcal{X}^\vee . This is indeed the case. The only problem with our current construction is that this is quite difficult to see. We will therefore use a slightly different version of this later on.

6.2.2 The naive \mathcal{X} -scattering diagram

Now consider $\mathcal{D}_{\mathcal{X}, S}^0$. We would like to take the lattice N from §4.7 to be trivial, P^{gp} to be M , and P to contain the span of the v_i 's. The problem is that the v_i 's often do not live in any strictly convex cone. In fact, one can easily see that they are contained in a convex cone if and only if S is acyclic (this cone may not be *strictly* convex, but it can be made strictly convex by some sequence of mutations). We will therefore need a trick in order to deal with the non-acyclic cluster varieties.

Let us first look at the acyclic situation and consider how things work out here. In fact, let us make the stronger assumption that p_2^* restricted to the lattice N_{u_f} of non-frozen vectors is injective. We state the next lemma very vaguely and refer to Theorem 1.33 of [GHKK] for the precise statement. Let $\mathcal{D}_{X, S} := \text{S}(\mathcal{D}_{\mathcal{X}, S}^0)$.

Lemma 6.2.2 ([GHKK], Theorem 1.33). *If $p_2^*|_{N_{u_f}}$ is injective, then $\mathcal{D}_{X, S}$ is mutation invariant. That is, μ_j acts on the scattering diagram in a way that takes $\mathcal{D}_{X, S}$ to $\mathcal{D}_{X, \mu_j(S)}$.*

Lemma 6.2.3 ([GHKK], Construction 1.38 and Lemma 2.9). *$\mathcal{D}_{X, S}$ contains the Langland's dual cluster complex \mathcal{C}^\vee (in the sense that the walls of $\mathcal{D}_{X, S}$ cut out the chambers of \mathcal{C}). Furthermore, \mathcal{C}^\vee forms a simplicial fan.*

Proof. Note that $\mathcal{D}_{X, S}^0$ cuts out the cone $\mathcal{C}_S \subset \mathcal{C}$. The mutation invariance of Lemma 6.2.2 thus implies that every $\mathcal{C}_{S'}$ is cut out by some walls of $\mathcal{D}_{X, S}$. This gives the first claim. We refer to [GHKK] for the second claim. \square

We now have that seeds for \mathcal{A}^\vee correspond certain cells in $\mathcal{D}_{X, S}$. The wall-crossing formula for passing between the cells is exactly the mutation formula for the corresponding seed tori. Thus, applying the construction of §4.4.1 to the cells of \mathcal{C}^\vee (with no need for modding out by any ideals)

produces exactly \mathcal{A}^\vee . To define the theta functions, we can pick a point Q in the interior of \mathcal{C}_S^\vee for some seed S for \mathcal{A}^\vee , and then as in 4.5 we define

$$\vartheta_q|_{\mathcal{A}_S^\vee} := \sum_{\gamma | \text{Ends}(\gamma) = (q, Q)} c_\gamma z^{m_\gamma}.$$

Unfortunately, there is no guarantee that ϑ_q will be a Laurent polynomial (as opposed to some infinite sum). We can deal with this by introducing a formal version of \mathcal{A}^\vee where infinite sums are allowed. Alternatively we can define $\text{can}(\mathcal{A}^\vee)$ to be the submodule of $\text{up}(\mathcal{A}^\vee) := \Gamma(\mathcal{A}^\vee, \mathcal{O}_{\mathcal{A}^\vee})$ generated by the theta functions which are in fact Laurent polynomials. [GHKK] shows that $\text{can}(\mathcal{A}^\vee)$ is in fact closed under multiplication, thus defining an algebra $\text{mid}(\mathcal{A}^\vee)$.

$\text{mid}(\mathcal{A}^\vee)$ always contains the cluster algebra and is contained in the upper cluster algebra (hence the name), and quite often it does equal the entire upper cluster algebra (in particular in all acyclic cases, since here the lower and upper cluster algebras are known to be equal).

Conjecture 6.2.4. $\text{mid}(\mathcal{A}^\vee) = \text{up}(\mathcal{A}^\vee)$ if and only if the $\text{up}(\mathcal{A}^\vee)$ is finitely generated.

6.3 The relation to the log Calabi-Yau surface constructions

Since we have seen that log Calabi-Yau surfaces U appear as fibers of cluster \mathcal{X} -varieties, it is natural to wonder what exactly the construction from Chapter 4 has to do with the cluster situation. We will explain this now, in arbitrary dimension.

6.3.1 The scattering diagram in U^{trop}

We first directly generalize the construction from §4.7.3 to our current situation. We choose a smooth compactification (Y, D) of U^{ft} with boundary divisors corresponding to points in $U(\mathbb{Z}^t)$. By a smooth compactification, I mean that D contains no singular points of Y , but U is allowed singularities.

Assume the compactification is such that we have a toric model $\pi : (Y, D) \rightarrow (\bar{Y}, \bar{D})$, with blowups being those corresponding to the non-frozen seed vectors—that is, \bar{Y} has $U_S(\mathbb{Z}^t)$ as its cocharacter lattice, and the fan Σ includes the rays generated by vectors (\cdot, e_i) , $i \in I \setminus F$. Note that \bar{Y} is smooth. If we let (\hat{Y}, \hat{D}) be the compactification of \mathcal{X} corresponding to this fan, then the blowups are along the loci $H_{i,Y} := \{1 + z^{e_i} = 0\} \cap \hat{D}_{(\cdot, e_i)} \cap Y$. Recall that these blowups are repeated d'_i times, and we keep only the last exceptional divisor, contracting the others to a singular locus. Also recall that $H_{i,Y}$ is a union of $\text{ind}_{\bar{N}_1}(e_i, \cdot)$ irreducible hypertori.

Now, as in §4.7.3, we define $\bar{\varphi} : U(\mathbb{R}^t) \rightarrow A_1(\bar{Y}, \mathbb{Z})$ to be an integral Σ -piecewise linear function with bending parameter $[C_\rho]$ along each codimension one wall ρ (C_ρ here is the curve corresponding to ρ). We define $P_0^{gp} = \pi^*(A_1(\bar{Y}, \mathbb{Z}))$, so $P^{gp} := A_1(Y, \mathbb{Z}) = P_0^{gp} \oplus \mathcal{E}$, with \mathcal{E} being the lattice generated by the exceptional curves (i.e., including a curve E_i for each $i \in I \setminus F$ equal to the exceptional curve over one of the points blown up by μ_i). Take $P := \pi^*(\text{NE}(\bar{Y})) \oplus \mathcal{E}$. Define $\varphi_\pi : U^{\text{trop}} \rightarrow P^{gp}$ by $\varphi_\pi(u) := (\pi^*(\bar{\varphi}(u)), 0)$.

Generalizing Equation 4.7, we now define our initial scattering diagram $\mathfrak{D}_{U,S}^0$ to be the scattering diagram in $M_{S,\mathbb{R}}$ with walls

$$\left\{ e_i^\perp, (1 + z^{\bar{\varphi}_\pi(v_i) - d'_i[E_i]})^{\text{ind}_{\bar{N}_1}((e_i, \cdot))} \mid i \in I \setminus F \right\}. \quad (6.3)$$

The $\text{ind}((e_i, \cdot))$ -term is to account for the number of irreducible hypertori being blown up in $D_{(\cdot, e_i)} \cap Y$. We will think of \mathfrak{D}_U^0 as living in $U(\mathbb{R}^t) \oplus P_{\mathbb{R}}^{gp}$, rather than just living in $U(\mathbb{R}^t)$ with some exponents in P^{gp} . Here, we replace walls with their preimages under the projection, which we will identify with p_2^* .² Our goal is to show that $\mathcal{D}_{U,S}^0 = \mathcal{D}_{\mathcal{A},S}^0$. This means that the construction of the mirror to \mathcal{A} is a direct generalization of the mirror constructions of Chapter 4!

Remark 6.3.1. Note that if we do not intersect e_i^\perp with $U(\mathbb{R}^t)$ above, this also defines a scattering diagram in $\mathcal{X}(\mathbb{R}^t)$ which this time does satisfy the convexity condition. Setting $z^{\varphi(v_i) - d'_i[E_i]} = 1$ for each $i \in I \setminus F$ recovers the naive $\mathcal{D}_{\mathcal{X},S}^0$ from Equation 6.1.

6.3.2 $K_2 \cong A_1(Y, \mathbb{Z})$

Let U be a fiber of some \mathcal{X}^{ft} , up to irrelevant loci. Let S be a seed for \mathcal{X} , and let $\mathcal{F} := \{f_i := p_2^*(e_i)\}_{i \in F}$. Assume the f_i 's are primitive and distinct.³ Let (Y, D) be a partial minimal model for U such that each D_i equals some D_{f_i} , and conversely each $D_{f_i} \subset D$. Let (\tilde{Y}, \tilde{D}) be a further compactification which is compact, nonsingular, and admits a toric model $\pi : (\tilde{Y}, \tilde{D}) \rightarrow (\bar{Y}, \bar{D})$ corresponding to S , as in the previous subsection. Let $(\tilde{D} \setminus D)^\perp \subseteq A_1(Y, \mathbb{Z})$ be the sublattice generated by classes of curves C such that $[C] \cdot [\tilde{D}_k] = 0$ for all $D_k \subseteq \tilde{D} \setminus D$ (note that this does not depend on the choice of (\tilde{Y}, \tilde{D})). For each $i \in I \setminus F$, let $E_i := \pi^{-1}(x_i)$ for some generic point $x_i \in H_i$. Let $\mathcal{E} \subset A_1(Y, \mathbb{Z})$ denote the sublattice generated by the $d'_i[E_i]$'s.

Proposition 6.3.2. *S determines an injection $\kappa : K_2 \hookrightarrow A_1(Y, \mathbb{Z})$ with image $(\pi^*(A_1(\bar{Y}, \mathbb{Z})) \oplus \mathcal{E}) \cap (\tilde{D} \setminus D)^\perp$. In particular, if Y is compact and smooth (so each $d'_i = 1$), then κ is an isomorphism. Dropping compactness, it is $(\tilde{D} \setminus D)^\perp$.⁴*

Proof. An element $r \in K_2$ corresponds to a relation $\sum_{i \in I} a_i v_i = 0 \in U_S^{\text{trop}}(\mathbb{Z})$, which corresponds to a the curve class $[C_r] \in A_1(\bar{Y}, \mathbb{Z})$ such that $[C_r] \cdot [\bar{D}_\rho] = \sum d'_j a_j$, where the sum is over all j such that $v_j \in \rho$. Assume U is generic. Define the map

$$\kappa : r \mapsto \pi^*[C_r] - \sum_{i \in I \setminus F} a_i d'_i [E_i]. \quad (6.4)$$

We easily see that this is indeed in $(\pi^*(A_1(\bar{Y}, \mathbb{Z})) \oplus \mathcal{E}) \cap (\tilde{D} \setminus D)^\perp$. For the inverse, the coefficient of e_i in $\kappa^{-1}([C])$ is

$$a_i = [C] \cdot [D_{v_i}] \quad (6.5)$$

for $i \in F$ and is determined by the intersection with $\pi^{-1}(H_i)$ for $i \in I \setminus F$.

If U is not generic (i.e., if two different H_i 's coincide), the map can be obtained by using the Gauss-Manin connection to parallel transport $\kappa(r)$ from a nearby fiber of λ . \square

Definition 6.3.3. We take P^{gp} to be the image of κ .

²This does introduce more integral points, but the theta functions corresponding to two integral points in the same fiber will only differ by the pullback of a monomial on the base of the mirror family.

³It may be a good idea to try and avoid this assumption, but several formulas would get more complicated, and I do not know of any geometric reason to study cases where this assumption fails. The [GHKK] construction does apply without this assumption, but it does not use our geometric motivation.

⁴Our argument is a generalization of the one in [GHK13a] for the two-dimensional case with no frozen vectors (their Theorem 5.5). The case with no frozen vectors is also closely related to their Theorem 4.1.

Example 6.3.4. Suppose Y is compact and smooth. Consider $e_i, e_j \in E$ with $i \in I \setminus F, j \in F$, $p_2 * (e_i) = d'_i f_j$, where $f_j = p_2^*(e_j)$ is primitive. Then $d'_i[E_i] = \kappa(d'_i e_j - e_i)$.

6.3.3 $\mathcal{A}_S^{\text{trop}}$ as $U_S^{\text{trop}} \oplus P_{\mathbb{R}}^{gp}$

We now have that $\mathcal{A}^{\text{trop}}$ is an P^{gp} -torsor over U^{trop} . Choose an integral linear section $s : \overline{N}_2 \otimes \mathbb{R} \rightarrow N_{\mathbb{R}}$ of p_2^* . This determines an identification $\mathcal{A}_S^{\text{trop}} \cong U_S^{\text{trop}} \oplus P_{\mathbb{R}}^{gp}$.

We now return to the assumptions of §4.7.3 about (Y, D) being a smooth compactification (at least smooth at points on the boundary) which admits a toric model corresponding to S . We want to describe a section $\varphi : U_S^{\text{trop}} \rightarrow \mathcal{A}_S^{\text{trop}}$ which agrees with the map φ_{π} from before. First note that if we view φ_{π} as a function to $P_{\mathbb{R}}^{gp}$, the choice of s only changes φ by a linear function, so the bending parameters (which are all we care about) are unchanged.

We must assume that every wall of $\mathcal{D}_{U,S}^0$ is a union of cones of Σ .⁵ We define

$$\varphi(f_i) = e_i,$$

and we extend linearly over the maximal cones of Σ .

Lemma 6.3.5. *Let ρ be a codimension 1 wall of Σ , corresponding to a curve class $[C] \in A_1(\overline{Y}, \mathbb{Z})$. Then φ has bending parameter $\pi^*([C])$ along ρ .*

Proof. Let f_0, \dots, f_{n-1} denote the primitive vectors generating ρ , cyclically ordered respecting to the orientation of U^{trop} . Let σ_0, σ_n denote the two maximal cones containing ρ (recall that \overline{Y} smooth implies each maximal cone is generated by a basis), and let f_0, f_n denote the other generators of these two cones, respectively. Assume the indexing is chosen so that the ordering respects the orientation. By Proposition 2.7.5, $f_n = -f_0 - \sum_{i=1}^{n-1} ([C] \cdot D_{f_i}) f_i$.

Recall Equation 2.1 defining the bending parameter b_{ρ} :

$$b_{\rho} m_{\rho} = \varphi^{\text{trop}}|_{\sigma_n} - \varphi^{\text{trop}}|_{\sigma_0},$$

where m_{ρ} is a primitive element of the dual space 0 along ρ and positive on σ_n , and by restriction, we really mean the linear extension of the restriction. Evaluating both sides at f_n , we have

$$\begin{aligned} b_{\rho} &= (e_n - \sum_j e_{nj}) + (e_0 - \sum_j e_{0j}) + \sum_{i=1}^{n-1} ([C] \cdot D_{f_i})(e_i) \\ &= \sum_{i \in F} ([C] \cdot D_{f_i}) e_i. \end{aligned}$$

Now applying the map κ from Equation 6.4, this clearly becomes $b_{\rho} = \pi^*([C])$, as desired. \square

Remark 6.3.6. We note that if one equips U^{trop} with the GHK integral linear structure, then the argument above shows that the bending parameter corresponding to ρ is the class of the curve corresponding to ρ in Y , rather than the pullback of the one in \overline{Y} . This is consistent with the definition of $\varphi_{\text{NE}(Y)}$ in §4.3. I do not take this approach because it is not clear how to define the scattering diagram for this integral linear structure.

⁵I suspect that this can be avoided if we correctly specify an integral linear structure on $\mathcal{A}(\mathbb{R}^t)$, but I have not yet figured this out. The assumption will not really limit us though, since we can always add frozen vectors to make the assumption satisfied and then remove them later on.

Theorem 6.3.7. *Using the above identification of $\mathcal{A}_S^{\text{trop}}$ with $U_S^{\text{trop}} \times (\pi^*(A_1(\overline{Y}, \mathbb{Z})) \oplus \mathcal{E})$, we have $\mathcal{D}_{U,S}^0 = \mathcal{D}_{\mathcal{A},S}^0$. Consequently, the [GHKK] construction of the mirror to \mathcal{A} generalizes the [GHK11] construction of the mirror to (Y, D) . In particular, the mirror to (Y, D) is \mathcal{X}^\vee .*

Proof. First recall the definitions of the scattering diagrams:

$$\begin{aligned} \mathfrak{D}_{U,S}^0 &= \left\{ e_i^\perp, (1 + z^{\widetilde{\varphi}_\pi(v_i) - d'_i[E_i]})^{\text{ind}_{\overline{N}_1}((e_i, \cdot))} \mid i \in I \setminus F \right\} \\ \mathcal{D}_{\mathcal{A},S}^0 &:= \{(e_i, \cdot)^\perp, (1 + z^{e_i})^{\text{ind}_M((e_i, \cdot))} \mid i \in I \setminus F\} \subset \mathcal{A}_S^{\text{trop}}. \end{aligned}$$

Also, recall that $\{v_i\}_{i \in F}$ is assumed to be a spanning set of primitive vectors in $U_S^{\text{trop}}(\mathbb{Z})$. For $\cdot \wedge \cdot$ the standard wedge form on $U^{\text{trop}}(\mathbb{Z})$, we can assume we have scaled so that $[e_i, e_j] = d_i d_j (v_i \wedge v_j)$, and so $\langle e_i, e_j \rangle = v_i \wedge v_j$. Consequently, $\{p_1^*(e_i)\}_{i \in F}$ consists of primitive vectors spanning $\overline{N}_1^{\text{sat}} \subset M$, and so $\overline{N}_1 = \overline{N}_1^{\text{sat}}$. Hence, $\text{ind}_M((e_i, \cdot)) = \text{ind}_{\overline{N}_1}((e_i, \cdot))$.

Since $p_2^*(a) \in e_i^\perp$ means exactly that $(e_i, a) = 0$, we see that $(p_2^*)^{-1}(e_i^\perp)$ is exactly $(e_i, \cdot)^\perp$, so the supports of the walls of the two scattering diagrams agree. Now it remains to check that $\kappa(e_i) = \widetilde{\varphi}_\pi(v_i) - [E_i]$. [THIS $\widetilde{\varphi}_\pi$ IS φ FROM BEFORE]. Lemma 6.3.5 tells us that for $p_2^*(e_i) = d'_i p_2^*(e_j)$, $i \in I \setminus F$, $j \in F$, we have $\widetilde{\varphi}_\pi(v_i) = d'_i e_j$. Now applying Example 6.3.4, we get $d'_i e_j - \kappa^{-1}(d'_i [E_i]) = e_i$, as desired. \square

6.3.4 The Relative Torus Action and Removing Frozen Vectors

For the constructions above, one often has to include extra frozen vectors in the seed data in order to satisfy the assumptions. Of course, if we want to remove the frozen vector e_j , $j \in F$, we can simply take the subspace $(e_j^*)^\perp \subset \mathcal{A}_S^{\text{trop}}$, and also take the obvious intersection of $\mathcal{D}_{\mathcal{A},S}^0$ with this subspace. This clearly recovers $\mathcal{D}_{\mathcal{A}_j, S_j}^0$, where we use the subscript j to indicate that we have removed e_j from the seed data.

Geometrically, we can interpret this removal frozen vectors as taking the quotient of the mirror family $\mathcal{X}^\vee \rightarrow T_{K_2^*}$ (or possibly a compactification, infinitesimal, or formal version of this) by the “relative torus action” defined in §5 of [GHK11].⁶ This means that for $j \in F$, we consider the action of the one parameter family $T_{\langle D_j \rangle} \subset T_{\text{Pic}(Y)}$ on the mirror family. For $a := \sum a_i e_i \in K_2$, z^a is an eigenfunction for the action of $T_{\langle D_j \rangle}$ with eigenvalue $\kappa(a) \cdot D_j$, which by Equation 6.5 is equal to a_j . Thus, monomials corresponding to points in $(e_j^*)^\perp \cap K_2$ give the ring of invariants for the action on $T_{K_2^*}$. Noting that the scattering automorphisms do not affect the value of e_j^* on the exponents of the monomials (because frozen vectors do not appear in the scattering functions), we get the equivariance of the action, and we also see that the theta functions are eigenfunctions for the action of $T_{\langle D_j \rangle}$ with ϑ_q having eigenvalue $(e_j^*)(q)$. In summary:

Theorem 6.3.8. *Let $\mathcal{X}^\vee \rightarrow B$ be the mirror family to \mathcal{A} corresponding to the seed S (B being $T_{K_2^*}$ or whatever compactified, infinitesimal, or formal subspace is appropriate). For any $j \in F$, the one parameter family $T_{\langle D_j \rangle}$ corresponding to the inclusion of $D_j := D_{p_2^*(e_j)}$ into K_2^* acts equivariantly on $\mathcal{X}^\vee \rightarrow B$. The theta functions are eigenfunctions with the eigenvalue of ϑ_q being $(e_j^*)(q)$ (interpreted with respect to the seed S). The quotient is the mirror to the cluster variety obtained by removing the frozen vector e_j from S to obtain a seed we call S_j . $\mathcal{D}_{\mathcal{A}_j, S_j}^0$ is obtained from $\mathcal{D}_{\mathcal{A}, S}^0$ by intersecting with $(e_j^*)^\perp$*

⁶The relative torus action in [GHK11] is the one generated by our $T_{\langle D_j \rangle}$'s for all $D_j \subset D$. Thus, the quotient by their relative torus action actually corresponds to removing all the frozen vectors.

6.4 Using Principal Coefficients

To get the mirror to \mathcal{X} , we could use the scattering diagram described above in Remark 6.3.1. We will see that the scattering diagram for $\mathcal{X}^{\text{prin}}$ is an extension of this, and then, as in [GHKK], we will use the scattering diagrams in $(\mathcal{X}^{\text{prin}})^{\text{trop}}$ and $(\mathcal{A}^{\text{prin}})^{\text{trop}}$ to better understand the mirrors to \mathcal{X} and \mathcal{A} respectively. In particular, we will be able to see that the mirror to \mathcal{A} (resp. \mathcal{X}) is in fact \mathcal{X}^\vee (resp. \mathcal{A}^\vee) as desired.

By Lemma 5.2.5, since $\langle \cdot, \cdot \rangle_{S^{\text{prin}}}$ is non-degenerate, each $\mu_j^{\mathcal{A}}$ preserves the centers of the other blowups. Thus, up to codimension 2, we

Chapter 7

Examples and Applications of Cluster Algebras

In this chapter we will briefly look at some applications of cluster algebras. This is a very small sample of the known applications, and even for these cases we do not go into much detail. The goal is simply to help motivate the study of cluster algebras and to give the interested reader some ideas of where to go with the theory.

7.1 Cluster structures on double Bruhat cells of semisimple groups

7.2 Cluster varieties and surfaces

This section is based closely on [FG06].

Let \bar{S} be a closed oriented surface of genus g with non-intersecting open discs D_1, \dots, D_n . $S := \bar{S} \setminus \bigcup D_i$. Let \hat{S} denote the data of S along with a finite (possibly empty) set of distinct marked boundary points $\{x_1, \dots, x_k\}$, considered up to isotopy. Let $\partial\hat{S}$ denote the boundary of S , minus the marked points. Let N be the number of connected components of $\partial\hat{S}$. We assume that \hat{S} is hyperbolic, meaning that either $g > 1$, or $g = 1$ with $N \geq 1$, or $g = 0$ with $N \geq 3$. Furthermore, we assume $N > 0$.

Let G be a group. A G -local system is the data of a G -principal bundle along with a flat connection. It is a standard fact that the moduli space $\mathcal{L}_{G,S}$ of G -local systems on S can be identified with $\text{Hom}(\pi_1(S), G) // G$, where the quotient is by the conjugation action of G (this is a GIT quotient, and the result is an algebraic stack).

Let \mathcal{B} be the flag variety parametrizing Borel subgroups of G (for G a subgroup of $\text{GL}_n(\mathbb{C})$, these are the subgroups which are isomorphic to the subgroup of upper-triangular matrices in some basis). If we choose a Borel subgroup B , we can identify $\mathcal{B} = G/B$. Let $U := [B, B]$ be the corresponding maximal unipotent subgroup (i.e., strictly upper-triangular matrices).

For a G -local system \mathcal{L} on S , G acting on the right, one defines the associated *flag bundle* $\mathcal{L}_{\mathcal{B}}$ and

principal affine bundle \mathcal{L}_A by

$$\begin{aligned}\mathcal{L}_B &:= \mathcal{L} \times_G \mathcal{B} = \mathcal{L}/B, \\ \mathcal{L}_A &:= \mathcal{L}/U.\end{aligned}$$

Definition 7.2.1. Let G be a split reductive group (usually with trivial center). A *framed G -local system* on \widehat{S} is a pair (\mathcal{L}, β) where \mathcal{L} is a G -local system on S and β is a flat section of $\mathcal{L}_B|_{\partial\widehat{S}}$. $\mathcal{X}_{G,\widehat{S}}$ is the moduli space of framed G -local systems on \widehat{S} .

Choose a point y_i on each component of $\partial\widehat{S}$. β may be viewed as a choice of flag F_{y_i} in \mathcal{L}_{y_i} for each i . If the component C_i containing y_i is a circle (i.e., a boundary component of S with no marked points), then F_{y_i} must be monodromy invariant about this circle.

Example 7.2.2. Let \widehat{S} be a disk with $k \geq 3$ marked points on the boundary. Since S is simply connected, any G -local system \mathcal{L} on S is trivial. $\mathcal{X}_{G,\widehat{S}}$ can then be identified with the space of ordered sets of k not necessarily distinct flags in G —i.e., $\mathcal{X}_{G,\widehat{S}} \cong G \backslash \mathcal{B}^k$, where G acts diagonally by conjugation. Note that the cyclic ordering of the k flags is part of the data of $\mathcal{X}_{G,\widehat{S}}$, but extending this cyclic ordering to an actual ordering requires a choice. Similarly, if $s_G = \text{Id}$ (e.g. if $|G|$ has odd order) then $\mathcal{A}_{G,\widehat{S}} \cong G \backslash \mathcal{A}^k$, where $\mathcal{A} := G/U$ is the principal affine space of G .

For example, if $G = \text{PGL}_2(\mathbb{C})$, then $\mathcal{B} \cong \mathbb{P}^1$, and $\mathcal{X}_{G,\widehat{S}}$ is a compactification of the moduli space $\mathcal{M}_{0,k}$ of k distinct ordered marked points on \mathbb{P}^1 , up to an automorphism of \mathbb{P}^1 . The corresponding cluster variety we will construct can be identified with a compactification $\mathcal{M}_{0,k}^{\text{cyc}}$ of $\mathcal{M}_{0,k}$ which is very closely related to the usual Deligne-Mumford compactification $\overline{\mathcal{M}}_{0,k}$ (see [FG11] for the details).

The definition of the \mathcal{A} -space is similar (using \mathcal{L}_A in place of \mathcal{L}_B), but with a complication involving keeping track of certain “twisting” data. If the center of G has odd order, this twisting goes away, and we can state the definition as follows:

Definition 7.2.3. Let G be a split simply connected semisimple algebraic group of odd order (i.e., type $A_{2k}, E_6, E_8, F_4, G_2$). A *decorated G -local system* on \widehat{S} is a pair (\mathcal{L}, α) where \mathcal{L} is a G -local system on S and α is a flat section of $\mathcal{L}_A|_{\partial\widehat{S}}$. $\mathcal{A}_{G,\widehat{S}}$ is the moduli space of decorated G -local systems on \widehat{S} .

If the order of the center is not odd, we must instead use decorated *twisted G -local systems*. Very briefly (cf. [FG06] §2 for more details), let w_0 be the longest element of the Weyl group of G . This admits a natural lift \overline{w}_0 to G [I THINK I WILL NEED THIS FOR THE DOUBLE BRUHAT CELL STUFF, SO MAYBE I SHOULD INCLUDE MORE DETAIL]. Let $s_G := \overline{w}_0^2$. s_G is central and $s_G^2 = 1$. For example, if $G = \text{SL}_m$, then $s_G = (-1)^{m-1} \text{Id}$.

Now let $T'S$ be the punctured tangent bundle of S : i.e., the tangent bundle minus the 0-section. Let γ denote a curve in a fiber of $T'S$ generating the fundamental group of the fiber. Let σ_S denote the corresponding class in $\pi_1(T'S)$. Let \mathbf{C}_i denote a small annulus in S with one of its boundary circles being a boundary circle C_i of S . Let \mathbf{C}'_i be the complement of the marked points in C_i . View \mathbf{C}_i as $C_i \times [0, 1]$, and let $\mathbf{C}'_i = C'_i \times \{1\} \subset \mathbf{C}_i$. We use a nonvanishing section of $T\mathbf{C}_i \times \{1\} \subset TS$ to view \mathbf{C}'_i as living in $T'S$.

Definition 7.2.4. Let G be a split simply connected semisimple algebraic group. A *twisted G -local system* on S is a G -local system on $T'S$ with monodromy s_G about σ_S . A *decorated twisted G -local system* on S is a twisted G -local system represented by \mathcal{L} on $T'S$ along with a choice of locally constant section α of $\mathcal{L}_A|_{\cup \mathbf{C}'_i}$. $\mathcal{A}_{G,\widehat{S}}$ is the moduli space of decorated twisted G -local systems (\mathcal{L}, α) on \widehat{S} .

7.2.1 Cluster structures

[FG06] describe cluster structures on $\mathcal{A}_{\mathrm{SL}_m, \widehat{S}}$ and $\mathcal{X}_{\mathrm{PGL}_m, \widehat{S}}$. They claim that cluster structures can be put on these spaces for other groups, but from what I understand there are still details that need worked out. The main ideas are supposed to be already captured in the relatively simple SL_2 and PGL_2 examples, so we will just cover these.

Definition 7.2.5. Choose a single distinguished point on each connected component of $\partial\widehat{S}$. An *ideal triangulation* T of \widehat{S} is a triangulation of S with all vertices at these distinguished points. We view boundary disks without marked points as punctures in \overline{S} , so the distinguished point on such a disk can instead be viewed as the point removed when making the puncture. Also, we allow degenerate triangles where an edge appears twice. An edge of T is called *internal* if it is not contained in $\partial\overline{S}$ (boundary curves connect two distinct distinguished points on the same boundary component—in particular, edges cannot be contained in punctures).

Note that a point $z \in \mathcal{X}_{\mathrm{PGL}_2(\mathbb{C})}$ determines a flag (in this case, a point in \mathbb{P}^1) at each vertex of T . Let e be an internal edge of T , contained in a quadrilateral Q_e . Let x_1, x_2, x_3, x_4 be the four vertices of Q_e , clockwise ordered, with x_1 being a vertex of e (it does not matter which vertex). Let p_1, \dots, p_4 be the corresponding points in \mathbb{P}^1 determined by a choice of $z \in \mathcal{X}_{\mathrm{PGL}_2(\mathbb{C})}$ —we can view them as living in the same \mathbb{P}^1 using parallel transport in Q with respect to the flat connection corresponding to z . We have a cluster variable $X_e : \mathcal{X}_{\mathrm{PGL}_2(\mathbb{C}), \widehat{S}} \rightarrow \mathbb{C}^*$ corresponding to e defined as the cross-ratio $X_e(z) = -\frac{(p_4-p_1)(p_2-p_3)}{(p_4-p_3)(p_2-p_1)}$ (i.e., (p_1, p_2, p_3, p_4) can be identified with $(0, -1, \infty, X_e(z))$).

Each choice of T corresponds to a seed for $\mathcal{X}_{\mathrm{PGL}_2(\mathbb{C}), \widehat{S}}$ (viewed as a cluster \mathcal{X} -variety)—the cluster variables for the corresponding seed are the X_e 's as above for the internal edges e . If E is the set of internal edges, then in the corresponding seed, $\langle e_1, e_2 \rangle$ is the number of triangles in T containing e_1 and e_2 with e_2 following e_1 in the counterclockwise order, minus the number of such triangles with e_1 following e_2 in the counterclockwise order (so this pairing is always an integer in $[-2, 2]$). In other words, we can view each internal edge as a vertex of the corresponding quiver with arrows going counterclockwise through the triangles. A mutation with respect to an edge e then corresponds to replacing e with the other diagonal of Q_e .

We thus obtain a cluster \mathcal{X} -variety equal to the subset of $\mathcal{X}_{\mathrm{PGL}_2(\mathbb{C}), \widehat{S}}$ where all the X_e 's for some triangulation T are \mathbb{C}^* -valued.

We can similarly give a nice geometric interpretation of the corresponding \mathcal{A} -coordinates which realizes the corresponding cluster \mathcal{A} -variety as a subspace of $\mathcal{A}_{\mathrm{SL}_2(\mathbb{C}), \widehat{S}}$. Namely, a point $w \in \mathcal{A}_{\mathrm{SL}_2(\mathbb{C}), \widehat{S}}$ determines a vector $v_i \in \mathrm{SL}_2(\mathbb{C})/U \cong \mathbb{C}^2$ at each vertex x_i of T . For an internal edge e of T with vertices x_1, x_2 , we associate the coordinate $A_e := v_1 \wedge v_2$. To compute this wedge product, we can parallel transport v_1 to x_2 along a path following the edges of Q_e in the clockwise direction. Suppose we were to instead define $A'_e := v_2 \wedge v_1$ with v_2 parallel transported to x_1 along Q_e in the clockwise direction. It appears at first that this $A'_e = -A_e$, but this is where the twisting comes in: to define A_e , we could have instead parallel transported v_2 to x_1 counterclockwise along Q_e . Let σ be the monodromy action on $\mathrm{SL}_2(\mathbb{C})/U$ corresponding to parallel transporting along a path following the edges of Q_e in the clockwise direction. Since this path is contractible, the contribution to the monodromy only comes from the fact that the path goes around the puncture (the removed 0-section) once when lifted to $T'S$. Thus, $\sigma = -\mathrm{Id}$, and so $A'_e = \sigma(v_2) \wedge v_1 = v_1 \wedge v_2$. So the \mathcal{A} -monomials are indeed well-defined.

Example 7.2.6. For the situations from Example 7.2.2, triangulations include $k-3$ edges. Choosing a “zig-zag” triangulation results in a type A_{k-3} -quiver.

Remark 7.2.7. We note without proof that the action of the mapping class group $\Gamma_{\widehat{S}}$ of \widehat{S} on \mathcal{A} and \mathcal{X} gives an embedding of $\Gamma_{\widehat{S}}$ into the cluster modular group of \mathcal{A} and \mathcal{X} .

7.2.2 Positive real points and the relation to Teichmüller spaces

Recall the notion of semifield-valued points from §5.3.3. The $\mathbb{R}_{>0}$ -valued points tend to have interesting geometric interpretations. We explore some of these here:

Example 7.2.8. Suppose \widehat{S} has no marked boundary points and $n > 0$ holes, so we can say $\widehat{S} = S$. Theorem 1.7 of [FG06] says

- $\mathcal{A}_{\mathrm{SL}_2(\mathbb{C}),S}(\mathbb{R}_{>0})$ is equal to Penner’s decorated Teichmüller space \mathcal{T}_S^d .
- Let \mathcal{T}_S^+ be the finite $(2^n : 1)$ cover of the classical Teichmüller space \mathcal{T}_S consisting of points $p \in \mathcal{T}_S$ together with a choice of orientation for each boundary component of S (technically, a choice of orientation for each boundary component which is non-cuspidal with respect to the hyperbolic structure corresponding to p —this cover is ramified on the boundary of \mathcal{T}_S). Then $\mathcal{X}_{\mathrm{PGL}_2(\mathbb{C}),S}(\mathbb{R}_{>0}) = \mathcal{T}_S^+$.
- Recall that we have a forgetful map $\pi : \mathcal{X}_{G,S} \hookrightarrow \mathcal{L}_{G,S}$. [FG06] defines $\mathcal{L}_{G,S}(\mathbb{R}_{>0})$ to be the image of $\pi|_{\mathcal{X}_{G,S}(\mathbb{R}_{>0})}$, and they identify $\mathcal{L}_{\mathrm{PGL}_2(\mathbb{C}),S}(\mathbb{R}_{>0})$ with \mathcal{T}_S .

The forgetful map $\pi : \mathcal{X}_{\mathrm{PGL}_2(\mathbb{C}),S} \rightarrow \mathcal{L}_{\mathrm{PGL}_2(\mathbb{C}),S}$ is also generically $2^n : 1$. For a generic point x in $\mathcal{L}_{\mathrm{PGL}_2(\mathbb{C}),S}$, the corresponding monodromy action on $G/B \cong \mathbb{P}^1$ has two invariant points (because a generic linear map on \mathbb{C}^2 has two eigenspaces) and the fiber over x is the set of choices of such an invariant point for each hole. The map is ramified when the monodromy around a hole is unipotent (which for positive real points does indeed correspond to the hole being cuspidal).

The space $\mathcal{L}_{G,S}$ has a Poisson structure where symplectic leaves correspond to specifying the traces of the monodromies around the holes. We have seen that \mathcal{X} -spaces also have a Poisson structure. The map π is Poisson.

The 2-form $\widetilde{\Omega}$ on \mathcal{A} generalizes the Weil-Petersson 2-form on decorated Teichmüller space. The image of the map $\mathcal{A} \rightarrow \mathcal{X}$ is equal to the quotient of \mathcal{A} by the null-foliation of $\widetilde{\Omega}$. It is also equal to the subset of \mathcal{X} where all monodromies are unipotent (i.e., the ramification locus of π).

[FG11] also describes how to compactify the \mathcal{X} -space and $\mathcal{X}(\mathbb{R}_{>0})$, and they relate this to Thurston’s compactification of Teichmüller space. I don’t plan to say any more about this.

Example 7.2.9. Let \mathcal{X} be the \mathcal{X} -space corresponding to the situation from Example 7.2.2 (k marked points on the boundary of a disk). Then $\mathcal{X}(\mathbb{R}_{>0})$ can be identified with the subspace of $\mathcal{M}_{0,k}$ consisting of points in $\mathbb{P}^1(\mathbb{R}) \cong S^1$ whose cyclic ordering with respect to the orientation of $\mathbb{P}^1(\mathbb{R})$ agrees with that of the corresponding components of $\partial\widehat{S}$.

7.2.3 Laminations and Canonical Coordinates

§12 of [FG06] identifies $\mathcal{A}_{\mathrm{SL}_2(\mathbb{C}),S}(\mathbb{Q}^t)$ with set of “rational \mathcal{A} -laminations” of S , and $\mathcal{X}_{\mathrm{PSL}_2(\mathbb{C}),S}(\mathbb{Z}^t)$ with the set of “integral \mathcal{X} -laminations.” These are defined as follows:

Definition 7.2.10. A rational (resp., integral) \mathcal{A} -lamination is a homotopy class of a finite collection of disjoint simple unoriented closed curves with rational (resp., integral) weights. The weights must be positive unless the curve surrounds a hole. Ignoring weight 0 curves or combining homotopic curves (adding their weights) produces equivalent laminations.

A rational (resp., integral) \mathcal{X} -lamination is a homotopy class of a finite collection of non-self-intersecting and pairwise non-intersecting curves which are either closed or connect two (not necessarily distinct) boundary components of S , along with positive rational (resp., integral) weights and a choice of orientation for any boundary component intersecting at least one of the curves. Curves retracting to boundary components can be removed, and homotopic curves can be combined as before.

[FG06] also shows that $\mathcal{A}_{\mathrm{SL}_2(\mathbb{C}),S}(\mathbb{Z}^t)$ can be identified with the half-integral \mathcal{A} -laminations (those where each weight is in $\frac{1}{2}\mathbb{Z}$) which satisfy a certain parity condition (in particular including all integral \mathcal{A} -laminations).

Fock and Goncharov then describe the theta functions on $\mathcal{A}_{\mathrm{SL}_2(\mathbb{C}),S}$ and $\mathcal{X}_{\mathrm{PSL}_2(\mathbb{C}),S}$ (up to a sign, which can be determined by imposing positivity) corresponding to these laminations as follows:

- Suppose q is a loop with weight k not surrounding a puncture.
 - Viewed as an \mathcal{A} -lamination, the corresponding $\vartheta_q : \mathcal{X}_{\mathrm{PSL}_2(\mathbb{C}),S} \rightarrow \mathbb{C}$ is given by taking the trace of the k^{th} power of the monodromy around q .
 - Viewed as an \mathcal{X} -lamination, with $\vartheta_q : \mathcal{A}_{\mathrm{SL}_2(\mathbb{C}),S} \rightarrow \mathbb{C}$ given by the trace of the k^{th} power of the monodromy around the lift of q to $T'S$.
- Suppose q is a closed curve surrounding a puncture p on S with weight k (an \mathcal{A} -lamination). Let $(\mathcal{L}, \beta) \in \mathcal{X}_{\mathrm{PSL}_2(\mathbb{C}),S}$. Then the choice of monodromy-invariant flag β at the puncture p determines a choice of eigenspace for the monodromy of \mathcal{L} around q . ϑ_q is then λ^k , where λ is the corresponding eigenvalue.
- If q is a weight k curve connecting two punctures (an \mathcal{X} -lamination), then ϑ_q is just A_q^k , where A_q is the cluster A -monomial defined as in §7.2.1 (i.e., when viewing q as an edge in some triangulation).
- The rest of the theta functions are determined by the following: For q_1 and q_2 non-intersecting classes of curves containing no common isotopy class, $\vartheta_{q_1+q_2} = \vartheta_{q_1}\vartheta_{q_2}$.

This deals with all integral \mathcal{A} -laminations, but for \mathcal{X} -laminations which are curves connecting two punctures, the orientation at each puncture must be the one induced by the orientation of the surface. This can be extended to all \mathcal{X} -laminations as follows: The group $(\mathbb{Z}/2\mathbb{Z})^n$ acts birationally on \mathcal{A} by products of the involutions taking the decoration v_p at a puncture p to $\alpha(p)v_p$, where $\alpha(p)$ is defined by saying that the monodromy around p is given in a basis v_p, v'_p by $M_p := \begin{pmatrix} 1 & 0 \\ \alpha(p) & 1 \end{pmatrix}$ (see [FG06], §12.6 for more details). $(\mathbb{Z}/2\mathbb{Z})^n$ also acts on the \mathcal{X} -laminations by changing the orientations at the punctures, and imposing that the action is equivariant determines the rest of the theta functions. This action is in the cluster modular group if and only if $n > 1$.

Exercise 7.2.11. Check that the four-punctured sphere and once-punctured torus really do correspond to the seeds in Example 5.3.2. Observe that the theta functions on the \mathcal{X} -space, as described above, are as we claimed in Exercise 4.5.5.

Appendix A

Some Algebraic Geometry Background

A.1 Divisors

I assume we are dealing with normal varieties over \mathbb{C} , but this may be weakened for much of what I say. A prime Weil divisor in a variety is a closed irreducible codimension 1 subvariety. A **Weil divisor** is a formal sum of prime Weil divisors. A **principal** Weil divisor is one given by the zeroes and poles of a rational function—that is, one of the form $(f) := \sum \text{val}_Y(f) \cdot Y$, where f is a rational function and the sum is over all closed irreducible codimension 1 subvarieties Y . val_Y here denotes the order of zero (or negative the order of a pole) along Y . Two Weil divisors are **linearly equivalent** if they differ by a principal Weil divisor.

A **Cartier divisor** is a global section of $\mathcal{M}^*/\mathcal{O}^*$, the sheaf of nonzero meromorphic functions modulo the sheaf of invertible holomorphic functions. In other words, choosing an affine open cover $\{U_i\}$, it is a rational function f_i on each U_i such that f_i/f_j is an invertible regular function on $U_i \cap U_j$. A **principal** Cartier divisor is one coming from a global rational function. Cartier divisors modulo linear equivalence on an irreducible variety X may be identified with $\text{Pic}(X)$, the group of “invertible sheaves” (i.e., line bundles) (without irreducibility this gives a subgroup of $\text{Pic}(X)$). The transition maps are given by the f_j/f_i 's.

Cartier divisors can be identified with a subset of the Weil divisors (the “locally principal” Weil divisors, i.e., those locally given by zeroes and poles of rational functions) using the zeroes and poles of the f_i 's. This inclusion clearly identifies principal Cartier divisors with principal Weil divisors, and the inclusion is an isomorphism if the variety is irreducible and “locally factorial” (local rings are UFD's). In particular, smooth varieties are locally factorial.

For $D = \sum a_i D_i$ a Cartier divisor expressed as a sum of prime Weil divisors, we can identify the global sections of the associate line bundle $\mathcal{O}(D)$ with the subspace of rational functions whose pole along D_i is at worst $-a_i$ for each i .

If $D_1 \sim D_2$ (\sim meaning linearly equivalent), then for any curve C , $C \cdot D_1 = C \cdot D_2$, where \cdot denotes the intersection form (if the intersections are transverse, this is the number of intersection points). Blowing up a smooth point on a surface gives an exceptional divisor of self-intersection (-1) ,

so conversely, self-intersection (-1) curves are the only ones on a surface that can be blown down to get a smooth point (other curves with negative self-intersection can be blown down to singular points).

Given a Cartier divisor $D \subset X$, let $|D|$ denote the space of effective divisors linearly equivalent to D . $|D|$ is called a linear system, and is isomorphic to $\text{Proj}(H^0(X, \mathcal{L}_D))$, where \mathcal{L}_D is the line bundle corresponding to D . b is called a *base point* of $|D|$ if b is contained in every element of $|D|$. If $|D|$ is basepoint free, then every $b \in X$ determines a hyperplane H in $|D|$ of divisors passing through b . $|H|$ corresponds to a point in the dual projective space $|D|^*$ (i.e., the space of hyperplanes in $H^0(X, \mathcal{L}_D)$ rather than the space of lines), and this determines a map $X \rightarrow |D|^*$. Choosing a basis for $H^0(X, \mathcal{L}_D)$ gives an identification with $\mathbb{P}^{\dim(|D|)}$. Note that hyperplane sections of the image of X in $|D|^*$ correspond to elements of $|D|$.

A.2 Formal schemes

Here I will very briefly introduce the idea of formal schemes, which serve as the algebro-geometric analogue of analytic neighborhoods. In algebraic geometry, the “value” of an element f in a ring R at a prime ideal $\mathfrak{p} \in \text{Spec}(R)$ is defined to be the image \bar{f} of f in the field of functions $\kappa(\mathfrak{p})$ of R/\mathfrak{p} . Alternatively, one can first localize $R_{\mathfrak{p}}$ (i.e., allow division by elements not in \mathfrak{p}), and then mod out by the image of \mathfrak{p} . More generally, if we are not dealing with affine schemes, we can replace this localization $R_{\mathfrak{p}}$ with the localization of the structure sheaf at the point \mathfrak{p} .

In particular, if \mathfrak{p} is maximal and $\text{Spec}(R)$ is a variety, then this is what one usually thinks of as the value of f at \mathfrak{p} , and the localization step can be skipped. For example, for $f \in \mathbb{k}[x]$, $f(a)$ is the image of f in $\mathbb{k}[x]/\langle x - a \rangle$. If we instead take $\mathbb{k}[x]/\langle x - a \rangle^2$, then the image of f is actually $f(a) + f'(a)x$. $\text{Spec} \mathbb{k}[x]/\langle x - a \rangle^2$ is therefore viewed as a “fat point,” which contains data about the tangent space to $\langle x = a \rangle$ instead of just the point itself.

We can similarly consider $\mathbb{k}[x]/\langle x - a \rangle^k$ for any $k \in \mathbb{Z}_{>0}$, with the image of f being the degree $(k - 1)$ Taylor polynomial for f . Taking the inverse limit (of the rings) with respect to k , we obtain $\mathbb{k}[[x - a]]$, the ring of formal (i.e., not necessarily converging) powers series centered at $x = a$. The ringed space whose underlying topological space is $\text{Spec} \mathbb{k}[x]/\langle x - a \rangle$ and whose structure sheaf is this inverse limit $\mathbb{k}[[x - a]]$ (or rather, the corresponding inverse limit of sheaves) is called the formal completion of $\text{Spec} \mathbb{k}[x]$ along $\langle x - a \rangle$, and we will denote it by $\text{Spf}(\mathbb{k}[[x - a]])$. Similarly if we replace $\mathbb{k}[x]$ with a more general Noetherian ring (or more generally, a Noetherian scheme) and $\langle x - a \rangle$ by some other ideal (or a sheaf of ideals).

When we construct a mirror family, it might only be a formal deformation of the large complex structure limit, meaning that it will be a family over a formal scheme which possibly has only one closed point, the fiber over which is the degenerate fiber (the large complex structure limit). Intuitively, one may think of this as being a family over an analytic open disk with singular fiber over the origin. Extending to an algebraic family will require a certain affineness assumption.

Bibliography

- [Aur] D. Auroux, *Mirror symmetry and T-duality in the complement of an anticanonical divisor*, arXiv:0706.3207v2.
- [BFZ05] A. Berenstein, S. Fomin, and A. Zelevinsky, *Cluster algebras. III. Upper bounds and double Bruhat cells*, Duke Math. J. **126** (2005), no. 1, 1–52.
- [BZ05] A. Berenstein and A. Zelevinsky, *Quantum cluster algebras*, Adv. Math. **195** (2005), no. 2, 405–455.
- [CdIOGP92] P. Candelas, X. de la Ossa, P. Green, and L. Parkes, *A pair of Calabi-Yau manifolds as an exactly soluble superconformal theory*, Essays on mirror manifolds, Int. Press, Hong Kong, 1992, pp. 31–95.
- [CJS11] D. Cox, Little J., and H. Schenck, *Toric varieties*, Graduate Studies in Mathematics, vol. 124, American Mathematical Society, Providence, RI, 2011.
- [Cox95] D. Cox, *The homogeneous coordinate ring of a toric variety*, J. Algebraic Geom. **4** (1995), no. 1, 17–50.
- [CPS] M. Carl, M. Pumperla, and B. Siebert, *A tropical view on Landau-Ginzburg models*, Preprint, 2011.
- [FG06] V. Fock and A. Goncharov, *Moduli spaces of local systems and higher Teichmüller theory*, Publ. Math. IHES **103** (2006), 1–212.
- [FG09] ———, *Cluster ensembles, quantization and the dilogarithm*, Ann. Sci.Éc. Norm. Sup. (4) **42** (2009), no. 6, 865–930.
- [FG11] ———, *Cluster \mathcal{X} -varieties at infinity*, arXiv:1104.0407.
- [Ful93] W. Fulton, *Introduction to toric varieties*, Annals of Mathematics Studies, vol. 131, Princeton University Press, Princeton, NJ, 1993.
- [FZ02] S. Fomin and A. Zelevinsky, *Cluster algebras I: Foundations*, J. Amer. Math. Soc. **15** (2002), 497–529.
- [GHK] M. Gross, P. Hacking, and S. Keel, *Mirror symmetry for log Calabi-Yau surfaces II*, (in preparation).
- [GHK11] ———, *Mirror symmetry for log Calabi-Yau surfaces I*, arXiv:1106.4977.

- [GHK13a] ———, *Birational Geometry of Cluster Algebras*, arXiv:1309.2573.
- [GHK13b] ———, *Moduli of surfaces with an anti-canonical cycle*, arXiv:1211.6367.
- [GHKK] M. Gross, P. Hacking, S. Keel, and M. Kontsevich, *Canonical bases for cluster algebras*, arXiv:1411.1394.
- [GLS08] C. Geiss, B. Leclerc, and J. Schröer, *Partial flag varieties and preprojective algebras*, Ann. Inst. Fourier (Grenoble) **58** (2008), no. 3, 825–876.
- [GPS09] M. Gross, R. Pandharipande, and B. Siebert, *The tropical vertex*, arXiv:0902.0779v2.
- [Gro11] M. Gross, *Tropical geometry and mirror symmetry*, CBMS Regional Conference Series in Mathematics, vol. 114, A.M.S., 2011.
- [Gro13] M. Gross, *Mirror Symmetry and the Strominger-Yau-Zaslow conjecture*, arXiv:1212.4220v2.
- [GS03] M. Gross and B. Siebert, *Affine manifolds, log structures, and mirror symmetry*, Turkish J. Math. **27** (2003), no. 1, 33–60.
- [GS06] ———, *Mirror symmetry via logarithmic degeneration data. I*, J. Differential Geom. **72** (2006), no. 2, 169–338.
- [GS11a] ———, *From real affine geometry to complex geometry*, Ann. of Math. (2) **174** (2011), no. 3, 1301–1428.
- [GS11b] ———, *An invitation to toric degenerations*, Surveys in differential geometry. Volume XVI. Geometry of special holonomy and related topics, Surv. Differ. Geom., vol. 16, Int. Press, Somerville, MA, 2011, pp. 43–78.
- [GSV05] M. Gekhtman, M. Shapiro, and A. Vainshtein, *Cluster algebras and Weil-Petersson forms*, Duke Math. J. **127** (2005), no. 2, 291–311.
- [GSV10] ———, *Cluster algebras and Poisson geometry*, Mathematical Surveys and Monographs, vol. 167, American Mathematical Society, Providence, RI, 2010.
- [Hit97] N. Hitchin, *The moduli space of special Lagrangian submanifolds*, Ann. Scuola Norm. Sup. Pisa Cl. Sci. (4) **25** (1997), no. 3-4, 503–515.
- [HKK⁺03] K. Hori, S. Katz, A. Klemm, R. Pandharipande, R. Thomas, C. Vafa, R. Vakil, and E. Zaslow, *Mirror symmetry*, Clay Mathematics Monographs, vol. 1, American Mathematical Society, Providence, RI; Clay Mathematics Institute, Cambridge, MA, 2003, With a preface by Vafa.
- [Kel12] B. Keller, *Cluster algebras and derived categories*, Derived categories in algebraic geometry, EMS Ser. Congr. Rep., Eur. Math. Soc., Zürich, 2012, pp. 123–183.
- [KM98] J. Kollár and S. Mori, *Birational geometry of algebraic varieties*, Cambridge Tracts in Mathematics, vol. 134, Cambridge University Press, Cambridge, 1998.

- [Kon95] M. Kontsevich, *Homological algebra of mirror symmetry*, Proceedings of the International Congress of Mathematicians, Vol. 1, 2 (Zürich, 1994), Birkhäuser, Basel, 1995, pp. 120–139.
- [KS06] M. Kontsevich and Y. Soibelman, *Affine structures and non-Archimedean analytic spaces*, The unity of mathematics, Progr. Math., vol. 244, Birkhäuser Boston, Boston, MA, 2006, pp. 321–385.
- [KS13] ———, *Wall-crossing structures in Donaldson-Thomas invariants, integrable systems and Mirror Symmetry*, arXiv:1303.3253v4.
- [Li02] J. Li, *A degeneration formula of GW-invariants*, J. Differential Geom. **60** (2002), no. 2, 199–293.
- [LS14] K. Lee and R. Schiffler, *Positivity for cluster algebras*, arXiv:1306.2415v3.
- [Lus90] G. Lusztig, *Canonical bases arising from quantized enveloping algebras*, J. Amer. Math. Soc. **3** (1990), no. 2, 447–498.
- [Lus94] ———, *Total positivity in reductive groups*, Lie theory and geometry, Progr. Math., vol. 123, Birkhäuser Boston, Boston, MA, 1994, pp. 531–568.
- [Lus10] G. Lusztig, *Introduction to quantum groups*, Modern Birkhäuser Classics, Birkhäuser/Springer, New York, 2010, Reprint of the 1994 edition.
- [Mana] T. Mandel, *Classification of rank 2 cluster varieties*, arXiv:1407.6241v2.
- [Manb] ———, *Tropical Theta Functions and Log Calabi-Yau Surfaces*, (in preparation).
- [Mik05] G. Mikhalkin, *Enumerative tropical algebraic geometry in \mathbb{R}^2* , J. Amer. Math. Soc. **18** (2005), 313–377.
- [NS06] T. Nishinou and B. Siebert, *Toric degenerations of toric varieties and tropical curves*, Duke Math. J. **135** (2006), 1–51.
- [Pen87] R. C. Penner, *The decorated Teichmüller space of punctured surfaces*, Comm. Math. Phys. **113** (1987), no. 2, 299–339.
- [Rud] H. Ruddat, *Mirror duality of Landau-Ginzburg models via Discrete Legendre Transforms*, arXiv:1204.5611v2.
- [Sco06] J. Scott, *Grassmannians and cluster algebras*, Proc. London Math. Soc. (3) **92** (2006), no. 2, 345–380.
- [Spe13] D. Speyer, *An Infinitely Generated Upper Cluster Algebra*, arXiv:1305.6867.
- [SYZ96] A. Strominger, S.-T. Yau, and E. Zaslow, *Mirror symmetry is T-duality*, Nuclear Phys. B **479** (1996), no. 1-2, 243–259.