Title
Numerics and stability for orbifolds with applications to symplectic embeddings

Author
Wormleighton, Ben

Publication Date
2020

Peer reviewed|Thesis/dissertation

Numerics and stability for orbifolds
with applications to symplectic embeddings

by

Ben Wormleighton

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Mathematics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor David Eisenbud, Chair
Professor Kenneth Ribet
Professor Vivek Shende
Professor Yasunori Nomura

Spring 2020

Numerics and stability for orbifolds
with applications to symplectic embeddings

Abstract

Numerics and stability for orbifolds
with applications to symplectic embeddings

by

Ben Wormleighton

Doctor of Philosophy in Mathematics

University of California, Berkeley

Professor David Eisenbud, Chair

This thesis studies the geometry of orbifolds – primarily via variation of GIT, derived category methods, and numerics – and develops connections of equivariant algebraic geometry with embedding problems in symplectic geometry, and with lattice point counting for rational polytopes. We also compile many aspects of the disparate toolkit required to rigorously study orbifolds.

Contents

Acknowledgments

First and foremost, I want to thank and offer everything expressed in this thesis to Jesus. It is because of His ongoing grace that I am able to live in the freedom and transformation that He freely offers. To my wife, Stephanie – thank you for your love, support, prayers, and deep care for my thriving and wellbeing. To my parents, thank you for supporting me in pursuing my curiosity wherever around the world it took me, and for believing in me every step of the way. To my parents-in-law, thank you for being such a significant part of my life in the US and for making me feel at home every time I'm with you.

I would like to thank the academic mentors in the Bay Area who have played a major role in helping me mature over the last five years. I am very grateful to my advisor for allowing me the freedom to explore while keeping me accountable to getting things written, and for sharing with me a panoramic sense of what a positive contribution to mathematics is; to Jenny Harrison for guiding me through difficult and complicated moments of my PhD, for fearlessly sharing her experiences, and for inspiring me to be an educator who truly cares; to Michael Hutchings for being willing to cross back-and-forth between algebraic and symplectic geometry with me, and for welcoming me into his research community; to Matthias Beck for his deeply human and empathetic eye for what really matters; to my first advisor David Nadler for graciously supporting me both while I was his student and afterwards; to Bernd Sturmfels for his consistent interest and excellent model of what being a boldly inclusive academic can look like; and to Vivek Shende for always being willing to talk about something intriguing.

I would also like to thank my mentors and collaborators from around the world who have offered plentiful support during my time in Berkeley. My work on the McKay correspondence owes a great deal to Alastair Craw, Yukari Ito – especially her generous hosting of me at Nagoya University during the summer of 2016 – and Tom Ducat; my work on polytopes and orbifold del Pezzo surfaces was born out of a very fruitful and enjoyable stay at Imperial College London where I learned much about the practice of research from Al Kasprzyk, Alessio Corti, Tom Coates, and Mohammad Akhtar; my ongoing exploration of the world of symplectic embeddings has been encouraged by Vinicius Ramos, who also kindly hosted me at IMPA during December 2019; and I'm grateful to Antoine Song for introducing me to the theory of minimal hypersurfaces. As will become evident, I owe much of my mathematical palate to my first mentor Miles Reid. His encouragement and support has lasted well past my time at Warwick, and I am certain that his influence will endure in the questions that I choose to pursue.

There are many graduate students at Berkeley I am delighted to thank for the ways they have enriched my social and mathematical life in the Bay. Đan Daniel, I will miss our enjoyably confusing lunches and mutual enthusiasm to talk about absolutely anything;

Eugene, Mike, Sander, and Alex: thank you for being some of the finest officemates to be found anywhere in Evans; I'm grateful to the 'Eisenbuddies', especially Ritvik, Chris, and Mengyuan, for many fun discussions; Julian, I'm really glad we got to work on a project together; Ruairidh, thank you for being an island of British sanity; and lastly thank you to my fellow Berkeley Connect mentors Ed, Michael, Catherine, and Nic whose real compassion for students I admire so much.

I'd like to thank some of my other excellent friends who are graduate students at other universities, especially Clara and Jonathan whose camaraderie and painful jokes I have very much appreciated throughout the job application process.

I want to thank the staff at Berkeley who have helped me in ever so many ways over the last five years. Barb, Vicky, Jon, Isabel, and Marsha – thank you for caring for us graduate students generously and unjudgingly; thanks to Jon and Sug Woo for each serving as head of graduate studies; and thanks to Antonio and Nikhil for supporting me during my time as a Berkeley Connect mentor. I also want to thank Professors Shende, Ribet, and Nomura for kindly being on my dissertation committee.

I'd like to thank my closest friends: David, Valentin, Sam, Sam, Sam (= Sam^3), Kevin, Garrison, and Chad. You are all men of such excellent character, and you've each contributed uniquely and powerfully to the last five years. Here's to many more years of calls, Skypes, and talking about nonsense.

Thank you also to the church communities in the US that I've had the joy of either being a part of or being connected to: the Ark, Emmaus Church Community, and Sanctuary SF. I want to especially thank some people who've been truly wonderful brothers, sisters, fathers, and mothers in Christ to me: Lauren, Bethany, David, Yining, Gary, Jason, and Cory.

And lastly, thank you, the reader, for making it through these acknowledgments and onto my thesis! I hope that what follows illustrates some of my mathematical process, and ultimately provokes many curiosities and questions of your own – this is my chief purpose of writing.

List of Figures

Chapter 1

Overview

Orbifolds are rich geometric objects that lie at the intersection of many fields, including algebraic and symplectic geometry, representation theory, and combinatorics. This thesis lives in this intersection: resolving questions about the birational geometry of important families of orbifolds by representation-theoretic and combinatorial means and, conversely, using the framework of orbifold geometry to shed new insight on aspects of symplectic embeddings and polytope combinatorics.

McKay correspondence

We first examine the local model for orbifolds – quotient singularities $\mathbb{C}^n/G$ for a finite subgroup $G \subseteq \mathrm{GL}_n(\mathbb{C})$ – through the McKay correspondence in dimensions two and three. The McKay correspondence is an expansive body of work encompassing many ways in which the representation theory of G influences the geometry of $\mathbb{C}^n/G$ and its resolutions.

When $G \subseteq \mathrm{SL}_n(\mathbb{C})$ the singularity $\mathbb{C}^n/G$ is Gorenstein, which allows many more techniques from birational geometry to be applied. In dimensions two and three there is a crepant resolution

$$G\text{-Hilb}\,\mathbb{C}^n \to \mathbb{C}^n/G$$

with a natural moduli description, called the G-Hilbert scheme. There is an entrancing story due to Gonzalez-Sprinberg–Verdier, Reid, Craw, Logvinenko, and Cautis–Craw–Logvinenko of how the irreducible representations of G control the exceptional fibre $E \subseteq G\text{-Hilb}\,\mathbb{C}^3$ that now goes by the name of 'Reid's recipe'. Here we are concerned with in some sense the complementary question: *how does representation theory help to relate different crepant resolutions of $\mathbb{C}^n/G$?*

In two dimensions this question is addressed by work of Kronheimer [53]. One can rephrase the moduli problem for $G\text{-Hilb}\,\mathbb{C}^n$ in any dimension to add a parameter or 'stability condition' θ living inside a large vector space to produce moduli spaces $\mathcal{M}_\theta$ of quiver representations. When $n = 2, 3$ and for generic θ, $\mathcal{M}_\theta$ is a crepant resolution of $\mathbb{C}^n/G$, and recovers the G-Hilbert scheme for certain choices of θ. For surface singularities, where there is a unique crepant resolution, Kronheimer showed that the space of stability conditions has a wall-and-chamber structure that can be identified with the Weyl chambers inside the Cartan subalgebra of a certain simple Lie algebra. Crossing walls corresponds to contracting and 'regrowing' various strata of the exceptional fibre.

In Chapter 3 following the author's work in [76] we develop an analog of Kronheimer's result for abelian subgroups of $\mathrm{SL}_3(\mathbb{C})$ using the interpretation of the McKay correspondence on the level of derived categories by Bridgeland–King–Reid [14], Craw–Ishii , [26], and Cautis–Craw–Logvinenko [17]. There is again a wall-and-chamber structure on the space of stability conditions. Denote by $\mathfrak{C}_0$ the chamber for the G-Hilbert scheme.

Theorem 1 (Theorem 3.9.1)**.** *There is an algorithm – called the* unlocking procedure *– to explicitly compute a set of inequalities defining $\mathfrak{C}_0$ from the data of Reid's recipe and the combinatorics of the exceptional fibre. Moreover, one can determine which of these inequalities are irredundant and so actually define walls of $\mathfrak{C}_0$. The birational type, unstable locus, and derived equivalence for the wall can be read from the wall equation.*

The full result that will be described in Chapter 3 is rather stronger than this statement: it produces an explicit list of walls and their birational properties for any finite abelian subgroup of $\mathrm{SL}_3(\mathbb{C})$ without running any computation. The largely combinatorial presentation of walls that this yields also serves to provide simpler proofs of results of Craw–Ishii and Nolla de Celis–Sekiya. There should be versions for this result for other Gorenstein singularities, such as for Gorenstein toric singularities via the theory of dimer models [11, 25, 47].

Numerics

The numerics, particularly the Hilbert function

$$\mathrm{hilb}_{(X,L)}(m) := \chi(X, L^{\otimes m})$$

for L a $\mathbb{Q}$-line bundle on X, when X is an orbifold are structured by Reid's *orbifold Riemann–Roch*. This family of results decomposes the Hilbert function for a polarised orbifold (X, L) as

$$\chi(X, L^{\otimes m}) = \text{initial term} + \sum_p q_p(L^{\otimes m})$$

where the initial term is a 'Riemann–Roch-like expression' in terms of $c_1(L)$ and various classes on X, the sum is over singular points p of X, and $q_p(L^{\otimes m})$ is some quasi-polynomial function in m. Following Reid we study the local contributions made by orbifold singularities to the Hilbert function. The first question we address is for del Pezzo surfaces with orbifold singularities; *to what extent can the anticanonical Hilbert function* $\mathrm{hilb}_{(X,-K_X)}(m)$ *capture the singular locus of* X*?* It is more convenient to package the Hilbert function by its generating function or its Hilbert series

$$\mathrm{Hilb}_{(X,-K_X)}(t) := \sum_{m\geq 0} \chi(X, -mK_X)t^m$$

where we have also converted to additive notation for the divisors $-mK_X$ as opposed to the multiplicative notation for the corresponding $\mathbb{Q}$-line bundles $(\omega_X^\vee)^{\otimes m}$. In this notation orbifold Riemann–Roch becomes

$$\mathrm{Hilb}_{(X,L)}(t) = \text{initial term} + \sum_p Q_p(L)$$

where the initial term is the generating function for a polynomial sequence, and $Q_p(L)$ is a particular rational function in t. We refer to the difference

$$\mathrm{Hilb}_{(X,-K_X)}(t) - \text{initial term} = \sum_p Q_p(-K_X)$$

as the 'total orbifold contribution' for $(X, -K_X)$. Up to two elementary number-theoretic conjectures (Conjectures 4.4.8 and 4.5.8), we have the following result.

Theorem 2 (Theorem 4.6.12)**.** *Fix a power series* $H(t) \in \mathbb{N}[[t]]$*. Either there are no orbifold del Pezzo surfaces with Hilbert series equal to* $H(t)$*, or*

- *the collection of orbifold del Pezzo surfaces with Hilbert series* $H(t)$ *naturally breaks up into finitely many families, which are indexed by lattice points in a polytope associated to* $H(t)$*,*
- *for each family, there is a set of rigid singularities* $\mathcal{RB}$ *that appear on every del Pezzo surface in that family,*
- *the singularities on different orbifold del Pezzo surfaces in the same family differ by adding 'cancelling tuples' to the singularities in* $\mathcal{RB}$*.*

Cancelling tuples are sets of singularities arising from a series of crepant toric blowups of a $\mathbb{Q}$-Gorenstein smoothable cyclic quotient singularity (named as 'T-singularities' and classified in two dimensions by Kollár–Shepherd-Barron). The moral of this result is that there is a surprising amount of structure in the collection of orbifold del Pezzo surfaces with a given Hilbert series.

The finiteness of the set of families with properties listed above has consequences for classification.

Theorem 3 (Theorem 4.7.2)**.** *For a given collection of $\mathbb{Q}$-Gorenstein rigid singularities $\mathcal{B}$ there exist constants $m > 0$ and $M > 0$ dependent only on and computable from $\mathcal{B}$ such that, for any orbifold del Pezzo surface X with singularities $\mathcal{B}$,*

$$m \leq K_X^2 \leq M$$

This allows us to show that certain power series cannot occur as the Hilbert series of orbifold del Pezzo surfaces.

When X is a toric orbifold, Hilbert functions for ample divisors correspond to lattice point counting functions for polytopes. We take the perspective that Reid's orbifold Riemann–Roch for polarised toric varieties should be viewed within the polytope algebra framework of McMullen's work on Ehrhart theory [58, 59]. Suppose Q is a simple rational polytope. Denote by Σ_Q its inner normal fan, and by (X_Q, D_Q) the associated polarised toric variety. X_Q is an orbifold as we assumed Q was simple. Suppose Q is 'integer-height': that is, for each edge $e \subseteq Q$ and choice of primitive normal n_e one has $\langle x, n_e\rangle \in \mathbb{Z}$ for all $x \in e$. By orbifold Riemann–Roch, there is then a decomposition

$$|Q \cap \mathbb{Z}^n| = \text{initial term} + \sum_{\sigma \in \Sigma_Q(n)} q_\sigma(Q) \tag{$\star$}$$

where we denote $q_\sigma(Q) = q_\sigma(D_Q)$. McMullen's work studies abstract decompositions

$$|Q \cap \mathbb{Z}^n| = \sum_{F \subseteq Q} \phi(N(F, Q)) \operatorname{vol}(F) \tag{$\dagger$}$$

where $N(F, Q)$ is the normal cone to a face $F \subseteq Q$, for some function ϕ on cones in $\mathbb{R}^n$. There are many choices of function ϕ that work. When Q is rational ϕ depends both on the normal cone to a face and the residue class of the affine span of the face under lattice translation.

This allows geometric tools from deformation and singularity theory to be leveraged to solve problems in rational Ehrhart theory, which studies the growth of lattice point counts in dilates of rational polytopes. Such functions are known to be quasi-polynomials. There is a natural candidate r_Q called the denominator of Q for the period of this quasi-polynomial. We study the problem of *quasi-period collapse*; seeking to characterise when the minimal period π_Q of this quasi-polynomial is smaller than r_Q. The decomposition ($\star$) is especially amenable to studying these problems since all the periodicity is contained in the orbifold contributions $q_\sigma(Q)$.

Mutations are a crucial part of an ongoing program exploring Fano mirror symmetry [20]. For a polytope P and a good choice of combinatorial data (w, F) relative to P, there exists a new polytope $\mu_{(w,F)}P$ called the mutation of P. In two dimensions, if two Fano polygons

P and P' are related by a sequence of mutations then there exists a $\mathbb{Q}$-Gorenstein deformation between the corresponding toric del Pezzo surfaces X_P and $X_{P'}$. Such a deformation preserves the anticanonical Hilbert series, hence the quasi-periods of the dual polytopes $P^\vee$ and $(P')^\vee$ agree. However it does not in general preserve the Gorenstein index, which is the geometric avatar of the denominator. This the fundamental observation that allows us to access new understanding of quasi-period collapse by geometric means.

Using this insight into the numerics of orbifold del Pezzo surfaces in combination with the work of Akhtar–Kaspzryk on polytope mutation in [3], the author and Kasprzyk evidenced the value of this approach in two dimensions. In the following we denote by ℓ_σ the 'local index' [34, Note 3.19] of an orbifold singularity $\sigma \in X_Q$.

Theorem 4 (Theorem 4.9.3)**.** *Let $P \subseteq \mathbb{R}^2$ be a Fano polygon. The quasi-period of the (rational) dual polygon $Q = P^\vee$ has*

$$\pi_Q \leq \operatorname{lcm}\{\ell_\sigma : \sigma \in \mathcal{RB}\}$$

where $\mathcal{RB}$ is the reduced basket of singularities for the orbifold del Pezzo surface X_Q. Furthermore, Q exhibits quasi-period collapse if there exists some singularity $\tau \in X_Q^{\mathrm{sing}}$ of local index not dividing $\operatorname{lcm}\{\ell_\sigma \mid \sigma \in \mathcal{RB}\}$.

This provides a geometric characterisation of quasi-period collapse for many two-dimensional polytopes, which can be used to construct and classify interesting families of examples of polygons with prescribed quasi-period collapse.

Symplectic embeddings

We also form connections between the numerics of $\mathbb{Q}$-line bundles on orbifolds and obstructions to symplectic embeddings between symplectic 4-manifolds.

An active body of work in symplectic geometry is concerned with constructing (ideally sharp) numerical obstructions to the existence of embeddings

$$\iota\colon (X,\omega) \to (X,\omega') \text{ such that } \iota^*\omega' = \omega$$

One of the most successful recent movements in this area was the development of *Embedded Contact Homology* (ECH) by Hutchings, originally to provide a symplectic model of Seiberg–Witten Floer homology. A family of optimisation problems in ECH were used by Hutchings and many coauthors to produce a weakly increasing sequence of real numbers $c_k(X,\omega)$ associated to a symplectic 4-manifold (X,ω) called *ECH capacities* such that

$$(X,\omega) \text{ symplectically embeds in } (X',\omega') \implies c_k(X,\omega) \leq c_k(X',\omega') \text{ for all } k \in \mathbb{Z}_{\geq 0}$$

We construct a family of optimisation problems in algebraic geometry that produce a sequence of invariants

$$c_k^{\mathrm{alg}}(Y, C)$$

where C is a curve in an orbifold Y, called *algebraic capacities* and show that for many important toric manifolds (X_Ω, ω) the ECH capacities $c_k(X, \omega)$ agree with the algebraic capacities of a suitable algebraic compactification of the open manifold $X_\Omega^\circ = X_\Omega \setminus \partial X_\Omega$.

Theorem 5 (Theorem 5.3.16). *Suppose Ω is a rational convex domain. Then*

$$c_k(X_\Omega, \omega) = c_k^{\mathrm{alg}}(Y_\Omega, D_\Omega)$$

where (Y_Ω, D_Ω) is a natural polarised toric surface associated to (X_Ω, ω).

From the numerics of $\mathbb{Q}$-divisors on orbifolds, we deduce several results regarding subtler embedding obstructions from the sub-leading asymptotics of ECH capacities, answering questions of Cristofaro-Gardiner and Hutchings. This connection opens up a great deal of potential for further fruitful application of birational algebro-geometric methods to symplectic embedding problems, and contrariwise. As an example, we formulate the following conjecture for an 'algebraic Weyl law' controlling the asymptotics of solutions to these algebro-geometric optimisation problems, and prove it in certain cases.

Conjecture 1 (Conjecture 5.5.2). *For (Y, A) a polarised $\mathbb{Q}$-factorial surface,*

$$\lim_{k \to \infty} \frac{c_k^{\mathrm{alg}}(Y, A)^2}{k} = 2A^2$$

Declaration

Chapter 2 consists of exposition original in presentation but not in content. The results of Chapter 3 are original to me having appeared in [76]. Chapter 4 is partly expository and partly based on my work in [77] and my joint work with A. M. Kasprzyk in [50]. The results of Chapter 5 are original to me following and slightly extending my paper [75].

Chapter 2

Introduction to orbifolds

2.1 Quotients

Taking quotients in a geometric category is notoriously challenging. For instance, it is well-known that arbitrary quotients of manifolds by group actions fail to be manifolds.

Example 2.1.1. *Let ξ be an irrational real number. Consider the action of the group $G = \mathbb{Z}^2$ on the (topological, smooth,...) manifold $M = \mathbb{R}$ given by*

$$(m, n) \cdot x = x + m + n\xi$$

This action has an uncountable number of orbits, each of which is dense in $\mathbb{R}$, and so the quotient as a topological space has the indiscrete topology and so is not a manifold (for instance, it is not Hausdorff.)

Similarly, taking quotients of varieties or schemes is a difficult business; the central problem here is that satisfying the desired mapping properties of a quotient may suppress the opportunity for interesting geometry. In the following we mostly follow [15].

Definition 2.1.2. *Given a category $\mathcal{C}$ and an object M with a G-action $\alpha\colon G \to \mathrm{Aut}(M)$, a* categorical quotient *of M by G is an object X in $\mathcal{C}$ with a morphism $\pi\colon M \to X$ such that:*

- *π is equivariant: $\pi \circ \alpha(g) = \pi \circ \alpha(h)$ for all $g, h \in G$,*
- *π is initial among G-equivariant morphisms in $\mathcal{C}$ whose domain is M; that is, if X' is an object in $\mathcal{C}$ with a G-action $\beta\colon G \to \mathrm{Aut}(X')$ and $f\colon M \to X'$ is a G-equivariant*

morphism, then there exists a unique morphism $\widetilde{f}$ *in* $\mathcal{C}$ *such that*

$$\begin{tikzcd} M \arrow[d, "\pi"'] \arrow[dr, "f"] & \\ X \arrow[r, "\widetilde{f}"'] & X' \end{tikzcd}$$

commutes.

Example 2.1.3. *If* $M = \mathbb{C}^n$ *with the* $\mathbb{C}^\times$*-action by scaling, regarded in the category of complex schemes, a categorical quotient is just a point since every* $\mathbb{C}^\times$*-equivariant morphism from* $\mathbb{C}^n$ *is constant.*

In this example, we see a natural and rich geometric situation being compressed too far. We restrict to the category of algebraic groups acting on varieties for the next definition.

Definition 2.1.4. *Given a scheme* M *with a* G*-action* $\alpha\colon G \to \mathrm{Aut}(M)$*, a* geometric quotient *of* M *by* G *is a scheme* X *with a morphism* $\pi\colon M \to X$ *such that:*

- *for each* $y \in Y$*,* $\pi^{-1}(y)$ *is a* G*-orbit,*
- X *has the quotient or pushforward topology:* $U \subseteq X$ *is open iff* $\pi^{-1}(U) \subseteq M$ *is open*
- $\mathcal{O}_X = \pi_*(\mathcal{O}_M^G)$.

Example 2.1.5. *Continuing Example 2.1.3,* $X = pt$ *is not a geometric quotient for the* $\mathbb{C}^\times$*-action on* $\mathbb{C}^n$ *since the preimage of the single point is not an orbit. However, restricting the action to* $M = \mathbb{C}^n \setminus \{0\}$ *produces a geometric quotient* $\mathbb{P}^{n-1}$*.*

Restricting the G-action to a suitable G-invariant open subset is a persistent theme in the treatment of quotients. The next theorem due to Rosenlicht indicates why this is the case.

Theorem 2.1.6 ([70, Theorem 2])**.** *For any* G*-action on a variety* M*, there is a* G*-invariant open subset* $M_0 \subseteq M$ *such that the* G*-action on* M_0 *has a geometric quotient.*

The issue is that this open subset is not canonical in any sense, and so different choices for M_0 yield different quotients. There are two responses to this problem:

- enlarge the category of schemes to include objects that play the role of good quotients in both senses described above (stacks),
- embrace the nonuniqueness of geometric quotients by studying them collectively.

We will primarily emphasise the second approach and see that many situations of geometric, algebraic, and combinatorial interest naturally arise by studying the ways these different quotients are related to each other, and to the action of G on M. We will also discuss some relations with the first approach, since the two are typically not mutually exclusive in applications.

2.2 Invariants

Suppose M is an affine variety or scheme over k with coordinate ring $k[M]$, and that G is a subgroup of $\mathrm{Aut}(M)$. Define the *invariant algebra* for G acting on M by

$$k[M]^G := \{f \in k[M] : f(g \cdot x) = f(x) \text{ for all } x \in M\}$$

The problem of understanding the invariants of such a group action has a long history (for example, at least back to [12].) We state one of the main theorems in this direction. Recall that G is reductive if every kG-module is semisimple.

Theorem 2.2.1 ([15, Theorem 1.24])**.** *Let M be an affine variety and let $G \subseteq \mathrm{Aut}(M)$ be a reductive group. Then*

- *$k[M]^G$ is finitely generated,*
- *Define the* affine quotient *of M by G*

$$M/\!/ G := \operatorname{Spec} k[M]^G$$

The morphism $M \to M/\!/ G$ induced by the inclusion $k[M]^G \subseteq k[M]$ is surjective, G-equivariant, and makes $M/\!/ G$ a categorical quotient in the category of affine varieties over k.

One can globalise this construction to arbitrary schemes using a G-invariant affine open cover.

2.3 Orbifold singularities

Suppose V is a vector space over k, and that $G \subseteq \mathrm{GL}(V)$ is a finite subgroup. We consider the affine quotient $V/\!/ G$.

Example 2.3.1. *Suppose $V = k^2$ for a field k of characteristic not 2, and $G = \langle -id \rangle$ is the group of order 2 generated by the negative of the identity matrix. If the dual basis to the standard basis e_1, e_2 on k^2 is denoted x, y, then the affine quotient of V by G is*

$$\operatorname{Spec} k[x,y]^G = \operatorname{Spec} k[x^2, xy, y^2] \cong (uw = v^2) \subseteq \mathbb{A}^3$$

This affine variety has one singularity, at the origin.

Just like in this example $V/\!/\,G$ is usually singular. We call singularities arising on such affine quotients *orbifold singularities* or *quotient singularities*. If the group G is finite cyclic, we call these singularities *cyclic quotient singularities*. Notice that we do not demand that these singularities are isolated.

Example 2.3.2. *Finite cyclic subgroups of $\mathrm{GL}_n(\mathbb{C})$ provide many valuable examples of orbifold singularities. We use the notation*

$$\frac{1}{r}(a_1, \dots, a_n)$$

to denote both the subgroup of $\mathrm{GL}_n(\mathbb{C})$ generated by the matrix

$$\begin{pmatrix} \varepsilon^{a_1} & & & \\ & \varepsilon^{a_2} & & \\ & & \ddots & \\ & & & \varepsilon^{a_n} \end{pmatrix}$$

where ε is a primitive rth root of unity, and the corresponding cyclic quotient singularity.

Definition 2.3.3. *Let X be a variety over k. We say that X is an* orbifold *or* has orbifold singularities *if its (potentially nonisolated) singularities are all locally isomorphic to orbifold singularities.*

2.4 GIT

Geometric Invariant Theory (GIT) is a theory that systematises Theorem 2.1.6 within a general framework; that is, it gives a general method of selecting invariant open subsets with geometric quotients. We draw from a combination of [15, 39, 66] for this subsection.

Suppose $M = \operatorname{Spec}(R)$ is a smooth affine variety with the action of a reductive algebraic group G. For a choice of character $\chi \in G^\vee$ define

$$M/\!/_\chi G := \operatorname{Proj} \bigoplus_{d \geq 0} R_{\chi^{\otimes d}}$$

where $R_{\chi^{\otimes d}} := \{f \in R : f(g \cdot x) = \chi(g)^d f(x)\}$. Elements of $R_{\chi^{\otimes d}}$ are called $\chi^{\otimes d}$-semi-invariant functions. This Proj is equal to

$$M/\!/_{\chi}G = (M \setminus \Delta_\chi)/\!/\, G =: M^{\mathrm{ss}}_{\chi}/\!/\, G$$

for some closed, G-invariant $\Delta_\chi \subseteq M$.

Let us analyse what is happening. By taking the Proj we are removing the locus Δ_χ of all points $p \in M$ such that every semi-invariant function f has $f(p) = 0$. Such points are called χ-*unstable*. All other points are called χ-*semistable*. This produces a G-equivariant morphism

$$\pi \colon M \setminus \Delta_\chi \to M/\!/_{\chi}G$$

by evaluating at each point. At this point we have obtained what is referred to in the literature as a 'good quotient'. However, it is not clear that distinct G-orbits in $M \setminus \Delta_\chi$ go to different points in $M/\!/_{\chi}G$. For this we require the notion of stable points.

Definition 2.4.1. *A point $p \in M$ is* χ-stable *if it is χ-semistable, if locally near p all G-orbits are closed (for instance, inside a G-invariant distinguished open set), and if* $\mathrm{Stab}_G(p)$ *is finite.*

The local condition on orbits implies that distinct orbits near x are separated by semi-invariant functions as desired. Denote by M^{s}_{χ} the (open) set of χ-stable points in M.

Theorem 2.4.2. *Consider the morphism $\pi \colon M^{ss} \to M/\!/_{\chi}G$. There exists an open set $X \subseteq M/\!/_{\chi}G$ such that $\pi^{-1}(X) = M^s_\chi$ and π restricts to a geometric quotient $M^s_\chi \to X$.*

Example 2.4.3. *Suppose $G = k^\times$ and let it act on $M = k^{n+1}$ by scaling. Suppose $a \in \mathbb{Z}_{>0}$ and $\chi(t) = t^a$. Then Δ_χ is given by the vanishing of the ideal generated by the χ-eigenvectors in the coordinate ring $\mathbb{C}[x_0, \ldots, x_n]$ of M. These are exactly the homogeneous polynomials of degree a and so $\Delta_\chi = \{0\}$ giving the familiar description of projective space (in its nth Veronese embedding). Note that for $a < 0$, the χ-eigenspace is trivial and so the quotient is empty. We suggest the reader contemplates what happens when $a = 0$.*

As this example illustrates, this construction is effective at capturing quotients that interest us but can backfire by producing empty quotients when some choices of χ yield no stable points.

Example 2.4.4. *Consider $G = \mathrm{GL}_r(k)$ acting on $M = \mathrm{Hom}(k^r, k^n)$ and let k be infinite. For any character χ, notice that χ-stable points of M must have rank r since otherwise it is* *easy to create infinite stabilisers. Moreover these low rank maps are also χ-unstable. Pick the character $\chi(A) = \det(A)^a$ for positive a. We can write such a map as*

$$X = \begin{pmatrix} a_{1,1} & a_{1,2} & \dots & a_{1,n} \\ \vdots & \vdots & \ddots & \vdots \\ a_{r-1,1} & a_{r-1,2} & \dots & a_{r-1,n} \\ 0 & 0 & \dots & 0 \end{pmatrix}$$

Note that such matrices are fixed by

$$A_t = \begin{pmatrix} 1 & 0 & \dots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \dots & 1 & 0 \\ 0 & 0 & \dots & 0 & t \end{pmatrix}$$

and so for semi-invariant f one has

$$f(X) = f(A_t \cdot X) = \chi(A_t) f(X) = t^a f(X)$$

hence $f(X) = 0$. For full rank $X \colon k^r \to k^n$ there is a nonzero $r \times r$ minor and the ath power of this defines a χ-semi-invariant function that is nonzero at X, hence evidencing that X is χ-semistable. Indeed, in this situation there are no non-stable semistable ('strictly semistable') points.

It follows that $M^{ss}_\chi/G = M^s_\chi/G$ parameterises r-dimensional subspaces of k^n up to coordinate change; that is, it is isomorphic to the Grassmannian $\mathrm{Gr}(r, n)$.

Note that the existence of semistable points requires that the kernel K of the action on X is sent to 1 by χ. For an argument over the real or complex numbers, one can see this as otherwise, if $|\chi(g)| < 1$ for some $g \in K$, $f(x) = f(g^m \cdot x) = \chi(g)^m f(x) \to 0$ as $m \to \infty$.

Observe that the GIT quotient also comes with an ample line bundle from the Proj construction. We can view this line bundle as coming from $\mathcal{O}_M$ on the affine variety M that we have equipped with a G-action via the character χ. It descends to an ample line bundle L on the GIT quotient, for which the graded ring is exactly

$$R(M/\!/_\chi G, L) = \bigoplus_{d \geq 0} R_{\chi^{\otimes d}}$$

As was clear from the examples above, the resulting quotient depended very strongly on the choice of character, but also that many different characters can yield the same quotient. We will explore this dependence further in §3 and sooner in §2.5. It was also the case that many characters existed for which the notions of stable and semistable agreed. We call such characters 'generic' and will indeed see that they are truly generic in many contexts. There are also many variations of the above (c.f. [66]) for GIT quotients of (quasi)projective varieties. In these contexts many authors prefer to emphasise the choice of equivariant line bundle on M instead of a character to parameterise their GIT quotients.

2.5 Toric orbifolds

Toric varieties will be a recurrent source of examples and results in the following. Philosophically, toric varieties are partial compactifications of algebraic tori by torus-invariant boundary strata. They are described combinatorially by cones, fans, and polytopes. This and much more is detailed in [23]. We outline some of the key constructions below. We will work over the complex numbers to prepare to make contact with the symplectic methods in later chapters, though much of the following carries over the arbitrary fields.

Affine toric varieties arise from cones

Let $N \cong \mathbb{Z}^n$ be a lattice and let $N_{\mathbb{R}} := N \otimes_{\mathbb{Z}} \mathbb{R}$ be the associated real vector space. A *cone* σ in $N_{\mathbb{R}}$ is a subset of the form

$$\mathrm{Cone}(S) := \{\sum_{v \in S} \lambda_v v : \lambda_v \geq 0, \text{all but finitely many } \lambda_v \text{ are zero}\}$$

Let $M = N^\vee := \mathrm{Hom}_{\mathbb{Z}}(N, \mathbb{Z})$ be the dual lattice to N, and $M_{\mathbb{R}} = M \otimes_{\mathbb{Z}} \mathbb{R}$ the dual vector space to $N_{\mathbb{R}}$. Define the *dual cone* to a cone $\sigma \subset N_{\mathbb{R}}$ to be

$$\sigma^\vee := \{v \in M_{\mathbb{R}} : \langle u, v\rangle \geq 0 \text{ for all } u \in \sigma\}$$

where $\langle \cdot, \cdot \rangle$ is the dual pairing $N_{\mathbb{R}} \times M_{\mathbb{R}} \to \mathbb{R}$. Suppose now that σ is a *rational polyhedral* cone: that there is a finite set of lattice points $S \subset N$ such that $\sigma = \mathrm{Cone}(S)$. Such a cone σ gives an affine toric variety U_σ as follows.

- **Input:** σ, a rational polyhedral cone
- Dualise to $\sigma^\vee$
- Take lattice points $\sigma^\vee \cap M$ to obtain a semigroup
- Take the semigroup algebra $\mathbb{C}[\sigma^\vee \cap M]$; this is a finitely generated $\mathbb{C}$-algebra
- **Output:** $U_\sigma := \mathrm{Spec}\, \mathbb{C}[\sigma^\vee \cap M]$.

Notice that the dense open torus arises from $\mathbb{C}[\sigma^\vee \cap M] \subset \mathbb{C}[M] \cong \mathbb{C}[\mathbb{Z}^n]$, which is the ring of Laurent polynomials, or the ring of functions for the torus $(\mathbb{C}^\times)^n$. The cone σ (or rather $\sigma^\vee$) is describing which functions on the torus extend to global functions on U_σ, which is equivalent to describing the variety. One can describe the torus inside U_σ intrinsically as

$$T_N := N \otimes_{\mathbb{Z}} \mathbb{C}^\times$$

In this presentation, a vector $m \in M$ gives a function $\chi^m : T_N \to \mathbb{C}$ via

$$\chi^m(n \otimes t) = t^{\langle m,n \rangle}$$

Example 2.5.1. *Take $N = \mathbb{Z}^2$ and let $\sigma = \mathrm{Cone}(e_1, e_2)$. The dual cone is $\sigma^\vee = \mathrm{Cone}(e^1, e^2)$ giving*

$$\sigma^\vee \cap M = \mathbb{Z}^2_{\geq 0} \text{ and } \mathbb{C}[\sigma^\vee \cap M] \cong \mathbb{C}[x, y]$$

Hence $U_\sigma \cong \mathbb{C}^2$. In this case, $\sigma^\vee$ prescribes that the only Laurent polynomials extending to all of U_σ are the polynomials.

Toric varieties arise from fans

To construct non-affine (in particular, compact) toric varieties we glue together affine toric varieties in an torus-equivariant way. The combinatorial avatar of this process is collecting cones together in a *fan*. To start with, a *face* of a cone σ is a subset of σ of the form $\sigma \cap (\langle m, \cdot \rangle = 0)$ for some $m \in \sigma^\vee$. The cones forming the boundary of σ are examples of faces, as is the vertex of the cone (the origin). A *fan* in $N_\mathbb{R}$ is a collection of cones $\Sigma = \{\sigma\}$ such that

- if $\tau \subset \sigma$ is a face, then $\tau \in \Sigma$
- for any two cones $\sigma_1, \sigma_2 \in \Sigma$, $\sigma_1 \cap \sigma_2$ is a face of each

A fan Σ produces a toric variety Y_Σ via gluing two affine pieces $U_{\sigma_1}, U_{\sigma_2}$ according to the (potentially zero-dimensional) face they have in common.

Example 2.5.2. *Take $N = \mathbb{Z}^2$ and Σ to be the fan containing the cones $\sigma_1 = \mathrm{Cone}(e_1, e_2), \sigma_2 = \mathrm{Cone}(e_1, -e_1 - e_2), \sigma_3 = \mathrm{Cone}(e_2, -e_1 - e_2)$ and their faces. The two-dimensional cones give three copies of $\mathbb{C}^2$ and the gluing prescribed by the faces makes this into $\mathbb{P}^2$. For example, σ_1 and σ_3 share the face $\mathrm{Cone}(e_2)$ that corresponds to the toric variety $\mathbb{C}^\times \times \mathbb{C}$. Gluing $\mathbb{C}^2$ to $\mathbb{C}^2$ along $\mathbb{C}^\times \times \mathbb{C}$ is familiar from the gluing construction of projective space.*

Compact toric varieties arise from polytopes

Suppose $P \subset N_{\mathbb{R}}$ is a lattice polytope. One can produce a fan Σ_P from P via

$$\Sigma_P := \{\mathrm{Cone}(S) : S \subseteq \mathrm{Vert}(P) \text{ such that all } u \in S \text{ share a face}\}$$

This is called the *face fan* of P and defines a toric variety $Y_P := Y_{\Sigma_P}$ that turns out to be compact.

A polytope $Q \subset M_{\mathbb{R}}$ also defines a toric variety V_Q. Let $L_Q = \#Q \cap M$ and define a map $\phi_Q : T_N \to \mathbb{P}^{L_Q-1}$ by $x \mapsto (\chi^m(x))_{m \in Q \cap M}$. The toric variety V_Q is defined to be the closure of the image of ϕ_Q in $\mathbb{P}^{L_Q-1}$. If we define the dual polytope

$$P^\vee := \{v \in M_{\mathbb{R}} : \langle u, v\rangle \geq -1\}$$

then the toric variety Y_P is also described abstractly as the variety $V_{kP^\vee}$ for large enough k, from which it is readily apparent that it is compact.

Example 2.5.3. *A polytope for $\mathbb{P}^2$ is the triangle with vertices $e_1, e_2, -e_1 - e_2$. The dual polytope is the triangle with vertices $2e^1 - e_2, -e_1 + 2e_2, -e^1 - e^2$. This has 10 lattice points and describes the third Veronese (or anticanonical) embedding of $\mathbb{P}^2$ in $\mathbb{P}^9$.*

In the V_Q presentation, one can interpret Q as the moment polytope for the compact torus action on V_Q by composing the map ϕ_Q with the moment map on $\mathbb{P}^{L_Q-1}$.

Polytopes arise from divisors

A (Weil) divisor on a normal variety is a formal $\mathbb{Z}$-linear combination of codimension one subvarieties. Divisors on a variety X up to an equivalence relation called rational equivalence form a group called the *class group* of X. For a toric variety X containing dense open torus T, the class group is generated by the components of the toric boundary $X \setminus T$. If $X = Y_\Sigma$ is given by a fan, these boundary components correspond to the rays of Σ. The set of rays is commonly denoted $\Sigma(1)$. Thus, every divisor on Y_Σ is rationally equivalent to one of the form

$$\sum_{\rho \in \Sigma(1)} a_\rho D_\rho$$

One can associate a polytope $P(D)$ to a divisor of this form as follows. Let u_ρ be the primitive lattice point lying on the ray ρ. Then set

$$P(D) := \{v \in M_{\mathbb{R}} : \langle u_\rho, v\rangle \geq -a_\rho \text{ for all } \rho \in \Sigma(1)\}$$

The hyperplanes defining the facets of $P(D)$ are given by $\langle u_\rho, \cdot\rangle = -a_\rho$ and so this construction of $P(D)$ taking in the data $(u_\rho, a_\rho)_{\rho\in\Sigma(1)}$ is often referred to as a 'facet presentation' for $P(D)$. Denote by $\mathcal{O}(D)$ the line bundle associated to a (Cartier) divisor D.

Lemma 2.5.4 ([23], Prop. 4.3.3)**.** *Let $D = \sum_\rho a_\rho D_\rho$. A basis of $H^0(\mathcal{O}(D))$ is in bijection with lattice points of $P(D)$. That is,*

$$\#P(D) \cap M = L_{P(D)} = h^0(\mathcal{O}(D))$$

Notice that there can be multiple facet presentations corresponding to the same divisor if some of the hyperplanes give redundant inequalities.

Divisors arise from support functions

Fix a fan Σ. The *support* $|\Sigma|$ of Σ is the union of the cones it contains. A *support function* on Σ is a function $\varphi : |\Sigma| \to \mathbb{R}$ such that $\varphi|_\sigma$ is linear for each $\sigma \in \Sigma$. An *integral* support function is a support function such that $\varphi(|\Sigma| \cap N) \subset \mathbb{Z}$. An integral support function φ produces a (Cartier) divisor D via

$$D = -\sum_{\rho\in\Sigma(1)} \varphi(u_\rho) D_\rho$$

and this process is actually reversible (so long as D is Cartier).

Singularities on toric varieties

There are several combinatorial criteria for controlling how wild the singularities on a toric variety can be. Most relevant to the present work is the following result.

Theorem 2.5.5 ([23])**.** *A toric variety X_Σ of dimension n is an orbifold iff every top-dimensional cone in Σ is generated by $n+1$ rays. Such fans are called* simplicial *as their top-dimensional pieces are cones over simplices.*

Let σ be an n-dimensional cone with $n+1$ rays. The key insight behind this result is that the lattice Λ generated by the minimal generators of rays in σ has finite index in N and the quotient $G = N/\Lambda$ is a finite abelian group that acts on the toric variety $\mathbb{C}^n$ obtained by replacing N by Λ as the lattice we use. The quotient $\mathbb{C}^n/G$ is exactly the toric variety U_σ (using N as lattice) and so we obtain the local description of X_Σ as an orbifold.

Toric varieties as quotients

There is an alternative construction of toric varieties as GIT quotients due to Cox [21]. Let Σ be a full-dimensional fan in $N_{\mathbb{R}}$ with $\Sigma(1)$ the set of rays in Σ. We define the *Cox ring* of $X = X_\Sigma$ to be

$$S(X) = \mathbb{C}[x_\rho : \rho \in \Sigma(1)]$$

A monomial $x_{\rho_1}^{a_1} \dots x_{\rho_s}^{a_s}$ here defines a Weil divisor $D = \sum a_i D_{\rho_i}$, leading us to often denote this monomial by x^D. We grade $S(X)$ by

$$\deg(D) = [D] \in A^1(X)$$

where $A^1(X)$ is the Chow group of divisors on X. Observe that by construction $x^D x^E = x^{D+E}$ and so this is indeed a grading. Denote the graded pieces by S_α for $\alpha \in A^1(X)$.

We construct an ideal of $S(X)$ called the *irrelevant ideal*: for a cone σ, create a monomial

$$x^\sigma = \prod_{\rho \notin \sigma(1)} x_\rho$$

Denote by B_Σ the ideal generated by these monomials, equivalently the ideal generated by x^σ for all maximal cones of σ. The variety Z_Σ of this ideal in $\operatorname{Spec} S(X) = \mathbb{C}^{\Sigma(1)}$ has codimension at least 2. We will see that X_Σ can be obtained as a categorical quotient of $\mathbb{C}^{\Sigma(1)} \setminus Z_\Sigma$, as a geometric quotient exactly when Σ is simplicial, and that this can all be manoeuvred into the framework of GIT.

The Chow group $A^1(X)$ fits into an exact sequence

$$M \to \mathbb{Z}^{\Sigma(1)} \to A^1(X) \to 0$$

Applying the Gale dual functor $\operatorname{Hom}_{\mathbb{Z}}(-, \mathbb{C}^\times)$ which is exact as $\mathbb{C}^\times$ is divisible,

$$T \leftarrow (\mathbb{C}^\times)^{\Sigma(1)} \leftarrow G_\Sigma \leftarrow 0$$

where T is the torus inside X_Σ and where G_Σ is by definition the Gale dual of $A^1(X)$. Note that $(\mathbb{C}^\times)^{\Sigma(1)}$ and hence G_Σ act naturally on $\operatorname{Spec} S(X) = \mathbb{C}^{\Sigma(1)}$ and also on $\mathbb{C}^{\Sigma(1)} \setminus Z_\Sigma$.

Theorem 2.5.6 ([21, Theorem 2.1])**.** *X_Σ is naturally isomorphic to the categorical quotient $(\mathbb{C}^{\Sigma(1)} \setminus Z_\Sigma)/G_\Sigma$. If Σ is simplicial, this is a geometric quotient.*

Example 2.5.7. *Consider the A_n singularity $\frac{1}{n+1}(1, n)$. This is the affine toric variety U corresponding to the cone σ with ray generators $(0, 1), (n + 1, -n)$. We can compute $A^1(U)$ from*

$$\mathbb{Z}^2 \to \mathbb{Z}^2 \to A^1(U) \to 0$$

where the map $\mathbb{Z}^2 \to \mathbb{Z}^2$ *is given by*

$$(a, b) \mapsto (a(n+1) - bn, b)$$

whose cokernel is isomorphic to $\mathbb{Z}/(n+1)$. *Hence* $G_\Sigma \cong \mathbb{Z}/(n+1)$, *which acts on* $\mathbb{C}^2$ *by the familiar action inside* $\mathrm{SL}_2(\mathbb{C})$.

Toric varieties as GIT quotients

The quotient construction of a toric variety

$$X = X_\Sigma = (\mathbb{C}^{\Sigma(1)} \setminus Z_\Sigma)/G_\Sigma$$

looks strikingly similar to a GIT quotient. Observe the G_Σ-action just comes from the $A^1(X)$-grading on $S(X)$. If we pick a character $\chi \in G_\Sigma^\vee = A^1(X)$ corresponding to an ample divisor on X we can compute directly that $\mathrm{rad}(B_\chi) = \mathrm{rad}(B_\Sigma)$ and so

$$\mathbb{C}^{\Sigma(1)} /\!/_\chi G_\Sigma = X_\Sigma$$

Indeed, as χ is an ample divisor on X, it has no basepoints on X nor does its lift to $\mathbb{C}^{\Sigma(1)}$, which is the linearisation provided by χ for $\mathcal{O}_{\mathbb{C}^{\Sigma(1)}}$.

Hence, every toric variety can be realised as a GIT quotient of affine space by an abelian algebraic group. Reversing this process, one can provide a GIT construction of toric varieties from fans that allows variation of GIT parameter to produce interesting geometry.

Fix a finite collection of primitive vectors $\mathcal{B} \subseteq N$. We regard these as ray generators for a fan Σ. Note that the Chow group of any X_Σ with precisely these rays is fixed, and so one obtains a 'combinatorial class group' $A^1(\mathcal{B})$ defined by the sequence

$$M \to \mathbb{Z}^{\mathcal{B}} \xrightarrow{\pi} A^1(\mathcal{B}) \to 0 \qquad (*)$$

This produces an algebraic group $G_{\mathcal{B}} = \mathrm{Hom}(A^1(\mathcal{B}), \mathbb{C}^\times)$ that acts on $\mathbb{A}^{\Sigma(1)}$ through the grading $\deg(x_\rho) = \pi(\rho)$. The character group of $G_{\mathcal{B}}$ is $A^1(\mathcal{B})$ and so a choice of class $\chi \in A^1(\mathcal{B})$ yields a GIT quotient

$$X = \mathbb{C}^{\mathcal{B}} /\!/_\chi G$$

This is toric since G is abelian and, when G is a torus and χ is generic, will be a smooth toric variety projective over the affine variety $\mathrm{Spec}\,(S_\chi)_0$, which is independent of χ. The sequence $(*)$ gives that $A^1(X) = A^1(\mathcal{B})$.

We rephrase Cox's theorem in this language.

Theorem 2.5.8. *Every toric variety with fan* Σ *having* $\Sigma(1) \subseteq \mathcal{B}$ *occurs as a GIT quotient* $\mathbb{A}^{\mathcal{B}} /\!/_\chi G$ *for some character* χ.

Note that that 'degenerate' characters inside $A^1(\mathcal{B})$ that are not ample for any X_Σ with $\Sigma(1) = \mathcal{B}$ can produce fans with fewer rays than $\mathcal{B}$ offers.

There are many choices of ample line bundle for a given toric variety X. The space of these is the interior of the nef cone $\mathrm{Nef}(X) \subseteq A^1(X)_\mathbb{R} = A^1(\mathcal{B})_\mathbb{R}$. For a different X, there will be a different nef cone in $A^1(\mathcal{B})_\mathbb{R}$. In this way, part of $A^1(\mathcal{B})_\mathbb{R}$ can be decomposed into cones that form a fan. Notice that ample divisor classes must be effective and so this fan won't be supported beyond the positive octant. In the toric literature this is called the 'secondary fan' [22, §7]. Within each chamber, the GIT quotient is constant as an abstract variety though its embedding changes with χ.

Example 2.5.9. *Let* $v_0 = (0, 0, -2), v_1 = (1, 1, 1), v_2 = (1, -1, 1), v_3 = (-1, -1, 1), v_4 = (-1, 1, 1) \in \mathbb{R}^3$*. These vectors satisfy* $v_1 + v_3 = v_2 + v_4$ *and* $v_0 = -(v_1 + v_3) = -(v_2 + v_4)$*. There are three complete toric varieties* X_Σ *with* $\Sigma(1) \subseteq \{v_0, v_1, v_2, v_3, v_4\} = \mathcal{B}$ *that, inside the slice* $(z = 1)$*, look like:*

Figure 2.1: Fans with $\Sigma(1) \subseteq \mathcal{B}$

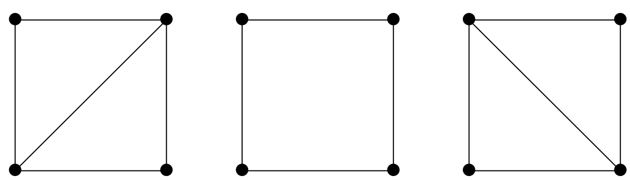

We compute the nef cones of each of these toric varieties $X_1 = X_{\Sigma_1}, X = X_\Sigma, X_2 = X_{\Sigma_2}$*. Recall from [23] that a divisor is nef iff its support function* ψ *is convex iff for every primitive collection* $\{w_1, \ldots, w_r\}$ *one has*

$$\psi(w_1 + \cdots + w_r) \geq \psi(w_1) + \cdots + \psi(w_r)$$

From [7] 'primitive collection' is a set of ray generators of Σ *that do not span a cone, but such that every proper subset does. The primitive collections for* Σ_1 *are* $\{v_0, v_1, v_3\}, \{v_2, v_4\}$ *and so, letting* $-a_i = \psi(v_i)$*, the inequalities defining* $\mathrm{Nef}(X_1)$ *are*

$$0 = \psi(v_0 + v_1 + v_3) \geq \psi(v_0) + \psi(v_1) + \psi(v_3) = -a_0 - a_1 - a_3 \Leftrightarrow a_0 + a_1 + a_3 \geq 0$$
$$-a_1 - a_3 = \psi(v_1 + v_3) = \psi(v_2 + v_4) \geq \psi(v_2) + \psi(v_4) = -a_2 - a_4 \Leftrightarrow a_2 + a_4 \geq a_1 + a_3$$

The exact sequence computing the class group of X_1 *is*

$$\mathbb{Z}^3 \to \mathbb{Z}^5 \to \mathbb{Z}^2$$

where the second map is $(a_0, \dots, a_4) \mapsto (a_0 + a_1 + a_3, a_0 + a_2 + a_4) =: (s, t)$*. In* $A^1(X_1)_{\mathbb{R}}$*, these inequalities are expressed simply as* $t \geq s \geq 0$*. The ample cone is* $t > s > 0$*.*

Of course $A^1(X_2) \cong A^1(X_1) \cong A^1(\mathcal{B})$*. The primitive collections for* X_2 *are* $\{v_0, v_2, v_4\}, \{v_1, v_3\}$*, which switch the roles of* a_1, a_3 *and* a_2, a_4 *(or* s *and* t*) in the inequalities. Hence, the nef cone of* X_2 *is given by* $s \geq t \geq 0$ *and the ample cone by* $s > t > 0$*.*

Consider the singular toric variety X*. The primitive collections are* $\{v_1, v_3\}$ *and* $\{v_2, v_4\}$ *giving inequalities*

$$a_1 + a_3 \geq 0 \text{ and } a_2 + a_4 \geq 0$$

that are actually equivalent because of the relation $v_1 + v_3 = v_2 + v_4$ *and the fact that they all share a cone. This produces the locus* $s = t \geq 0$*.*

Figure 2.2: Secondary fan for $\mathcal{B}$

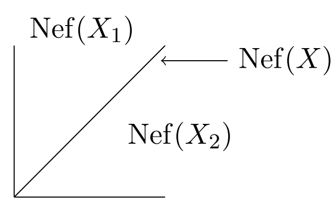

There are clear maps $X_1 \to X \leftarrow X_2$*, which are resolutions of* X*. These resolutions have a single rational curve as exceptional fibre corresponding to the additional two dimensional cone from the diagonal of the square at height* 1*. One obtains a birational map* $X_1 \to X_2$*, which is exactly the flop in this curve.* X_1 *and* X_2 *are the minimal resolutions of* X *- isomorphic as varieties, but not as* X*-varieties. This is the exemplar of wall-crossing geometry: crossing a 'wall' – a cone of characters yielding a 'degenerate' GIT quotient – produced a correspondence* $X_1 \to X \leftarrow X_2$ *realising a classical birational surgery.*

2.6 Quiver varieties

A quiver is a directed graph. More formally, a quiver is a pair $Q = (Q_0, Q_1)$ of sets with two maps $h, t \colon Q_1 \to Q_0$ taking an element of Q_1, an arrow in the quiver, to its head vertex and its tail vertex respectively.

Example 2.6.1. *Consider the quiver* $Q = \bullet \rightrightarrows \bullet$*. This can be encoded formally as* $Q = (\{1, 2\}, \{a, b\})$ *with* $h(a) = h(b) = 2$ *and* $t(a) = t(b) = 1$*.*

Definition 2.6.2. *A* representation *of* Q *(over a fieldk) is a* Q_0*-graded kvector space* $V = \bigoplus_{i \in Q_0} V_i$ *with linear maps* $f_\alpha \colon V_i \to V_j$ *for each arrow* $i \xrightarrow{\alpha} j$ *in* Q_1*. The* dimension vector *of V is $\underline{\dim}V := (\dim V_i)_{i\in Q_0}$. We often write a representation as either just V or as a pair $(V,(f_\alpha)_{\alpha\in Q_1})$. A graded linear map $T = (T_i)_{i\in Q_0}\colon V \to W$ between two representations $V = (V,(f_\alpha))$ and $(W,(g_\alpha))$ of Q is a homomorphism of representations if*

$$T_h(\alpha) \circ f_\alpha = g_\alpha \circ T_{t(\alpha)}$$

for all $\alpha \in Q_1$.

We obtain an abelian category $\mathrm{rep}_k(Q)$ of representations of Q over k.

Fix a dimension vector $\underline{\mathrm{d}}$ for Q. A *stability condition* for Q is an element θ of

$$\Theta_{\underline{\mathrm{d}}} := \{\vartheta \in \mathrm{Hom}_{\mathbb{Z}}(\mathbb{Z}^{Q_0^G}, \mathbb{R}) : \vartheta(\underline{\mathrm{d}}) = 0\}$$

We say that a representation V of dimension vector $\underline{\mathrm{d}}$ is *θ-stable* if $\theta(U) := \theta(\underline{\dim}U) > 0$ for all $0 \neq U \subsetneq V$. This is a combinatorialisation of usual GIT stability for quivers, as follows.

Define for $V = \bigoplus_{i\in Q_0} V_i$

$$\mathcal{R}(Q,V) = \bigoplus_{\alpha\in Q_1} \mathrm{Hom}(V_{t(\alpha)}, V_{h(\alpha)})$$

The graded automorphism group acting on $\mathcal{R}(Q,V)$ is

$$\mathrm{GL}(V) := \prod_{i\in Q_0} \mathrm{GL}(V_i)$$

We can view $\theta = (\theta_i)_{i\in Q_0} \in \Theta_{\underline{\dim}(V)}$ as a character for $\mathrm{GL}(V)$ via

$$\chi_\theta(A_i)_{i\in Q_0} = \prod_{i\in Q_0} \det(A_i)^{\theta_i}$$

and hence consider the GIT quotient

$$\mathcal{M}_\theta(Q,V) := \mathcal{R}(Q,V)/\!/_\theta\, \mathrm{GL}(V)$$

Observe that the condition $\theta(\underline{\dim}(V)) = 0$ corresponds to the fact that the scalar subgroup $k^\times \subseteq \mathrm{GL}(V)$ acts trivially on $\mathcal{R}(Q,V)$.

It follows from an argument of King [51, Prop. 3.1] that GIT stability for χ_θ corresponds to the combinatorial notion of θ-stability introduced above. He uses this to show the following.

Theorem 2.6.3 ([51, Prop. 4.3 & Prop. 5.3])**.** *There is a fine moduli space $\mathcal{M}_\theta(Q,\underline{\mathrm{d}})$ parameterising θ-stable representations of Q with dimension vector $\underline{\mathrm{d}}$. Moreover, this moduli space is projective.*

It is often the case that the representations appearing naturally when taking quotients come with relations between the maps f_α. The natural place to express relations in terms of arrows is the *path algebra*

$$kQ := k\langle Q_1\rangle / I_{\text{path}}$$

where I_{path} is generated by

$$\alpha\beta = 0 \text{ if } t(\alpha) \neq h(\beta)$$

In other words, elements of kQ are linear combinations of oriented paths in Q. This clearly furnishes an equivalence of categories

$$\operatorname{rep}(Q) \simeq kQ\text{-mod}$$

Relations between the maps f_α composing a representation of Q can be expressed by a two-sided ideal $I \subseteq kQ$. Indeed, the category $\operatorname{rep}(Q, I)$ of representations of Q such that the f_α obey all relations in I is equivalent to the category kQ/I-mod. There is an analog to Theorem 2.6.3 though we lose the properness of the moduli space.

Theorem 2.6.4. *There is a fine moduli space $\mathcal{M}_\theta(Q, I, \underline{\mathrm{d}})$ parameterising θ-stable representations of Q satisfying relations I with dimension vector $\underline{\mathrm{d}}$. Moreover, this moduli space is quasi-projective.*

This arises from the same GIT construction using the $\mathrm{GL}(V)$-action on the representation space

$$\mathcal{R}(Q, I) := \{(f_\alpha)_{\alpha\in Q_1} \in \bigoplus_{\alpha\in Q_1} \operatorname{Hom}(V_{t(\alpha)}, V_{h(\alpha)}) : f_\gamma = 0 \text{ for all } \gamma \in I\}$$

where for $\gamma \in kQ$ we define f_γ by defining $f_{\alpha_1\ldots\alpha_r} = f_{\alpha_r} \circ \cdots \circ f_{\alpha_1}$ and extending linearly.

An important source of ideals in path algebras is potentials. Define kQ_{cyc} to be the vector subspace of kQ generated by oriented cycles in Q. We call elements of kQ_{cyc} *potentials*. A potential $W \in kQ_{\text{cyc}}$ produces an ideal of kQ via the following procedure. Define the derivative of a cycle $\gamma = \alpha_1 \ldots \alpha_r$ with respect to an arrow β by

$$\partial_\beta\gamma = \begin{cases} \alpha_1 \ldots \alpha_{i-1}\alpha_i \ldots \alpha_r & \beta = \alpha_i, \\ 0 & \beta \neq \alpha_i \text{ for all } i \end{cases}$$

One can extend the derivative linearly to define $\partial_\beta W$ for any potential W. With this definition, the ideal $J(W)$ corresponding to W, called the *Jacobian ideal* of W, is the two-sided ideal of kQ generated by $\{\partial_\beta W : \beta \in Q_1\}$. The quotient $\mathcal{P}(Q, W) := kQ/J(W)$ is called the *Jacobian algebra* of (Q, W).

Nakajima quiver varieties

There is a symplectic version of these moduli spaces when the quiver Q and relations I take a certain form, and when $k = \mathbb{C}$. Let $Q = (Q_0, Q_1)$ be a quiver. Define the 'double quiver' $Q^{\mathrm{dbl}} = (Q_0, Q_1 \amalg Q_1^*)$ to be the quiver with vertices Q_0 and with two arrows α, α^* for every arrow α in Q such that α links the same two vertices as it does in Q, and $h(\alpha^*) = t(\alpha), t(\alpha^*) = h(\alpha)$. Let $U = \bigoplus_{i \in Q_0} U_i$ and define the 'U-extended double quiver' $\widetilde{Q}^{\mathrm{dbl}}$ to be quiver with vertices $\widetilde{Q}_0 = Q_0 \amalg \{\infty\}$ and arrows $Q_1 = Q_1 \amalg Q_1^* \amalg Q_1^\infty$ where

$$Q_1^\infty = \{a_{ij}, b_{ij}\}_{i \in Q_0, j \in \{1, \ldots, \dim U_i\}}$$

with $h(a_{ij}) = i$ and $t(a_{ij}) = \infty$, and $h(b_{ij}) = \infty$ and $t(b_{ij}) = i$.

Example 2.6.5. *Let Q be the A_n quiver $\bullet \to \cdots \to \bullet$ with n vertices. Q^{dbl} is*

$$\begin{tikzcd}
\bullet \arrow[r, shift left] & \bullet \arrow[l, shift left] \arrow[r, shift left] & \cdots \arrow[l, shift left] \arrow[r, shift left] & \bullet \arrow[l, shift left] \arrow[r, shift left] & \bullet \arrow[l, shift left]
\end{tikzcd}$$

If $U_i = k$ when i is one of the outermost vertices and $U_i = 0$ otherwise, we obtain the following quiver as $\widetilde{Q}^{dbl}$

$$\begin{tikzcd}
& & \bullet \arrow[dll, bend right] \arrow[drr, bend left] & & \\
\bullet \arrow[urr, bend left] \arrow[r, shift left] & \bullet \arrow[l, shift left] \arrow[r, shift left] & \cdots \arrow[l, shift left] \arrow[r, shift left] & \bullet \arrow[l, shift left] \arrow[r, shift left] & \bullet \arrow[l, shift left] \arrow[ull, bend right]
\end{tikzcd}$$

that is recognisable as the double of the affine $\widetilde{A}_n$ quiver.

We call U a 'framing' for Q. Our treatment has already differed from the original work of Nakajima [61]; we are using a reformulation due to Crawley-Boevey [29]. The notes of Ginzburg [39] reconcile the two. We define

$$\begin{aligned}
\widetilde{\mathcal{R}}(Q, U, V) &= \bigoplus_{\alpha \in Q_1} \mathrm{Hom}(V_{t(\alpha)}, V_{h(\alpha)}) \oplus \mathrm{Hom}(V_{t(\alpha^*)}, V_{h(\alpha^*)}) \oplus \bigoplus_{i \in Q_0} \bigoplus_{i=1}^{\dim U_i} V_i \oplus V_i^* \\
&= \bigoplus_{\alpha \in Q_1} \mathrm{Hom}(V_{t(\alpha)}, V_{h(\alpha)}) \oplus \mathrm{Hom}(V_{h(\alpha)}, V_{t(\alpha)}) \oplus \bigoplus_{i \in Q_0} \bigoplus_{i=1}^{\dim U_i} V_i \oplus V_i^*
\end{aligned}$$

This is the representation space $\mathcal{R}(\widetilde{Q}^{\mathrm{dbl}}, \widetilde{V})$ where $\widetilde{V}$ is the $\widetilde{Q}_0$-graded vector space

$$\widetilde{V} = \bigoplus_{i \in Q_0} V_i \oplus \mathbb{C}$$

where $\deg \mathbb{C} = \infty$. It inherits a holomorphic symplectic structure via the pairing

$$\operatorname{Hom}(V_{t(\alpha)}, V_{h(\alpha)}) \times \operatorname{Hom}(V_{h(\alpha)}, V_{t(\alpha)}) \to \mathbb{C}, (f_\alpha, f_{\alpha^*}) \mapsto \operatorname{tr}(f_{\alpha^*} f_\alpha)$$

$\widetilde{\mathcal{R}}(Q, U, V)$ has a $\mathrm{GL}(V)$-action where the one dimensional vector space at ∞ is left alone. This action is Hamiltonian, and so we obtain a moment map

$$\mu \colon \widetilde{\mathcal{R}}(Q, U, V) \to \mathfrak{gl}(V), (f_\alpha) \cup (f_{\alpha^*}) \cup (a_{ij}, b_{ij}) \mapsto \left(\prod_{t(\alpha)=i} f_{\alpha^*} f_\alpha - \prod_{h(\alpha)=i} f_\alpha f_{\alpha^*} + \prod_{j=1}^{\dim U_i} b_{ij} a_{ij} \right)_{i \in Q_0}$$

conflating $\mathfrak{gl}(V)$ and $\mathfrak{gl}(V)^\vee$. Choosing $\lambda \in \mathfrak{gl}(V)$ gives a level set $\mu^{-1}(\lambda)$ whose Hamiltonian reduction we denote by

$$\mathfrak{M}_\lambda(Q, U, V) := \mu^{-1}(\lambda) / \operatorname{GL}(V)$$

The quotient here is simply the affine quotient on the level of varieties. Observe that taking the preimage of the moment map is imposing a particular set of relations on representations in $\widetilde{\mathcal{R}}(Q, U, V)$ however in general this symplectic construction is not a special case of the previous GIT construction since there in the symplectic picture there is no action at the vertex ∞. We will see some special cases later where they do actually agree. Nakajima's vision [61] for these varieties was to encode the relations for Kac–Moody algebras in terms of their homology. When $\dim(V_i) = 1$ for all $i \in Q_0$, $\mathfrak{M}_\lambda(Q, U, V)$ is a 'hypertoric variety' – the quaternionic analog of a complex toric variety – whose study has many valuable applications in combinatorics [67].

2.7 Bridgeland stability

In contrast to the previous types of stability that have been discussed, Bridgeland stability requires no explicitly equivariant context. We follow [13, 54]. Let $\mathcal{A}$ be an abelian category. The Grothendieck group of $\mathcal{A}$ is the abelian group

$$K(\mathcal{A}) := \mathbb{Z} \cdot \operatorname{Obj} \mathcal{A} / \sim$$

where $E_1 + E_3 \sim E_2$ whenever $0 \to E_1 \to E_2 \to E_3 \to 0$ is an exact sequence. In other words, K-theory trivialises extensions in $\mathcal{A}$. A *stability function* on $\mathcal{A}$ is a group homomorphism

$$Z \colon K(\mathcal{A}) \to \mathbb{C}$$

such that for all nonzero $E \in K(\mathcal{A})$ one has $Z(E) \in \mathbb{H} \cup \mathbb{R}_{<0}$, where $\mathbb{H} := \{a + ib : b > 0\}$ is the upper halfplane. Define the Z-slope of E to be

$$\phi_Z(E) := \frac{1}{\pi} \arg(Z(E))$$

It is clear from this definition why we require the image to be constrained.

Example 2.7.1. *Let Q be a quiver and consider the abelian category $\mathcal{A} = \mathrm{rep}(Q)$. If Q has no loops or oriented cycles, then*

$$K(\mathcal{A}) = \mathbb{Z} \cdot \{S_i : i \in Q_0\}$$

where S_i are the vertex simple modules with dimension vector $(\delta_{ij})_{j \in Q_0}$. The reason for this is easy to see from a small example. Take $Q = \bullet \to \bullet$ with vertex simple modules

$$S_1 = \mathbb{C} \to 0 \text{ and } S_2 = 0 \to \mathbb{C}$$

Take a general representation of the form $V_1 \xrightarrow{f} V_2$. Split this into

$$\ker f \oplus V_1' \to \operatorname{im} f \oplus V_2' = (\ker f \to 0) \oplus (V_1' \xrightarrow{\sim} \operatorname{im} f) \oplus (0 \to V_2')$$

The bookends are clearly multiples of vertex simples, so it suffices to express $\mathbb{C} \xrightarrow{1} \mathbb{C}$ as an extension of vertex simples. This is achieved by the sequence

$$0 \to (0 \to \mathbb{C}) \to (\mathbb{C} \to \mathbb{C}) \to (\mathbb{C} \to 0) \to 0$$

In this situation where $K(\mathcal{A})$ is freely generated one can choose $z_i \in \mathbb{H} \cup \mathbb{R}_{<0}$ and define $Z(S_i) = z_i$ to obtain a stability function.

Example 2.7.2. *The original motivating example for stability functions comes from vector bundles on algebraic curves. Consider the stability function given by $Z(E) = -\deg(E) + i\,\mathrm{rk}(E)$ on the category Vec_X of vector bundles on a smooth projective curve X. The associated slope is a rescaling of $\deg(E)/\mathrm{rk}(E)$, which is recognisable as the usual slope for vector bundles on curves.*

Definition 2.7.3. *A nonzero object $E \in \mathcal{A}$ is Z-stable if for any nontrivial $0 \neq F \subsetneq E$ one has $\phi_Z(F) < \phi_Z(E)$. E is Z-*semistable *if such $F \subsetneq E$ have $\phi_Z(F) \leq \phi_Z(E)$.*

Definition 2.7.4. *A* stability condition *for $\mathcal{A}$ is a stability function $Z \colon \mathcal{A} \to \mathbb{C}$ such that each $E \in \mathcal{A}$ has a filtration*

$$0 = E_0 \subsetneq \cdots \subsetneq E_{n-1} \subsetneq E_n = E$$

with quotients $Q_i = E_i/E_{i-1}$ satisfying:

- *each Q_i is semistable,*
- *$\phi_Z(Q_i) > \phi_Z(Q_{i+1})$.*

Such a filtration is called a Harder-Narasimhan filtration*.*

Example 2.7.5. *Let Q be a quiver without loops or oriented cycles. Let $\theta \in \mathrm{Hom}_{\mathbb{Z}}(\mathbb{Z}^{Q_0}, \mathbb{R})$. As in §2.6 define $\theta(E) := \theta(\underline{\dim}(E))$ for a representation E of Q. A stability condition for* $\mathrm{rep}(Q)$ *is given by*

$$Z_\theta(E) = \theta(E) + i \sum_{j \in Q_0} \dim E_j$$

This arises from the previous construction by setting $Z(S_j) = \theta(j) + i$. Let us observe that Z_θ-stability is nearly the same as θ-stability. If E is Z_θ-stable and $\theta(E) = 0$, then a subrepresentation $F \subset E$ has slope $\phi_{Z_\theta}(F) < \phi_{Z_\theta}(E) = 1/2$, which means that the real part of $Z_\theta(F)$ is larger than zero or, equivalently, that $\theta(F) > 0$. Note the requirement for $\theta(E) = 0$ comes from the fact that we are considering all of $\mathrm{rep}(Q)$ *and not just those of fixed dimension vector in order to study filtrations.*

There is a modification of this definition for stability conditions on general triangulated categories such as the derived category $D^b(X)$ of a variety X: these are *Bridgeland stability conditions.*

Chapter 3

Stability for orbifolds

3.1 Introduction to the McKay correspondence

Let $G \subset \mathrm{SL}_n(\mathbb{C})$ be a finite subgroup. When $n = 2$ there is a famous ADE classification of such subgroups that matches the classification of Du Val or modality zero singularities by taking a subgroup G to the quotient singularity $0 \in \mathbb{C}^2/G$. This observation and the surrounding deep interactions of the geometry of $\mathbb{C}^2/G$ and its resolutions, and the representation theory of G are known as the two-dimensional McKay correspondence [5, 8, 37, 48, 49, 57, 69]. In this case, the unique minimal or crepant resolution has a modular interpretation as the G-Hilbert scheme $G\text{-Hilb}\,\mathbb{C}^2$.

The moduli space $G\text{-Hilb}\,M$ for M a variety and $G \subseteq \mathrm{Aut}(M)$ a finite subgroup parameterises G-clusters in M: zero-dimensional G-invariant subschemes M of $\mathbb{C}^2$ with $H^0(\mathcal{O}_Z) \cong \mathbb{C}[G]$ as G-modules. One should think of G-clusters as being 'scheme-theoretic group orbits', or degenerations of free group orbits with the reduced scheme structure. This was generalised to three dimensions for finite abelian subgroups of $\mathrm{SL}_3(\mathbb{C})$ by Nakamura [62] who showed that $G\text{-Hilb}\,\mathbb{C}^3$ is a crepant resolution of $\mathbb{C}^3/G$ and then to all subgroups by the celebrated work of Bridgeland–King–Reid [14]. They moreover established an equivalence of categories

$$D^b(G\text{-Hilb}\,\mathbb{C}^3) \simeq D^b_G(\mathbb{C}^3) \tag{3.1}$$

which also holds if $G\text{-Hilb}\,\mathbb{C}^3$ is replaced by any projective crepant resolution of $\mathbb{C}^3/G$.

3.2 Abelian McKay correspondence

Let $G \subset \mathrm{SL}_3(\mathbb{C})$ be a finite abelian subgroup. We will assume that G is cyclic, however most of what follows carries over to the non-cyclic case. We will denote by $\frac{1}{r}(a, b, c)$ the cyclic subgroup of $\mathrm{SL}_3(\mathbb{C})$ generated by the matrix

$$g = \begin{pmatrix} \varepsilon^a & & \\ & \varepsilon^b & \\ & & \varepsilon^c \end{pmatrix}$$

where ε is a primitive rth root of unity and $a+b+c \equiv 0 \bmod r$.

From the work of Craw [24] one can reinterpret G-Hilb as a moduli space of quiver representations. The quiver in question is the *McKay quiver* with vertices indexed by irreducible representations of G and the number of arrows between ρ and ρ' defined to be

$$\dim \mathrm{Hom}_G(\rho' \otimes \rho_{\mathrm{std}}, \rho)$$

where ρ_{std} is the standard representation of G acting on $\mathbb{C}^3$. We choose the dimension vector $\underline{\mathrm{d}} = (\dim \rho)_\rho$ and a stability parameter $\theta \in \Theta$ as defined above. Locally to this chapter we define $\mathcal{M}_\theta := \mathcal{M}_\theta(Q, \underline{\mathrm{d}})$ to be the fine moduli space of θ-stable representations of the McKay quiver with dimension vector $\underline{\mathrm{d}}$ subject to certain relations coming from the associated preprojective algebra. When $\theta(\rho) > 0$ for all ρ one has that $\mathcal{M}_\theta(Q, \underline{\mathrm{d}}) = G\text{-Hilb}\,\mathbb{C}^3$.

It is apparent from [14] that their smoothness result and equivalence of categories (3.1) holds for any generic θ and so one obtains many resolutions $\mathcal{M}_\theta(Q, \underline{\mathrm{d}})$ and corresponding equivalences of categories

$$\Phi_\theta : D^b(\mathcal{M}_\theta) \to D^b_G(\mathbb{C}^3)$$

These equivalences are Fourier-Mukai transforms coming from the universal family of $\mathcal{M}_\theta$. Consider the diagram

$$\begin{tikzcd} & \mathcal{Z}_\theta \arrow[dl] \arrow[dr] & \\ \mathcal{M}_\theta \arrow[dr] & & \mathbb{C}^3 \arrow[dl] \\ & \mathbb{C}^3/G & \end{tikzcd}$$

where $\mathcal{Z}_\theta \to \mathcal{M}_\theta$ is the universal family. The equivalence Φ_θ is the Fourier-Mukai transform with kernel $\mathcal{O}_{\mathcal{Z}_\theta}$ as described in [26, §2.4].

As alluded to in Chapter 2, the stability space Θ has a wall-and-chamber structure such that any θ, θ' from the same open chamber $\mathfrak{C} \subseteq \Theta$ produce isomorphic moduli spaces: $\mathcal{M}_\theta \cong \mathcal{M}_{\theta'}$. For simplicity we denote by $\mathcal{M}_{\mathfrak{C}}$ and $\Phi_{\mathfrak{C}}$ the moduli space and equivalence of categories for any generic $\theta \in \mathfrak{C}$.

When G is abelian, each resolution $\mathcal{M}_{\mathfrak{C}}$ is toric. Fix the lattice $N = \mathbb{Z}^3 + \mathbb{Z} \cdot (\frac{a}{r}, \frac{b}{r}, \frac{c}{r})$. The singularity $\mathbb{C}^3/G$ is described by the cone $\sigma = \mathrm{Cone}(e_1, e_2, e_3)$ inside

$$N_{\mathbb{R}} := N \otimes_Z \mathbb{R} = \mathbb{R}^3_{\langle x_1, x_2, x_3 \rangle}$$

and crepant resolutions correspond to triangulations of the cone face $\sigma \cap (x_1 + x_2 + x_3 = 1)$ such that the vertices of each triangle lies in N, and each triangle is smooth (its vertices form a $\mathbb{Z}$-basis of N). In figures we will always draw only the cone face to produce two-dimensional pictures.

3.3 Reid's recipe

Let us focus on the case $\mathcal{M}_{\mathfrak{C}} = G\text{-Hilb}\,\mathbb{C}^3$. We will denote the universal G-cluster by $\mathcal{Z}$ and the chamber of Θ corresponding to G-Hilb by $\mathfrak{C}_0$. Craw-Reid [28] present an entertaining algorithm to construct the triangulation for G-Hilb that, after commenting on some of the salient features, we will use without comment.

We call the triangulation for G-Hilb the *Craw-Reid triangulation*. It divides the junior simplex into so-called 'regular triangles' of equal side length that fall into one of two cases:

- *corner triangles*, which have one of the vertices e_1, e_2, e_3 of the junior simplex as a vertex
- *meeting of champions*, for which none of the vertices of the junior simplex are vertices

Craw–Reid show that there is at most one meeting of champions triangle (possibly of side length zero, in which case it is a point). After dividing the junior simplex into such triangles, one subdivides them further into smooth triangles as depicted in Figure 3.1: the resulting unimodal triangulation describes the resolution G-Hilb.

Figure 3.1: A regular triangle and its triangulation

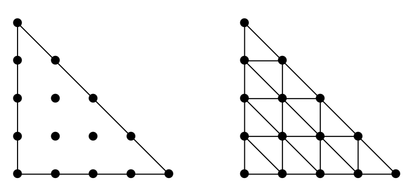

In early versions of the McKay correspondence [69] one of the chief aims was to supply a bijection from irreducible characters of G to a basis of cohomology on a crepant resolution. This was explicitly computed for G-Hilb by Craw [24] when G is abelian using 'Reid's recipe':

a labelling of exceptional subvarieties by characters of G. Reid's recipe is one of the main tools we will use to compute walls and so we will describe it in some detail.

An exceptional curve C in G-Hilb corresponds to an edge in the Craw-Reid triangulation, which in turn corresponds to a two-dimensional cone in the fan for G-Hilb. A primitive normal vector (α, β, γ) to this cone defines a G-invariant ratio of monomials

$$x^\alpha y^\beta z^\gamma = m_1/m_2$$

where x, y, z are eigencoordinates on $\mathbb{C}^3$ for G. Mark the curve C with the character by which G acts on m_1 (or m_2). We define the *χ-chain* to be the collection of all exceptional curves (or edges in the Craw-Reid triangulation) marked by the character χ. We say that a triangle in the Craw-Reid triangulation is a *χ-triangle* if one of its edges is marked with the character χ.

After marking all curves, there is a procedure for labelling the compact exceptional divisors, or interior vertices of the triangulation. Let D be such a divisor corresponding to a vertex v. There are three cases:

- v is trivalent: $D \cong \mathbb{P}^2$ and the three exceptional curves in D are all marked with the same character χ. **Mark D with $\chi^{\otimes 2}$.**

- v is 4- or 5-valent, or 6-valent and not inside a regular triangle: D is a Hirzebruch surface blown up in valency $-\,4$ points. There are two pairs of exceptional curves in D each marked with the same character χ and χ'. **Mark D with $\chi \otimes \chi'$.**

- v is 6-valent and lies in the interior of a regular triangle: D is a del Pezzo surface of degree 6, and there are three pairs of exceptional curves each marked with the same character χ, χ', χ''. D has two G-invariant maps to $\mathbb{P}^2$, **mark D by the two characters arising from the monomials constituting these two maps**. These two characters ϕ_1, ϕ_2 satisfy

$$\chi \otimes \chi' \otimes \chi'' = \phi_1 \otimes \phi_2$$

For more detail see [24, Lemmas 3.1-3.4]. We will frequently refer to a curve or a divisor marked with a character χ as a χ-curve or a χ-divisor.

Example 3.3.1. *In Figure 3.2 with $G = \frac{1}{30}(25, 2, 3)$, the leftmost curve marked with the character 20 has normal $(-2, 25, 0)$ giving the G-invariant ratio y^{25}/x^2. G acts on the numerator and denominator by the character $\varepsilon \mapsto \varepsilon^{20}$, hence the marking. The divisor marked with 23 incident to the previous curve marked with 20 has two pairs of curves with characters 20 and 3 and a fifth curve with character 15. Thus, the divisor is correctly marked by $20 + 3 = 23$.*

We refer to divisors of the first two types - that is, all divisors not isomorphic to a del Pezzo surfaces of degree 6 - as *Hirzebruch divisors*, and to divisors isomorphic to a del Pezzo surface of degree 6 as *del Pezzo divisors*. We ask the reader to have grace on the slight abuse of terminology as $\mathbb{P}^2$ is also a del Pezzo surface. For a character χ marking a curve, we denote by $\mathrm{Hirz}(\chi)$ the set of characters marking Hirzebruch divisors in the interior of the χ-chain and by $\mathrm{dP}(\chi)$ the set of characters marking del Pezzo divisors in the interior of the χ-chain. We will often say 'along the χ-chain' in place of 'in the interior of the χ-chain'.

3.4 Walls for G-Hilb

By definition, the vertices of the McKay quiver biject with the irreducible representations $\mathrm{Irr}(G)$ of G and so one can conflate the stability space Θ with a quotient of the representation ring of G (tensored with $\mathbb{R}$). As θ varies, one obtains many different crepant resolutions of $\mathbb{C}^3/G$; in the case that G is abelian, Craw-Ishii [26] show that all projective crepant resolutions arise in this way. The stability space Θ has a wall-and-chamber structure such that the moduli space $\mathcal{M}_\theta(Q, \underline{\mathrm{d}})$ is constant so long as θ remains inside a given chamber. We denote the moduli space $\mathcal{M}_{\mathfrak{C}} := \mathcal{M}_\theta(Q, \underline{\mathrm{d}})$ for any generic θ in a chamber $\mathfrak{C}$. Denote the chamber corresponding to $G\text{-Hilb}\,\mathbb{C}^3$ by $\mathfrak{C}_0$. The positive octant

$$\Theta^+ := \{\theta \in \Theta : \theta(\rho) > 0 \text{ for all nontrivial } \rho \in \mathrm{Irr}(G)\}$$

lies inside $\mathfrak{C}_0$ however it is not usually equal to it. The primary purpose of this chapter is to provide explicit combinatorial inequalities defining $\mathfrak{C}_0$ and identify precisely which of these define walls of $\mathfrak{C}_0$. We remark that such equations were computed for a group of order 11 in [26, Example 9.6]. For alternative interpretations of this wall-and-chamber structure in related contexts, see [63].

[26, Theorem 9.5] gives an abstract description of such inequalities, however making calculations or deducing general statements from it are difficult tasks. One can view some of the results herein as a combinatorialisation of this theorem, which turn out to be very amenable to applications. To briefly outline the context and notation of [26] that we will also use, for a chamber $\mathfrak{C} \subseteq \Theta$ the equivalence from (3.1) induces an isomorphism $\varphi_{\mathfrak{C}} : K_0(\mathcal{M}_{\mathfrak{C}}) \to K_G(\mathbb{C}^3) = \mathrm{Rep}(G)$. Here $K_0(\mathcal{M}_{\mathfrak{C}})$ denotes the K-group of sheaves supported on the exceptional fibre of $\mathcal{M}_{\mathfrak{C}} \to \mathbb{C}^3/G$. Walls in Θ are cut out by hyperplanes $(\sum_i \alpha_i \cdot \theta(\chi_i) = 0)$ for some characters $\chi_i \in \mathrm{Irr}(G)$ and integers $\alpha_i \in \mathbb{Z}$. The inequalities in [26] have three different forms, each coming from exceptional subvarieties. Firstly, each exceptional curve $C \subseteq G\text{-Hilb}\,\mathbb{C}^3$ gives an inequality of the form

$$\theta(\varphi_{\mathfrak{C}_0}(\mathcal{O}_C)) > 0$$

The characters appearing in of these inequalities are packaged in collections of monomials associated to exceptional curves that were named by Nakamura in a different context as *G-igsaw pieces* in [62]. Our first result is to pin down which characters lie in G-igsaw pieces.

Theorem 3.4.1. *There is a combinatorial procedure that we call the* ***unlocking procedure*** *for computing the characters appearing in a G-igsaw piece for an exceptional curve in* $G\text{-Hilb}\,\mathbb{C}^3$.

To briefly illustrate how the procedure works, we consider the example of $G = \frac{1}{30}(25, 2, 3)$. The triangulation for G-Hilb is shown in Figure 3.2 along with Reid's recipe.

Figure 3.2: G-Hilb and Reid's recipe for $G = \frac{1}{30}(25, 2, 3)$

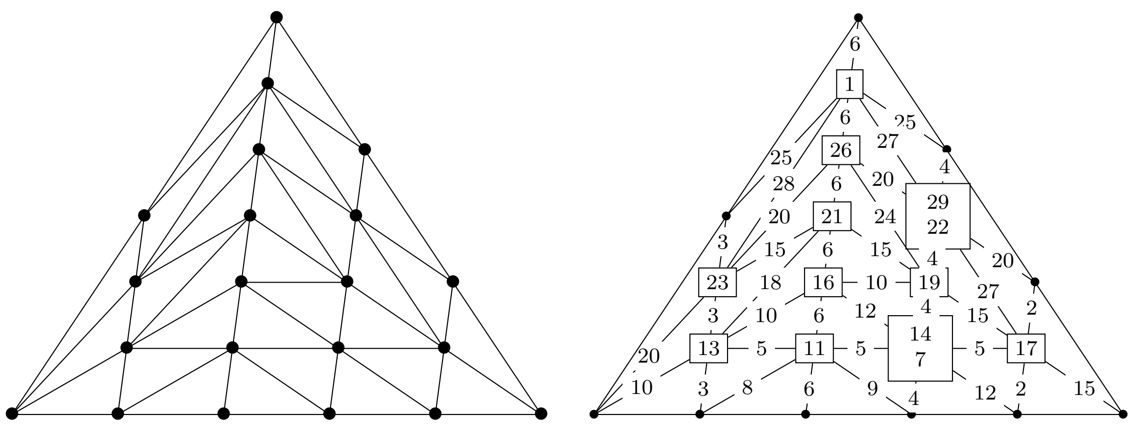

We will demonstrate the unlocking procedure for the curve C shown on the left side of Figure 3.3 marked with the character 5; that is, the character taking $g \mapsto \varepsilon^5$. On the right side of Figure 3.3 we illustrate the unlocking procedure. Roughly, we consider all the curves (or edges) marked with 5, add one character marking each divisor containing two such curves (or vertices between two edges marked with 5), and finally add the characters appearing in G-igsaw pieces for certain curves cohabiting a divisor with a curve marked with 5. In this case, the only such extra curve is marked with 9 and the G-igsaw piece for this curve consists just of the character 9 itself. Some recursion will be necessary in general to compute the smaller G-igsaw pieces of such curves. It follows that the G-igsaw piece for C has characters $5, 9, 11, 14$. Observe that this G-igsaw piece only picked one of the two characters $7, 14$ marking a divisor containing two 5-curves. We will elaborate later on how the unlocking procedure identifies which of the two characters should be added.

Walls inside Θ are of various types denoted `0-III` depending on the birational geometry of the moduli spaces near the wall following [74]. Walls of Type `I` correspond to flops induced

Figure 3.3: Unlocking for a 5-curve

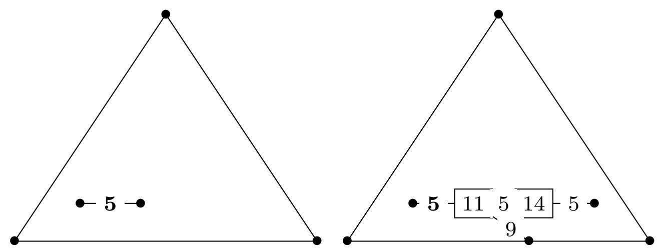

by curves. By [26, Theorem 9.12], every flop in a single exceptional $(-1,-1)$-curve can be realised by a wall-crossing of Type I directly from $\mathfrak{C}_0$, which is very much not true for other resolutions; see [26, Example 9.13]. These walls are the easiest to compute. We denote the set of characters appearing in a G-igsaw piece for an exceptional curve C by G-ig(C). The following result is implied by [26, Cor. 5.2 + Prop. 9.7 + Theorem 9.12] and Theorem 3.4.1.

Proposition 3.4.2 (Prop. 3.8.2)**.** *Suppose $C \subseteq G\text{-Hilb}\,\mathbb{A}^3$ is an exceptional $(-1,-1)$-curve marked with character χ by Reid's recipe. Then, the necessary inequality corresponding to C that defines a Type* I *wall of $\mathfrak{C}_0$ is given by*

$$\theta(\varphi_{\mathfrak{C}_0}(\mathcal{O}_C)) = \sum_{\chi \in \text{G-ig}(C)} \theta(\chi) > 0$$

where G-ig(C) *is computed by the unlocking procedure.*

Walls of Type III arise from exceptional curves corresponding to certain 'boundary' edges in the triangulation for G-Hilb. The inequalities potentially defining such walls are computed by the following result, which is a consequence of [26, Cor. 5.2] and Theorem 3.4.1.

Proposition 3.4.3. *Suppose $C \subseteq G\text{-Hilb}\,\mathbb{A}^3$ is an exceptional boundary curve marked with character χ by Reid's recipe. Then, the inequality corresponding to C is given by*

$$\theta(\varphi_{\mathfrak{C}_0}(\mathcal{O}_C)) = \sum_{\chi \in \text{G-ig}(C)} \deg(\mathcal{R}_\chi|_C)\theta(\chi) > 0$$

where G-ig(C) *is computed by the unlocking procedure, and where $\mathcal{R}_\chi$ is the tautological line bundle for χ.*

We can also use the unlocking procedure to compute inequalities that do not come from exceptional curves. The other two kinds of inequality come from exceptional divisors. For each character ψ marking a divisor, we obtain an inequality $\theta(\psi) > 0$. The second kind of inequality coming from divisors is more complicated.

Proposition 3.4.4 (Prop. 3.8.1)**.** *Suppose D' is a (not necessarily prime) exceptional divisor in G-Hilb $\mathbb{C}^3$. Then any $\theta \in \mathfrak{C}_0$ satisfies*

$$\theta(\varphi_{\mathfrak{C}_0}(\omega_{D'}^{\vee})) = \sum_{C \subseteq D'} \sum_{\chi \in \text{G-ig}(C)} \theta(\chi) > 0$$

where C ranges over exceptional curves inside D'.

It was shown in [26, Prop. 3.8] that there are no Type `II` walls in Θ. However, it is still interesting to compute the inequalities $\theta(\varphi_{\mathfrak{C}}(\mathcal{O}_C)) > 0$ for $(1, -3)$-curves that would induce contractions of this type.

Proposition 3.4.5 (Prop. 3.8.3)**.** *Suppose $C \subseteq G$-Hilb $\mathbb{A}^3$ is an exceptional $(1, -3)$-curve marked with character χ by Reid's recipe. Then, the inequality corresponding to C is given by*

$$\theta(\varphi_{\mathfrak{C}_0}(\mathcal{O}_C)) = 2 \cdot \theta(\chi^{\otimes 2}) + \sum_{\chi \in \text{G-ig}(C) \setminus \{\chi^{\otimes 2}\}} \theta(\chi) = 0$$

where G-ig(C) *is computed by the unlocking procedure.*

As a result of Prop. 3.4.4 and Prop. 3.4.5 we can immediately deduce the conclusion [26, Prop. 3.8] for $\mathfrak{C}_0$.

Corollary 3.4.6 (Cor. 3.8.4)**.** *$\mathfrak{C}_0$ has no Type `II` walls.*

We can similarly reprove [26, Theorem 9.12] by combinatorial means.

Proposition 3.4.7 (Prop. 3.8.5)**.** *Each flop in a $(-1, -1)$-curve in G-Hilb $\mathbb{C}^3$ is induced by a wall-crossing from $\mathfrak{C}_0$.*

We can use these formulae to show exactly which inequalities are necessary to define $\mathfrak{C}_0$.

Theorem 3.4.8 (Theorem 3.9.1)**.** *Suppose $G \subseteq \mathrm{SL}_3(\mathbb{C})$ is a finite abelian subgroup. The walls of the chamber $\mathfrak{C}_0$ for G-Hilb $\mathbb{C}^3$ and their types are as follows:*

- *a Type `I` wall for each exceptional $(-1, -1)$-curve,*
- *a Type `III` wall for each generalised long side,*
- *a Type `0` wall for each irreducible exceptional divisor,*
- *the remaining walls are of Type `0` coming from divisors as in Prop. 3.4.4. We discuss which of these are necessary and how to reconstruct the divisor D' in §3.9.*

We will define the term 'generalised long side' in Definition 3.8.13, which is an entirely combinatorial notion.

The remainder of this chapter is devoted to proving Theorem 3.9.1 and its consequences.

3.5 G-igsaw pieces

Consider the G-clusters at torus-fixed points of G-Hilb, or triangles in the Craw-Reid triangulation. The ideal defining such a cluster is a monomial ideal and one can draw a Newton polygon in the hexagonal lattice $\mathbb{Z}^3/\mathbb{Z}\cdot(1,1,1)$ to illustrate the monomial basis. An example of a torus-fixed G-cluster for the group $G = \frac{1}{6}(1,2,3)$ is shown in Figure 3.4. Notice that there is exactly one monomial in each character space for G as desired.

Figure 3.4: A torus-fixed G-cluster for $G = \frac{1}{6}(1,2,3)$

$$\begin{array}{ccccc} & y & & & \\ yz & & 1 & & x \\ & z & & xz & \end{array}$$

The monomial ideal in $\mathbb{C}[x,y,z]$ defining this cluster is

$$\langle x^2, y^2, z^2, xy\rangle$$

Torus-fixed G-clusters for adjacent triangles separated by an exceptional curve C differ by taking a subset of the monomials basing one G-cluster and moving them to other monomials in the same character space; that is, multiplying by G-invariant ratios of monomials. This process was studied in [62] and called a *G-igsaw transformation*. Let χ be the character marking C. The subset of monomials in one of the two torus-invariant G-clusters partaking in the G-igsaw transformation is called a *G-igsaw piece* for C. There is a single monomial that divides all others in the G-igsaw piece, and this is the monomial in the G-cluster in the character space for χ.

Example 3.5.1. *We continue the example of $G = \frac{1}{6}(1,2,3)$. The Craw-Reid triangulation and Reid's recipe for this group is shown in Figure 3.5. The triangle labelled by $\star$ is the triangle corresponding to the G-cluster from Figure 3.4.*

Passing through the 4-curve C adjacent to the triangle $\star$ performs a G-igsaw transformation with G-igsaw piece centred on the monomial with character 4, which in this case is xz. The G-igsaw transformation switches xz for y^2 - since the G-invariant ratio for C is xz/y^2 - producing the new G-cluster

$$\begin{array}{cccc} y^2 & & & \\ & y & & \\ yz & & 1 & x \\ & z & & \end{array}$$

Figure 3.5: Triangulation and Reid's recipe for $\frac{1}{6}(1,2,3)$

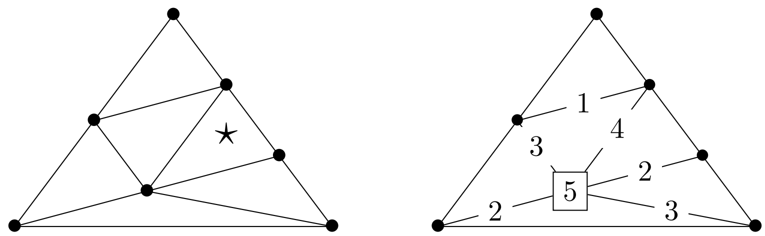

If we pass through the 2-curve bordering $\star$ then the G-igsaw piece contains the monomials y, yz and produces the G-cluster

$$\begin{array}{cccccc} & 1 & & x & & x^2 \\ z & & xz & & x^2z & \end{array}$$

As the two G-igsaw pieces for a given curve are related by multiplying by G-invariant ratios, it is clear that they each have the same set of characters represented by their monomials. We denote the set of characters in either G-igsaw piece for a curve C by G-ig(C). For convenience we will also denote by $\chi(m)$ the character by which G acts on a monomial m.

3.6 Tautological bundles

The sheaf $\mathcal{R} = \pi_* \mathcal{O}_{\mathcal{Z}}$ is locally free with fibre $H^0(\mathcal{O}_Z)$ above $Z \in G\text{-Hilb}\,\mathbb{C}^3$. It splits into eigensheaves

$$\mathcal{R} = \bigoplus_{\chi \in \operatorname{Irr} G} \mathcal{R}_\chi$$

and these summands are called *tautological line bundles.* Since G is abelian, the $\mathcal{R}_\chi$ are line bundles. [24] gives relations between these line bundles in K-theory, which translate to divisibility relations between eigenmonomials. For a triangle τ in the Craw-Reid triangulation, denote the monomial generating $\mathcal{R}_\chi|_{U_\tau}$ by $r_{\chi,\tau}$. We usually omit reference to τ so long as the context is clear.

Theorem 3.6.1 ([24, Theorem 6.1])**.** *The relations between (generators of) tautological line bundles are described by Reid's recipe in the following way.*

- *If three lines marked with the same character χ meet at a vertex marked with $\psi = \chi^{\otimes 2}$ then*

$$r_\chi^2 = r_\psi$$

- *If four or five or six lines consisting of two pairs marked by characters χ, χ' and zero or one or two extra lines marked with further characters meet at a vertex marked with $\psi = \chi \otimes \chi'$ then*
$$r_\chi \cdot r_{\chi'} = r_\psi$$
- *If six lines consisting of three pairs marked by characters χ, χ', χ'' meet at a vertex marked with ϕ, ϕ' then*
$$r_\chi \cdot r_{\chi'} \cdot r_{\chi''} = r_\phi \cdot r_{\phi'}$$

The claim is that these relations hold and generate all relations between tautological bundles. We will make heavy use of these divisibility relations between eigenmonomials to study G-igsaw pieces for exceptional curves.

As alluded to, the work of Craw-Cautis-Logvinenko [17] categorifies Reid's recipe via the tautological bundles. Many of the constructions in [17, §3-4] resemble constructions made in §3.7 below, however the computations they make are for the inverse equivalence of (3.1) to that utilised in [26] and here. It would be of interest to make a more detailed comparison.

Evident from [17,26] and below, characters marking a divisor or a single curve are special. They are termed 'essential characters' and have been further examined in [27, 73].

3.7 Computing characters in G-igsaw pieces

Our motivating question for this section is the following: *let C be a χ-curve, what are the characters that appear in a G-igsaw piece for C?* As we shall see, the answer depends somewhat on how C sits inside G-Hilb, though it is still completely combinatorial.

Monomials for divisors

We will begin by proving results for $(-1,-1)$-curves (or those lying in the interior of a regular triangle), starting with the following results relating the characters marking divisors along the χ-chain to G-igsaw pieces for χ-curves.

Lemma 3.7.1. *Suppose C is a χ-curve. Then* G-ig(C) *includes the character for each Hirzebruch divisor along the χ-chain.*

Proof. Theorem 3.6.1 implies that $r_\chi \mid r_\psi$ for each $\psi \in \mathrm{Hirz}(\chi)$ on every triangle. Hence, any G-igsaw piece featuring r_χ - such as a G-igsaw piece for C - will also feature each r_ψ and so $\psi \in$ G-ig(C). □

Lemma 3.7.2. *Suppose C is a $(-1,-1)$-curve inside a regular triangle Δ marked with χ. Then the G-igsaw piece for C includes exactly one of each pair of characters marking a del Pezzo divisor inside Δ that is along the χ-chain.*

Proof. Consider the local picture deduced from the proof of [24] Theorem 6.1 shown in Figure 3.6 for eigenmonomials near a vertex v inside Δ. We assume that Δ is a corner triangle with e_3 as vertex and one side coming from a ray out of e_1; the meeting of champions case is similar. Here ϕ_1, ϕ_2 denote the characters marking the del Pezzo divisor at v, and a, b, c, d, e, f are positive integers coming from the edges in the Craw-Reid triangulation defining out Δ. More precisely, the two sides incident to e_3 have the ratios $x^d : y^b$ and $y^e : x^c$ marking them, and the side coming from a ray out of e_1 has ratio $z^f : y^c$. $r = f$ is the side length of the regular triangle.

Figure 3.6: Generators for tautological bundles near v

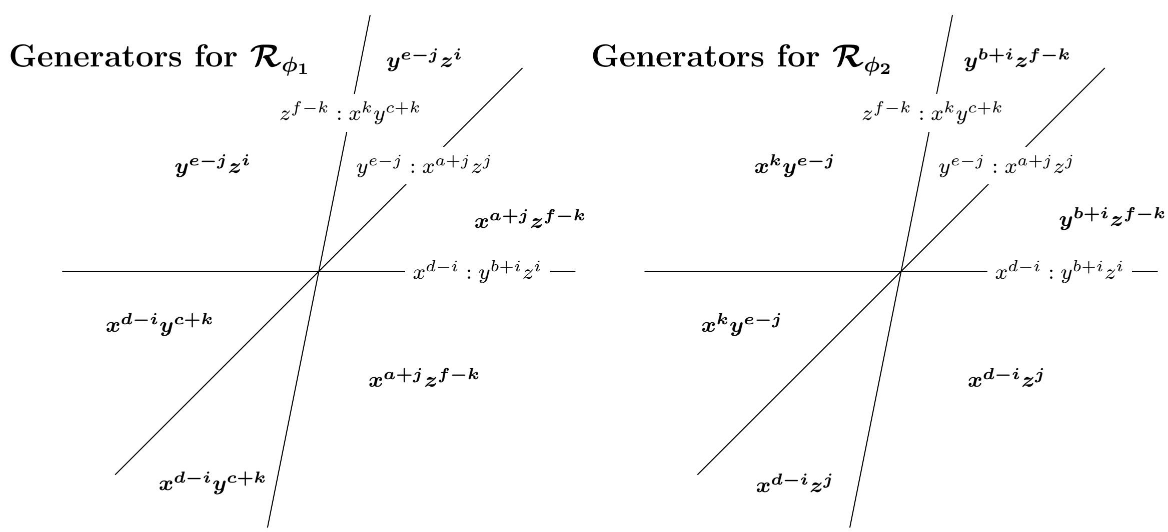

Suppose C is a curve marked with $x^{d-i} : y^{b+i}z^i$ - not necessarily incident to v - and hence that $\chi = \chi(x^{d-i})$. We will consider only this case as the same analysis goes through for each of the curves marked with the other two ratios. The computation of eigenmonomials in Figure 3.6 implies the divisibility relations shown in Figure 3.7. The χ-chain is dashed for emphasis.

Indeed, some of the divisibility relations are clear; for example in the lower half of the diagrams when $r_\chi = x^{d-i}$ and $r_{\phi_1} = x^{d-i}y^{c+k}$ or $r_{\phi_2} = x^{d-i}z^j$. The remaining claims are:

$$x^{d-i} \nmid x^k y^{e-j} \qquad x^{d-i} \nmid x^{a+j}z^j \qquad y^{b+i}z^i \mid y^{e-j}z^i \qquad y^{b+i}z^i \mid y^{b+i}z^{f-k}$$

Figure 3.7: Divisibility relations near v

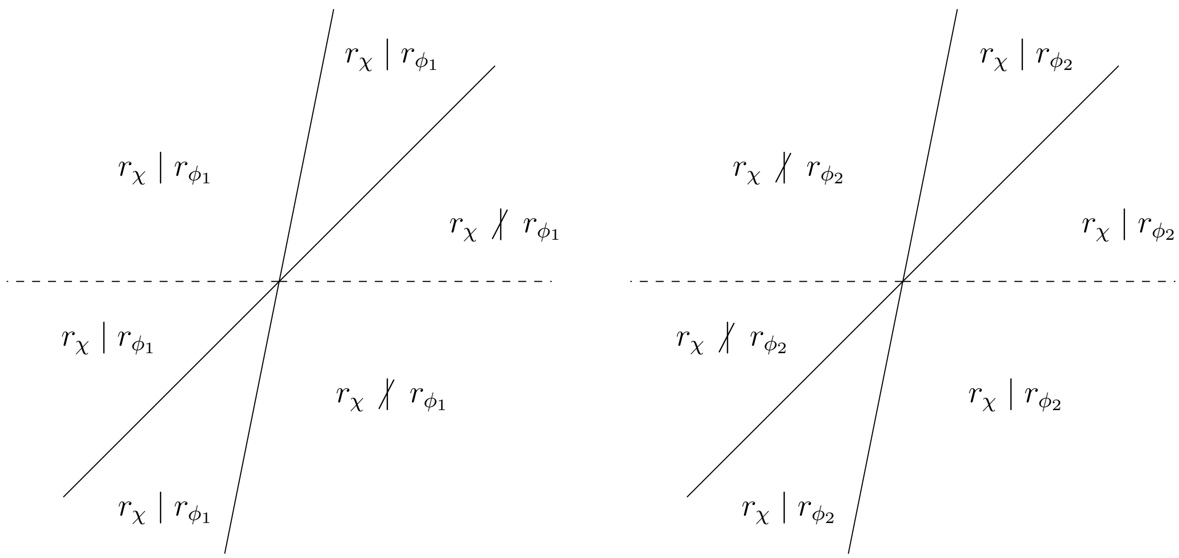

which are equivalent to

$$d-i>k \qquad d-i>a+j \qquad b+i\le e-j \qquad i\le f-k$$

respectively. Observe that the monomial x^{d-i} doesn't appear in any G-graphs for triangles above the χ-chain in the diagram but $x^{a+j}z^{f-k}$ does. Hence $x^{d-i} \nmid x^{a+j}z^{f-k}$ by the convexity of monomial bases. Similarly one sees that $x^{d-i} \nmid x^k y^{e-j}$. From [28, Prop. 3.1]

$$i+j+k = r\pm 1 \tag{3.2}$$

$$d-a = e-b-c = f = r \tag{3.3}$$

with the $\pm$ depending on whether the χ-triangle we are using to compute a G-igsaw piece for C is 'up' or 'down' (see [28, §3.2]). Notice also that since v is in the interior of a regular triangle, each of $i,j,k\ge 1$. From (3.3),

$$b+i = e-c-j-k\pm 1 \le e-j$$

and so $y^{b+i}z^i \mid y^{e-j}z^i$. From (3.2),

$$f-k = i+j\mp 1\ge i$$

and so $y^{b+i}z^i \mid y^{b+i}z^{f-k}$. It follows that r_χ divides r_{ϕ_1} and does not divide r_{ϕ_2} for every del Pezzo divisor 'to the right' of C in the orientation of Figure 3.7, and that r_χ divides r_{ϕ_2} but not r_{ϕ_1} for every del Pezzo divisor 'to the left' of C, which establishes the lemma. $\square$

This also expands on how the position of a curve determines which of the characters marking a del Pezzo surface along the χ-chain makes it into the G-igsaw piece, as alluded to while calculating walls for the example $G = \frac{1}{30}(25, 2, 3)$ in §2.

Lemma 3.7.3. *Suppose C is a $(-1,-1)$-curve marked with χ. Then* G-ig(C) *contains exactly one of each pair of characters marking a del Pezzo divisor along the χ-chain.*

Proof. Note that the only generalisation of Lemma 3.7.2 in this claim is that its conclusion also holds for del Pezzo divisors along the χ-chain but in a different regular triangle to C. This follows since at least one of the monomials in the ratio marking C (and the part of the χ-chain inside Δ) still marks the χ-chain after it passes into a new regular triangle. □

Lemmas 3.7.1-3.7.3 give an effective way of finding characters in G-ig(C). However, this will turn out to not supply all characters in G-ig(C). We will transition into discussion of the recursive procedure for filling in the remaining characters, and of the methods we will use to verify that all characters have been located. We start with a lemma of Craw-Ishii.

Lemma 3.7.4 ([26, Lemma 9.1])**.** *A character χ marks a torus-invariant compact divisor $D \subseteq G\text{-Hilb}\,\mathbb{C}^3$ iff r_χ is in the socle of every torus-invariant G-cluster in D.*

Select a $(-1,-1)$-curve C marked with χ. This lies in two del Pezzo divisors from the endpoints of the corresponding line segment. From Lemma 3.7.2 r_χ divides exactly two of the monomials in the character spaces labelling these two divisors. Suppose τ is a χ-triangle neighbouring C. By the shape of the ratios in Figure 3.6 we can assume that r_χ is not a power of a single variable. The Unique Valley Lemma [62, Lemma 3.3] of Nakamura implies that r_χ divides exactly two elements of the socle of the torus-invariant G-cluster Z_τ corresponding to τ. Lemma 3.7.4 implies that the elements in the socle of Z_τ that r_χ divides correspond exactly to these two characters labelling the neighbouring del Pezzo divisors. These are the outermost monomials in the G-igsaw piece for C on τ, so that knowing them will allow us to count how many characters appear in G-ig(C).

Recursive procedure: 'unlocking'

We will describe the recursive procedure to compute G-igsaw pieces using only the data of Reid's recipe.

Input: an exceptional $(-1,-1)$-curve C marked with a character χ.
Let $S = \{\chi\}$.

dP for each del Pezzo surface along the χ-chain, add one of the two characters marking it to S

H1 for each Hirzebruch divisor along the χ-chain, add the character marking it to S

For each Hirzebruch divisor D along the χ-chain, define a set of curves $\mathcal{C}_\chi(D)$ by:

- if D is a boundary vertex on either side of which the χ-chain consists of boundary edges, then $\mathcal{C}_\chi(D) = \emptyset$,
- if D is a boundary vertex of a regular triangle Δ that is not in the previous case, then $\mathcal{C}_\chi(D)$ consists of all curves strictly inside Δ incident to D that are not marked with χ.

H2 for each Hirzebruch divisor D along the χ-chain, add the characters from the G-igsaw pieces for curves in $\mathcal{C}_\chi(D)$ to S.

Observe that to compute the characters on the G-igsaw piece for these curves combinatorially, one may need to iteratively apply the procedure until reaching a curve where the characters in the G-igsaw piece can be read off immediately (see below for a description of such curves).

Output: G-ig$(C) = S$.

We call this the *unlocking procedure* as passing through a Hirzebruch divisor 'unlocks' the simpler G-igsaw puzzles for the relevant curves incident to it that one can recursively solve and then feed into the G-igsaw piece for C. It can be visualised as a flow through the triangulation emanating from the curve C with preferred paths defining its tributaries. We will first prove the validity of the unlocking procedure for curves inside regular triangles (i.e. those able to define flops, or $(-1,-1)$-curves) before justifying the procedure for the other exceptional curves.

Type `Iy` curves

In order to study G-igsaw pieces for curves inside regular triangles our treatment of the three kinds of curve marked with different ratios as in Figure 3.6 will now diverge. We now consider the edges marked with ratios of the form $y^{e-j} : x^{a+j}z^j$. We say that such curves are of **Type `Iy`**. The analysis from Lemma 3.7.2 gives a precise description of the socle of the nearby torus-invariant G-clusters - depicted in Figure 3.8 - and hence the G-igsaw pieces for χ-curves.

Figure 3.8: Generators of eigenspaces along a $\chi(x^{a+j}z^j)$-chain inside a regular triangle

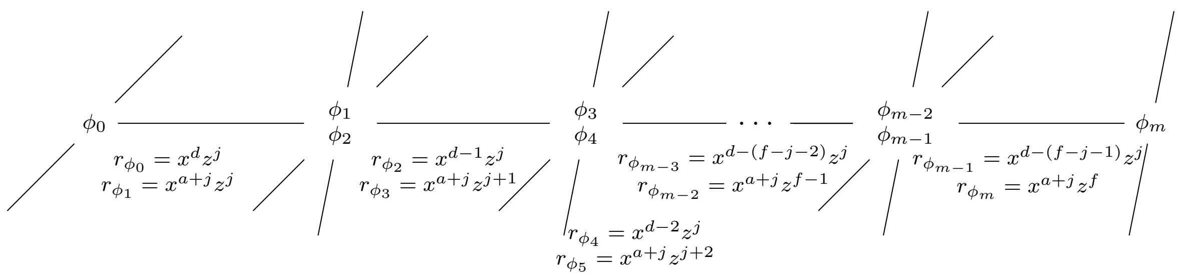

Lemma 3.7.5. *The G-igsaw piece for a χ-curve of Type* **Iy** *on a χ-triangle chosen so that in the coordinates used above $r_\chi = x^{a+j}z^j$ is*

$$\begin{array}{ccccccc} & & & r_\chi & xr_\chi & \dots & x^{f-i-j-1}r_\chi \\ & & zr_\chi & & & & \\ & \iddots & & & & & \\ z^i r_\chi & & & & & & \end{array}$$

where the curve corresponds to the ith line segment from the bottom edge of the triangle. Moreover, the χ-chain does not continue outside of this regular triangle. In particular, $\mathrm{Hirz}(\chi) = \emptyset$*.*

Proof. The calculation of the G-igsaw piece follows immediately from the description of the eigenmonomials in Figure 3.9. The χ-chain cannot continue outside of this regular triangle since neither r_{ϕ_0} nor r_{ϕ_m} are divisible by r_χ and so Theorem 3.6.1 implies that there cannot be two edges marked with χ incident to either boundary vertex. $\square$

Notice that this means that there are $f-j-1$ characters to account for, excluding χ. But this is exactly the number of del Pezzo surfaces along the χ-chain, each of which contributes one character that depends on how far along the chain the curve is.

Corollary 3.7.6. *For a χ-curve C of Type* **Iy**, $\text{G-ig}(C)$ *consists exactly of χ and precisely one character from each del Pezzo divisor along the χ-chain.*

Observe that this is a situation in which there is no recursion necessary since $\mathrm{Hirz}(\chi) = \emptyset$. This is one of the base cases that we will reduce to.

Type Ix curves

Suppose now that C is a χ-curve inside a corner triangle with e_3 as a vertex that is marked with the ratio $x^{d-i} : y^{b+i}z^i$. We say that such curves are of **Type Ix**. [24, Lemma 3.4] yields the identities in Figure 3.9 for eigenmonomials on triangles neighbouring the χ-chain, which allow us to completely describe G-igsaw pieces inside regular triangles. In the following we continue the notation of Figure 3.6 and let $\kappa = r - (i+1)$.

Figure 3.9: Generators of eigenspaces along a χ-chain inside a regular triangle

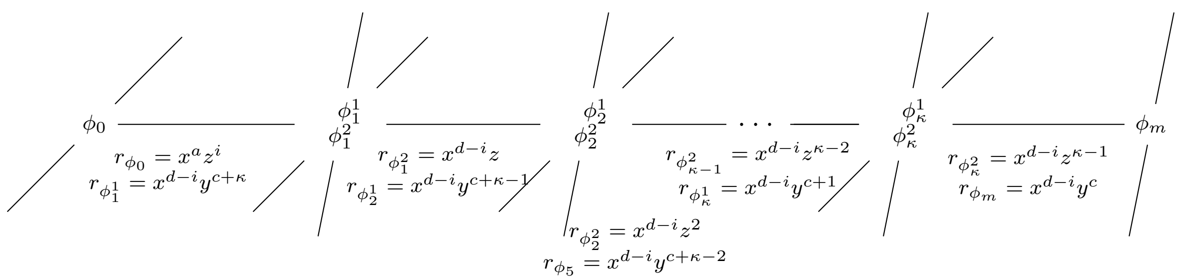

Lemma 3.7.7. *The G-igsaw piece for a χ-curve C of Type Ix on a χ-triangle chosen so that in the coordinates used above $r_\chi = x^{d-i}$ is*

$$\begin{array}{ccccc}
y^{c+k-1}r_\chi & & & & \\
 & \ddots & & & \\
 & & yr_\chi & & \\
 & & & r_\chi & \\
 & & zr_\chi & & \\
 & ⋰ & & & \\
z^j r_\chi & & & &
\end{array}$$

where C corresponds to the $(j+1)$th line segment from the left edge of the triangle, and $i+j+k=r$. Moreover, the χ-chain continues to the right and does not continue to the left of Figure 3.9.

Proof. The same argument as for Lemma 3.7.5 applies, except that r_χ does divide r_{ϕ_m} and so by Theorem 3.6.1 the χ-chain must continue past the rightmost vertex. □

Notice that the only characters in any such G-igsaw piece that are unaccounted for by divisors along the χ-chain in the same regular triangle are those for the monomials

$$yr_\chi, \dots, y^c r_\chi$$

though $y^c r_\chi = r_{\phi_m}$, which we have seen corresponds to a Hirzebruch divisor appearing along the χ-chain.

Lemma 3.7.8. *Suppose C is a χ-curve of Type* `Ix` *such that the χ-chain continues into a boundary edge of a corner triangle with e_2 as a vertex. Then* G-ig(C) *consists of χ, one character from every del Pezzo divisor along the χ-chain, and the characters marking Hirzebruch divisors along the χ-chain.*

This follows since the e_2-corner triangle has side length c and so there are exactly c Hirzebruch divisors along the boundary part of the χ-chain that contribute the remaining c characters to the G-igsaw piece. We say that the curves from Lemma 3.7.8 are of **Type** `Ixb`. This is the other base case to which the unlocking procedure reduces.

Lemma 3.7.9. *Suppose C is a χ-curve of Type* `Ix` *inside an e_2-corner triangle and suppose that the χ-chain continues into an e_3-corner triangle (not necessarily the boundary). Then* G-ig(C) *consists of χ, one character from each del Pezzo surface along the χ-chain, the character marking the Hirzebruch divisor D between the two regular triangles, and the characters from the G-igsaw piece of the* `Iy` *curve also incident to D inside the e_3-corner triangle.*

Proof. Let C' be the Type `Iy` curve incident to D in the e_3-corner triangle. Denote its character by χ'. From Lemma 3.7.5 the characters in the G-igsaw piece for C' are χ' and one character from each del Pezzo divisor along the χ'-chain inside this regular triangle. From examining the situation explicitly, on the lower χ-triangle neighbouring C one has $r_\chi = x^{d-i}$ and $r_{\chi'} = x^{d-i}z^i$ so that $r_\chi \mid r_{\chi'}$ near C. Also, one can see that the zone where $r_{\chi'}$ divides one character from each del Pezzo divisor includes this basic triangle containing C and so these divisibility relations remain. Hence, the G-igsaw piece for C' is contained in the G-igsaw piece for C. The divisibility relations are depicted in Figure 3.10.

Suppose the ratio marking the common edge of the two corner triangles is $z^f : y^c$. From Lemma 3.7.7 noting the change in notation coming from using an e_2- instead of an e_3-corner triangle, the G-igsaw piece for C is missing f characters from the χ-chain to the right. Continuing the adapted notation, we let the χ-chain enter the e_3-corner triangle at height $d - i$ so that there are $f - i$ new characters along the χ-chain corresponding to the del Pezzo divisors along the χ-chain and the boundary Hirzebruch divisor D. There are $f - (f - i) - 1 = i - 1$ divisors along the χ'-chain, making a contribution of i characters in total including χ' itself. Thus these account for all of the f missing characters. $\square$

Note that this vindicates the unlocking procedure for such curves, where only one recursion was required to unlock the single Type `Iy` curve. The final case to consider is when the

Figure 3.10: Unlocking for a Type `Ix` curve in an e_3-corner triangle

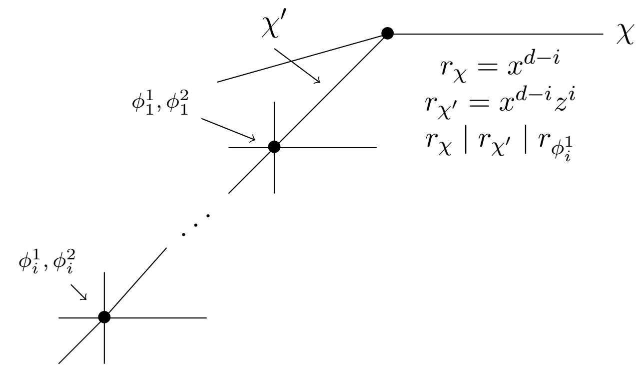

χ-chain merges into an e_1-corner triangle, whence it becomes a different kind of curve that we will treat separately after the present case. We are still able to analyse it however.

Suppose the χ-chain passes through n e_1-corner triangles before entering an e_2-corner triangle. Note that the curves in this last triangle are either boundary or `Iy` curves and so this is the final triangle the χ-chain passes through. Note further that, as an `Iy` curve cannot continue out of a regular triangle, the χ-chain must feed into the boundary of the final e_2-corner triangle.

Let the ratio $x^{d_m} : y^{b_m}$ mark the edge opposite e_1 for the mth e_1-corner triangle Δ_m from the left and so Δ_m has side length d_m. Suppose the χ-chain enters Δ_m at height i_m. This means that the χ-chain picks up $d_m - i_m$ divisors from del Pezzo divisors and a single Hirzebruch divisor inside Δ_m. From analysing local divisibility relations as above, it is clear that r_χ divides all of the monomials in the G-igsaw pieces for the Type `Ix` curve incident to the χ-chain and the leftmost Hirzebruch divisor inside each of these regular triangles. See Figure 3.11 for a schematic. We denote $D_m := \sum_{q=1}^m d_q$ and $BD_m := \sum_{q=1}^m (b_q + d_q)$.

By computing the characters on the nearby del Pezzo divisor, one can tell that these Type `Ix` curves each have $b_m + i_m$ characters in their G-igsaw pieces, making the total number of characters they contribute to the G-igsaw piece of C

$$\sum_{q=1}^{n}(d_q - i_q + b_q + i_q) = \sum_{m=1}^{n}(b_q + d_q)$$

From the equations (3.3) the ratios marking the edges from e_1 for the e_1-corner triangles are of the form

$$z^{f+\sum_{q=1}^m d_q} : y^{c-\sum_{q=1}^m (b_q+d_q)} \text{ for } m = 0, \ldots, n$$

with the last edge marked by $z^{f+\sum_{q=1}^n d_q} : y^{c-\sum_{q=1}^n (b_q+d_q)}$. In particular, this means that the e_2-corner triangle has side length $c - \sum_{q=1}^n (b_q + d_q)$ and so the final part of the χ-chain

Figure 3.11: Unlocking for a Type `Ix` curve in a series of e_1-corner triangles

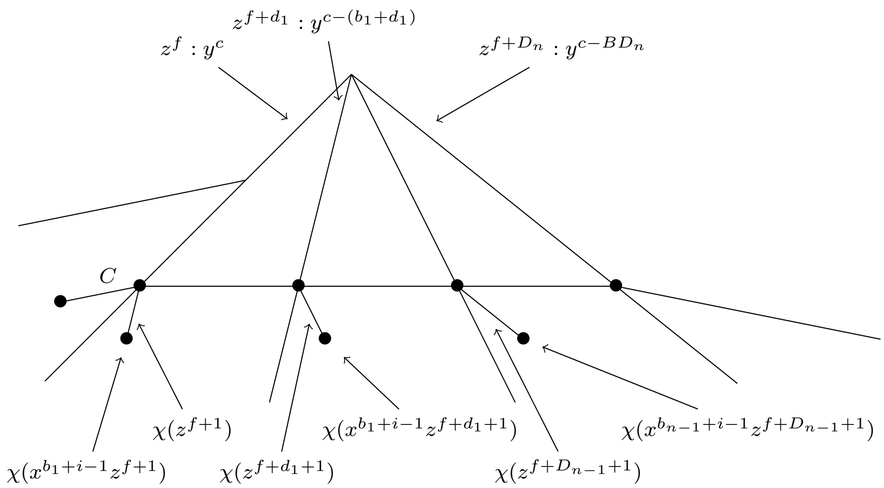

contributes $c - \sum_{q=1}^{n}(b_q + d_q)$ characters to G-ig(C). Thus, in total we have

$$\sum_{q=1}^{n}(b_q + d_q) + c - \sum_{q=1}^{n}(b_q + d_q) = c$$

characters, which is exactly the number that are not accounted for by del Pezzo divisors in the e_3-corner triangle that C inhabits by Lemma 3.7.7. This completes the proof of validity of the unlocking procedure for curves of Type `Ix`.

Type `Iz` curves

The final type of curve occurring inside regular triangles is Type `Iz`: the curves marked by ratios of the form $z^{f-k} : x^k y^{c+k}$ in the coordinates we have been using for an e_3-corner triangle. We repeat the G-igsaw analysis for these curves, represented in Figure 3.12 with the $\chi = \chi(z^{f-k})$-chain dashed.

As in all previous cases, exactly one character marking each incident del Pezzo surface has a monomial divisible by r_χ and so we can pin down the socle and hence the G-igsaw piece for such a curve.

Lemma 3.7.10. *The G-igsaw piece for a $(-1,-1)$-curve marked with χ on a χ-triangle*

Figure 3.12: Divisibility relations near v

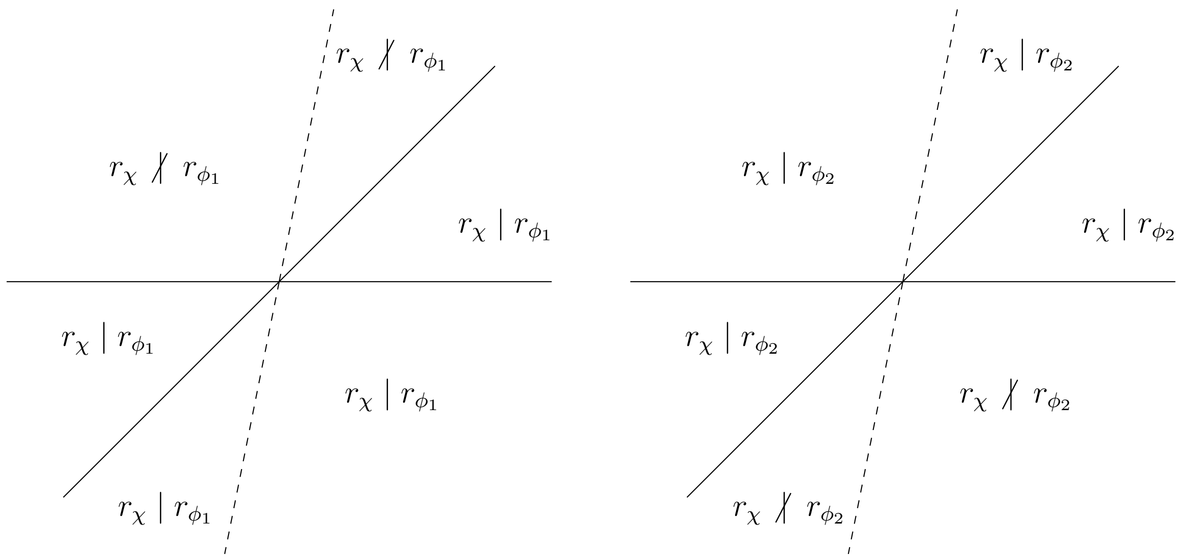

chosen so that in the coordinates used above $r_\chi = z^{f-k}$ is

$$\begin{array}{cccccc} y^{b+i-1}r_\chi & & & & & \\ & \ddots & & & & \\ & & yr_\chi & & & \\ & & & r_\chi & xr_\chi & \dots & x^{d-i-k}r_\chi \end{array}$$

where the curve corresponds to the ith line segment from the bottom edge of the triangle.

This means that there are $b+d-k$ characters in the G-igsaw piece for such a `Iz` curve. We shift notation to match the setup of the final case for Type `Ix` curves shown in Figure 3.11. In particular, we assume our Type `Iz` curve C lies in an e_1-corner triangle. Suppose it lies in the mth triangle from the left. From considering local divisibility relations near Hirzebruch divisors along the χ-chain this implies that C unlocks $m-1$ Type `Iy` curves to the left and $n-m$ Type `Ix` curves to the right. From the calculations for Type `Ix` curves, the $n-m$ Type `Ix` curves each feature $b_q + i_q$ characters in their G-igsaw pieces. From a similar calculation, one can verify that the Type `Iy` curves contain i_q characters in their G-igsaw pieces. These unlocked curves thus contribute

$$\sum_{q=1}^{m-1} i_q + \sum_{q=m+1}^{n} (b_q + i_q) = \sum_{q=m+1}^{n} b_q + \sum_{q=1}^{n} i_q - i_m$$

characters to G-ig(C). The part of the χ-chain in the e_3-corner triangle studied in the previous case contributes $f - i_0$ characters, and the part in the e_2-corner triangle contributes $c - \sum_{q=1}^{n}(b_q + d_q)$. If $i_0 \neq 0$ then we unlock another `Iy` curve with i_0 characters appearing in its G-igsaw piece. If $i_0 = 0$ then the χ-chain continues along the boundary of an e_3-corner triangle, contributing f characters. In either case there are f characters coming from the e_3-corner triangle. Lastly, there are $\sum_{q=1}^{n}(d_q - i_q)$ del Pezzo and Hirzebruch divisors along the part of the χ-chain inside e_1-corner triangles, giving in total

$$\underbrace{f}_{e_3\text{-corner}} + \underbrace{\sum_{q=m+1}^{n} b_q + \sum_{q=1}^{n} i_q - i_m}_{\text{unlocked curves}} + \underbrace{\sum_{q=1}^{n}(d_q - i_q)}_{e_1\text{-corner}} + \underbrace{c - \sum_{q=1}^{n}(b_q + d_q)}_{e_2\text{-corner}} = f + c - \sum_{q=1}^{m} b_q - i_m$$

characters. Compare to the quantity $b + d - k$ in Lemma 3.7.10, which in these coordinates is

$$c - \sum_{q=1}^{m}(b_q + d_q) + f + \sum_{q=1}^{m} d_q - i_m = f + c - \sum_{q=1}^{m} b_q - i_m$$

showing that every character in G-ig(C) is accounted for.

G-igsaw pieces for other curves

The procedure described above also works to compute G-igsaw pieces for curves not found in the interior of regular corner triangles. Firstly, direct computations of divisibility relations show that unlocking procedure as described above carries over verbatim to curves inside or whose chains pass through a meeting of champions triangle. There is a more significant expansion required for curves corresponding to boundary edges of regular triangles.

Input: An exceptional curve C corresponding to a boundary edge of a regular triangle. Let $S = \{\chi\}$.

H1 for each Hirzebruch divisor along the χ-chain, add the character marking it to S.

For each Hirzebruch divisor D along the χ-chain, define a set of curves $\mathcal{C}_\chi(D)$ to consist of the curves contained in D whose chain terminates at D or whose corresponding edges are along 'broken chains' at D. A broken chain at D is a ρ-chain for some character ρ such that D is contained in the interior of the chain and the two ρ-curves incident to D are marked with different ratios. Pictorally, this means that the edges corresponding to these curves have different slopes. We say that a chain passing through D is 'straight' at D if the two curves incident to D in the chain are marked with the same ratio; that is, the corresponding edges have the same slope. These situations are shown in Figure 3.13. For some examples in the case $G = \frac{1}{30}(25, 2, 3)$ depicted in Figures 3.2 and 3.18, the 15-chain is broken at the divisor D_{21} marked with 21 but straight at the divisors D_{19} and D_{17} marked with 19 and 17 respectively, and the 5-chain is straight at all the divisors it contains. It follows that, in this example,

$$\mathcal{C}_6(D_{11}) = \{C_8, C_9\} \text{ and } \mathcal{C}_6(D_{21}) = \{C^1_{15}, C^2_{15}, C_18\}$$

where C_ρ is the curve incident to the relevant divisor marked with ρ, and C^1_{15}, C^2_{15} are the two 15-curves incident to D_{21}.

We say that a divisor D is 'ahead' of a boundary curve C if the edge corresponding to C lies between the vertex for D and the vertex that the straight line containing C emanates from. For example, when $G = \frac{1}{30}(25, 2, 3)$, D_{19} and D_{21} are ahead of the 15-curve C contained in D_{17} and D_{19}, whereas D_{17} is not ahead of C.

H2 For each Hirzebruch divisor D ahead of C along the χ-chain, add the characters from the G-igsaw pieces for curves in $\mathcal{C}_\chi(D)$ to S.

Output: G-ig$(C) = S$.

Figure 3.13: Straight and broken chains

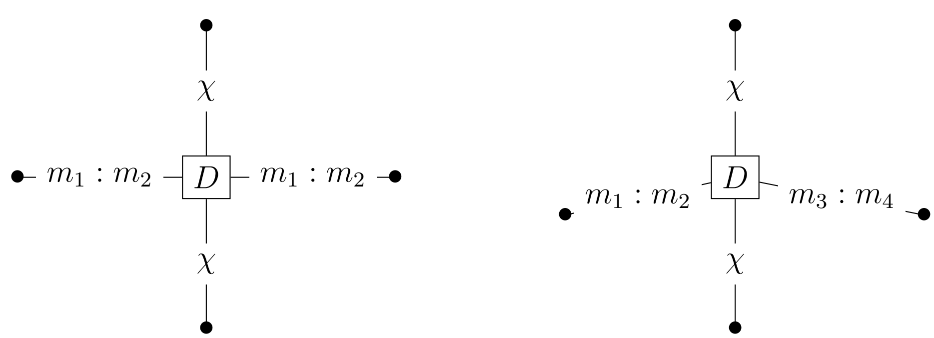

One proves that this procedure is valid in an analogous way to the procedure for curves of Types `Ix-Iz` using neighbouring divisors to compute the socle and hence the G-igsaw piece for C and then testing local divisibility relations to evidence that all these characters come from the subvarieties in the procedure. We will sketch some new elements of the proof below.

Proof. Choose coordinates so that C lies along the boundary of an e_1-corner triangle. Consider the two boundary curves C and C' shown in Figure 3.14.

One can verify using local divisibility relations that the only difference between the G-igsaw piece for C and for C' is that the latter loses the characters in the G-igsaw pieces for

Figure 3.14: Two boundary curves

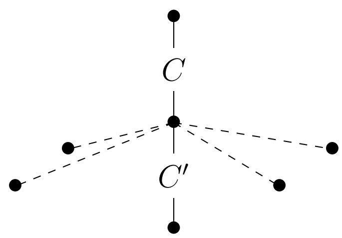

the curves on one side of the χ-chain. It hence suffices to just compute the G-igsaw piece for the curve in the χ-chain incident to e_1.

Suppose that D is a Hirzebruch divisor along the χ-chain. If D is at the boundary of two e_1-corner triangles or an e_1-corner triangle and a meeting of champions - as shown in Figure 3.15 - then one can check that r_χ divides the G-igsaw pieces for the Type `Iy` curves C_3 and C_4.

Figure 3.15: D bordering two e_1-corner triangles or meeting of champions

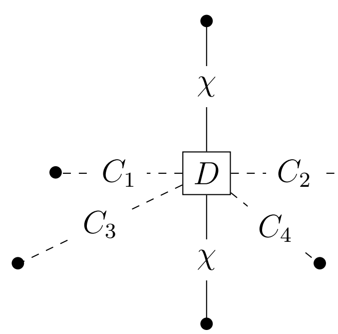

Suppose now that D borders an e_2- and an e_3-corner triangle, or an e_1-corner triangle and an e_3-corner triangle. We illustrate this situation in Figure 3.16, along with some of the ratios marking curves.

The same argument as in the previous case gives that r_χ divides the G-igsaw pieces for C_3 and C_4.

To treat the remaining two curves C_1 and C_2 in each case, we use a generalised form of [28, §3.3.2]: an edge ℓ continues in a straight line past a boundary edge ℓ_0 if and only if the ratio marking ℓ features any common variables x, y, z raised to a strictly lower exponent than in the ratio marking ℓ_0. One can verify this by a case-by-case analysis using as its base the original result from [28]. This implies that r_χ divides the G-igsaw pieces for ‘broken

Figure 3.16: D bordering an e_1- or an e_2-corner triangle and an e_3-corner triangle

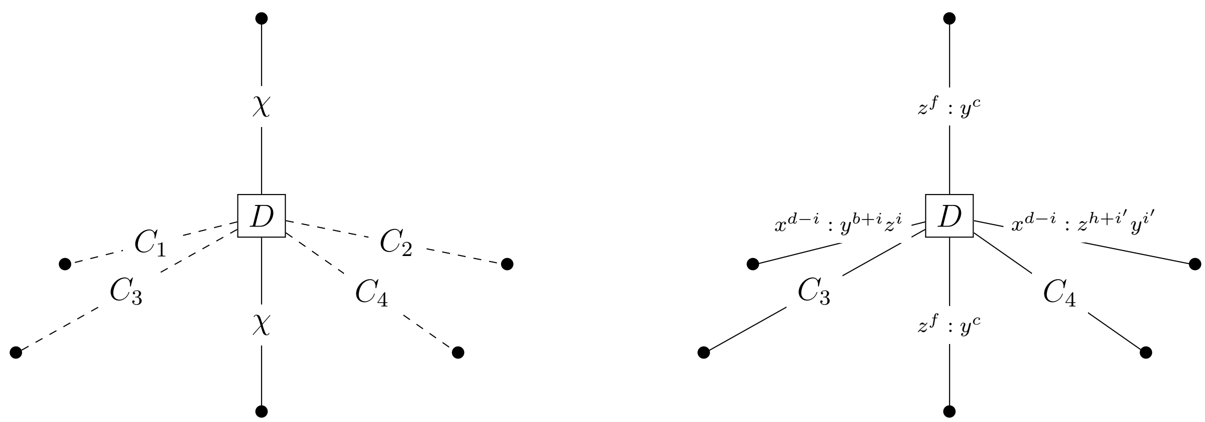

edges' that do not continue in a straight line past the χ-chain and that it does not divide any monomials in the G-igsaw pieces for 'straight edges' that do continue past the χ-chain.

Variations of the arguments above work just as well for the cases not depicted when some of the edges incident to D are also boundary edges of regular triangles. Counting up all these monomials and comparing them with a socle calculation shows that these are all the characters in the G-igsaw piece for C, which validates the unlocking procedure for boundary curves. $\square$

As an example use case, if G-Hilb has a meeting of champions of side length 0 with the three champions marked with a character χ then for any curve C along the χ-chain the characters in the G-igsaw piece are given by the unlocking procedure applied to the branch of the χ-chain that C lies on, combined with all the characters from (Hirzebruch) divisors along the other two branches of the χ-chain. We will see an example of this in §3.7.

Observe that the unlocking procedure for these curves directly generalises the unlocking procedure for $(-1,-1)$-curves in the sense that the construction of $\mathcal{C}_\chi(D)$ in §3.7 agrees with the construction via broken chains here.

Example: $G = \frac{1}{30}(25, 2, 3)$

We will illustrate the unlocking procedure for G-Hilb in the case that $G = \frac{1}{30}(25, 2, 3)$. In the figures below, dashed lines are edges within a regular triangle and undashed lines are the result of the first stage of the Craw-Reid triangulation.

Figure 3.17: G-Hilb for $G = \frac{1}{30}(25, 2, 3)$

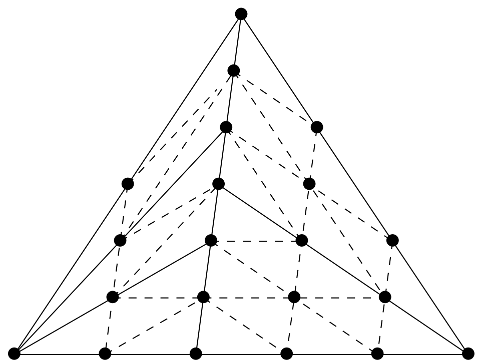

Figure 3.18: Reid's recipe for $G = \frac{1}{30}(25, 2, 3)$

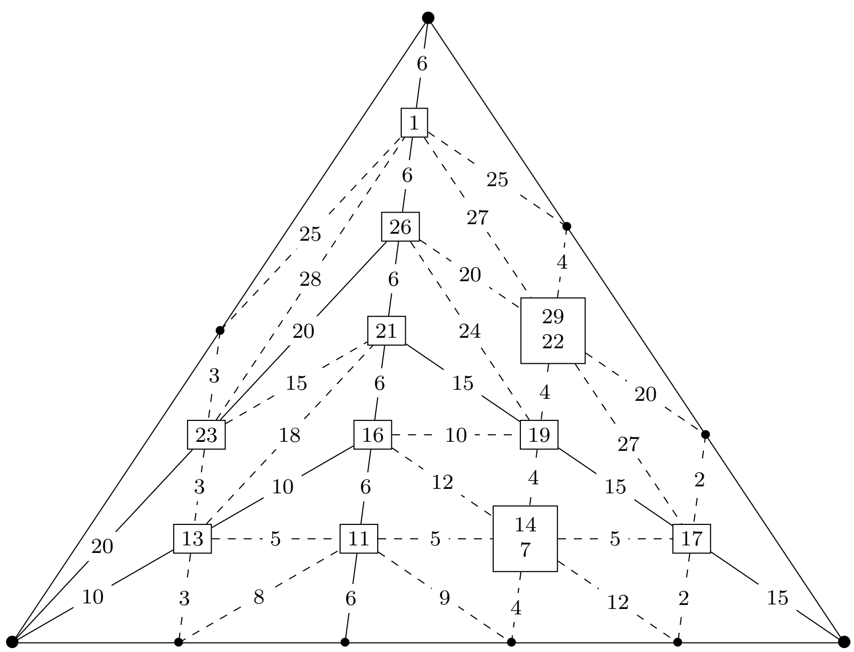

Consider the 15-curve C_{15} shown in Figure 3.19. This curve is of Type `Ixb` since after the $(-1, -1)$-curve on the left side it feeds into a boundary edge of a regular triangle. This gives

$$\text{G-ig}(C_{15}) = \{15, 17, 19, 21\}$$

Figure 3.19: Unlocking for a 15-curve

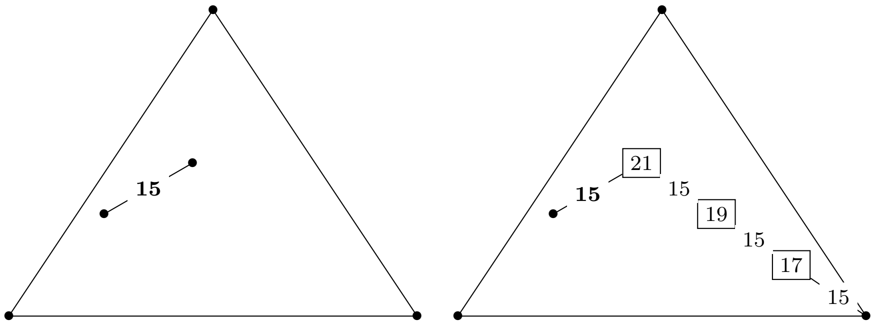

Consider the 5-curve C_5 shown in Figure 3.20. It passes into the right side of the junior simplex, unlocking the 9-curve of Type `Iy` and giving

$$\text{G-ig}(C_5) = \{5, 9, 11, 14\}$$

Figure 3.20: Unlocking for a 5-curve

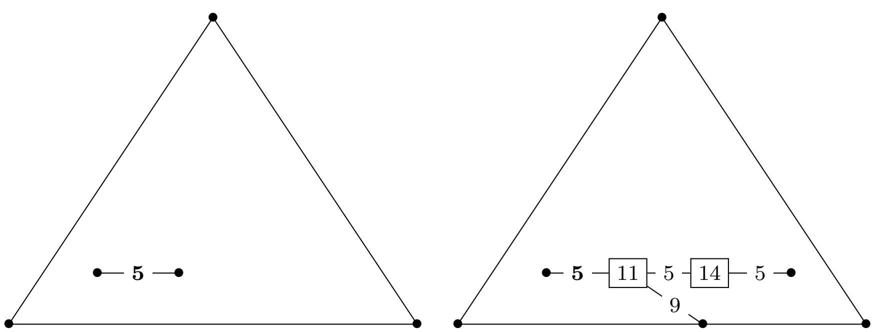

Consider the 2-curve C_2 shown in Figure 3.21. This is a curve of Type `Iz`. We first get the character 17 marking the divisor on the 2-chain, unlocking the 27-chain. The 27-chain contains a del Pezzo divisor contributing the character 22 in this case. Hence

$$\text{G-ig}(C_2) = \{2, 17, 22, 27\}$$

Figure 3.21: Unlocking for a 2-curve

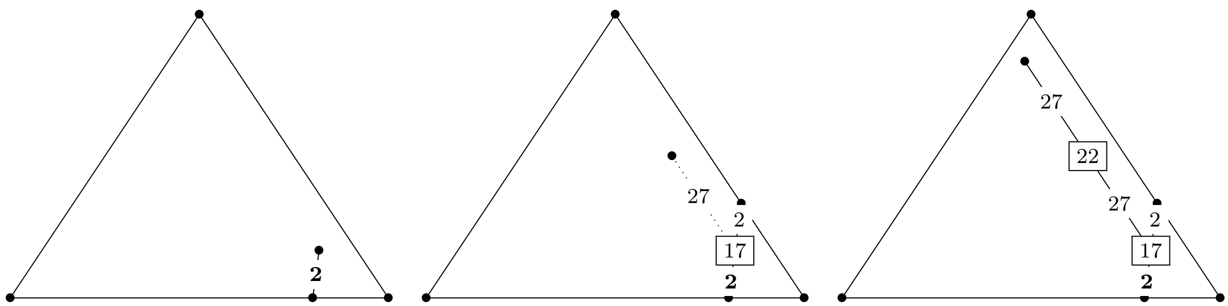

Lastly, we will consider the boundary 15-curve C'_{15} shown in Figure 3.22.

Figure 3.22: Unlocking for a boundary 15-curve

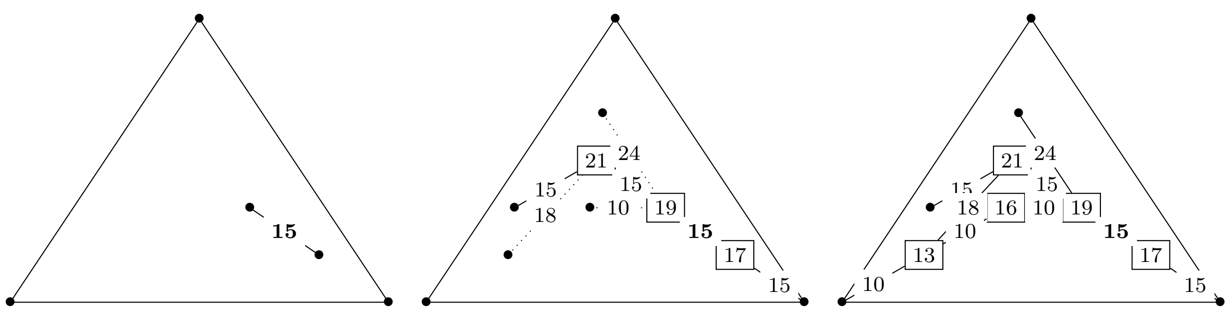

At the first step we include the 15-chain and the curves of Type `Ix` and `Iy` unlocked by it. These curves are marked with characters $10, 24, 18$. The 18-curve and the 24-curve are of Type `Iy` and only contribute their own character to the G-igsaw piece. The 10-curve is of Type `Ixb` and so we add the Hirzebruch divisors along the 10-chain. As a result

$$\text{G-ig}(C'_{15}) = \{10, 13, 15, 16, 17, 18, 19, 21, 24\}$$

Example: $G = \frac{1}{35}(1, 3, 31)$

We will use the example of $G = \frac{1}{35}(1, 3, 31)$ to illustrate a phenomenon implicit, but less clear in the long side picture. The triangulation for G-Hilb is shown in Figure 3.23. Reid's recipe is found in Figure 3.24.

Figure 3.23: G-Hilb for $G = \frac{1}{35}(1, 3, 31)$

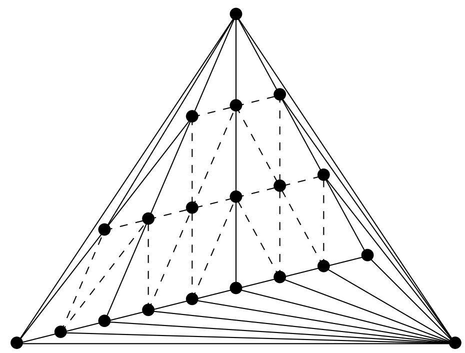

Figure 3.24: Reid's recipe for $G = \frac{1}{35}(1, 3, 31)$

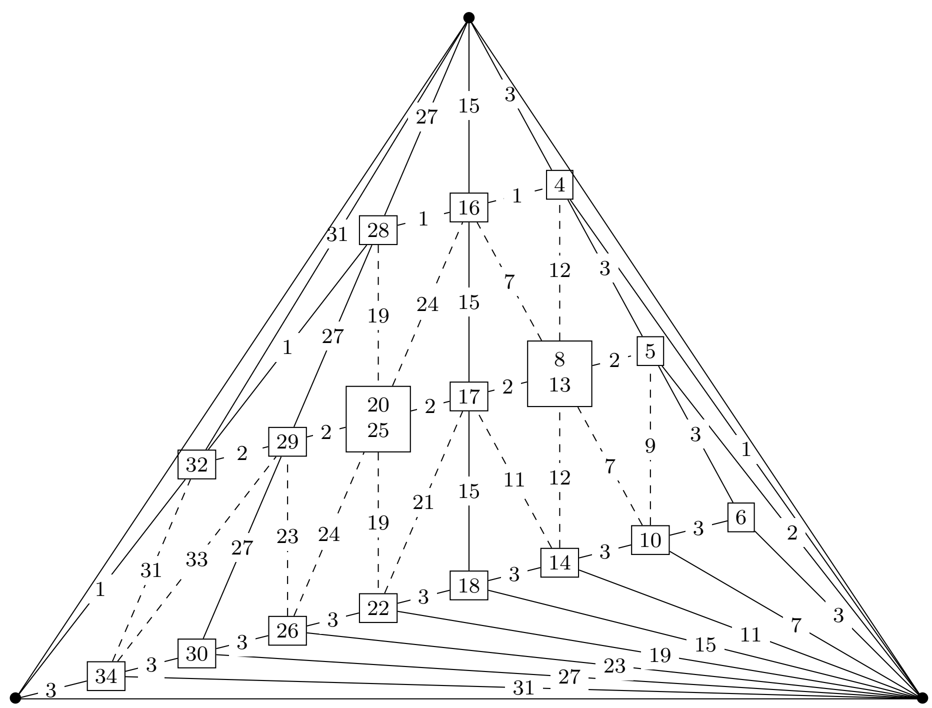

Consider the 3-curve C_3 incident to e_1. The unlocking procedure for this curve is shown in Figure 3.25 giving

$$\text{G-ig}(C_3) = \left\{ \begin{array}{c} 1, 2, 3, 4, 5, 6, 8, 9, 10, 12, 13, 14, 16, 17, 18, \\ 20, 21, 22, 24, 25, 26, 28, 29, 30, 32, 33, 34 \end{array} \right\}$$

Notice that every chain meeting the 3-chain in a vertex is broken there. Repeating for the next 3-curve along the chain produces the same unlocking sequence except that the topmost part including the 1-chain and the 12-chain are not included, capturing that the monomials in the corresponding character spaces are no longer divisible by r_3 there.

Figure 3.25: Unlocking for a 3-curve

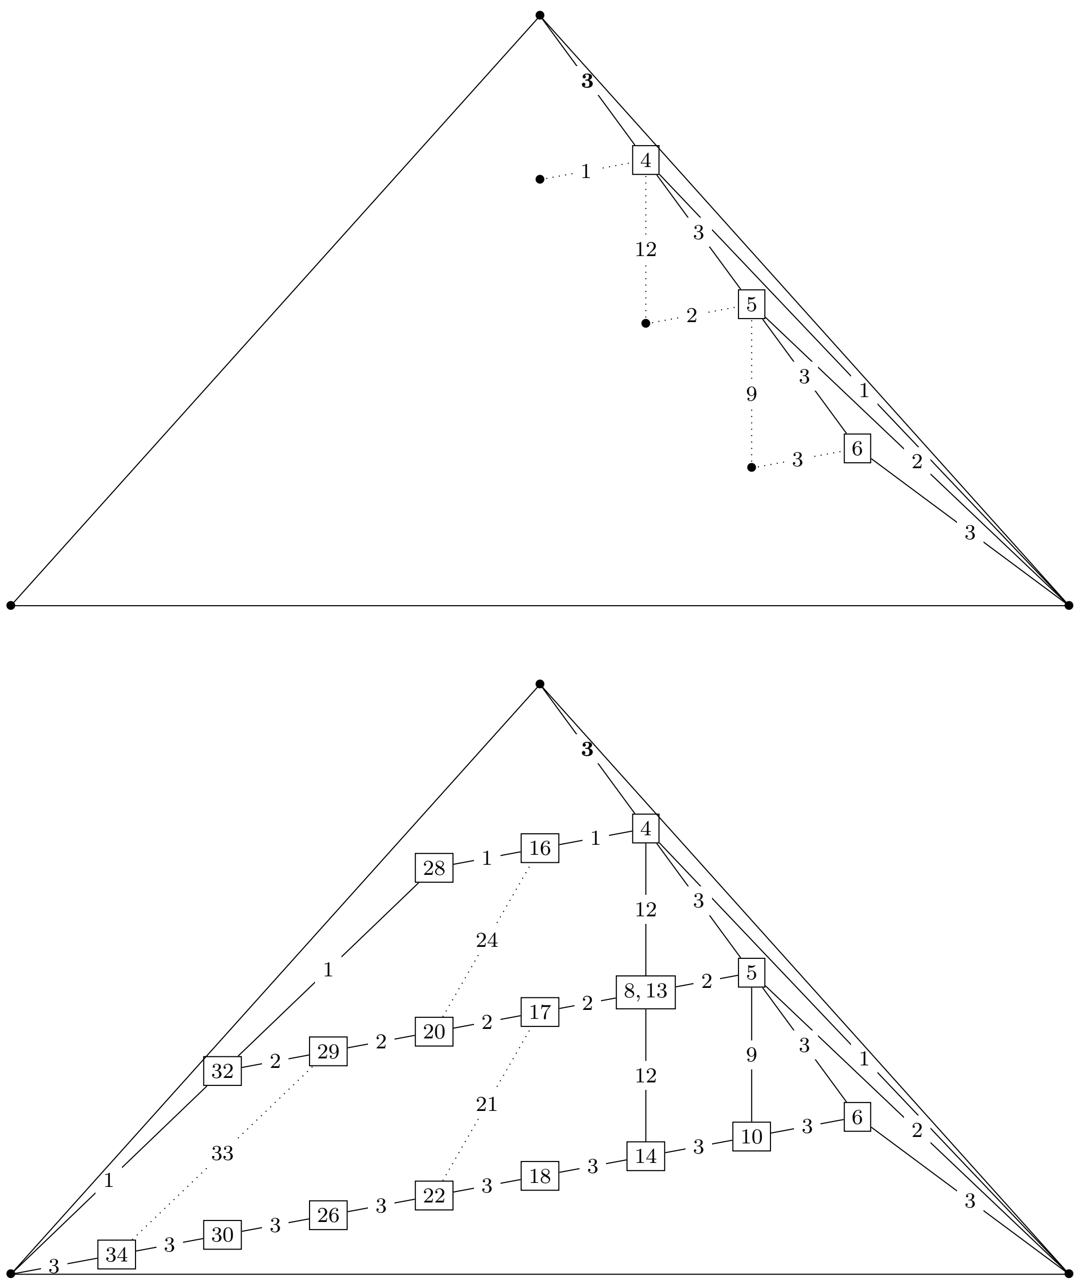

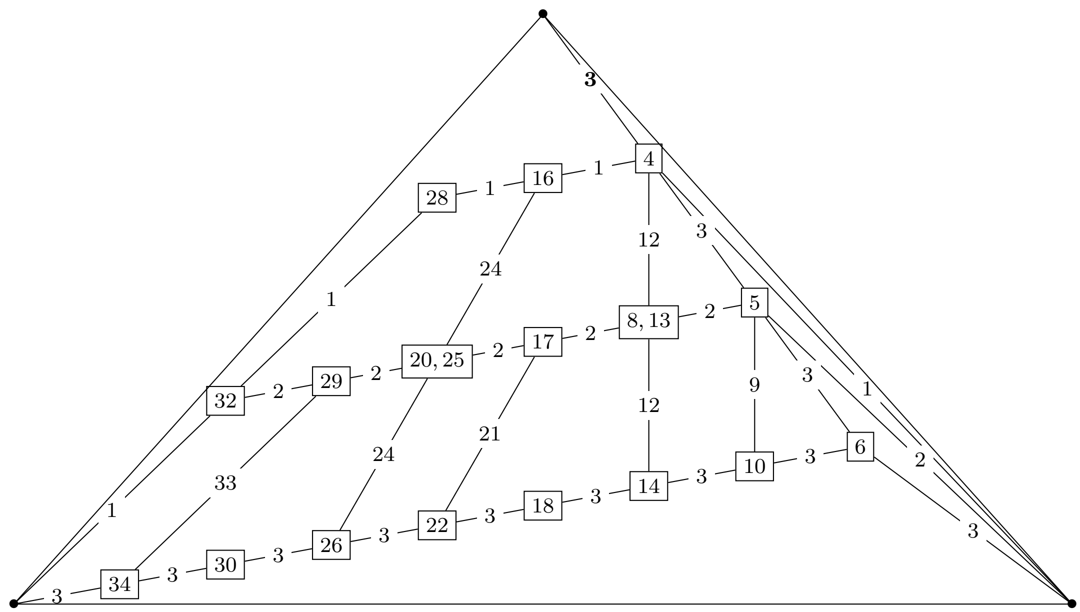

3.8 Computing the walls of $\mathfrak{C}_0$

Inequalities from curves

In [26, §9] Craw-Ishii provide an abstract description of some sufficient inequalities to carve out the chamber $\mathfrak{C}_0$. These inequalities arise from the interpretation of $\mathfrak{C}_0$ mapping to the ample cone of G-Hilb. From Kleiman's criterion, one obtains inequalities

$$\theta(\varphi_{\mathfrak{C}_0}(\mathcal{O}_C)) > 0$$

for all exceptional curves $C \subseteq G$-Hilb. If this inequality is necessary to define $\mathfrak{C}_0$, the geometry of C determines the type of the wall as follows:

- If C is a $(-1,-1)$-curve - that is, it corresponds to an interior edge inside a regular triangle - then $(\theta(\varphi_{\mathfrak{C}_0}(\mathcal{O}_C)) = 0) \cap \overline{\mathfrak{C}}_0$ is a Type `I` wall.
- If C is a $(1,-3)$-curve - that is, it corresponds to one of the edges incident to a trivalent vertex - then $(\theta(\varphi_{\mathfrak{C}_0}(\mathcal{O}_C)) = 0) \cap \overline{\mathfrak{C}}_0$ is a Type `II` wall.
- If C is contained in a Hirzebruch divisor but it is not in either of the previous cases, then $(\theta(\varphi_{\mathfrak{C}_0}(\mathcal{O}_C)) = 0) \cap \overline{\mathfrak{C}}_0$ is a Type `III` wall.

As discussed, there are no Type `II` walls and so the inequalities from the second case cannot be necessary. One can express the inequality $\theta(\varphi_{\mathfrak{C}}(\mathcal{O}_C)) > 0$ abstractly via [26, Cor. 5.2], a consequence of which is

$$\theta(\varphi_{\mathfrak{C}}(\mathcal{O}_C) = \sum_{\rho} \deg(\mathcal{R}_\rho|_C)\theta(\rho)$$

Any character ρ not in G-ig(C) has $\mathcal{R}_\rho|_C = \mathcal{O}_C$ and so it doesn't appear in the sum above. It follows that

$$\theta(\varphi_{\mathfrak{C}}(\mathcal{O}_C)) = \sum_{\rho \in G\text{-ig}(C)} \deg(\mathcal{R}_\rho|_C)\theta(\rho)$$

which empowers the unlocking procedure to compute and interpret these inequalities.

Inequalities from divisors

To complete the classification, the walls of Type `0` are obtained from divisors. Suppose F is a G-equivariant coherent sheaf on $\mathbb{C}^3$ with $H^0(F) = \mathbb{C}[G]$ as G-modules; in the language of [26] F is a *G-constellation*, which generalises the notion of G-cluster. Suppose $\mathfrak{w}$ is a Type `0` wall of $\mathfrak{C}_0$ with unstable locus D. Let $\theta \in \mathfrak{w}$ and suppose that $S \subseteq F$ is a nontrivial subsheaf such that $\theta(S) = 0$. That is, F is θ-destabilised by S. [26, Cor. 4.6 + Remark 4.7] says that either S or $Q = F/S$ is rigid and so defines a constant family on D. It follows that D is exactly the divisor parameterising such rigid sub- or quotient sheaves. From [26, Cor. 5.6 + Theorem 9.5] the equation of a type `0` wall $\mathfrak{w}$ is one of the following:

- if $\mathfrak{w}$ comes from a divisor D parameterising rigid subsheaves, then $D = D_\psi$ is irreducible and the inequality defining the wall is $\theta(\varphi_{\mathfrak{C}_0}(\mathcal{R}_\psi^{-1}|_D) = \theta(\psi) > 0$.
- if $\mathfrak{w}$ comes from a divisor D' parameterising rigid quotient sheaves, then D' is connected but potentially reducible and the inequality defining the wall is $\theta(\omega_{D'}) = \theta(Q) < 0$ where Q is the representation defined by a quotient sheaf in D'.

We can be more precise in the second case. For the representation Q to be constant across D, it means that $\mathcal{R}_\rho|_{D'}$ is trivial for all $\rho \subseteq Q$. Equivalently, all torus-invariant G-clusters in D' share the same eigenmonomial r_ρ for each $\rho \subseteq Q$ or, also equivalently, $\rho \notin$ G-ig(C) for any $C \subseteq D'$.

Proposition 3.8.1. *Suppose D' is a (possibly reducible) divisor in G-*Hilb $\mathbb{C}^3$*. Then the inequality for the rigid quotient parameterised by D' is*

$$\theta(\varphi_{\mathfrak{C}_o}(\omega_{D'}^\vee)) = \sum_{C \subseteq D'} \sum_{\chi \in \text{G-ig}(C)} \theta(\chi) > 0$$

Proof. This follows since reversing the inequality $\theta(Q) < 0$ gives

$$\theta(\varphi_{\mathfrak{C}_0}(\omega_{D'})) = \theta(Q) < 0 \Longrightarrow \theta(\varphi_{\mathfrak{C}_0}(\omega_{D'}^{\vee})) = \theta(\mathbb{C}[G]/Q) > 0$$

and $\mathbb{C}[G]/Q$ contains exactly the characters in the statement of the proposition. □

Type I walls

We know from [26, Theorem 9.12] that all flops in a single $(-1,-1)$-curve C are achieved by a wall-crossing from $\mathfrak{C}_0$. Moreover, we have $\deg(\mathcal{R}_\rho|_C) = 1$ for all $\rho \in$ G-ig(C) from [26, Cor. 6.3]. The unlocking procedure hence gives a combinatorial way of writing down the equations of these walls.

Proposition 3.8.2. *Suppose $C \subseteq G$-Hilb $\mathbb{A}^3$ is an exceptional $(-1,-1)$-curve marked with character χ by Reid's recipe. Then, the Type I wall corresponding to C is given by*

$$\theta(\varphi_{\mathfrak{C}_0}(\mathcal{O}_C)) = \sum_{\chi \in \text{G-ig}(C)} \theta(\chi) = 0$$

where G-ig(C) *is computed by the unlocking procedure.*

No Type II walls

Proposition 3.8.3. *Suppose $C \subseteq G$-Hilb $\mathbb{A}^3$ is an exceptional $(1,-3)$-curve marked with character χ by Reid's recipe. Then, the inequality corresponding to C is given by*

$$\theta(\varphi_{\mathfrak{C}_0}(\mathcal{O}_C)) = 2 \cdot \theta(\chi^{\otimes 2}) + \sum_{\chi \in \text{G-ig}(C) \setminus \{\chi^{\otimes 2}\}} \theta(\chi) = 0$$

where G-ig(C) *is computed by the unlocking procedure.*

Proof. Notice that such a curve C lies inside the exceptional $\mathbb{P}^2$ in the meeting of champions case when the meeting of champions triangle has side length 0. Thus the $\mathbb{P}^2$ is marked with $\chi^{\otimes 2}$ and lies in the socle of any torus-invariant G-cluster. From Theorem 3.6.1 $r_{\chi^{\otimes 2}} = r_\chi^2$ and so r_χ^2 is the furthest character from r_χ in the G-igsaw piece in some direction. Note that

$$\deg(\mathcal{R}_\rho|_C) = \min\{k : r_\chi^k \mid r_\rho\}$$

and so all the characters in G-ig(C) appear with multiplicity 1 except for r_χ^2, which appears with multiplicity 2. This gives the required formula. □

As a result we can immediately deduce the conclusion of [26, Prop. 3.8] for $\mathfrak{C}_0$.

Corollary 3.8.4. *$\mathfrak{C}_0$ has no Type* `II` *walls.*

Proof. Suppose C is an exceptional $(1,-3)$-curve marked with χ. From Prop. 3.8.3 a G-igsaw piece for C consists of χ, χ^2, and the characters marking the (Hirzebruch) divisors along the χ-chain. Let D' be the exceptional $\mathbb{P}^2$ containing C. Consider the inequality for rigid quotients parameterised by D': from Prop. 3.8.1 the characters appearing in this inequality are exactly the characters in the G-igsaw pieces of all three χ-curves converging at D'. These are

$$\{\chi, \chi^{\otimes 2}\} \cup \mathrm{Hirz}(\chi)$$

which are exactly the characters appearing in the inequality for C. However, the inequality for rigid quotients parameterised by D' has multiplicities all equal to 1. When combined with the inequality $\theta(\chi^{\otimes 2}) > 0$ coming from rigid subsheaves parameterised by D' this implies that the inequality $\varphi_{\mathfrak{C}_0}(\mathcal{O}_C) > 0$ is redundant. □

All flops in $(-1,-1)$-curves

Using Prop. 3.8.2 and the unlocking procedure one can show directly that every $(-1,-1)$-curve produces a necessary inequality, recovering [26, Theorem 9.12] by purely combinatorial means.

Proposition 3.8.5. *Suppose C is an exceptional $(-1,-1)$ curve inside G-Hilb $\mathbb{C}^3$. Then the inequality $\theta(\varphi_{\mathfrak{C}_0}(\mathcal{O}_C)) > 0$ is necessary and so defines a wall of $\mathfrak{C}_0$.*

Proof. Suppose C is marked with χ. From the unlocking procedure we can write the inequality corresponding to C in the form

$$\theta(\chi) + \sum_i \theta(\psi_i) + \theta(\rho_1) + \sum_i \theta(\psi_i^1) + \cdots + \theta(\rho_m) + \sum_i \theta(\psi_i^m) > 0 \tag{3.4}$$

where ρ_j are the characters marking curves C_j unlocked by C and ψ_i^j are the characters in the G-igsaw piece for C_j. Note that curves unlocked by C cannot continue on both sides of the χ-chain, since they meet the χ-chain at a Hirzebruch divisor found at the intersection of the χ-chain and an edge of a regular triangle, where only two chains can continue. The inequality for the $(-1,-1)$-curve C_j is

$$\theta(\varphi_{\mathfrak{C}_0}(\mathcal{O}_{C_j})) = \theta(\rho_j) + \sum_i \theta(\psi_i^j) > 0$$

In order to express (3.4) in terms of other inequalities, we must have an inequality featuring the character χ. These can only arise from other χ-curves or divisors parameterising rigid quotients not featuring χ. Other χ-curves will feature at least one different character in their G-igsaw piece compared to G-ig(C): indeed, other curves in the same regular triangle will feature a different collection of del Pezzo divisors, curves in other regular triangles will either feature different del Pezzo divisors or unlock different curves, and χ-curves along a boundary edge will have different unlocking behaviour. In particular, the inequalities from these curves will not be summands of the inequality (3.4). Inequalities from rigid quotients not containing χ will also not be summands of (3.4) since the unlocking procedure implies that there are no divisors D_ρ along the χ-chain for which all characters marking curves incident to D_ρ are represented in G-ig(C). It follows that (3.4) is necessary. $\square$

Irredundant inequalities - examples

The aim of these final sections is to precisely describe all the walls of $\mathfrak{C}_0$. We start with an accessible example.

Example 3.8.6. *Consider $G = \frac{1}{6}(1, 2, 3)$. G-Hilb and Reid's recipe are shown in Figure 3.26. We compute the inequalities coming from curves and divisors that define $\mathfrak{C}_0$ via the*

Figure 3.26: G-Hilb and Reid's recipe for $\frac{1}{6}(1, 2, 3)$

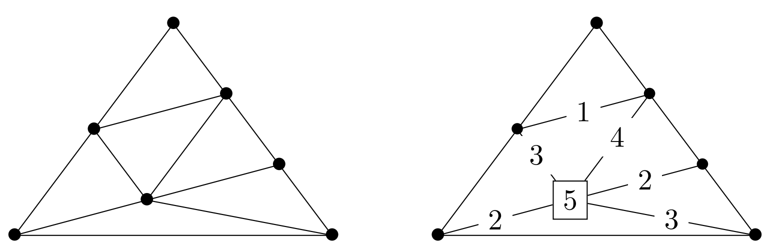

unlocking procedure.

$$\theta(\chi_1) > 0 \tag{A_1}$$
$$\theta(\chi_2) + \theta(\chi_5) > 0 \tag{A_2}$$
$$\theta(\chi_2) + \theta(\chi_3) + 2\theta(\chi_4) + 2\theta(\chi_5) > 0 \tag{B_2}$$
$$\theta(\chi_3) + \theta(\chi_5) > 0 \tag{A_3}$$
$$\theta(\chi_3) + \theta(\chi_4) + \theta(\chi_5) > 0 \tag{B_3}$$
$$\theta(\chi_4) > 0 \tag{A_4}$$
$$\theta(\chi_5) > 0 \tag{A_5}$$
$$\theta(\chi_2) + \theta(\chi_3) + \theta(\chi_4) + \theta(\chi_5) > 0 \tag{B_5}$$

(A_1) is from the curve marked with the essential character 1*. Similarly for (A_4). We then have two inequalities (A_2) and (B_2) coming from the two* 2*-curves, and two (A_3) and (B_3) from the two* 3*-curves. The* 5*-divisor gives two inequalities (A_5) and (B_5) for rigid subsheaves and quotients it parameterises.*

We can see that (B_2) is redundant by expressing it as a combination of (A_2), (A_3) and (A_4). Similarly, (B_3) can be expressed in terms of (A_3) and (A_4). No further reductions are possible, and so the walls of $\mathfrak{C}_0$ (with their types) in this example are:

$$\theta(\chi_1) = 0 \qquad \texttt{(I)}$$

$$\theta(\chi_2) + \theta(\chi_5) = 0 \qquad \texttt{(III)}$$

$$\theta(\chi_3) + \theta(\chi_5) = 0 \qquad \texttt{(I)}$$

$$\theta(\chi_4) = 0 \qquad \texttt{(I)}$$

$$\theta(\chi_5) = 0 \qquad \texttt{(0)}$$

$$\theta(\chi_2) + \theta(\chi_3) + \theta(\chi_4) + \theta(\chi_5) = 0 \qquad \texttt{(0)}$$

Example 3.8.7. *We continue with a more detailed example for $G = \frac{1}{30}(25, 2, 3)$. Continuing the calculations in §3.7, we find that the inequalities from curves in G-Hilb are:*

$$\theta_2 + \theta_{27} + \theta_{22} + \theta_{17} > 0 \qquad (A_2)$$

$$\theta_2 + \theta_5 + \theta_8 + \theta_{11} + \theta_{14} > 0 \qquad (B_2)$$

$$\theta_3 + \theta_{13} + \theta_{18} + \theta_{23} + \theta_{28} > 0 \qquad (A_3)$$

$$\theta_3 + \theta_5 + \theta_7 + \theta_9 + \theta_{11} + \theta_{13} + \theta_{23} + \theta_{28} > 0 \qquad (B_3)$$

$$\theta_3 + \theta_5 + \theta_7 + \theta_9 + \theta_{11} + \theta_{13} + \theta_{15} + \theta_{17} + \theta_{19} + \theta_{21} > 0 \qquad (C_3)$$

$$\theta_4 + \theta_{29} + \theta_{24} + \theta_{19} + \theta_{14} > 0 \qquad (A_4)$$

$$\theta_4 + \theta_7 + \theta_{29} + \theta_{24} + \theta_{19} > 0 \qquad (B_4)$$

$$\theta_4 + \theta_7 + \theta_{10} + \theta_{13} + \theta_{16} + \theta_{19} + \theta_{29} > 0 \qquad (C_4)$$

$$\theta_4 + \theta_7 + \theta_{10} + \theta_{13} + \theta_{16} + \theta_{19} + \theta_{22} > 0 \qquad (D_4)$$

$$\theta_5 + \theta_7 + \theta_9 + \theta_{11} > 0 \qquad (A_5)$$

$$\theta_5 + \theta_7 + \theta_8 + \theta_{11} > 0 \qquad (B_5)$$

$$\theta_5 + \theta_8 + \theta_{11} + \theta_{14} > 0 \qquad (C_5)$$

$$\begin{matrix}\theta_6+\theta_8+\theta_9+\theta_{10}+\theta_{11}+\theta_{13}+2\theta_{12}+2\theta_{14}+2\theta_{16}+2\theta_{15}+2\theta_{17}+2\theta_{19}+3\theta_{18}\\ +3\theta_{20}+3\theta_{22}+3\theta_{21}+3\theta_{23}+3\theta_{25}+4\theta_{24}+4\theta_{26}+4\theta_{28}+4\theta_{27}+4\theta_{29}+4\theta_1\end{matrix}>0 \quad (\mathbf{A_6})$$

$$\begin{matrix}\theta_6+\theta_8+\theta_{10}+2\theta_{12}+2\theta_{14}+2\theta_{16}+3\theta_{18}+\theta_9+\theta_{11}+\theta_{13}+2\theta_{15}+2\theta_{17}\\ +2\theta_{19}+3\theta_{21}+\theta_1+4\theta_{26}++2\theta_{16}+3\theta_{23}+3\theta_{20}+3\theta_{22}+4\theta_{24}+4\theta_{26}\end{matrix}>0 \quad (\mathbf{B_6})$$

$$\begin{matrix}\theta_6+\theta_8+\theta_{10}+2\theta_{12}+2\theta_{14}+2\theta_{16}+3\theta_{18}+\theta_9+\theta_{11}+\theta_{13}\\ +2\theta_{15}+2\theta_{17}+2\theta_{19}+3\theta_{21}+\theta_1+\theta_{26}+\theta_{21}+\theta_{16}\end{matrix}>0 \quad (\mathbf{C_6})$$

$$\theta_6+\theta_8+\theta_{10}+2\theta_{12}+2\theta_{14}+2\theta_{16}+\theta_1+\theta_{26}+\theta_{21}+\theta_{16}+\theta_9+\theta_{11}+\theta_{13}>0 \quad (\mathbf{D_6})$$

$$\theta_6+\theta_1+\theta_{26}+\theta_{21}+\theta_{16}+\theta_{11}+\theta_8+\theta_9>0 \quad (\mathbf{E_6})$$

$$\theta_6+\theta_1+\theta_{26}+\theta_{21}+\theta_{16}+\theta_{11}>0 \quad (\mathbf{F_6})$$

$$\theta_8>0 \quad (A_8)$$

$$\theta_9>0 \quad (A_9)$$

$$\theta_{10}+\theta_{13}+\theta_{16}>0 \quad (A_{10})$$

$$\theta_{10}+\theta_{12}+\theta_{14}+\theta_{16}+\theta_{18}+\theta_5+\theta_7+\theta_9+\theta_{11}+\theta_{13}>0 \quad (\mathbf{B_{10}})$$

$$\theta_{10}+\theta_{13}+\theta_{12}+\theta_{14}+\theta_{16}>0 \quad (\mathbf{C_{10}})$$

$$\theta_{12}+\theta_7>0 \quad (A_{12})$$

$$\theta_{12}+\theta_{14}>0 \quad (B_{12})$$

$$\theta_{15}+\theta_{17}+\theta_{19}+\theta_{21}>0 \quad (A_{15})$$

$$\theta_{15}+\theta_{17}+\theta_{19}+\theta_{18}+\theta_{21}>0 \quad (\mathbf{B_{15}})$$

$$\theta_{15}+\theta_{17}+\theta_{18}+\theta_{21}+\theta_{24}+\theta_{10}+\theta_{13}+\theta_{16}+\theta_{19}>0 \quad (\mathbf{C_{15}})$$

$$\theta_{15}+\theta_{18}+\theta_{21}+\theta_{24}+\theta_{27}+\theta_{10}+\theta_{13}+\theta_{16}+\theta_{19}+\theta_{22}+\theta_5+\theta_8+\theta_{11}+\theta_{14}+\theta_{17}>0 \quad (\mathbf{D_{15}})$$

$$\theta_{18} > 0 \tag{A_{18}}$$

$$\theta_{20} + \theta_{23} + \theta_{26} + \theta_{29} > 0 \tag{A_{20}}$$

$$\theta_{20} + \theta_{22} + \theta_{23} + \theta_{26} > 0 \tag{B_{20}}$$

$$\theta_{20} + \theta_{23} + \theta_{22} + \theta_{24} + \theta_{26} > 0 \tag{$\mathbf{C_{20}}$}$$

$$\theta_{20} + \theta_{15} + \theta_{17} + \theta_{19} + \theta_{21} + \theta_{22} + \theta_{24} + \theta_{26} + \theta_{28} > 0 \tag{$\mathbf{D_{20}}$}$$

$$\theta_{24} > 0 \tag{A_{24}}$$

$$\theta_{25} + \theta_{27} + \theta_{29} + \theta_1 > 0 \tag{A_{25}}$$

$$\theta_{25} + \theta_{28} + \theta_1 > 0 \tag{B_{25}}$$

$$\theta_{27} + \theta_{22} > 0 \tag{A_{27}}$$

$$\theta_{27} + \theta_{29} > 0 \tag{B_{27}}$$

$$\theta_{28} > 0 \tag{A_{28}}$$

The bolded inequalities correspond to curves C with $\mathcal{N}_C$ not of type $(-1,-1)$. We know by [26, Theorem 9.12] that the other inequalities are necessary and define Type `I` *walls of $\mathfrak{C}_0$. The inequalities from divisors parameterising rigid subsheaves are:*

$$\theta_1 > 0 \tag{A_1}$$

$$\theta_7 > 0 \tag{A_7}$$

$$\theta_{11} > 0 \tag{A_{11}}$$

$$\theta_{13} > 0 \tag{A_{13}}$$

$$\theta_{14} > 0 \tag{A_{14}}$$

$$\theta_{16} > 0 \tag{A_{16}}$$

$$\theta_{17} > 0 \tag{A_{17}}$$

$$\theta_{19} > 0 \tag{A_{19}}$$

$$\theta_{21} > 0 \tag{A_{21}}$$

$$\theta_{22} > 0 \tag{A_{22}}$$

$$\theta_{23} > 0 \tag{A_{23}}$$

$$\theta_{26} > 0 \tag{A_{26}}$$

$$\theta_{29} > 0 \tag{A_{29}}$$

We record the redundancies for the bold (or potentially redundant) inequalities.

$$\begin{matrix} (\mathbf{F_6}) + (A_8) + (A_9) + (A_{10}) + (B_{12}) + (A_{15}) \\ +(A_{18}) + (A_{20}) + (A_{24}) + (A_{25}) + (B_{27}) + (A_{28}) \end{matrix} \implies (\mathbf{A_6})$$

$$(\mathbf{F_6}) + (A_8) + (A_9) + (A_{10}) + (B_{12}) + (A_{15}) + (A_{18}) + (A_{20}) + (A_{24}) \implies (\mathbf{B_6})$$

$$(\mathbf{F_6}) + (A_8) + (A_9) + (A_{10}) + (B_{12}) + (A_{15}) + (A_{18}) \implies (\mathbf{C_6})$$

$$(\mathbf{F_6}) + (A_8) + (A_9) + (A_{10}) + (B_{12}) \implies (\mathbf{D_6})$$

$$(\mathbf{F_6}) + (A_8) + (A_9) \implies (\mathbf{E_6})$$

$$(A_5) + (B_{12}) + (A_{18}) \implies (\mathbf{B_{10}})$$

$$(A_{10}) + (B_{12}) \implies (\mathbf{C_{10}})$$

$$(A_{15}) + (A_{18}) \implies (\mathbf{B_{15}})$$

$$(A_{15}) + (A_{18}) + (A_{10}) + (A_{24}) \implies (\mathbf{C_{15}})$$

$$(A_{15}) + (A_{18}) + (A_{10}) + (A_{24}) + (A_{27}) + (C_5) \implies (\mathbf{D_{15}})$$

$$(B_{20}) + (A_{24}) \implies (\mathbf{C_{20}})$$

$$(A_{15}) + (B_{20}) + (A_{24}) \implies (\mathbf{D_{20}})$$

We have killed off the inequalities from all curves except for the $(-1,-1)$*-curves and one curve* $(\mathbf{F_6})$ *from the long side.*

Irredundant inequalities from curves

Observe that the vast majority of inequalities in Examples 3.8.6-3.8.7 define walls of Type $\mathtt{I}$. We should be unsurprised by the cancellation of all except one bolded inequality in Example 3.8.7 due to the following result from [26].

Lemma 3.8.8 ([26, Corollaries 6.3 & 6.5])**.** *Suppose* $w = (\sum \alpha_i \theta_i = 0)$ *is a Type* $\mathtt{I}$ *or* $\mathtt{III}$ *wall of* $\mathfrak{C}_0$*. Then all* $\alpha_i \in \{0, 1\}$*.*

Chambers other than $\mathfrak{C}_0$ can have coefficients $\alpha_i = -1$, however since the trivial representation appears in no G-igsaw piece we can exclude this possibility.

Corollary 3.8.9. *Suppose* $G\text{-Hilb}\,\mathbb{C}^3$ *has a meeting of champions of side length* 0*. Then the inequality for any curve along one of the three champions is redundant.*

Proof. Suppose χ is the character marking each of the champions. Then, by Theorem 3.6.1, $r_\chi^2 = r_{\chi^2}$ globally on G-Hilb and so $\deg(\mathcal{R}_{\chi^2}|_C) = 2$ for all χ-curves C. It follows from Lemma 3.8.8 that none of these inequalities can be strict. □

We can also show this directly via unlocking. This reproves Cor. 3.8.4.

Lemma 3.8.10. *Suppose C is a χ-curve. If the unlocking procedure for C doesn't unlock a curve or divisor marked with χ^2 then all the coefficients in the inequality $\theta(\varphi_{\mathfrak{C}_0}(\mathcal{O}_C)) > 0$ are equal to 0 or 1.*

Proof. This is because if some ρ has $\deg(\mathcal{R}_\rho|_C) \geq 2$ then $r_\chi^2 \mid r_\rho$ and so r_χ^2 must feature in the G-igsaw piece for C and is hence equal to r_{χ^2} near C. □

Lemma 3.8.11. *Suppose a curve C_0 unlocks a curve C_1 of character ρ. Let $\psi \in$ G-ig(C_1). If C is a curve that unlocks C_0, then $\deg(\mathcal{R}_\psi|_C) \geq \deg(\mathcal{R}_\rho|_C)$.*

Proof. As used previously, $\deg(\mathcal{R}_\rho|_C) = \max\{k \in \mathbb{Z}_{\geq 0} : r_\chi^k \mid r_\rho\}$. From this formulation, clearly if $r_\rho \mid r_\psi$ then $\deg(\mathcal{R}_\psi|_C) \geq \deg(\mathcal{R}_\rho|_C)$, but this is the case by definition of G-igsaw piece. □

We say that an inequality $\sum_i \alpha_i \theta(\chi_i) > 0$ with nonnegative coefficients is a *summand* of another inequality $\sum_j \beta_j \theta(\rho_j) > 0$ with nonnegative coefficients if the difference $\sum_i \alpha_i \theta(\chi_i) - \sum_j \beta_j \theta(\rho_j)$ also has nonnegative coefficients in the basis $\operatorname{Irr} G$. If an inequality coming from curves or divisors decomposes into other inequalities as summands, then it is redundant and does not define a wall of $\mathfrak{C}_0$.

Lemma 3.8.12. *Suppose C is a curve on the boundary of a regular triangle marked with a character χ. Suppose the χ-chain contains a $(-1,-1)$-curve. Then the inequality $\theta(\varphi_{\mathfrak{C}_0}(\mathcal{O}_C)) > 0$ is redundant.*

Proof. Suppose C is marked with character χ. Let C_0 be the first $(-1,-1)$-curve in the χ-chain moving inwards from C. Then the G-igsaw piece for C consists of exactly the characters in the G-igsaw piece for C_0 along with the characters in the G-igsaw pieces for any curves $C_1, \ldots, C_n$ unlocked by C at Hirzebruch divisors before C_0. Let the character marking C_i be χ_i. The inequality for C decomposes as

$$\theta(\varphi_{\mathfrak{C}_0}(\mathcal{O}_C)) = \sum_{\rho \in \text{G-ig}(C_0)} \alpha_\rho \theta(\rho) + \sum_{i=1}^{n} \sum_{\rho \in \text{G-ig}(C_i)} \beta_\rho^i \theta(\rho) \tag{3.5}$$

where α_ρ and β_ρ^i are nonnegative multiplicities given by the appropriate calculation of $\deg(\mathcal{R}_\rho|_?)$. Note that $\alpha_\chi = 1$. One can thus write

$$\begin{aligned}\theta(\varphi_{\mathfrak{C}_0}(\mathcal{O}_C)) = \theta(\varphi_{\mathfrak{C}_0}(\mathcal{O}_{C_0})) &+ \sum_{\rho \in \text{G-ig}(C_0)} (\alpha_\rho - 1)\theta(\rho) \\ &+ \sum_{i=1}^{m} \left(\beta_{\chi_i}^i \theta(\varphi_{\mathfrak{C}_0}(\mathcal{O}_{C_i})) + \sum_{\rho \in \text{G-ig}(C_i)} (\beta_\rho^i - \beta_{\chi_i}^i)\theta(\rho) \right)\end{aligned}$$

From Lemma 3.8.11, $\alpha_\rho - 1$ and $\beta_\rho - \beta_{\chi'}$ are both nonnegative. If all the remaining ρ in these sums with nonzero coefficients are characters marking divisors then one can express each term $\gamma_\rho\theta(\rho) = \gamma_\rho\theta(\varphi_{\mathfrak{C}_0}(\mathcal{R}_\psi^{-1}|_D))$ for some divisor D, thus evidencing that (3.5) is redundant. Suppose instead that some $\rho = \rho_1$ marks a curve unlocked by C_0 or some C_i. We assume the latter; the former is treated identically. Denote this new curve by $C_{i,1}$. Then

$$\begin{aligned}\sum_{\rho\in\text{G-ig}(C_i)} &(\beta^i_\rho - \beta^i_{\chi_i})\theta(\rho) \\ &= (\beta^i_{\rho_1} - \beta^i_{\chi_i})\theta(\varphi_{\mathfrak{C}_0}(\mathcal{O}_{C_{i,1}})) + \sum_{\rho\in\text{G-ig}(C_{i,1})}(\beta^i_\rho - \beta^i_{\rho_1})\theta(\rho) + \sum_{\rho\notin\text{G-ig}(C_{i,1})}(\beta^i_\rho - \beta^i_{\chi_i})\theta(\rho)\end{aligned}$$

where again each coefficient is nonnegative by Lemma 3.8.11 applied to $C_{i,1}$. Observe that there are strictly fewer nonzero coefficients in this expression than before, since at the least we removed the term for ρ_1. Continuing in this way for each character appearing that marks a curve, we can reduce to the situation where the only characters with nonzero coefficients in the error term are those that mark divisors. At that point we have already seen how to express the error term in terms of inequalities coming from divisors, and so we have shown that (3.5) is redundant. □

Definition 3.8.13. *Let χ be a character marking a curve in G-Hilb. We say that the χ-chain is a* generalised long side *if it starts and ends on the boundary of the junior simplex, and all the edges along the χ-chain are boundary edges of regular triangles. We exclude the lines meeting at a trivalent vertex if there is a meeting of champions of side length* 0 *from this definition.*

For example, any long side is a generalised long side. The 15-chain for $\frac{1}{35}(1,3,31)$ is a generalised long side as can be seen in Figure 3.24.

Example 3.8.14. *We compute the inequalities for curves along the* 15*-chain in G-Hilb for $G = \frac{1}{35}(1,3,31)$. From the unlocking procedure or computing G-igsaw pieces directly, the inequalities for the* 15*-curves starting from* e_1 *and moving downwards are*

$$\theta_{15} + \theta_{18} + \theta_{21} + \theta_{24} + \theta_7 + \theta_{10} + \theta_{13} + \theta_{16} + \theta_{11} + \theta_{14} + \theta_{17} + \theta_{20} > 0 \qquad (A_{15})$$

$$\theta_{15} + \theta_{18} + \theta_{21} + \theta_{16} + \theta_{11} + \theta_{14} + \theta_{17} > 0 \qquad (B_{15})$$

$$\theta_{15} + \theta_{16} + \theta_{17} + \theta_{18} > 0 \qquad (C_{15})$$

$$\theta_{15} + \theta_{16} + \theta_{17} + \theta_{18} > 0 \qquad (D_{15})$$

Clearly (C_{15}) and (D_{15}) depend on each other; the inequality is the same since they are fibres of the $\mathbb{P}^1$-bundle structure on the Hirzebruch surface marked with 18*, and so contracting* *one must contract the other. We consider some of the additional inequalities coming from* $(-1,-1)$*-curves:*

$$\theta_7 + \theta_{10} + \theta_{13} > 0 \tag{A_7}$$
$$\theta_{11} + \theta_{14} > 0 \tag{A_{11}}$$
$$\theta_{21} > 0 \tag{A_{21}}$$
$$\theta_{24} + \theta_{20} > 0 \tag{A_{24}}$$

We can deduce

$$(C_{15}) + (A_7) + (A_{11}) + (A_{21}) + (A_{24}) \Longrightarrow (A_{15})$$
$$(C_{15}) + (A_{11}) + (A_{21}) \Longrightarrow (B_{15})$$

so that (A_{15}) *and* (B_{15}) *are redundant.*

Consider a generalised long side marked with character χ. Recall that each χ-chain consists of potentially several straight line segments. We call a curve in the χ-chain *final* if it is the furthest curve along the χ-chain away from a vertex along such a line segment. For example, for $G = \frac{1}{35}(1, 3, 31)$, the dashed curves in Figure 3.27 are final.

Figure 3.27: Final curves for $G = \frac{1}{35}(1, 3, 31)$

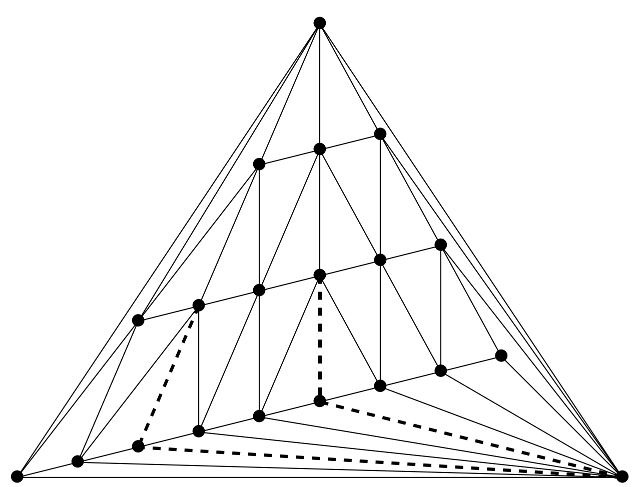

Final curves not along a long side are also those contained in an exceptional Hirzebruch surface (with no blowups) or, equivalently, those corresponding to edges incident to a 4-valent vertex. There can be at most two final curves for each generalised long side, with exactly one when the generalised long side is actually a long side.

Lemma 3.8.15. *Suppose χ is a character marking a curve and that the χ-chain is a generalised long side. Then, the inequality for each non-final curve C in the χ-chain is is redundant. The final curves all produce the same inequality:*

$$\theta(\chi) + \sum_{\psi \in \mathrm{Hirz}(\chi)} \theta(\psi) > 0$$

which is a necessary inequality defining a Type `III` *wall of $\mathfrak{C}_0$.*

Proof. First, the inequality for a final χ-curve C features only the Hirzebruch divisors along the χ-chain by the unlocking procedure. It has all nonzero coefficients equal to 1 for the following reason. χ^2 cannot mark a Hirzebruch divisor along the χ-chain because to do so one would require another chain, say with character ρ, to cross the χ-chain and have $\chi \otimes \rho = \chi^2$. Of course, this would mean that $\rho = \chi$, but chains do not self-intersect. Hence, χ^2 does not appear in the G-igsaw piece for C and so all multiplicities must be equal to 1 by Lemma 3.8.10. This is clearly a necessary inequality, as χ is the only character in the inequality coming from a curve and there is no divisor that contains only χ-curves - in contrast to the case of a trivalent vertex.

To see that the other inequalities coming from curves along a generalised long side are redundant, we will decompose these inequalities similarly to before. Let C be such a curve and write

$$\theta(\varphi_{\mathfrak{C}_0}(\mathcal{O}_C)) = \theta(\chi) + \sum_{\psi \in \mathrm{Hirz}(C)} \alpha_\psi \theta(\psi) + \sum_{i=1}^{n} \sum_{\rho \in \text{G-ig}(C_i)} \beta^i_\rho \theta(\rho)$$

where $C_1, \ldots, C_n$ are the curves unlocked by C. By exactly the same methods as in the proof of Lemma 3.8.12, one can express the final term as a sum of inequalities from curves and divisors. The first two terms are equal to

$$\theta(\chi) + \sum_{\psi \in \mathrm{Hirz}(C)} \alpha_\psi \theta(\psi) = \theta(\varphi_{\mathfrak{C}_0}(\mathcal{O}_{C'})) + \sum_{\psi \in \mathrm{Hirz}(\chi)} (\alpha_\psi - 1)\theta(\varphi_{\mathfrak{C}_0}(\mathcal{R}_\psi^{-1}|_{D_\psi}))$$

where C' is a final χ-curve and D_ψ is the divisor marked with ψ. Of course $\alpha_\psi \geq 1$ and so we have shown that the inequality from C is redundant. $\square$

We consider the example $G = \frac{1}{25}(1, 3, 21)$, which has a meeting of champions of side length 2.

Example 3.8.16. *We show the triangulation for G-Hilb and Reid's recipe for $G = \frac{1}{25}(1, 3, 21)$ in Figures 3.28-3.29. Observe that of the three champions, the* 3*-chain and* 9*-chain are generalised long sides whilst the* 1*-chain contains a* $(-1, -1)$*-curve. We hence obtain two Type*

III walls from the champions and another for the 21-chain, with inequalities

$$\theta_3 + \theta_8 + \theta_{12} + \theta_{16} + \theta_{20} + \theta_{24} > 0 \qquad (F_3)$$

$$\theta_9 + \theta_{10} + \theta_{11} + \theta_{12} > 0 \qquad (C_9)$$

$$\theta_{21} + \theta_{22} + \theta_{23} + \theta_{24} > 0 \qquad (C_{21})$$

Figure 3.28: G-Hilb for $G = \frac{1}{25}(1, 3, 21)$

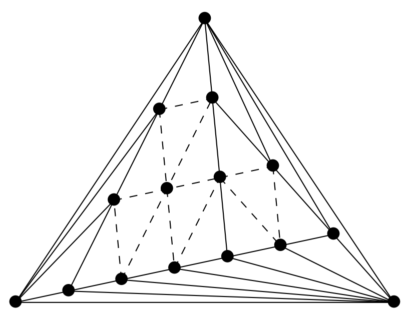

Figure 3.29: Reid's recipe for $G = \frac{1}{25}(1, 3, 21)$

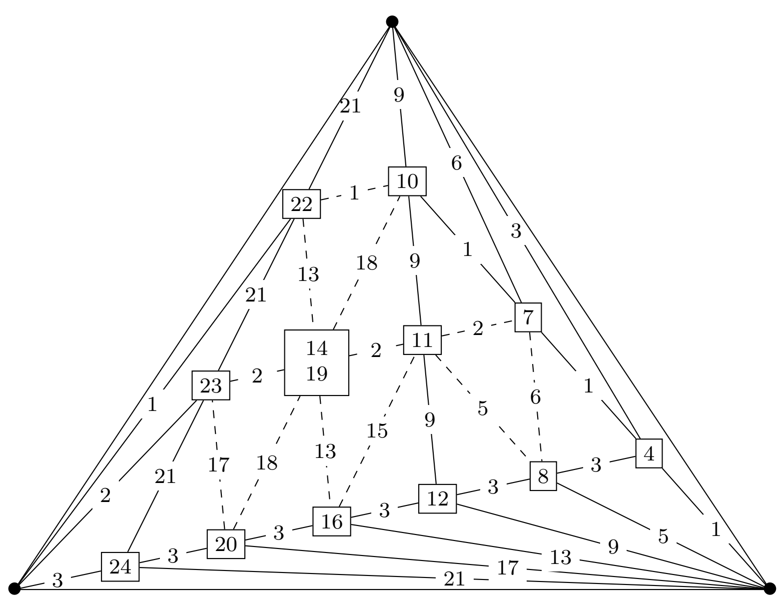

3.9 Summary

We compile the main results - Cor. 3.8.4, Prop. 3.8.5, Lemma 3.8.12, Lemma 3.8.15 - of this section.

Theorem 3.9.1. *Suppose $G \subseteq \mathrm{SL}_3(\mathbb{C})$ is a finite abelian subgroup. The walls of the chamber $\mathfrak{C}_0$ for G-$\mathrm{Hilb}\,\mathbb{C}^3$ and their types are as follows:*

- *a Type `I` wall for each exceptional $(-1,-1)$-curve,*
- *a Type `III` wall for each generalised long side,*
- *a Type `0` wall for each irreducible exceptional divisor,*
- *each remaining wall is of Type `0` and comes from a divisor parameterising a rigid quotient.*

Prop. 3.8.1 describes how to recover the unstable locus or the corresponding reducible divisor D' for each wall of Type `0` from a rigid quotient. Let $\mathfrak{w}$ be a wall of $\mathfrak{C}_0$. Denote by $E(\mathfrak{w})$ the set of edges in the Craw-Reid triangulation corresponding to curves C for which all characters in G-ig(C) appear in the equation of the wall. The desired divisor D' inducing $\mathfrak{w}$ is then the union of the divisors corresponding to vertices for which all incident edges are in $E(\mathfrak{w})$. We observe that the unlocking procedure allows the check of which walls from rigid quotients are necessary to be performed combinatorially.

Chapter 4

Numerics for orbifolds

Having studied some aspects of the birational geometry of orbifold singularities via stability, we now attend to studying numerical invariants of quotient singularities and the implications this has for the global geometry of projective orbifolds.

4.1 Orbifold Riemann-Roch

The canonical reference for numerical invariants of quotient singularities is [68], which we imitate substantially in the following.

Suppose $\mathbb{C}^n/\mu_r$ is a quotient singularity with quotient map $\pi\colon \mathbb{A}^n \to \mathbb{A}^n/\mu_r$. μ_r acts on $\pi_*\mathcal{O}_{\mathbb{A}^n}$, decomposing it into eigensheaves L_i consisting of all regular functions satisfying $\varepsilon \cdot f = \varepsilon^i f$.

Now let X be an orbifold. Let $Q \in X$ be a singularity; we say that Q is a cyclic quotient singularity of type $\frac{1}{r}(a_1, \dots, a_n)_i$ if there exists an affine chart $Q \in U \subseteq X$ in which x is the only singular point and $U \cong \mathbb{A}^n/\mu_r$ with the weights of actions being $a_1, \dots, a_n$, and locally $\mathcal{O}_X(D) \cong L_i$.

Theorem 4.1.1 ([68, Theorem 8.5]). *Suppose μ_r acts on a smooth projective variety Y with N fixed points at which μ_r acts with weights $a_1, \dots, a_n$. Let $\pi\colon Y \to X$ be the quotient map, and let $\mathcal{L}_i$ be the ith eigensheaf of $\pi_*\mathcal{O}_Y$ as above. Then*

$$\chi(Y, \mathcal{L}_i) = \frac{1}{r}\chi(\mathcal{O}_Y) + N \cdot \sigma_i(\frac{1}{r}(a_1, \dots, a_n)) \tag{4.1}$$

where

$$\sigma_i(\frac{1}{r}(a_1, \dots, a_n)) = \frac{1}{r}\sum_{\varepsilon \in \mu_r} \frac{\varepsilon^i}{\prod_{i=1}^n (1 - \varepsilon^{a_i})}$$

neglecting any terms where $1 - \varepsilon^{a_j} = 0$ for some j.

This is the local situation: the appearance of a singularity of type $\frac{1}{r}(a_1,\ldots,a_n)_i$ contributes to the Euler characteristic via the quantity $\sigma_i(\frac{1}{r}(a_1,\ldots,a_n))$, which is commonly called a *Dedekind sum*. Dedekind sums are intriguing quantities much studied in number theory and combinatorics - for example, [78] and the survey [9] - that frequently recur in studying the cohomology of toric varieties [65].

Since the quantities $\sigma_i(\frac{1}{r}(a_1,\ldots,a_n))$ are completely local to the singularity, with multiple singularities one simply adds the various contributions to obtain the Euler characteristic. We thus have:

Theorem 4.1.2 ([68, Cor. 8.6]). *Suppose X is an orbifold and that $L = \mathcal{O}_X(D)$ is a divisorial sheaf on X. Then*

$$\chi(X, L) = \text{initial term} + \sum_p q_p(D)$$

where the initial term is a 'Riemann–Roch' type formula involving D and various classes on X, and

$$q_p(D) = \sigma_i(\frac{1}{r}(a_1,\ldots,a_n)) - \sigma_0(\frac{1}{r}(a_1,\ldots,a_n))$$

when p is a singular point of type $\frac{1}{r}(a_1,\ldots,a_n)$ and $\mathcal{O}_X(D)$ is locally isomorphic to L_i near p.

The terms $\sigma_i - \sigma_0$ come from eliminating $\chi(\mathcal{O}_Y)$ from (4.1). A precise description of the initial term is given in [68, §8.6]

4.2 Numerics of quotient singularities

Properties of cyclic quotient singularities

In this subsection, we will review some of the deformation theory of cyclic quotient singularities. Recall that the singularity $\frac{1}{r}(1,a)$ is the affine toric variety associated to the cone

$$\sigma_{r,a} = \mathrm{Cone}(e_2, re_1 - ae_2) \subseteq N_{\mathbb{R}}$$

The lattice height of such a cone - the lattice distance between the origin and the line segment joining the two primitive ray generators of the cone; the *edge* of the cone - is called the *local index* [34, Note 3.19] of the cone and can be calculated as in [4] to be

$$\ell_{\sigma_{r,a}} = \frac{r}{\gcd(r, a+1)}$$

The width of a cone is the number of lattice line segments along the edge of the cone. Equivalently, it's one less than the number of lattice points along the edge, which can be computed to be

$$\mathrm{width}(\sigma_{r,a}) = \gcd(r, a+1)$$

We will often conflate a singularity and its corresponding cone in $N_{\mathbb{R}}$. Recall that an isolated cyclic quotient singularity is a T-singularity if it is $\mathbb{Q}$-Gorenstein smoothable. These were classified by Kollár and Shepherd-Barron:

Lemma 4.2.1 ([52, Prop. 3.10]). *An isolated cyclic quotient singularity is a T-singularity if and only if it takes the form*

$$\frac{1}{dn^2}(1, dnc - 1)$$

for some c with $\gcd(n, c) = 1$.

Notice that the local index of a T-singularity written as in this classification is $\ell_\sigma = n$ and that there are dn lattice points lying along the edge of the cone. On the level of cones, a T-singularity thus looks like:

Figure 4.1: T-singularity

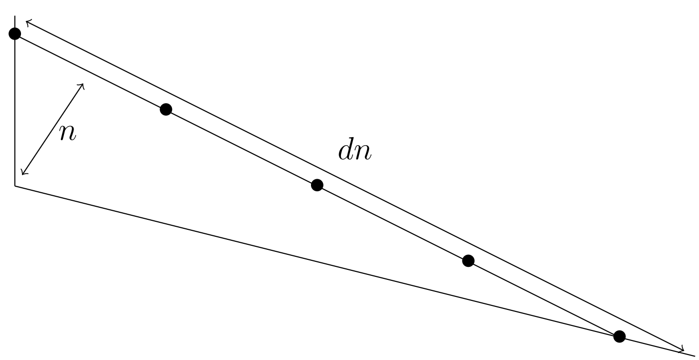

Prosaically, the cone of a T-singularity has width a multiple of its height. A T-singularity of width equal to local index - as thin as possible - is called an elementary T-singularity.

4.3 Polytope mutation

Crucial to the formulation of mirror symmetry for Fano varieties is the notion of polytope mutation developed in [1]. This is a combinatorial procedure that produces from a polytope P, a 'weight vector' $w \in M$, and an appropriate choice of convex subset $F \subseteq P$ another polytope $\mu_{w,F}P$ whose toric variety is a deformation of the toric variety associated to P. Applied locally to cones, it enables one to realise the smoothing of Kollár and Shepherd-Barron for T-singularities.

Let $N \cong \mathbb{Z}^d$ be a rank d lattice and set $N_{\mathbb{Q}} := N \otimes_{\mathbb{Z}} \mathbb{Q}$. Let $P \subseteq N_{\mathbb{Q}}$ be a lattice polytope. We require – and will assume for the remainder of this subsection – that P satisfies the following two conditions:

(a) P is of maximum dimension in N, $\dim(P) = d$;

(b) the origin is contained in the strict interior of P, $\mathbf{0} \in P^{\circ}$.

Condition (b) is not especially stringent, and can be satisfied by any polytope with $P^{\circ} \cap N \neq \emptyset$ by lattice translation. It is, however, an essential requirement in what follows.

Let $M := \mathrm{Hom}(N, \mathbb{Z}) \cong \mathbb{Z}^d$ denote the dual lattice. Given a polytope $P \subseteq N_{\mathbb{Q}}$, the dual polyhedron is defined by

$$P^* := \{u \in M_{\mathbb{Q}} \mid u(v) \geq -1 \text{ for all } v \in P\} \subseteq M_{\mathbb{Q}}.$$

Condition (b) gives that P^* is a (typically rational) polytope.

Following [2, §3], let $w \in M$ be a primitive lattice vector. Then $w : N \to \mathbb{Z}$ determines a height function (or grading) which naturally extends to $N_{\mathbb{Q}} \to \mathbb{Q}$. We call $w(v)$ the *height* of $v \in N_{\mathbb{Q}}$. We denote the set of all points of height h by $H_{w,h}$, and write

$$w_h(P) := \mathrm{conv}(H_{w,h} \cap P \cap N) \subseteq N_{\mathbb{Q}}$$

for the (possibly empty) convex hull of lattice points in P at height h.

Definition 4.3.1. *A* factor *of $P \subseteq N_{\mathbb{Q}}$ with respect to $w \in M$ is a lattice polytope $F \subseteq w^{\perp}$ such that for every negative integer $h \in \mathbb{Z}_{<0}$ there exists a (possibly empty) lattice polytope $R_h \subseteq N_{\mathbb{Q}}$ such that*

$$H_{w,h} \cap \mathrm{vert}(P) \subseteq R_h + |h|\, F \subseteq w_h(P).$$

Here '+' denotes Minkowski sum, and we define $\emptyset + Q = \emptyset$ for every lattice polytope Q.

Definition 4.3.2. *Let $P \subseteq N_{\mathbb{Q}}$ be a lattice polytope with $w \in M$ and $F \subseteq N_{\mathbb{Q}}$ as above. The* mutation *of P with respect to the data (w, F) is the lattice polytope*

$$\mu_{(w,F)}(P) := \mathrm{conv}\left(\bigcup_{h \in \mathbb{Z}_{<0}} R_h \cup \bigcup_{h \in \mathbb{Z}_{\geq 0}} (w_h(P) + hF) \right) \subseteq N_{\mathbb{Q}}.$$

It is shown in [2, Prop. 1] that, for fixed data (w, F), any choice of $\{R_h\}$ satisfying Definition 4.3.1 gives $\mathrm{GL}_d(\mathbb{Z})$-equivalent mutations. Since we regard lattice polytopes as being defined only up to $\mathrm{GL}_d(\mathbb{Z})$-equivalence, this means that mutation is well-defined. One can readily see that translating the factor F by some lattice point $v \in w^{\perp} \cap N$ gives isomorphic mutations: $\mu_{(w,F+v)}(P) \cong \mu_{(w,F)}(P)$. In particular if $\dim(F) = 0$ then $\mu_{(w,F)}(P) \cong P$. Finally, we note that mutation is always invertible [2, Lemma 2]: if $Q := \mu_{(w,F)}(P)$ then $P = \mu_{(-w,F)}(Q)$.

Remark 4.3.3. *Informally, mutation corresponds to the following operation on slices $w_h(P)$ of P: at height h one Minkowski adds or "subtracts" $|h|$ copies of F, depending on the sign of h. Definition 4.3.1 ensures that the concept of Minkowski subtraction makes sense.*

For a T-singularity with cone $\sigma_{dn^2,dnc-1}$, mutation with respect to the factor given by the cone edge and the weight vector given by the inward normal to the cone edge has the effect of removing a line segment of length $\ell_\sigma = n$ from the edge of $\sigma_{dn^2,dnc-1}$.

Figure 4.2: Mutation of a T-singularity

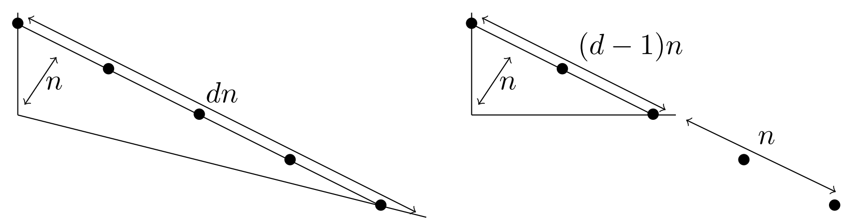

The result is that $\sigma_{dn^2,dnc-1}$ mutates to $\sigma_{(d-1)n^2,(d-1)nc-1}$. Successive mutation reduces it to the empty cone; in other words, the singularity is smoothed. The same procedure applies to a general singularity $\frac{1}{r}(1, a)$ yet may not terminate in the empty cone. Indeed, by definition the only cones for which this process will reduce to the empty cone are the T-singularities. Let $\sigma_{r,a}$ be the cone for $\frac{1}{r}(1, a)$ and let its local index be ℓ. In general, one will be able to remove subcones of the form $\sigma_{d\ell^2,d\ell c-1}$ from $\sigma_{r,a}$ by mutation and hence can deform to a cone with width equal to the residue of the width of $\sigma_{r,a}$; in particular, smaller than ℓ. This cone is the *residue* $\mathrm{res}(\sigma_{r,a})$ of $\sigma_{r,a}$. A singularity is called *residual* if it is an isolated cyclic quotient singularity with width less than its local index. In other words, no more line segments can be removed from its edge by mutation.

Lemma 4.3.4. *An isolated cyclic quotient singularity is a residual singularity if and only if it can be written in the form*

$$\frac{1}{k\ell}(1, kc - 1)$$

where ℓ is the local index, $0 < k < \ell$, $\gcd(\ell, c) = 1$, and $\gcd(\ell, kc - 1) = 1$.

Proof. Let $\frac{1}{r}(1, a)$ be a residual singularity of local index ℓ. Then, letting $k = \gcd(r, a + 1)$, $r = k\ell$ and $a + 1 = kc$ for some c with $\gcd(\ell, c) = 1$. The final coprimality condition is necessary for the singularity to be isolated. The width of $\frac{1}{r}(1, a)$ is k and so the condition $0 < k < \ell$ is equivalent to the singularity being residual. □

The quantity c is called the *slope* of the singularity (or of the cone).

Orbifold del Pezzo surfaces

We focus in this subsection on orbifold del Pezzo surfaces: two dimensional Fano varieties with only isolated cyclic quotient singularities. The *basket* of singularities on a orbifold del Pezzo surface X is the multiset of types of singularities appearing on X.

Example 4.3.5. *The weighted projective space $\mathbb{P}(a, b, c)$ is an orbifold del Pezzo surface, with basket*

$$\mathcal{B} = \{\frac{1}{a}(b, c), \frac{1}{b}(c, a), \frac{1}{c}(a, b)\}$$

Hilbert series

In [3] Ahktar–Kasprzyk use orbifold Riemann–Roch to produce the following formula for the anticanonical Hilbert series of an orbifold del Pezzo surface.

Lemma 4.3.6 ([3, Cor. 3.5])**.** *Suppose X is an orbifold del Pezzo surface with basket $\mathcal{B}$. Then*

$$\mathrm{Hilb}_X(t) = \frac{1 + (K_X^2 - 2)t + t^2}{(1 - t)^3} + \sum_{\sigma \in \mathcal{B}} Q_\sigma$$

where

$$Q_{\sigma_{r,a}} = \frac{1}{1 - t^r} \sum_{i=1}^{r} (\delta_{r,a,(a+1)i} - \delta_{r,a,0}) t^{i-1}$$

where $\delta_{r,a,i}$ is a Dedekind sum *defined as*

$$\delta_{r,a,i} = \frac{1}{r} \sum_{\xi \in \mu_r \setminus \{1\}} \frac{\xi^i}{(1 - \xi)(1 - \xi^a)}$$

Notice that each singularity σ in $\mathcal{B}$ makes two contributions to the Hilbert series: a local contribution in the orbifold correction terms Q_σ and a global contribution to the degree. The latter is captured by:

Lemma 4.3.7 ([3, Prop. 3.3]). *The degree of an orbifold del Pezzo surface with basket $\mathcal{B}$ is given by*

$$K_X^2 = 12 - n - \sum_{\sigma \in \mathcal{B}} A_\sigma$$

where n is the Euler number of the smooth locus of X, and

$$A_{\sigma_{r,a}} = m + 1 - \sum_{i=1}^{m} d_i^2 b_i + 2 \sum_{i=1}^{m} d_i d_{i+1}$$

with $[b_1, \ldots, b_m]$ the Hirzebruch-Jung continued fraction expansion of $\frac{r}{a+1}$ and d_i the discrepancy at the ith component of the exceptional fibre of the minimal resolution of $\sigma_{r,a}$.

Example 4.3.8. *A T-singularity τ of the form $\frac{1}{dn^2}(1, dnc - 1)$ has $A_\tau = d$ and $Q_\tau = 0$.*

It is worth noting that these formulae are explicit enough to actually compute all the possible combinations of singularities that would produce a given power series. However, they allow no structural description of this collection of possible baskets, which turns out to be very geometric.

4.4 Shattering and the hyperplane sum

Much of the present chapter is geared towards finding and investigating high amounts of structure in the collection of quotient singularities, especially at a fixed local index. We introduce two inverse operations - shattering and the hyperplane sum - that decompose or combine singularities of a fixed local index. The singularities that are indecomposable with respect to these operations contain the cohomological information of all quotient singularities and hence allow the problems considered here to be simplified.

Shattering

Suppose σ is a cone and that v is a primitive lattice point on the edge of σ. The fan obtained by inserting a ray through v and hence dividing σ into two new cones is a crepant blowup Y of X_σ. Recall that a map $f : Y \to X$ of Gorenstein schemes is crepant if $f^* K_X = K_Y$; that is, it has no discrepancy. In the toric situation above this occurs exactly when the ray generator v lies on the edge of σ. This toric variety Y has two affine pieces corresponding to the two dimensional cones, which are each cyclic quotient singularities. The local indices of these new singularities will be the same as the local index of the original cone since v was primitive. This procedure can be repeated to produce a collection of cones with edges lying along a common hyperplane. We call this operation *shattering*, and the cones obtained from

σ by this process *shards* of σ. Below we show $\frac{1}{9}(1,2)$ shattering into two shards corresponding to the singularities $\frac{1}{6}(1,1)$ and $\frac{1}{3}(1,1)$.

Figure 4.3: Shattering a T-singularity

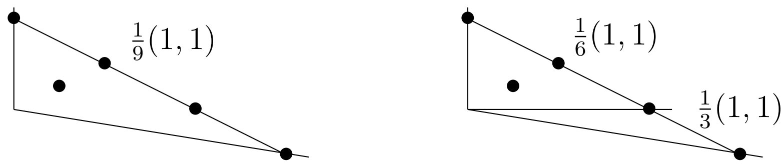

Note that the blowup corresponding to the ray through the other lattice point in the interior of the cone edge would not be crepant as this lattice point is not primitive.

The hyperplane sum

Definition 4.4.1. *The* hyperplane sum *of two cones* $\sigma_1 = \mathrm{Cone}(u, v)$ *and* $\sigma_2 = \mathrm{Cone}(v, w)$ *with ray generators clockwise ordered such that:*

- *the vectors* $v - u$ *and* $w - v$ *are parallel,*
- *both cones are of the same local index,*

to be $\sigma_1 * \sigma_2 := \mathrm{Cone}(u, w)$*, which is the usual Minkowski sum when it is defined.*

Equivalently σ_1 and σ_2 have hyperplane sum defined to be $\sigma_1+\sigma_2$ only when this cone also has the same local index as σ_1 and σ_2. One can see that this process is inverse to shattering. Continuing the conflation of cones and singularities, we will say that the hyperplane sum $\frac{1}{r}(1,a) * \frac{1}{s}(1,b)$ of two singularities is defined if there are two cones σ_1 and σ_2 corresponding to the singularities that are as in Definition 4.4.1. In this case, the hyperplane sum is defined to be the singularity corresponding to the resulting cone $\sigma_1 * \sigma_2$. Thus,

$$\frac{1}{6}(1,1) * \frac{1}{3}(1,1) = \frac{1}{9}(1,2)$$

Notice that this is well-defined as the singularity defined by a cone is unchanged by the action of $\mathrm{GL}_2(\mathbb{Z})$ on N but that it is noncommutative: the hyperplane sum $\sigma_2 * \sigma_1$ in general won't even be defined if $\sigma_1 * \sigma_2$ is.

Geometrically, the hyperplane sum is a crepant blowdown contracting the torus-invariant curve corresponding to the ray through v. From this is follows that:

Lemma 4.4.2. *(Additivity) Let $\sigma_1 * \cdots * \sigma_n = \sigma$. Then*

$$Q_{\sigma_1} + \cdots + Q_{\sigma_n} = Q_\sigma \text{ and } A_{\sigma_1} + \cdots + A_{\sigma_n} = A_\sigma$$

Corollary 4.4.3. *Let $\sigma_1 * \cdots * \sigma_n = \tau$, a T-singularity. Then*

$$Q_{\sigma_1} + \cdots + Q_{\sigma_n} = 0 \text{ and } A_{\sigma_1} + \cdots + A_{\sigma_n} = A_\tau = \frac{\mathrm{width}(\tau)}{\ell_\tau}$$

Proof of Lemma 4.4.2. Let X be a toric del Pezzo surface whose fan contains the cone σ and with all other cones smooth. Let Y be the toric variety associated to the fan that agrees with X everywhere outside of σ and there has σ replaced by $\sigma_1, \ldots, \sigma_n$. The natural map $f : Y \to X$ is given by $n-1$ crepant blowups in torus-fixed points according to the cones $\sigma_1, \ldots, \sigma_n$. Since f is surjective, $h^0(-dK_X) = h^0(-d(f^*K_X)) = h^0(-dK_Y)$. Thus the Hilbert series of X agrees with that of Y.

Note that $K_X^2 = f^*K_X^2 = K_Y^2$ and so the initial term of $\mathrm{Hilb}_X(t)$ and $\mathrm{Hilb}_Y(t)$ agree. Hence the orbifold correction terms Q_ς satisfy

$$Q_\sigma = \sum_{\varsigma \in \mathcal{B}_X} Q_\varsigma = \sum_{\varsigma \in \mathcal{B}_Y} Q_\varsigma = \sum_{i=1}^{n} Q_{\sigma_i}$$

as desired. Suppose that there are m smooth cones in the fans of X and Y. The degree formula above gives that

$$12 - m - A_\sigma = K_X^2 = K_Y^2 = 12 - m - \sum_{i=1}^{n} A_{\sigma_i}$$

and so the degree correction terms also agree. The corollaries follow from the lemma combined with Example 4.3.8. $\square$

In the sense of Cor. 4.4.3 T-singularities are negligible from the perspective of orbifold contributions to Hilbert series, though they still contribute to the degree. Say that a residual singularity σ is *hyperplane inverse* (or just 'inverse' if the context is clear) to another residual singularity σ' if the hyperplane sum $\sigma * \sigma'$ is defined and equal to an elementary T-singularity. By explicit calculation one finds that:

Lemma 4.4.4. *The hyperplane inverse of $\sigma : \frac{1}{kn}(1, kc-1)$ is*

$$\sigma^{-1} : \frac{1}{n(n-k)}(1, (n-k)(n-c)-1).$$

Notice that $(\sigma^{-1})^{-1}$ is isomorphic to σ. Let the 'dual singularity' $\overline{\frac{1}{r}(1,a)} := \frac{1}{r}(1,\overline{a})$ with $a\overline{a} \equiv 1 \bmod r$ and call a cyclic quotient singularity σ *self-dual* if $\overline{\sigma} = \sigma$. The effect of dualising a singularity reverses its gluing behaviour in that two singularities σ and σ' are hyperplane summable if and only if $\overline{\sigma}'$ and $\overline{\sigma}$ are. This follows as the element of $\mathrm{GL}_2(\mathbb{Z})$ taking σ to $\overline{\sigma}$ is orientation-reversing. The hyperplane sum is thus not actually defined on isomorphism classes of quotient singularities, but requires a choice of root of unity in addition.

Indecomposable singularities and maximal shatterings

Viewed on the level of cones, one can split a given residual singularity into 'indecomposable' hyperplane summands: those that cannot be shattered further. These are exactly the cones containing no primitive lattice points along the interior of their edge, which are not necessarily of width 1.

Definition 4.4.5. *A cyclic quotient singularity σ is* indecomposable *if a (or every) cone corresponding to it has no primitive lattice points lying on the interior of its edge.*

Note that an indecomposable singularity is in particular a residual singularity of local index at least 3. Define the *residual quiver* at local index ℓ to be the quiver with a vertex for every indecomposable singularity (distinguishing dual singularities) of local index ℓ and with an arrow drawn between σ and σ' if and only if σ and σ' are hyperplane summable. The following quiver is the result of applying the construction in the case of local index 5.

$$\begin{tikzcd}
\frac{1}{5}(1,1) \arrow[r] & \frac{1}{5}(1,2) \arrow[d] \\
\frac{1}{5}(1,3) \arrow[u] & \frac{1}{10}(1,1) \arrow[l]
\end{tikzcd}$$

Lemma 4.4.6. *For given ℓ the residual quiver at local index ℓ is a cycle of length $\phi(\ell)$, where ϕ is Euler's totient function. Moreover, there is exactly one indecomposable singularity of every slope $c \in (\mathbb{Z}/\ell\mathbb{Z})^\times$.*

Before proving this lemma, observe that for any two singularities σ_1 and σ_2, though they may not be hyperplane summable, there will always be a unique singularity σ_g of smallest width such that $\sigma_1 * \sigma_g$ and $\sigma_g * \sigma_2$ are well-defined. This is because, picking a cone for σ_1, one can always move a cone representing σ_2 to have edge lying on the hyperplane given by the cone edge of σ_1 by the action of $\mathrm{GL}_2(\mathbb{Z})$. This cone σ_g is called the *gluing cone* of σ_1 and σ_2.

Proof. If $\gcd(\ell, c-1) = 1$ then the singularity of slope c and width 1 is indecomposable. Suppose $\gcd(\ell, c-1) \neq 1$ but $\gcd(\ell, 2c-1) = 1$. Then the residual singularity with slope c and width 2 contains only non-primitive interior edge lattice points and so is also indecomposable. Continuing, if $\gcd(\ell, c-1) \neq 1$ and $\gcd(\ell, 2c-1) \neq 1$ but $\gcd(\ell, 3c-1) = 1$ then the residual singularity with slope c and width 3 will be indecomposable. For a given slope c one can pick $w = 2c^{-1} \in \mathbb{Z}/\ell\mathbb{Z}$ as a width for which $\gcd(\ell, wc-1) = 1$ since $wc \equiv 1 \bmod \ell$. Hence this process exhausts all $c \in (\mathbb{Z}/\ell\mathbb{Z})^{\times}$. Clearly no other values of c are possible and each level with fixed width excludes all previous levels, giving that the number of vertices of the residual quiver is equal to the number of choices of c, which is $\phi(\ell)$. There is a unique indecomposable singularity that can be glued onto a given singularity by extending its cone edge. It follows that every vertex is linked to exactly two others. This says that the quiver is a collection of cycles. However given any two indecompables σ_1 and σ_2 there is a gluing cone joining them. Maximally shattering this gluing cone gives a path from σ_1 to σ_2 in the residual quiver; therefore the quiver is connected and so forms a single cycle. $\square$

Conjecture #1

In this section we will state the first conjecture that the main results of this paper will depend upon. We will fix notation:

- $\mathcal{R}_\ell$ is the set of indecomposable singularities of local index ℓ
- $\mathrm{Res}(\ell)$ is the set of residual singularities of local index ℓ
- σ will denote a residual singularity of local index ℓ

From the general theory of Hilbert series for Gorenstein schemes - see, for example, [6], [16], or [60] - applied to this particular situation one can write

$$Q_\sigma = \frac{\delta_0 + \delta_1 t + \cdots + \delta_{\ell-1} t^{\ell-1}}{\ell(1-t^\ell)} \tag{4.2}$$

for $\delta_i \in \mathbb{Z}$. Define the *δ-vector* of σ to be $\delta(\sigma) := (\delta_0, \delta_1, \ldots, \delta_{\ell-1})$. From Lemma 4.3.6

$$\delta_i = \ell(\delta_{r,a,(a+1)(i+1)} - \delta_{r,a,0})$$

The δ-vector has the properties:

- $\delta(\sigma)$ is palindromic
- $\delta_0 = \delta_{\ell-1} = 0$.

Because of the second property, we will abbreviate the δ-vector to omit the first and last terms. For example, the δ-vector of $\frac{1}{5}(1,1)$ with

$$Q_{\frac{1}{5}(1,1)} = \frac{t^3 - 2t^2 + t}{5(1-t^5)}$$

is $(1,-2,1)$. We now prove these two properties.

Proof. Suppose $\sigma = \frac{1}{r}(1,a)$ with local index ℓ. We start by proving that Q_σ can be written in the form $(\star)$. It suffices that the numerator in the expression of Q_σ in Lemma 4.3.6 with denominator $1-t^r$ is divisible by $1 + t^\ell + t^{2\ell} + \cdots + t^{r-\ell}$, or equivalently that

$$\delta_{r,a,(a+1)(\ell+i)} = \delta_{r,a,(a+1)i}$$

This follows immediately from noting that $(a+1)\ell \equiv 0 \bmod r$ and that the final argument of such a Dedekind sum is well-defined modulo r. To prove the first property - that $\delta_{\ell-1-i} = \delta_i$ - observe that it suffices that the first ℓ terms of the numerator of Q_σ in the expression of Lemma 4.3.6 are palindromic, or that

$$\delta_{r,a,(a+1)(\ell-i)} = \delta_{r,a,(a+1)(i+1)}$$

Computing directly using the fact that $(a+1)\ell \equiv 0 \bmod r$,

$$\delta_{r,a,(a+1)(\ell-i)} = \sum \frac{\varepsilon^{(a+1)(\ell-i)}}{(1-\varepsilon)(1-\varepsilon^a)} = \sum \frac{\varepsilon^{-(a+1)i}}{(1-\varepsilon)(1-\varepsilon^a)}$$

Multiplying by $\varepsilon^{-(a+1)}$ in numerator and denominator yields

$$\sum \frac{\varepsilon^{-(a+1)(i+1)}}{\varepsilon^{-(a+1)}(1-\varepsilon)(1-\varepsilon^a)} = \sum \frac{\varepsilon^{-(a+1)(i+1)}}{(\varepsilon^{-1}-1)(\varepsilon^{-a}-1)} = \sum \frac{\varepsilon^{(a+1)(i+1)}}{(1-\varepsilon)(1-\varepsilon^a)} = \delta_{r,a,(a+1)(i+1)}$$

using the bijection $\varepsilon \mapsto \varepsilon^{-1}$ on the rth roots of unity. The second property now follows from the equality

$$\delta_{r,a,(a+1)\ell} = \delta_{r,a,0}$$

since $(a+1)\ell \equiv 0 \bmod r$ and so the $(\ell-1)$th coefficient $\delta_{\ell-1} = \delta_{r,a,(a+1)\ell} - \delta_{r,a,0} = 0$. $\square$

Lemma 4.4.7. *An isolated cyclic quotient singularity τ is a T-singularity if and only if $Q_\tau = 0$.*

Proof. The only if implication follows from Example 4.3.8. It suffices to show that every residual singularity makes a nonzero contribution to the Hilbert series. This follows using the shattering in [3] decomposing a cone into T-cones and a single residual cone. For the forward implication, let $\frac{1}{r}(1,a)$ be a residual singularity and consider the weighted projective plane $X = \mathbb{P}(1,a,r)$. This has three affine pieces isomorphic to $\mathbb{A}^2$, $\frac{1}{r}(1,a)$ and $\frac{1}{a}(1,r)$ and thus its Hilbert series is

$$\mathrm{Hilb}_X(t) = \frac{1+(K_X^2-2)t+t^2}{(1-t)^3} + Q_{\sigma_1} + Q_{\sigma_2}$$

where $\sigma_1 : \frac{1}{r}(1,a)$ and $\sigma_2 : \frac{1}{a}(1,r)$. Let these have local indices ℓ_1 and ℓ_2 - which are coprime - and write

$$\delta(\sigma_i) = (\delta^i_j)_{j=1,\dots,\ell_i-2}$$

One can compute the t-coefficient of the Hilbert series to be

$$h^0(-K_X) = 1 + K_X^2 + \frac{\delta^1_1}{\ell_1} + \frac{\delta^2_1}{\ell_2} \tag{4.3}$$

As the dimension of a vector space, this must be an integer. Recall that the degree of X is $(1+a+r)^2/ar$, which has the same fractional part as

$$\frac{1+a}{r} + \frac{1+r}{a} + \frac{1+a+r}{ar} \tag{4.4}$$

Consider the the residues of (4.4) mod $\mathbb{Z}\cdot\frac{1}{\ell_1}$ and $\mathbb{Z}\cdot\frac{1}{\ell_2}$. Suppose the residue of $(*)$ mod $\mathbb{Z}\cdot\frac{1}{\ell_2}$ is zero. Then

$$\frac{1+a}{r} + \frac{1+a+r}{ar} \equiv 0 \operatorname{mod} \mathbb{Z}\cdot\frac{1}{\ell_2}$$

as $(1+r)/a$ has denominator ℓ_2 in lowest terms. Combining fractions, this requires in particular that r divides $(1+a)^2$ as r is coprime to ℓ_2. Let k be the width of σ_1 so that $r = k\ell_1$ and $1+a = kc$ for some c coprime to ℓ_1. For $k\ell_1$ to divide k^2c^2 one must have that ℓ_1 divides k, but this is contrary to the definition of residual singularity. It follows that both residues are nonzero and so, for (†) to be an integer, δ^1_1 and δ^2_1 must be nonzero. $\square$

This proof actually shows that the first coefficient of the δ-vector of a residual singularity is nonzero. Notice that this argument would fail for a T-singularity where, by definition, the local index ℓ_1 divides the width k.

Given a semigroup S and a set R, let the formal semigroup consisting of S-linear combinations of elements of R be denoted by $S\langle R\rangle$. If $S = \mathbb{Z}$ this is just the formal lattice generated by R. Define the *δ-lattice* for local index ℓ to be the sublattice

$$\Delta(\ell) := \mathbb{Z}\langle \delta(\sigma) : \sigma \in \mathrm{Res}(\ell)\rangle \subseteq \mathbb{Z}^{\ell-2}$$

generated by all the δ-vectors of residual singularities (equivalently, indecomposable singularities) of local index ℓ.

Given a list of residuals $\mathcal{T} = (\sigma_1, \dots, \sigma_n)$ there is a unique expression

$$(\sigma_1^1, \dots, \sigma_1^{m_1}, \dots, \sigma_n^1, \dots, \sigma_n^{m_n})$$

where $\sigma_i^j \in \mathcal{R}_\ell$ and $\sigma_i^1 * \dots * \sigma_i^{m_i} = \sigma_i$. This tuple is the *maximal shattering* of $\mathcal{T}$ denoted by $\rho(\mathcal{T})$. Combinatorially the cones corresponding to each σ_i have been shattered as much as possible to decompose into a hyperplane sum in terms of $\mathcal{R}_\ell$. Note that by definition of $\mathcal{R}_\ell$ this is maximal exactly in this sense. Of course

$$\sum_i Q_{\sigma_i} = \sum_{i,j} Q_{\sigma_i^j} \text{ and } \sum_i A_{\sigma_i} = \sum_{i,j} A_{\sigma_i^j}$$

by the additivity in Lemma 4.4.2. From the characterisation in terms of non-primitive lattice points, the maximal shattering of any singularity is unique. After being linearly extended ρ defines a surjective monoid homomorphism $\mathbb{N}[\mathrm{Res}(\ell)] \to \mathbb{N}[\mathcal{R}_\ell]$ which is left-inverse to the inclusion $\mathbb{N}[\mathcal{R}_\ell] \to \mathbb{N}[\mathrm{Res}(\ell)]$. Consider the map $\widetilde{\Phi}_\ell$ completing the diagram

$$\begin{tikzcd} \mathbb{N}[\mathrm{Res}(\ell)] \arrow[d, "\rho"'] \arrow[dr, "\Phi_\ell"] & \\ \mathbb{N}[\mathcal{R}_\ell] \arrow[r, "\widetilde{\Phi}_\ell"'] & \Delta(\ell) \end{tikzcd}$$

where $\Delta(\ell)$ is the lattice of δ-vectors of orbifold contributions of local index ℓ as above. $\widetilde{\Phi}_\ell$ exists and is unique since, geometrically, ρ applies a collection of crepant blowups that preserve the orbifold contributions Q_σ and hence their δ-vectors. We make use of semigroups of the form $\mathbb{N}\langle R\rangle$ to record the (nonnegative) quantities of each singularity inside a basket. These maps all extend to lattice homomorphisms

$$\Phi_\ell^{\mathbb{Z}} : \mathbb{Z}\langle \mathrm{Res}(\ell)\rangle \to \Delta(\ell), \widetilde{\Phi}_\ell^{\mathbb{Z}} : \mathbb{Z}\langle \mathcal{R}_\ell\rangle \to \Delta(\ell), \rho^{\mathbb{Z}} : \mathbb{Z}\langle \mathrm{Res}(\ell)\rangle \to \mathbb{Z}\langle \mathcal{R}_\ell\rangle$$

With the objective of studying relations between orbifold contributions we assume the following conjecture, which has been verified up to local index 34 in Sage. It has echoes of the $\frac{1}{2}\phi(r)$ in [68] §5.9 as well as of many other results across the study of Dedekind sums. The reader can add the caviate $\ell \leq 34$ on any subsequent results making use of this conjecture.

Conjecture 4.4.8. $\operatorname{rank} \Delta(\ell) = \frac{1}{2}\phi(\ell)$.

4.5 Cancelling tuples

Note that any collection of singularities in the kernel of Φ_ℓ contributes zero in orbifold correction terms to the Hilbert series.

Definition 4.5.1. *A* cancelling tuple *is a finite collection of residual singularities* $\sigma_1, \ldots, \sigma_n$ *such that* $\sum_{i=1}^n Q_{\sigma_i} = 0$.

Equivalently, $\sum_{i=1}^n \sigma_i \in \ker \Phi_\ell$. If $\sigma_1, \ldots, \sigma_n$ is a cancelling tuple then $\rho(\sigma_1), \ldots, \rho(\sigma_n)$ is also a cancelling tuple. Hence $\sum_{i=1}^n \rho(\sigma_i) \in \ker \widetilde{\Phi}_\ell$. A cancelling tuple can be decomposed nonuniquely into *minimal* cancelling tuples: cancelling tuples that contain no smaller cancelling tuples.

Lemma 4.5.2. *An elementary T-singularity of local index ℓ is composed of exactly one of every indecomposable at local index ℓ glued in the cyclic order prescribed by the residual quiver. Hence, distinguishing dual singularities, there are exactly $\phi(\ell)$ T-singularities at local index ℓ parameterised by choosing a starting point in the residual quiver.*

Proof. Write $\tau = \sigma_1 * \cdots * \sigma_m$ with each σ_i an indecomposable singularity. Since a T-singularity can be glued to itself by mutation, σ_m must be the indecomposable singularity immediately preceding σ_1 in the residual quiver. Thus, since one must follow the cycle around the quiver in order to glue indecomposable singularities, τ is a circuit of the quiver beginning at σ_1 and ending at the previous vertex σ_m. Suppose σ_1 appears again as one of the σ_i for $i > 1$. Then $\tau = (\sigma_1 * \cdots * \sigma_m) * (\sigma_1 * \cdots * \sigma_m) * \cdots * (\sigma_1 * \cdots * \sigma_m)$ so that the cycle can end with σ_m. Suppose there are p cycles of the residual quiver in this decomposition of τ. The cone $\sigma_1 * \cdots * \sigma_m$ has $p \sum_{i=1}^m Q_{\sigma_i} = Q_\tau = 0$ and so $\sum_{i=1}^m Q_{\sigma_i} = 0$. Thus by Lemma 4.4.7 $\sigma_1 * \cdots * \sigma_m$ is also a T-singularity. This contradicts the fact that τ is an elementary T-singularity unless there is only a single circuit of the quiver. $\square$

Corollary 4.5.3. *The widths of all indecomposable singularities of local index ℓ sum to ℓ.*

Let $\sigma_1, \ldots, \sigma_n$ be the indecomposable singularities of local index ℓ listed in cyclic order according to the residual quiver. From the lemma above $\Phi_\ell(\sum_{i=1}^n \sigma_i) = 0$ and so $\sigma_1, \ldots, \sigma_n$ form a cancelling tuple. This arises from maximally shattering an elementary T-cone. Hence, $\operatorname{rank} \ker \widetilde{\Phi}_\ell^{\mathbb{Z}} \geq 1$ as there is at least one cancelling tuple of every local index, which is formed of indecomposable singularities. The objective of the rest of this section is to prove, assuming the conjecture on the rank of $\Delta(\ell)$, that:

Lemma 4.5.4. *All of the minimal cancelling tuples consisting of singularities of local index ℓ arise from shattering a T-cone in some way.*

Clearly shattering a T-cone does produce a cancelling tuple, however the converse is more subtle.

$\widetilde{\Phi}_\ell$ surjects onto the δ-lattice $\Delta(\ell) \cong \mathbb{Z}^{\frac{1}{2}\phi(\ell)}$ as Φ_ℓ does and so $\operatorname{rank} \widetilde{\Phi}_\ell^{\mathbb{Z}} = \frac{1}{2}\phi(\ell)$. We can identify isomorphic singularities in $\mathcal{R}_\ell$: this has the effect of conflating the singularities σ

and $\bar{\sigma}$. These two singularities have the same orbifold contribution and so $\widetilde{\Phi}_\ell$ passes to the quotient to give a surjection $\widetilde{\Phi}_\ell^{\mathbb{Z}} : \mathbb{Z}\langle \mathcal{R}_\ell/\cong\rangle \to \Delta(\ell)$. The rank of $\mathbb{Z}\langle \mathcal{R}/\cong\rangle$ is $\phi(\ell) - \frac{1}{2}S(\ell)$ where $S(\ell)$ is the number of non-self dual residuals of local index ℓ contained in the generating set $\mathcal{R}_\ell$, which is seen by noting that one of each pair $\sigma, \bar{\sigma}$ of non-self dual residuals are exactly those that are removed by quotienting out by isomorphism.

To prove Lemma 4.5.4 it suffices that $\operatorname{rank}\ker \widetilde{\Phi}_\ell^{\mathbb{Z}} = 1$ since, as seen, there is already a cancelling tuple obtained from cycling around the residual quiver and so if the kernel of $\widetilde{\Phi}_\ell^{\mathbb{Z}}$ is cyclic then this special cancelling tuple must generate it. This follows because the coordinate vector of this cancelling tuple in the standard basis of $\mathbb{Z}\langle \mathcal{R}_\ell/\cong\rangle$ is primitive as it contains a 1 as an entry corresponding to the single occurence of the self-dual indecomposable singularity $\frac{1}{\ell}(1,1)$ when ℓ is odd or $\frac{1}{2\ell}(1,1)$ when ℓ is even. Rank-nullity then informs us that Lemma 4.5.4 is equivalent to

$$S(\ell) = \phi(\ell) - 2$$

or, equivalently, that there are exactly two self-dual singularities contained in $\mathcal{R}_\ell$ for any ℓ.

A self-dual residual of width w and slope c is one for which $(wc-1)^2 \equiv 1 \bmod \ell w$ or, equivalently, $wc \equiv 2 \bmod \ell$. Suppose ℓ is odd. There is then exactly one self-dual residual of width 1 and 2 given by the equation $w = 2\bar{c}$ as 2 is invertible modulo ℓ. If ℓ is even then at width $w = 2$ one can solve for invertible c obtaining $c = 1$. Indeed $c \equiv 1 \bmod \ell$ is needed but this satisfies the coprimality conditions.

Lemma 4.5.5. *There are at most two self-dual indecomposables at any local index.*

Proof. Consider the residual quiver $\mathcal{Q}(\ell)$ for local index ℓ, which is a $\phi(\ell)$-cycle. It carries an involution given by $\iota : \sigma \mapsto \bar{\sigma}$ which reverses the direction of the arrows. ι hence fixes at most 2 vertices, which correspond to self-dual indecomposables by definition. $\square$

There are actually exactly two self-dual residuals at every local index; explicitly these are

$$\begin{cases} \frac{1}{\ell}(1,1), \quad \frac{1}{2\ell}(1,1) & \text{if } \ell \text{ is odd,} \\ \frac{1}{2\ell}(1,1), \quad \frac{1}{2\ell}(1,\ell+1) & \text{if } \ell \equiv 0,4 \bmod 8, \\ \frac{1}{2\ell}(1,1), \quad \frac{1}{4\ell}(1,\ell+1) & \text{if } \ell \equiv 2 \bmod 8, \\ \frac{1}{2\ell}(1,1), \quad \frac{1}{4\ell}(1,3\ell+1) & \text{if } \ell \equiv 6 \bmod 8. \end{cases}$$

as can be verified by some modular arithmetic. This proves Lemma 4.5.4 subject to Conjecture 4.4.8.

Corollary 4.5.6. *One can order the generating set $\mathcal{R}_\ell/\cong$ of $\mathbb{Z}\langle\mathcal{R}_\ell/\cong\rangle$ such that in those coordinates* $\ker\widetilde{\Phi}^{\mathbb{Z}}_\ell = \mathbb{Z}\cdot(1,2,\ldots,2,1)$ *where the ones correspond to the two self-dual indecomposables and the twos identify the $\frac{1}{2}\phi(\ell)-1$ non-self dual pairs.*

It also follows that the residual quiver always takes the form

$$\begin{tikzcd}
 & & \cdots \arrow[dr] & & \\
 & \sigma_1 \arrow[ur] & & \sigma_m \arrow[dr] & \\
\sigma^1_{\mathrm{sd}} \arrow[ur] & & & & \sigma^2_{\mathrm{sd}} \arrow[dl] \\
 & \bar{\sigma}_1 \arrow[ul] & & \bar{\sigma}_m \arrow[dl] & \\
 & & \cdots \arrow[ul] & &
\end{tikzcd}$$

where the σ^i_{sd} are the two self-dual indecomposables, and the $\sigma_i, \bar{\sigma}_i$ are the non-self dual pairs of indecomposables; so $m = \frac{1}{2}\phi(\ell) - 1$. Observe that if one maximally shatters a non-elementary T-cone τ then one must obtain a non-minimal cancelling tuple: the result will consist of a minimal cancelling tuple for each elementary T-cone inside τ. Because the degree contribution A_τ of a T-singularity τ is equal to its width divided by its local index, it follows that:

Corollary 4.5.7. *If $\sum Q_{\sigma_i} = 0$ then $\sum A\sigma_i \in \mathbb{N}$ and the second sum is zero iff the list of σ_is is empty.*

Conjecture #2

Another natural question is whether or not minimal cancelling tuples can involve singularities of different local indices. This is unresolved but serves to sharpen later results and so is stated as a conjecture.

Conjecture 4.5.8. *Suppose $\sum_{i=1}^n Q_{\sigma_i} = 0$. Then $\ell_{\sigma_i} = \ell_{\sigma_j}$ for all i, j.*

Remark 4.5.9. *Let $\mathrm{CQS}(\ell)$ be the set of cyclic quotient singularities of local index ℓ and consider the map $\rho : \mathrm{CQS}(\ell) \to \mathbb{N}[\mathcal{R}_\ell] \to \mathbb{N}[\mathcal{R}_\ell/\cong]$. This is a surjection onto a free semigroup of rank $\frac{1}{2}\phi(\ell)+1$. The T-singularities form a ray given in coordinates as above by $\mathbb{N}\cdot(1,2,\ldots,2,1)$ within this semigroup that exactly consists of the $\mathbb{Q}$-Gorenstein smoothable singularities. We are curious about any analogous combinatorics in higher dimensions, which may lead to or confirm a suitable definition of 'residual singularity' there.*

4.6 Recovering baskets from Hilbert series

Decomposing baskets

We will start by decomposing a basket of singularities into two pieces - one containing only cancelling tuples, and one that is actually detectable by the Hilbert series.

Definition 4.6.1. *Let X be an orbifold del Pezzo surface with basket $\mathcal{B}_X$.*

- *An* invisible basket *in $\mathcal{B}_X$ is a subset $\mathcal{IB} \subseteq \mathcal{B}_X$ such that $\sum_{\sigma \in \mathcal{IB}} Q_\sigma = 0$ and that no nonempty subcollection $\emptyset \neq T \subseteq \mathcal{B} \setminus \mathcal{IB}$ has that $\sum_{\sigma \in T} Q_\sigma = 0$.*

- *The collection $\mathcal{B}_X \setminus \mathcal{IB}$ is called the* reduced basket *for $\mathcal{IB}$ in $\mathcal{B}_X$. It will be denoted by $\mathcal{RB}$.*

Equivalently, call a multiset $\mathcal{S}$ of singularities *invisible* if $\sum_{\sigma \in \mathcal{S}} Q_\sigma = 0$. An invisible basket for X is a maximal invisible submultiset $\mathcal{IB} \subseteq \mathcal{B}_X$.

Definition 4.6.2. *Let $\mathcal{B}$ be a multiset of singularities. Set $\mathcal{B}(\ell) := \{\sigma \in \mathcal{B} : \ell_\sigma = \ell\}$ to be the ℓth* piece *of $\mathcal{B}$. Define $\mathcal{RB}(\ell)$ and $\mathcal{IB}(\ell)$ similarly.*

The orbifold correction terms of a Hilbert series provide data only on the level of a reduced basket as an invisible basket is by definition invisible to it. The extent to which a series determines a reduced basket is discussed below. Conjecture 4.5.8 implies the following:

> *Suppose X is an orbifold del Pezzo surface with basket $\mathcal{B}$ featuring singularities of local indices $\ell_1, .., \ell_N$. Then the decomposition*
>
> $$\mathrm{Hilb}_X(t) = \frac{1 + (K_X^2 - 2)t + t^2}{(1-t)^3} + \sum_{\sigma \in \mathcal{B}(\ell_1)} Q_\sigma + \cdots + \sum_{\sigma \in \mathcal{B}(\ell_N)} Q_\sigma$$
>
> *is unique in that it corresponds to grouping terms with a common denominator of the form $1 - t^{\ell_i}$.*

Consequently, write $Q_{\mathcal{B}}(\ell) := \sum_{\sigma \in \mathcal{B}(\ell)} Q_\sigma$ for the ℓth part of the decomposition of the orbifold contribution from $\mathcal{B}$. An easy fact independent of this conjecture is that the initial term can be identified from the Hilbert series as a whole as the only part with a triple pole at 1. Note that the order of vanishing at 1 cannot be diminished by the numerator since $K_X^2 > 0$ implies that $1 + (K_X^2 - 2)t + t^2$ cannot have 1 as a root.

Corollary 4.6.3. *The degree of an orbifold del Pezzo surface is determined by its Hilbert series.*

Convex geometry and reduced baskets

We situate the problem of computing the possible reduced baskets for an orbifold del Pezzo surface with a given Hilbert series or total orbifold contribution in the setting of convex geometry where it is most easily visualised. This will be accompanied by an example for local index 5. Recall that $\mathrm{Res}(\ell)$ is the set of indecomposable singularities of local index ℓ.

Denote by $\mathrm{Res}^+(\ell)$ a choice of one residual singularity of each hyperplane inverse pair $\{\sigma, \sigma^{-1}\}$. There is a map

$$\varphi : \mathbb{Z}\langle \mathrm{Res}^+(\ell)\rangle \to \mathbb{N}\langle \mathrm{Res}(\ell)\rangle$$

given by interpreting interpreting a linear combination $v_1\sigma_1 + \cdots + v_r\sigma_r \in \mathbb{Z}\langle \mathrm{Res}^+(\ell)\rangle$ as a basket of singularities by taking v_i-copies of σ_i if $v_i \geq 0$ or $(-v_i)$ copies of σ_i^{-1} if $v_i < 0$. Observe that by construction the image of φ consists of all baskets containing no cancelling pairs and hence it must contain every possibility for a reduced basket. Composing with the map Φ_ℓ to $\Delta(\ell)$ gives the linear map $\Phi_\ell^+ : \mathbb{Z}\langle \mathrm{Res}^+(\ell)\rangle \to \Delta(\ell)$ associating to σ the δ-vector of its orbifold contribution Q_σ. Note that this is well-defined as $Q_{\sigma^{-1}} = -Q_\sigma$. The elements of $\ker \Phi_\ell^+$ correspond to cancelling m-tuples with $m > 2$.

Example 4.6.4. *At local index* 5 *there are eight residual singularities falling into the following inverse pairs:*

$$\{\frac{1}{5}(1,1), \frac{1}{20}(1,11)\}, \{\frac{1}{5}(1,2), \frac{1}{20}(1,3)\}, \{\frac{1}{10}(1,1), \frac{1}{15}(1,11)\}, \{\frac{1}{10}(1,3), \frac{1}{15}(1,2)\}$$

Making the choice

$$\mathrm{Res}^+(5) = \{\sigma_1 = \frac{1}{5}(1,1), \sigma_2 = \frac{1}{20}(1,3), \sigma_3 = \frac{1}{10}(1,1), \sigma_4 = \frac{1}{15}(1,2)\}$$

gives orbifold contributions with δ-vectors

$$q_1 = (1,-2,1), q_2 = (2,1,2), q_3 = (3,4,3), q_4 = (1,3,1)$$

Their span is a two dimensional lattice since $q_1 + q_4 = q_2$ and $q_1 + 2q_4 = q_3$. These relations define the cancelling tuples

$$\frac{1}{5}(1,1), \frac{1}{5}(1,2), \frac{1}{15}(1,2) \quad \textit{and} \quad \frac{1}{5}(1,1), \frac{1}{15}(1,11), \frac{1}{15}(1,2), \frac{1}{15}(1,2).$$

Now suppose that $\delta \in \Delta(\ell)$ is the δ-vector of some rational function

$$Q = \frac{\delta_1 t + \delta_2 t^2 + \cdots + \delta_{\ell-2} t^{\ell-2}}{\ell(1-t^\ell)}$$

that could be the total orbifold contribution of some basket $\mathcal{B}$ of singularities of local index ℓ. Suppose $\mathcal{B}_0 \in \mathbb{Z}\langle \mathrm{Res}^+(\ell)\rangle$ satisfies

$$\Phi_\ell^+(\mathcal{B}_0) = \delta$$

That is, it is a particular solution to the problem of finding a basket producing the given total orbifold contribution. By definition, any other basket with this property will differ as an element of $\mathbb{Z}\langle \mathrm{Res}^+(\ell)\rangle$ by a cancelling tuple or, equivalently, an element of $\ker \Phi_\ell^+$.

Example 4.6.5. *Consider the orbifold contribution*

$$Q = \frac{2t^3 + t^2 + 2t}{5(1 - t^5)}$$

with δ-vector $(2, 1, 2)$. A particular basket producing this orbifold contribution is

$$\mathcal{B}_0 = \{\frac{1}{5}(1, 1), \frac{1}{15}(1, 2)\}$$

corresponding to the vector $(1, 0, 0, 1) \in \mathbb{Z}^4$ in the coordinates of Example 4.6.4. The set of baskets containing no cancelling pairs with this total orbifold contribution is, in coordinates,

$$\mathcal{B}_0 + \ker \Phi_\ell^+ = \{(1 + \lambda + \mu, -\lambda, -\mu, 1 + \lambda + 2\mu) : \lambda, \mu \in \mathbb{Z}\}$$

In order to find all of the reduced baskets - those not containing any cancelling tuples - that produce a given total orbifold contribution, one must exclude all baskets containing cancelling tuples. To this end, define the *signature* of a vector $v \in \mathbb{Z}^n$ to be

$$\mathrm{sgn}(v) := (\mathrm{sgn}(v_1), \ldots, \mathrm{sgn}(v_n))$$

where sgn is the usual sign function satisfying $\mathrm{sgn}(0) = 0$. Define

$$L_v := \bigoplus_{v_i \neq 0} \mathbb{N} \cdot \mathrm{sgn}(v_i) e_i \oplus \bigoplus_{v_i = 0} \mathbb{Z} \cdot e_i \text{ and } S(v) := v + L_v$$

A vector $u \in \mathbb{Z}^n$ is said to *feature* in another vector $v \in \mathbb{Z}^n$ if $u \in S(v)$. Note that $S(v)$ is a smooth affine rational polyhedral cone. Inside the lattice $\mathbb{Z}\langle \mathrm{Res}^+(\ell)\rangle$ using as coordinates the distinguished basis $\mathrm{Res}^+(\ell)$, the cone $S(v)$ consists of the baskets containing v since allowing no sign changes corresponds to the property that no singularities appearing in v can be removed in moving to a basket found in $S(v)$.

If v is a cancelling tuple, no reduced baskets can lie in $S(v)$ as all baskets there will all contain the cancelling tuple v. In particular, there can only be finitely many solutions along any affine ray of the form $\{u + \lambda v : \lambda \geq 0\}$ parallel to $\ker \Phi_\ell^+$, since eventually the cancelling tuple $v \in \ker \Phi_\ell^+$ will feature in $u + \lambda v$ for $\lambda \gg 0$. This shows:

Lemma 4.6.6. *There only finitely many reduced baskets along each affine ray parallel to* $\ker \Phi_\ell^+$*.*

Working with a given δ-vector $\delta \in \Delta(\ell)$ and particular choice of basket $\mathcal{B}_0$ whose total orbifold contribution has this δ-vector, this means that the reduced baskets producing this total orbifold contribution biject with lattice points in the complement of the union as v ranges over all cancelling tuples of the convex rational polyhedra

$$K_{v,\delta} := S(v) \cap (\mathcal{B}_0 + \ker \Phi_\ell^+) = S(v) \cap (\Phi_\ell^+)^{-1}(\delta)$$

inside $\mathcal{B}_0 + \ker \Phi_\ell^+$. More concisely, the reduced baskets producing an orbifold contribution with δ-vector δ biject with lattice points inside

$$(\Phi_\ell^+)^{-1}(\delta) \setminus \bigcup_{v \in \ker \Phi_\ell^+} K_{v,\delta}$$

Example 4.6.7. *Returning to Example 4.6.5 and the* δ*-vector* $(2,1,2) \in \Delta(5)$*, since* $\ker \Phi_\ell^+ \cong \mathbb{Z}^2$ *one can sketch the intersection* $K_{v,\delta}$ *of the cones* $S(v)$ *with* $\mathcal{B}_0 + \ker \Phi_\ell^+$*. Here some of the resulting polyhedra* $K_{v,\delta}$ *are drawn, enough for the purposes at hand.*

Figure 4.4: Possible reduced baskets for $(2,1,2)$

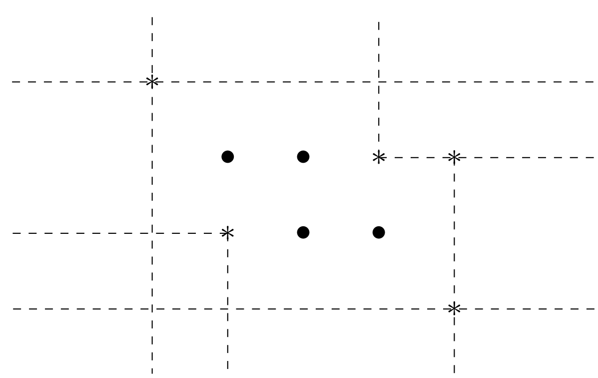

In this example the polyhedra $K_{v,\delta}$ *exclude a cobounded set and so there is only a finite number of possible reduced baskets for the* δ*-vector* $(2,1,2)$*. That is, there are only finitely many reduced baskets giving rise to the given orbifold contribution. They can be seen from this to be*

$$(1,0,0,1) \quad (1,-1,1,0) \quad (0,0,1,-1) \quad (0,1,0,0)$$

$$\text{or } \frac{1}{5}(1,1), \frac{1}{15}(1,2); \quad \frac{1}{5}(1,1), \frac{1}{5}(1,1), \frac{1}{5}(1,2), \frac{1}{10}(1,1); \quad \frac{1}{10}(1,1), \frac{1}{10}(1,3); \quad \frac{1}{20}(1,3).$$

They all have total degree contribution $\sum_{\sigma \in \mathcal{RB}} A_\sigma = -8/5$*.*

Proof of main result

We return to the situation of general but fixed local index ℓ in order to generalise the phenomena found in the example above. We apply Conjecture 4.4.8 to show that there are only a finite number of choices for reduced basket given a particular total orbifold contribution.

Theorem 4.6.8. *There are only finitely many reduced baskets for a given total orbifold contribution.*

Proof. From previous discussion around Conjecture 4.5.8 it suffices that we consider the case of a single local index ℓ. Recall the maximal shattering map $\rho : \mathbb{N}\langle \mathrm{Res}(\ell)\rangle \to \mathbb{N}\langle \mathcal{R}_\ell\rangle$. Let v be the image of an elementary T-singularity under ρ. As seen previously, in special coordinates this is the vector $(1, 2, \dots, 2, 1)$.

Suppose that $\mathcal{B}$ and $\mathcal{B}'$ are two baskets with the same total orbifold contribution. Then, subject to Conjecture 4.4.8, their maximal shatterings $\rho(\mathcal{B})$ and $\rho(\mathcal{B}')$ must differ by a vector of the form $\lambda \cdot v$ in coordinates as above. Notice also that there are only finitely many baskets $\mathcal{B} \in \mathbb{N}\langle \mathrm{Res}(\ell)\rangle$ with a given maximal shattering $\rho(\mathcal{B}) = \mathcal{T}$ as there are only finitely many ways to glue together the finite number of singularities in $\mathcal{T}$. Hence, if for any given $\mathcal{T} \in \mathbb{N}\langle \mathcal{R}_\ell\rangle$ there is some $\lambda \gg 0$ such that any $\mathcal{B} \in \mathbb{N}\langle \mathrm{Res}(\ell)\rangle$ with $\rho(\mathcal{B}) = \mathcal{T} + \lambda \cdot v$ contains a cancelling tuple, then the result would be shown. The reduced baskets with the total orbifold contribution for $\mathcal{T}$ would then correspond to the preimages under ρ of $\mathcal{T} + \mu \cdot v$ with $0 \le \mu < \lambda$ and that contain no cancelling tuples.

This statement is equivalent to the following, which we will actually prove: for each $\mathcal{T} \in \mathbb{N}\langle \mathcal{R}_\ell\rangle$ and for each $N \in \mathbb{N}$ there is $\lambda \gg 0$ such that every $\mathcal{B} \in \mathbb{N}\langle \mathrm{Res}(\ell)\rangle$ with $\rho(\mathcal{B}) = \mathcal{T} + \lambda \cdot v$ contains at least N cancelling tuples.

We proceed by induction on $|\mathcal{T}| = n$. If $n = 0$, then one can construct arbitrarily many cancelling tuples inside baskets whose maximal shattering is of the form $\lambda \cdot v$ as follows. Since the indecomposable singularities in a maximal shattering of the form $\lambda \cdot v$ coalesce to form a T-singularity τ of width λ, a basket $\mathcal{B}$ of size p with this maximal shattering corresponds to a choice of $p - 1$ lattice points on the edge of τ, for which the corresponding crepant blowups produce the singularities in $\mathcal{B}$. These lattice points must be width less than ℓ apart in order for the corresponding singularities to be residual. Notice that $p - 1 \ge \lambda$ as one has to choose a lattice point inside each of the elementary T-singularities constituting τ. If $\lambda = \ell + 1$ then at least two of the lattice points must differ by a cone of width a multiple of ℓ, which is hence a T-singularity. The cones subtended by the lattice points between these two lattice points thus form a cancelling tuple. Repeating the process starting from the end of this T-singularity, one can produce N cancelling tuples by setting $\lambda = N(\ell + 1)$.

Suppose the statement is true for all $\mathcal{T}$ of size $n-1$. Let $\mathcal{T}$ have size n and choose $\sigma_0 \in \mathcal{T}$. Then $\mathcal{T}_0 = \mathcal{T} \setminus \{\sigma_0\}$ has size $n - 1$ and so there is λ such that every $\mathcal{B}$ with $\rho(\mathcal{B}) = \mathcal{T}_0 + \lambda \cdot v$

contains at least $N+1$ cancelling tuples. By adding the single singularity σ_0 back in to such a basket $\mathcal{B}$, one can reduce the number of cancelling tuples at most by 1. Thus, this same λ has the property that every $\mathcal{B}$ with $\rho(\mathcal{B}) = \mathcal{T} + \lambda \cdot v$ contains at least N cancelling tuples. $\square$

Using the recursion for λ from the induction, one obtains:

Corollary 4.6.9. *For a collection of indecomposable singularities $\mathcal{T}$ of size n, every basket $\mathcal{B}$ with maximal shattering $\rho(\mathcal{B}) = \mathcal{T} + \lambda \cdot v$ contains at least N cancelling tuples when $\lambda \geq (N-n)(\ell+1)$.*

Though the proof above relies on Conjecture 4.4.8, it seems plausible that there is a direct convex geometric proof bypassing this dependency, the most general form of which could look like the following. Recall that an affine ray is a subset of $\mathbb{R}^n$ of the form $\{u + \lambda v : \lambda \geq 0\}$ for some $u, v \in \mathbb{R}^n$ and for some $\lambda_0 \in \mathbb{R}_{\geq 0}$. It is rational if u and v can be chosen to be lattice points.

Conjecture 4.6.10. *Suppose $K = \bigcup_{i=1}^{N} K_i \subseteq \mathbb{R}^n$ is the union of finitely many affine convex rational polyhedra and suppose that $\mathbb{R}^n \setminus K$ contains infinitely many lattice points. Then $\mathbb{R}^n \setminus K$ contains an affine rational ray.*

In combination with Lemma 4.6.6, by choosing K_i appropriately amongst the $K_{v,\delta}$ one obtains the same result but independently of Conjecture 4.4.8. The pertinent set

$$\mathrm{RBod}(\ell, \delta) := (\Phi_\ell^+)^{-1}(\delta) \setminus \bigcup_{v \in \ker \Phi_\ell^+} K_{v,\delta}$$

is then a bounded subset of $\mathbb{Z}\langle \mathrm{Res}^+(\ell) \rangle$ whose lattice points correspond to the finite number of reduced baskets whose total orbifold contribution has δ-vector δ. We call this subset the *reduced body* for the given total orbifold contribution.

Define the *width* of a collection of singularities to be the sum of the widths of the singularities. If the collection arises from shattering a T-singularity τ then its width equals the width of τ, which is also related to its total contribution to the degree via Cor. 4.4.3.

As noted in Cor. 4.6.3 the degree of an orbifold del Pezzo surface X can be read off from its Hilbert series. It has a decomposition

$$K_X^2 = 12 - (\mathcal{RK}_X^2(\ell_1) + \cdots + \mathcal{RK}^2(\ell_N) + \mathcal{IK}^2)$$

where $\mathcal{RK}^2(\ell) := \sum_{\sigma \in \mathcal{RB}(\ell)} A_\sigma$ is the degree contribution from the ℓth piece of the reduced basket and $\mathcal{IK}^2$ is the (nonnegative integral) contribution from the *extended invisible basket*

$$\widehat{\mathcal{IB}} := \mathcal{IB} \cup \{T\text{-singularities on } X\}$$

with

$$\mathcal{IK}^2 = \sum_{\sigma \in \widehat{\mathcal{IB}}} A_\sigma = \sum_{\sigma \in \widehat{\mathcal{IB}}} \frac{\mathrm{width}(\sigma)}{\ell_\sigma}$$

from Cor. 4.4.3. Notice that the Euler number term from Lemma 4.3.7 is accounted for in $\mathcal{IK}_X^2$ by including the T-singularities. Knowing a reduced basket for X then prescribes the degree contribution of the corresponding invisible basket. Notice that the definitions of these quantities only rely on the basket of singularities on X with a choice of invisible/reduced basket and so it m es sense to speak of each of them as associated to just a basket with a choice of invisible/reduced basket.

Example 4.6.11. *For the total orbifold contribution from Example 4.6.7 with δ-vector $(2,1,2) \in \Delta(5)$ the four possible reduced baskets all have $\mathcal{RK}^2 = -8/5$. Suppose one fixes a Hilbert series $H(t)$ with this total orbifold contribution from which one can read the degree K^2. Any invisible basket must have total degree contribution $\mathcal{IK}^2 = 12 - \mathcal{RK}^2 - K_X^2 = \frac{68}{5} - K_X^2$, which is indeed integral since $K_X^2 \equiv \frac{3}{5} \bmod \mathbb{Z}$ for any surface with this Hilbert series.*

We now collect the results of the paper subject to Conjectures 4.4.8 and 4.5.8, and established by the results and discussion of the last two sections.

Theorem 4.6.12. *Fix a power series $H(t) \in \mathbb{N}[[t]]$. Either there are no orbifold del Pezzo surfaces with Hilbert series equal to $H(t)$, or*

- *the reduced basket of an orbifold del Pezzo surface with Hilbert series equal to $H(t)$ is one of a finite number of possibilities, which are determined by the orbifold correction part of $H(t)$ and in bijection with the lattice points of the associated reduced body.*
- *the basket of such an orbifold del Pezzo surface with reduced basket $\mathcal{RB}$ is given by adding to the singularities in $\mathcal{RB}$ a collection of cancelling tuples – which are obtained by shattering T-singularities – whose total degree contribution is determined by $\mathcal{RB}$ and $H(t)$.*

Example 4.6.13. *For the total orbifold contribution*

$$\frac{8t^3 - t^2 + 8t}{5(1 - t^5)}$$

there are 18 possible reduced baskets which all lie in the affine plane $(5,0,0,3) + \ker \Phi$. Unlike in Example 4.6.7, the reduced body contains multiple lattice points along a single ray: $(5,0,0,3) + \lambda(1,0,-1,2)$, which have differing degree contributions. This shows that, in general, the total degree contribution from the reduced basket depends on more than simply the orbifold contribution.

To discuss an application of this result, recall that the Gorenstein index ℓ_X of a Fano variety X is the smallest positive integer m such that mK_X is Cartier. For orbifold del Pezzo surfaces, this is the lowest common multiple of the local indices of all the singularities on X.

Example 4.6.14. *Returning to the total orbifold contribution of Example 4.6.11 with δ-vector $(2,1,2) \in \Delta(5)$ it follows from the above that a basket for an orbifold del Pezzo surface X of Gorenstein index 5 with this orbifold contribution can contain at most 82 singularities: 4 at most from a reduced basket and then at most $5+1=6$ in a minimal cancelling tuple from the discussion of the $n=0$ case in the proof of Theorem 4.6.8, of which there can be at most $13 = \lfloor \frac{68}{5} \rfloor$ from the Fano condition $K_X^2 > 0$. If X is allowed to have Gorenstein index 5ℓ then imitating this calculation yields that the number of singularities on X can be at most*

$$4 + 13(5\ell + 1) = 65\ell + 17$$

This example generalises easily using Theorem 4.6.12 to the following corollary.

Corollary 4.6.15. *Fix a total orbifold contribution $Q \in \mathbb{Q}(t)$ and a positive integer ℓ_*. There exists a number $N(Q, \ell_*)$ dependent only on and computable from Q and ℓ_* such that any orbifold del Pezzo surface with total orbifold contribution Q and Gorenstein index bounded above by ℓ_* has at most $N(Q, \ell_*)$ singularities.*

As in the example, it is straightforward to compute $N(Q, \ell_*)$ using the proof of Theorem 4.6.8 once the reduced bodies for the ℓth pieces of Q have been found.

4.7 Degree bounds

Let $H(t) \in \mathbb{N}[[t]]$. If $H(t)$ is the Hilbert series of an orbifold del Pezzo surface X, then X must have degree given by the formula in Lemma 4.3.6. Denote this number by K_H^2. As seen in the previous section, choosing a reduced basket to capture the total orbifold contribution of $H(t)$ enforces a choice of the degree contribution of a corresponding invisible basket.

More precisely, for a choice of reduced basket $\mathcal{RB}$ and invisible basket $\mathcal{IB}$ one requires

$$K_H^2 = 12 - \mathcal{RK}^2 - \sum_{\sigma \in \widehat{\mathcal{IB}}} \frac{\mathrm{width}(\sigma)}{\ell_\sigma}$$

where $\mathcal{RK}^2 := \sum_{\sigma \in \mathcal{RB}} A_\sigma$. Since the last term is nonnegative, one must have $K_H^2 + \mathcal{RK}^2 \leq 12$. There are three cases:

- if $K_H^2 + \mathcal{RK}^2 < 12$ then there are infinitely many possibilities for the extended invisible basket of an orbifold del Pezzo surface with Hilbert series $H(t)$

- if $K_H^2 + \mathcal{RK}^2 = 12$, then there is only one possible extended invisible basket for an orbifold del Pezzo surface with Hilbert series $H(t)$ containing $\mathcal{RB}$, which is the empty set since there is no more flexibility in the degree allowing one to add T-singularities or cancelling tuples

- If $K_H^2 + \mathcal{RK}^2 > 12$ then there are no possible extended invisible baskets for orbifold del Pezzo surfaces surfaces with Hilbert series $H(t)$.

Testing across all reduced baskets gives the following non-existence result.

Corollary 4.7.1. *With notation as above, if $K_H^2 + \mathcal{RK}^2 > 12$ for all reduced baskets associated to $H(t)$ then there are no orbifold del Pezzo surfaces with Hilbert series $H(t)$.*

For example there are no orbifold del Pezzo surfaces with Hilbert series

$$\frac{1 + mt + t^2}{(1-t)^3}, m \geq 11$$

since the only reduced basket for this power series is the empty set with $\mathcal{RK}^2 = 0$ and so $K_H^2 + \mathcal{RK}^2 = m + 2 > 12$. Observe that in addition there are no toric orbifold del Pezzo surfaces with Hilbert series

$$\frac{1 + 10t + t^2}{(1-t)^3}$$

since a projective toric surface has at least three (possibly smooth) affine pieces or singularities corresponding to the faces of its polygon and so must feature at least one cancelling tuple.

In general, as there are only finitely many reduced baskets for a given total orbifold contribution, the discussion above along with the Fano condition $K_X^2 > 0$ show how to produce bounds on the degree of any orbifold del Pezzo surface with that total orbifold contribution.

Theorem 4.7.2. *For a given collection of residual singularities $\mathcal{B}$ there exist constants $m > 0$ and $M > 0$ dependent only on and computable from $\mathcal{B}$ such that, for any orbifold del Pezzo surface X with $\mathcal{B}_X = \mathcal{B}$,*

$$m \leq K_X^2 \leq M$$

4.8 A different perspective

Broadly speaking, the approach taken so far in this paper has been deformation-theoretic: we have sought to classify the possible collections of singularities that could correspond to baskets of singularities (which are by construction deformation classes of singularities) on an orbifold del Pezzo surface with a given Hilbert series. An alternative perspective one could take is to allow not just deformations but also crepant blowups. From this perspective, residual singularities are replaced by indecomposable singularities as the appropriate notion of 'rigid' singularities.

To briefly explore this perspective, define two invariants $\Psi_\ell(X)$ and $\widetilde{\Psi}_\ell(X)$ of an orbifold del Pezzo surface X by

$$\Psi_\ell(X) := \rho(\mathcal{B}_X(\ell)) \in \mathbb{N}\langle\mathcal{R}_\ell\rangle \text{ and } \widetilde{\Psi}_\ell(X) := \rho(\widetilde{\mathcal{B}}_X(\ell)) \in \mathbb{N}\langle\mathcal{R}_\ell\rangle$$

That is, $\Psi_\ell(X)$ is the collection of indecomposable singularities obtained by maximally shattering all the residues of singularities of local index ℓ on X, and $\widetilde{\Psi}_\ell(X)$ is the collection of indecomposable singularities obtained by maximally shattering all singularities on X. The former allows both deformations and crepant blowups to be taken; the latter allows no deformation, for example, it recognises the T-singularities on X. In this language, Theorem 4.6.12 becomes the following.

Theorem 4.8.1. *Fix a power series $H(t) \in \mathbb{N}[[t]]$. There exists an affine ray $\varrho_{H(t)} \subseteq \mathbb{N}\langle\mathcal{R}_\ell\rangle$ such that any orbifold del Pezzo surface X with Hilbert series $H(t)$ has the property that $\Psi_\ell(X)$ lies on $\varrho_{H(t)}$. Moreover the slope of $\varrho_{H(t)}$ is independent of $H(t)$ and corresponds to a maximally shattered elementary T-singularity of local index ℓ. The same is true for $\widetilde{\Psi}_\ell(X)$ with the same ray $\varrho_{H(t)}$ (possibly truncated).*

The theorem also holds if the Hilbert series $H(t)$ is replaced by a total orbifold contribution $Q \in \mathbb{Q}(t)$.

4.9 Hilbert functions of toric orbifolds

As discussed in Chapter 2 the Hilbert function of ample divisors on toric varieties are given by counting lattice points in polytopes. We develop this perspective in this section, and state results of the author and A. M. Kasprzyk from [50] using techniques from orbifold geometry (primarily deformation theory) to resolve questions in polytope combinatorics.

Ehrhart theory

Let M be a lattice and suppose $Q \subseteq M \otimes_{\mathbb{Z}} \mathbb{R} =: M_{\mathbb{R}}$ is a convex polytope; that is, the convex hull of finitely many points in $M_{\mathbb{R}}$. We will assume that our polytopes are full-dimensional, so that $\dim Q = \operatorname{rank} M$. We will often select $M = \mathbb{Z}^d$. If the vertices of Q are elements of

M, we say that Q is a *lattice polytope*. If its vertices are elements of $M_{\mathbb{Q}} := M \otimes_{\mathbb{Z}} \mathbb{Q}$ then we say that Q is a *rational polytope*.

Let

$$\mathrm{ehr}_Q(k) := |kQ \cap M|$$

be the function $\mathbb{Z}_{\geq 0} \to \mathbb{Z}$ counting the number of lattice points in dilates kQ of Q. We call this function the *Ehrhart function* of Q. The values of the Ehrhart polynomial of Q form a generating function

$$\mathrm{Ehr}_P(t) := \sum_{k \geq 0} \mathrm{ehr}_Q(k) t^k$$

called the *Ehrhart series* of Q.

Ehrhart [38] showed that when Q is a lattice polytope ehr_Q can be expressed as a polynomial of degree $\dim Q$

$$\mathrm{ehr}_Q(k) = c_d k^d + \ldots + c_1 k + c_0$$

which we call the *Ehrhart polynomial* of Q. The leading coefficient c_d is given by $\mathrm{vol}\, Q/d!$, c_{d-1} is equal to $\mathrm{vol}(\partial Q)/2(d-1)!$ and $c_0 = 1$. Here $\mathrm{vol}(\cdot)$ denotes the normalised volume, and ∂Q denotes the boundary of Q. For example, if Q is two-dimensional (that is, Q is a lattice polygon) we obtain

$$\mathrm{ehr}_Q(k) = \frac{\mathrm{vol}(Q)}{2} k^2 + \frac{|\partial Q \cap M|}{2} k + 1$$

Setting $k = 1$ in this expression recovers Pick's Theorem [64].

When P is a rational polytope the situation is more interesting. A *quasi-polynomial* with *period* $s \in \mathbb{Z}_{>0}$ is a function $q : \mathbb{Z} \to \mathbb{Q}$ defined by polynomials $q_0, q_1, \ldots, q_{s-1}$ such that

$$q(k) = q_i(k) \qquad \text{when } k \equiv i \bmod s$$

The degree of q is the largest degree of any monomial appearing in the q_i. The minimum period of q is called the *quasi-period* of q, and necessarily divides any other period. Ehrhart showed that ehr_Q is given by a quasi-polynomial of degree d, which we call the *Ehrhart quasi-polynomial* of Q.

Quasi-period collapse

Let π_Q denote the quasi-period of ehr_Q. The smallest positive integer $r_Q \in \mathbb{Z}_{>0}$ such that $r_Q Q$ is a lattice polytope is called the *denominator* of Q. It is certainly the case that ehr_Q is r_Q-periodic, however it is perhaps surprising that the quasi-period of ehr_Q does not always equal r_Q; this phenomenon is called *quasi-period collapse*.

Example 4.9.1 (Quasi-period collapse). *Consider the triangle*

$$Q := \mathrm{Conv}\{(5,-1),(-1,-1),(-1,1/2)\}$$

with denominator $r_Q = 2$. This has

$$\mathrm{ehr}_Q(k) = \frac{9}{2}k^2 + \frac{9}{2}k + 1$$

and so $\pi_Q = 1$, which is smaller than r_Q.

Quasi-period collapse is poorly understood, although it occurs in many contexts. For example, de Loera–McAllister [35, 36] consider polytopes arising naturally in the study of Lie algebras (Gel'fand–Tsetlin polytopes, and polytopes determined by Clebsch–Gordan coefficients) that exhibit quasi-period collapse. In dimension two McAllister–Woods [55] show that there exist rational polygons with r_Q arbitrarily large but with $\pi_Q = 1$ (see also [50]). Haase–McAllister [42] give a constructive view of this phenomena in terms of $\mathrm{GL}_d(\mathbb{Z})$-*scissor congruence*; here a polytope is partitioned into pieces that are individually modified via $\mathrm{GL}_d(\mathbb{Z})$ transformation and lattice translation, then reassembled to give a new polytope which by construction has equal Ehrhart quasi-polynomial but different denominator.

Quasi-period collapse for dual-Fano polygons

Definition 4.9.2. *A lattice polytope $P \subseteq N_\mathbb{R}$ is* Fano *if its vertices are primitive lattice points, and the origin lies within the interior of P. We say that a rational polytope $Q \subseteq M_\mathbb{R}$ is* dual-Fano *if Q is dual to a Fano polytope.*

The toric variety associated to a dual-Fano polytope via the inner normal fan (or equivalently from a Fano polytope via the spanning fan) is a toric Fano variety. A (dual-)Fano polygon thus defines a toric del Pezzo surface, which must have at worst orbifold singularities.

The Hilbert series of orbifold del Pezzo surfaces were studied previously in this chapter. We partitioned $\mathcal{B}$ into two pieces: a reduced basket and an invisible basket, for which the latter – along with any T-singularities – is not detectable by the Hilbert series. From our viewpoint it is this invisibility that causes quasi-period collapse.

Theorem 4.9.3. *Let $Q \subseteq M_\mathbb{R}$ be a dual-Fano polygon. Let $\mathcal{RB}$ be a reduced basket for X_Q with corresponding invisible basket $\mathcal{IB}$. Let the set of T-singularities on X_Q be denoted $\mathcal{T}$. The quasi-period of Q is bounded by*

$$\pi_Q \leq \mathrm{lcm}\{\ell_\sigma \mid \sigma \in \mathcal{RB}\}$$

Furthermore, Q exhibits quasi-period collapse if there exists some $\tau \in \mathcal{IB} \cup \mathcal{T}$ of local index not dividing $\operatorname{lcm} \ell_\sigma \mid \sigma \in \mathcal{RB}$*. Moreover, the quasi-period collapse is measured by* $\mathcal{IB}$*:*

$$r_Q = \operatorname{lcm}\{\pi_Q\} \cup \{\ell_\sigma \mid \sigma \in \mathcal{IB} \cup \mathcal{T}\}$$

We expect that the inequality in the first part of the theorem is an equality; this follows from stronger versions of the conjectures earlier in this chapter. This result empowers the construction of interesting families of examples – for example, many rational triangles with quasi-period 1 – and (currently only partial) classification results.

Example 4.9.4. *This result can be used to construct an infinite family of dual-Fano triangles corresponding to fake weighted projective planes (cyclic quotients of weighted projective planes) with arbitrary denominator ℓ and quasi-period 1. Let ℓ, w, c be integers with $0 < w, c < \ell$ and $\gcd(\ell, wc-1) = 1$. Let $\gcd(\ell, w) = g$, $w = gh$, and $\ell = gk$. Then, there is a fake weighted projective plane of the form*

$$\mathbb{P}(h, k-h, k)/(\mathbb{Z}/\ell g)$$

with singularities

$$\frac{1}{2\ell}(1, 2c-1), \quad \frac{1}{\ell(\ell-2)}(1, (\ell-2)(\ell-c)-1), \quad \frac{1}{\ell^2}(1, \ell(c-1)-1)$$

whose Hilbert function is a polynomial. The same techniques can be used to produce another infinite family of the form

$$\mathbb{P}(1, \ell-1, \ell)/(\mathbb{Z}/\ell)$$

with singularities

$$\frac{1}{\ell}(1, c-1), \quad \frac{1}{\ell(\ell-1)}(1, (\ell-1)(\ell-c)-1), \quad \frac{1}{\ell^2}(1, \ell(c-2)-1)$$

and again a polynomial Hilbert function. These examples are complementary to the dual-Fano triangles of qG-smoothable weighted projective spaces

$$\mathbb{P}(a^2, b^2, c^2)$$

for (a, b, c) a Markov triple satisfying

$$a^2 + b^2 + c^2 = 3abc$$

as classified by Hacking–Prokhorov [43] that also have polynomial Hilbert function.

Chapter 5

Applications to symplectic embeddings

The final chapter of this thesis addresses the application of numerics on orbifolds to symplectic embedding problems. We will also introduce elements of algebraic positivity, or the numerical asymptotics of divisors on orbifolds, as they also become highly relevant to this story. For the purposes of this chapter, a *symplectic manifold* is a smooth manifold potentially with boundary or corners equipped with a nondegenerate 2-form ω.

5.1 Symplectic embeddings and symplectic capacities

A smooth map $\iota\colon X_1 \to X_2$ between symplectic manifolds (X_1, ω_1) and (X_2, ω_2) is a *symplectic embedding* if it is a smooth embedding such that $\iota^*\omega_2 = \omega_1$.

Symplectic capacities are numerical invariants measuring obstructions to embedding one symplectic manifold into another. Perhaps the simplest such obstruction is the volume; a symplectic manifold (X_1, ω_1) can be embedded in another symplectic manifold (X_2, ω_2) only if $\mathrm{vol}(X_1, \omega_1) \leq \mathrm{vol}(X_2, \omega_2)$. A more sophisticated obstruction is the Gromov width: the supremum of the radii of balls that can symplectically embed into the given symplectic manifold. As Gromov's nonsqueezing theorem [40] illustrates, this is a nontrivial and interesting invariant even for simple submanifolds of $\mathbb{R}^n$.

There are many different capacities in past and current usage - see [19] and the numerous references therein for an overview - that were invented in order to answer more sophisticated embedding questions about symplectic 4-manifolds. In this paper we will focus on *Embedded Contact Homology* or *ECH capacities*, which were developed by Hutchings in [44] and have since been studied by many authors in, for example, [18], [31], [30], [33]. To an exact symplectic 4-manifold X with contact-type boundary they associate an increasing sequence of real numbers

$$c_k(X) \text{ for } k \in \mathbb{Z}_{\geq 0}$$

One of their early successes was studying embeddings of ellipsoids where the ellipsoid with symplectic radii a, b

$$E(a,b) := \{(x,y) \in \mathbb{C}^2 : \frac{|x^2|}{\pi a} + \frac{|y^2|}{\pi b} < 1\}$$

embeds into $E(c,d)$ iff $c_k(E(a,b)) \leq c_k(E(c,d))$ for all k. Moreover, $c_k(E(a,b))$ was computed to be the kth largest number of the form $am + bn$ for $m, n \in \mathbb{Z}_{\geq 0}$. We will introduce ECH capacities more formally in §5.2.

A particular type of symplectic manifold that ECH capacities provide an attractive means of studying is toric domains. Consider the moment map

$$\mu : \mathbb{C}^2 \to \mathbb{R}^2$$

for the 2-torus action on $\mathbb{C}^2$. Given a region $\Omega \subseteq \mathbb{R}^2$, $X_\Omega := \mu^{-1}(\Omega)$ is a toric symplectic 4-manifold potentially with boundary. If the domain Ω is a certain kind of convex polygon with two edges lying on the coordinate axes, X_Ω is called a *convex toric domain*. We omit mention of the symplectic form since we will always take the induced form from $\mathbb{C}^2$. Such symplectic manifolds are exact with contact-type boundary. The work of Cristofano-Gardiner–Choi [30] provides a somewhat combinatorial formula for the ECH capacities of such spaces in terms of lattice paths and lattice point counts, which we will make heavy use of.

5.2 ECH capacities

Combinatorial definitions

ECH is formally defined in terms of contact geometry. It is constructed explicitly in [44] however there is a combinatorial rephrasing of ECH in the case of toric domains that is most applicable to the situation at hand, which is how we will primarily present it here. This material comes from [18, 30, 44].

Suppose $\Omega \subseteq \mathbb{R}^2$ is any polygon. Define the Ω-*length* of a vector v to be

$$\ell_\Omega(v) := v \times p_v$$

where p_v is a boundary point of Ω such that Ω is contained in the right halfplane bounded by the line spanned by v translated to contain p_v. Here $\times$ means the cross product $u \times v = \det(u \mid v)$. Define the Ω-length of a piecewise linear path Λ to be

$$\ell_\Omega(\Lambda) = \sum \ell_\Omega(v_i)$$

where the sum ranges over the edge vectors v_i of Λ. Notice that, from a local calculation, one has

$$\ell_\Omega(\partial\Omega) = 2\operatorname{Vol}(\Omega)$$

Definition 5.2.1. *A* convex domain *is a convex region $\Omega \subseteq \mathbb{R}^2$ whose boundary consists of*

- *a line segment between the origin and a point $(a, 0)$ on the positive horizontal axis*
- *a line segment between the origin and a point $(0, b)$ on the positive vertical axis*
- *the graph of a convex piecewise linear function $f : [0, a] \to [0, b]$*

We say that a convex domain is a convex lattice domain if the points $(a, 0)$ and $(0, b)$ are lattice points and if the function f is piecewise linear such that each vertex is a lattice point. In other words, a convex lattice domain is a convex domain that is also a lattice polygon. Convex rational domains are defined similarly.

We call the corresponding symplectic manifold $X_\Omega = \mu^{-1}(\Omega)$ a *convex toric domain* if Ω is a convex domain, or a *convex toric lattice domain* if Ω is a convex lattice domain. One can also repeat these definitions with convex replaced by *concave*.

Following [30] - which built on [18,56] - the *weight sequence* associated to a convex lattice domain Ω is a sequence $w(\Omega)$ of numbers defined as follows. Let Δ_a be the convex hull of the points $(0, 0), (a, 0), (0, a)$. Let c be the smallest number such that $\Omega \subseteq \Delta_c$. Equivalently, c is the radius of the smallest ball in $\mathbb{C}^2$ containing X_Ω. The two components of the complement $\Delta_c \backslash \Omega$ are affine equivalent to two concave domains Ω_2 and Ω_3. There is a recursive definition weight sequences for concave domains as follows. Consider the concave domain Ω_2. Let b_1 be the largest real number such that $\Delta_{b_1} \subseteq \Omega_2$. The complement of Δ_{b_1} in Ω_2 consists of two (possibly empty) concave domains and so one can recurse to obtain a multiset of numbers $w(\Omega_2) := \{b_1, b_2, \dots\}$. We define

$$w(\Omega) := (c; w(\Omega_2); w(\Omega_3))$$

Example 5.2.2. *Let $\Omega = \operatorname{Conv}((0, 0), (0, 2a), (a, a), (a, 0))$ for some $a \in \mathbb{Z}_{>0}$. Here $c = 2a$, leaving a single concave region Ω_3 illustrated in the third figure, which is affine equivalent to Δ_a. The weight sequence for Ω is hence $(2a; \emptyset; a)$.*

Construction

Using the constructions above, we define ECH capacities combinatorially.

Figure 5.1: Example of weight sequence

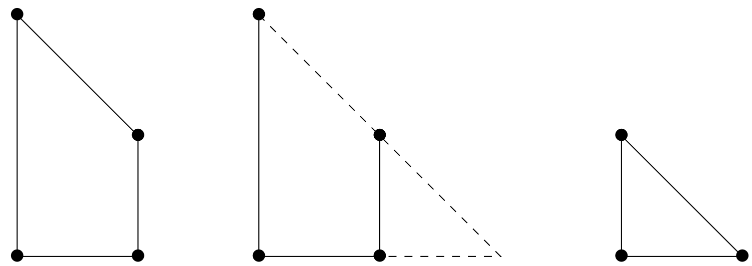

Definition 5.2.3. *A* convex lattice path *is a piecewise linear path starting on the positive vertical axis and ending on the positive horizontal axis such that its vertices are lattice points.*

After adding the pieces along the coordinate axes, convex lattice paths are exactly boundaries of convex lattice domains. For a polygon Λ we denote by L_Λ the number of lattice points enclosed by Λ, including on those its boundary. We state a result of Cristofaro-Gardiner but using the perspective of convex lattice domains instead of convex lattice paths.

Theorem 5.2.4 ([30, Cor. 8.5])**.** *Let Ω be a convex domain. Then*

$$c_k(X_\Omega) = \min\{\ell_\Omega(\partial\Lambda) : L_\Lambda = k+1\}$$

where the minimum is taken over convex lattice domains Λ.

Define the *cap function* of a symplectic 4-manifold X with contact-type boundary to be

$$\begin{aligned} \mathrm{cap}_X(r) &:= \#\{k \in \mathbb{Z}_{\geq 0} : c_k(X) \leq r\} \\ &= 1 + \max\{k \in \mathbb{Z}_{\geq 0} : c_k(X) \leq r\} \end{aligned}$$

for $r \in \mathbb{Z}_{\geq 0}$. In certain situations - such as ellipsoids with integral radii - the cap function recovers all of the ECH capacities.

Corollary 5.2.5. *If Ω is a convex domain, then*

$$\mathrm{cap}_{X_\Omega}(r) = \max\{L_\Lambda : \ell_\Omega(\partial\Lambda) \leq r\}$$

where the maximum ranges over convex lattice domains Λ.

Proof. After including the zeroth capacity, one has

$$\begin{aligned} \mathrm{cap}_{X_\Omega}(r) &= \#\{k : \exists\Lambda \text{ with } \ell_\Omega(\Lambda) \leq r \text{ and } L_\Lambda = k+1\} \\ &= 1 + \max\{k : \exists\Lambda \text{ with } \ell_\Omega(\Lambda) \leq r \text{ and } L_\Lambda = k+1\} \\ &= \max\{L_\Lambda : \ell_\Omega(\Lambda) \leq r\} \end{aligned}$$

as required. $\square$

ECH capacities and weight sequences

The weight sequence $w(\Omega)$ contains all the information required to compute $c_k(X_\Omega)$.

Lemma 5.2.6. *Suppose* $w(\Omega) = (c; a_1, \ldots, a_s; b_1, \ldots, b_t)$. *Then*

$$c_k(X_\Omega) = \min\{c_{k+k_2+k_3}(B(c)) - c_{k_2}(\amalg_{i=1}^s B(a_i)) - c_{k_3}(\amalg_{j=1}^t B(b_j)) : k_2, k_3 \in \mathbb{Z}_{\geq 0}\}$$

This follows from [30, Cor. A.5] combined with [18, Theorem 1.4].

Key properties of ECH capacities

ECH capacities have the following properties recorded in [18], which we will use throughout the chapter:

- **Monotonicity:** If (X, ω) embeds into (X', ω') then $c_k(X, \omega) \leq c_k(X', \omega')$ for all k
- **Disjoint union:** If $(X, \omega) = \amalg_{i=1}^n (X_i, \omega_i)$ then
$$c_k(X, \omega) = \max_{\sum k_i = k} \sum_{i=1}^n c_{k_i}(X_i, \omega_i) \text{ for all } k$$
- **Conformality:** For each k and $\lambda \in \mathbb{R}^+$, $c_k(X, \lambda\omega) = \lambda c_k(X, \omega)$

Asymptotics of ECH capacities

Asymptotically, capacities return the volume constraint for symplectic embeddings.

Theorem 5.2.7 ([31, Theorem 1.1])**.** *Suppose* (X, ω) *is a compact symplectic* 4*-manifold, then*

$$\lim_{k\to\infty} \frac{c_k(X, \omega)^2}{k} = 4\,\mathrm{Vol}(X, \omega)$$

When $X = X_\Omega$ *is a convex toric domain, this limit also equals* $4\,\mathrm{Vol}(\Omega)$.

This is known as the 'Weyl law' for ECH capacities.

5.3 Algebraic formulation of ECH capacities

There is a purely algebro-geometric framework that we will establish in which one can recast ECH capacities for convex toric domains arising from rational polygons. This also applies for a different class of toric domains called *free convex toric domains* that are defined in

§5.3. For a rational convex domain $\Omega \subseteq \mathbb{R}^2$ (or an irrational convex domain with rational slopes) one can associate the toric surface $Y_{\Sigma(\Omega)}$ associated to the inner normal fan of Ω. In this setting Ω defines a divisor D_Ω on $Y_{\Sigma(\Omega)}$.

Theorem 5.3.1. *(Theorem 5.3.16 + Theorem 5.3.17 + Theorem 5.3.20) Let Ω be any rational convex domain or a free rational convex toric domain. Then*

$$c_k(X_\Omega) = \min\{D \cdot D_\Omega : h^0(Y_{\Sigma(\Omega)}, D) \geq k+1\}$$
$$\mathrm{cap}_{X_\Omega}(r) = \max\{h^0(Y_{\Sigma(\Omega)}, D) : D \cdot D_\Omega \leq r\}$$

where both extrema range over all nef $\mathbb{Q}$- or $\mathbb{R}$-divisors on $Y_{\Sigma(\Omega)}$.

It is irrelevant whether we work with $\mathbb{Q}$- or $\mathbb{R}$-divisors. For the special case evaluating the cap function at $r\ell_\Omega(\partial\Omega) =: \lambda r$, we have

$$\mathrm{cap}_{X_\Omega}(\lambda r) = \max\{h^0(D) : (D - rD_\Omega) \cdot D_\Omega \leq 0\}$$

We can also state some of the results of this chapter in purely combinatorial terms. We will later describe a pseudonorm ℓ_Ω dependent on Ω called the Ω-*length*, which is central to the combinatorialisation of ECH capacities. For a polygon Λ, define its Ω-*perimeter* $\ell_\Omega(\partial\Lambda)$ to be the sum of the Ω-lengths of the line segments composing its boundary $\partial\Lambda$.

Theorem 5.3.2. *(Cor. 5.4.11) Suppose Ω is a tightly constrained convex lattice domain with lower bound*[1] *$r_0 = 0$ and let $\lambda = \ell_\Omega(\partial\Omega)$. Then $r\Omega$ contains the most lattice points of any convex lattice domain of Ω-perimeter at most $r\lambda$ for all $r \in \mathbb{Z}_{\geq 0}$.*

Ω-stretching

Consider a convex domain $\Omega \subseteq \mathbb{R}^2$. For a polygon Λ define $S_\Omega\Lambda$ to be the polygon with edges parallel to the edges of Ω by placing an edge of slope v_i at the point or points at which v_i is tangent to Λ, using corners if necessary. For example,

We call the resulting polygon $S_\Omega\Lambda$ the Ω-*stretching* of Λ. The following lemma is due to Michael Hutchings.

Lemma 5.3.3. $\ell_\Omega(\partial\Lambda) = \ell_\Omega(\partial S_\Omega\Lambda)$.

Proof. Let $p_1, \ldots, p_k$ denote the vertices of Ω. Let q_i be a point on $\partial\Lambda$ such that a tangent vector to $\partial\Lambda$ at q_i is parallel to the vector $p_i - p_{i-1}$. Then by definition, the Ω-length of $\partial\Lambda$ is

$$\sum_i p_i \times (q_{i+1} - q_i)$$

[1] For example, these assumptions are met if one of the weights of Ω is equal to 1, and we conjecture that they are met whenever the gcd of the weights is 1.

Figure 5.2: Example of Ω-stretching

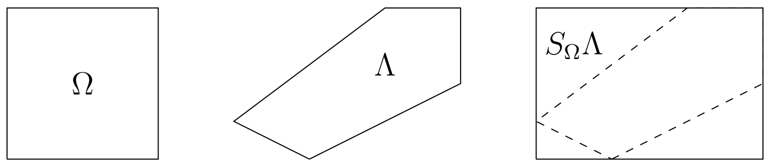

Notice that the same points q_i still satisfy the requirements for computing the Ω-length of $\partial S_\Omega \Lambda$, so that the nothing changes in the expression of $\ell_\Omega(\partial S_\Omega\Lambda)$ from that for $\ell_\Omega(\partial\Lambda)$. $\square$

The effect of Ω-stretching is to produce a polygon of the same Ω-length but with edges parallel to the edges of Ω.

Slope polytopes

Let Ω be a rational convex domain. Denote its set of edges by $\mathrm{Edge}(\Omega)$. An *edge-orientation* $\mathfrak{o}$ of Δ is an orientation of each of its edges in such a way that the boundary of Δ is an oriented cycle. A polygon with an edge-orientation is called *edge-oriented*. Given an edge-oriented convex lattice domain Ω, define the *slope* v_e of an edge $e \in \mathrm{Edge}(\Omega)$ to be the primitive lattice vector in the direction of the oriented edge. That is, of e has endpoints e_- and e_+ with orientation making e_- the tail and e_+ the head, v_e is the primitive ray generator of the ray $\mathbb{R}_{\geq 0} \cdot (e_+ - e_-)$.

Definition 5.3.4. *The* slope polytope *of an edge-oriented convex lattice domain Ω is the lattice polytope*

$$\mathrm{Sl}(\Omega) := \mathrm{Conv}(v_e : e \in \mathrm{Edge}(\Omega))$$

This produces a compact toric variety $Y_{\mathrm{Sl}(\Omega)}$ (via the spanning fan) on which the algebraic geometry side of the story will take place. We will actually work with a blowup of this toric variety, which we will denote by $\widetilde{Y}_{\mathrm{Sl}(\Omega)}$.

This blowup is obtained by creating a new fan by inserting rays through any slopes v_e that are not vertices of $\mathrm{Sl}(\Omega)$. For example, suppose that Ω has slopes $-e_1, e_2, e_1, e_1 - e_2, e_1 - 2e_2$. The slope polytope only has vertices $-e_1, e_2, e_1, e_1 - 2e_2$ and so one extra ray has to be added for $e_1 - e_2$. This is demonstrated pictorally below.

We will denote the resulting fan for the blowup by $\widetilde{\Sigma}_{\mathrm{Sl}(\Omega)}$. Observe that this fan is in some sense a rotation of the inner normal fan of Ω after picking bases, though they naturally

Figure 5.3: Blowup of $Y_{\mathrm{Sl}(\Omega)}$

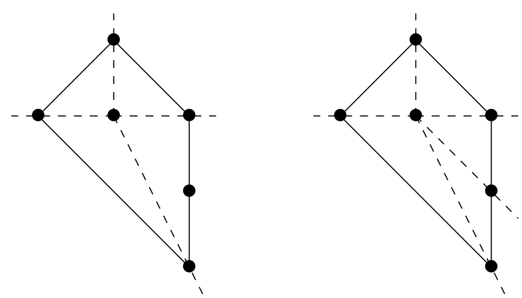

Fan for $Y_{\mathrm{Sl}(\Omega)}$ Fan for $\widetilde{Y}_{\mathrm{Sl}(\Omega)}$

live in dual lattices. At the end of this section we will provide an alternative version of the content below phrased in terms of the inner normal fan instead of the slope polytope. It can be favourable to use each of these perspectives at different times.

Balance divisors

As above, let Ω be a rational convex domain oriented clockwise with slope polytope $\mathrm{Sl}(\Omega)$. We will subsequently always assume that Ω has this orientation. Define the *Ω-length* of a vector $v \in \mathbb{R}^2$ to be

$$\ell_\Omega(v) := v \times p_v$$

where p_v is a boundary point of Ω such that the halfplane $p_v + \{u \in \mathbb{R}^2 : u \times v \geq 0\}$ contains Ω. Recall that the two-dimensional cross product $u \times v$ of two vectors u and v is defined to be the determinant of the matrix with u and v as first and second columns respectively.

Lemma 5.3.5. *Suppose Ω is lattice (resp. rational). The Ω-length is an integral (resp. rational) support function for the fan $\widetilde{\Sigma}_{\mathrm{Sl}(\Omega)}$.*

Proof. Suppose v, v' are adjacent slopes in Ω. The Ω-length applied to any vector $w \in \mathrm{Cone}(v, v') = \sigma$ is given by

$$\ell_\Omega(w) = p \times w$$

where p is the vertex shared between the two edges of slopes v and v' respectively. This is linear on the cone σ, which features in $\Sigma_{\mathrm{Sl}(\Omega)}$ by definition and describes all full-dimensional cones in $\Sigma_{\mathrm{Sl}(\Omega)}$ as v, v' range over adjacent slopes. ℓ_Ω is clearly integral on integral vectors when the vertices of Ω are lattice points, and similarly for the rational case. $\square$

Definition 5.3.6. *The* balance divisor *for Ω is the $\mathbb{Q}$-Cartier divisor D_Ω associated with the support function $-\ell_\Omega$. Notice the change in sign.*

Corollary 5.3.7. *The coefficients of D_Ω as a Weil divisor are*

$$a_v = \ell_\Omega(v)$$

for a (primitive) slope vector v of Ω.

Corollary 5.3.8. *D_Ω is ample.*

Proof. It is a straightforward check that $-\ell_\Omega$ is a strictly convex function, which corresponds to D_Ω being ample. □

Lemma 5.3.9. *The polytope for D_Ω is the result of rotating Ω by 90° anticlockwise around the origin.*

Proof. Denote by $\Omega^\perp$ the rotated version of Ω. The edges of Ω are by construction orthogonal to the rays of $\widetilde{\Sigma}_{\mathrm{Sl}(\Omega)}$ and so there is a facet presentation of $\Omega^\perp$ coming from this fan or, equivalently, a divisor D on $\widetilde{Y}_{\mathrm{Sl}(\Omega)}$. Order the slopes $v_1, \ldots, v_s$ with corresponding toric boundary divisors $D_1, \ldots, D_s$. It suffices that the coefficient a_i of D along D_i is the same as the corresponding coefficient in D_Ω. We will now compute this directly. The edge e_i of $P^\perp$ with slope v_i is carved out by the orthogonal hyperplanes to v_{i-1}, v_i, v_{i+1}. Suppose that v_{i-1}, v_{i+1} form a $\mathbb{Z}$-basis for $\mathbb{Z}^2$. They are independent over $\mathbb{Q}$ and the case when they are not a $\mathbb{Z}$-basis is similar. By a change of coordinates, suppose $v_{i-1} = (1,0), v_{i+1} = (0,-1)$ and $v_i = (\alpha, \beta)$. Then the endpoints of the edge in $\Omega^\perp$ corresponding to v_i are

$$\left(-a_{i-1}, \frac{\alpha a_{i-1} - a_i}{\beta}\right) \text{ and } \left(-\frac{\beta a_{i+1} + a_i}{\alpha}, a_{i+1}\right)$$

After rotating back, the Ω-length of v_i is then

$$\ell_\Omega(v_i) = \begin{vmatrix} \alpha & \beta \\ a_{i+1} & \frac{\beta a_{i+1} + a_i}{\alpha} \end{vmatrix} = a_i$$

which is the same as the corresponding coefficient in D_Ω. □

Corollary 5.3.10. *$L_{r\Omega} = h^0(rD_\Omega)$.*

Notice that there are many choices of Ω with the same slope polytope $\mathrm{Sl}(\Omega)$ and so to reflect the choice of Ω an extra choice has to be made in the geometry. This choice is a polarisation, where $\widetilde{Y}_{\mathrm{Sl}(\Omega)}$ is polarised by the ample divisor D_Ω. The same proof actually shows:

Corollary 5.3.11. *Let Λ be a polygon with all edges parallel to edges of Ω. Denote by $\Lambda^\perp$ the $90°$ anticlockwise rotation of Λ about the origin. The coefficients of a divisor D_Λ on $\widetilde{Y}_{\mathrm{Sl}(\Omega)}$ with polygon $\Lambda^\perp$ are*

$$D_\Lambda = \sum \ell_\Lambda(v) D_v$$

with notation as above.

When Ω is lattice, D_Ω is Cartier. Cartier divisors can also be characterised by their Cartier data, which has a toric version found in §4.2 of [23]. To this end, let $(a,b)^\perp := (-b,a)$. This has the property that $-u \cdot v^\perp = u \times v$.

Corollary 5.3.12. *The Cartier data for D_Λ is $m_{\sigma_i} = q_i^\perp$, where q_i is the vertex in common between the edges of slopes v_i, v_{i+1}, the vertices in $\Sigma(\Omega)$ bounding σ_i.*

Proof. As seen, $v_i \times q_i = a_i$ and so $v_i \cdot q_i^\perp = -a_i$. □

The balance divisor also captures the Ω-length by how it intersects other divisors. We will prove the following lemma in toric geometry to progress towards this.

Lemma 5.3.13. *Let X_Σ be a projective toric surface. An $\mathbb{R}$-divisor D on X_Σ is nef iff $D \cdot D_\rho$ equals the lattice length of the edge of $P(D)$ corresponding to ρ for each ray $\rho \in \Sigma(1)$.*

Proof. The if part is clear by the toric Kleiman condition. For the converse, observe that if D is ample then there is a unique facet presentation of $\Lambda^\perp := P(D)$ as every slope is represented by an edge in $P(D)$. This means that D must be equal to

$$\sum \ell_\Lambda(u_\rho) D_\rho$$

adapting notation from Cor. 5.3.11 and the result follows from the proof of that corollary. If D is nef, then it must be the case that some of the inequalities in the facet presentation are only just redundant: that is, none of the hyperplanes have empty intersection with $P(D)$, but some might only intersect at a vertex. This follows as the interior of the nef cone is the ample cone, or from the description of nef and ample divisors in [8] Theorem 2.15 or [23] Theorem 6.4.9. It suffices to show that $D \cdot D_\rho = 0$ for any ρ giving a redundant hyperplane (that is, an edge of length 0) but this follows from a direct calculation using [23] Prop. 6.4.4. □

Suppose that Λ is a polygon with edges parallel to the edges of Ω. As discussed above, there is a facet presentation of $\Lambda^\perp$ and so there is a nef divisor D_Λ on $\widetilde{Y}_{\mathrm{Sl}(\Omega)}$ with this as its polygon.

Lemma 5.3.14. $\ell_\Omega(\partial\Lambda) = D_\Lambda \cdot D_\Omega$.

Proof. From Lemma 5.3.13, the lattice length of the edge of slope v_i in Λ is $D_\Lambda \cdot D_i$. The Ω-length of the edge is thus $(D_\Omega \cdot D_i) \cdot \ell_\Omega(v_i)$. Summing all these up gives the Ω-perimeter as

$$\ell_\Omega(\partial\Lambda) = \sum (D_\Lambda \cdot D_i) \cdot \ell_\Omega(v_i) = D_\Lambda \cdot \sum \ell_\Omega(v_i) D_i = D_\Lambda \cdot D_\Omega$$

as required. □

Corollary 5.3.15. $\ell_\Omega(\partial\Lambda) = \ell_\Lambda(\partial\Omega)$.

Proof of Theorem 5.3.1

We are now in a position to convert the definition of ECH capacities and cap functions into purely algebro-geometric language.

Theorem 5.3.16. *Suppose Ω is a rational convex domain. Then*

$$c_k(X_\Omega) = \min_D \{D \cdot D_\Omega : h^0(\widetilde{Y}_{\mathrm{Sl}(\Omega)}, D) \geq k+1\}$$
$$\mathrm{cap}_{X_\Omega}(r) = \max_D \{h^0(\widetilde{Y}_{\mathrm{Sl}(\Omega)}, D) : D \cdot D_\Omega \leq r\}$$

where both extrema range over all nef $\mathbb{Q}$- or $\mathbb{R}$-divisors D on $\widetilde{Y}_{\mathrm{Sl}(\Omega)}$.

Proof. Since intersection with D_Ω describes the Ω-length and the number of lattice points enclosed equals h^0, the only thing to check is that the extrema ranging over nef $\mathbb{Z}$-, $\mathbb{Q}$- or $\mathbb{R}$-divisors is equivalent to ranging over convex lattice paths. We will focus on the real case from which it will eventually be apparent why the minima are achieved by integral nef divisors. We use nef divisors to ensure that each 'edge length' $D \cdot D_i$ is nonnegative. Note that the two equalities in the theorem are equivalent and so we will focus only on the first. For convenience denote

$$c_k^{\mathrm{alg}}(\widetilde{Y}_{\mathrm{Sl}(\Omega)}, D_\Omega) = \inf\{D \cdot D_\Omega : h^0(\widetilde{Y}_{\mathrm{Sl}(\Omega)}, D) \geq k+1\}$$

Note that a minimum really is attained. Indeed, pick a nef $\mathbb{R}$-divisor $D_\star$ with at least $k+1$ global sections. Then $c_k^{\mathrm{alg}}(\widetilde{Y}_{\mathrm{Sl}(\Omega)}) \leq D_\star \cdot D_\Omega$ and the infimum is the same if we take it over all nef $\mathbb{R}$-divisors with $h^0(\widetilde{Y}_{\mathrm{Sl}(\Omega)}, D) \geq k+1$ and $D \cdot D_\Omega \leq D_\star \cdot D_\Omega$. Observe that this extra condition places an upper bound on each of the (nonnegative) lattice lengths of edges of the polygon $P(D)$ for such D. This infimum thus takes place over a compact region inside the (closed) nef cone and is therefore realised by some divisor.

Suppose that $D = D_\Lambda$ realises this minimum. Its (rotated) polygon Λ must have a lattice point on every edge as otherwise one could perturb the coefficient in the facet presentation for an edge with no lattice point to obtain a divisor with the same number of global sections but smaller intersection with D_Ω. Notice that this implies that D is a $\mathbb{Q}$-divisor. Let Λ' be the convex hull of all lattice points in Λ. Note that $k'+1 = L_\Lambda = L_{\Lambda'}$ for some $k' \geq k$. Then $S_\Omega \Lambda' = \Lambda$ by construction (as we assumed that Λ has a lattice point on each edge) and so by Lemma 5.3.3 and Lemma 5.3.14 we have $\ell_\Omega(\partial\Lambda') = \ell_\Omega(\partial S_\Omega \Lambda) = \ell_\Omega(\partial\Lambda) = D \cdot D_\Omega$. Now we will show that, potentially after translation, $\partial\Lambda'$ is a convex lattice path in the sense of Definition 5.2.3.

Λ has two distinguished (possible length 0) edges of slopes $-e_1$ and e_2 by construction of $\mathrm{Sl}(\Omega)$ that meet at a point p_0. For these edges to each contain a lattice point, they must each be subsets of affine lines of the form $(y = \beta)$ and $(x = \alpha)$ respectively for some $\alpha, \beta \in \mathbb{Z}$. Hence $p_0 = (\alpha, \beta) \in \mathbb{Z}^2$ is a lattice point. We can thus use this lattice point to translate Λ back to the origin without changing the pairing with D_Ω (the Ω-length) or the dimension of global sections. By convexity Λ' thus also contains two adjacent edges with slopes $-e_1$ and e_2. Since Λ has slopes parallel to the slopes of Ω and is convex, the boundary of Λ forms a convex rational path in the sense of Definition 5.2.3. It follows that the boundary of Λ' forms a convex lattice path and hence features in the minimum of Theorem 5.2.4 giving the combinatorial formula for $c_{k'}(X_\Omega)$. Consequently,

$$c_k(X_\Omega) \leq c_{k'}(X_\Omega) \leq \ell_\Omega(\partial\Lambda') = \ell_\Omega(\partial\Lambda) = D \cdot D_\Omega = c_k^{\mathrm{alg}}(\widetilde{Y}_{\mathrm{Sl}(\Omega)}, D_\Omega)$$

For the converse inequality, suppose that Λ is a lattice polygon whose boundary $\partial\Lambda$ is a convex lattice path realising the minimum of Theorem 5.2.4. That is, $c_k(X_\Omega) = \ell_\Omega(\partial\Lambda)$ and $L_\Lambda = k+1$. Then $\Xi = S_\Omega\Lambda$ is a rational polygon with edges parallel to the edges of Ω, which hence defines a nef $\mathbb{Q}$-divisor D_Ξ on $\widetilde{Y}_{\mathrm{Sl}(\Omega)}$. Now, using Lemma 5.3.3 and Lemma 5.3.14,

$$c_k(X_\Omega) = \ell_\Omega(\partial\Lambda) = \ell_\Omega(\partial S_\Omega\Lambda) = D_\Xi \cdot D_\Omega$$

Notice that $S_\Omega\Lambda$ contains at least as many lattice points as Λ and so $h^0(\widetilde{Y}_{\mathrm{Sl}(\Omega)}, D_\Xi) \geq k+1$ giving

$$c_k^{\mathrm{alg}}(\widetilde{Y}_{\mathrm{Sl}(\Omega), D_\Omega}) \leq D_\Xi \cdot D_\Omega = c_k(X_\Omega)$$

which supplies the converse inequality. □

Notice that $c_k^{\mathrm{alg}}(\widetilde{Y}_{\mathrm{Sl}(\Omega)}, D_\Omega)$ uses $h^0 \geq k+1$ instead of equality (as in the original optimisation problem for ECH capacities in Theorem 5.2.4) because there might not be divisors on $\widetilde{Y}_{\mathrm{Sl}(\Omega)}$ with $k+1$ sections; for example, there are no divisors D on $\mathbb{P}^2$ with $h^0(\mathbb{P}^2, D) = 2$. Combinatorially, this comes from the fact that the lattice paths in Definition 5.2.3 are allowed any rational slopes whereas the paths coming from divisors in Theorem 5.3.16 must have edges parallel to edges of Ω.

Reformulation in terms of the inner normal fan

There is another fan one can associate to a polytope P now living in $M_{\mathbb{R}}$ called the *inner normal fan* $\Sigma(P)$, which consists of cones in $N_{\mathbb{R}}$. For a polygon $P \subseteq \mathbb{R}^2$, this is the fan with rays generated by inward-pointing normals to each of the faces and with all two-dimensional cones between them included. Observe that, after picking a basis as we implicitly did above, the fan $\widetilde{\Sigma}_{\mathrm{Sl}(\Omega)}$ for the blowup $\widetilde{Y}_{\mathrm{Sl}(\Omega)}$ of $Y_{\mathrm{Sl}(\Omega)}$ is the $90°$ anticlockwise rotation of $\Sigma(\Omega)$: taking slopes is dual to taking normals.

Completely analogously, we obtain a toric variety $Y_{\Sigma(Q)}$ that is isomorphic to the previous toric variety $\widetilde{Y}_{\mathrm{Sl}(\Omega)}$ with an ample divisor D_Ω whose coefficient along the divisor D_ρ is $\ell_\Omega(v)$, where ρ is the ray generated by a normal to the edge of slope v. $(Y_{\Sigma(\Omega)}, D_\Omega)$ has the same intersection theoretic and cohomological properties as the pair $(\widetilde{Y}_{\mathrm{Sl}(\Omega)}, D_\Omega)$ and so the results of the previous subsections exactly cross over to this setting.

Theorem 5.3.17. *Suppose X_Ω is a rational convex toric domain. Then*

$$c_k(X_\Omega) = \min\{D \cdot D_\Omega : h^0(Y_{\Sigma(\Omega)}, D) \geq k+1\}$$
$$\mathrm{cap}_{X_\Omega}(r) = \max\{h^0(Y_{\Sigma(\Omega)}, D) : D \cdot D_\Omega \leq r\}$$

where both extrema range over all nef $\mathbb{Q}$- or $\mathbb{R}$-divisors on $Y_{\Sigma(\Omega)}$.

We remark that the advantage of the inner normal fan in this context is its familiarity as a standard object of toric algebraic geometry, however the approach via slope polytopes is quite pleasing and may have better duality properties. For the sake of familiarity and consistency, we will continue to use $\Sigma(\Omega)$ instead of $\widetilde{\Sigma}_{\mathrm{Sl}(\Omega)}$ for the remainder of the chapter.

Free convex toric domains

One can also consider the situation when $\Omega \subseteq \mathbb{R}^2$ is a convex body that doesn't intersect the coordinate axes, which is where fibres of the moment map decrease in dimension and pick up nontrivial isotropy. We call such X_Ω *free convex toric domains.* This was one of the situations originally considered in [44]. There is an analogous theorem there to Theorem 5.2.4. To state it, we define for such Ω a new pseudonorm $\ell_\Omega^{v_\star}$ depending on a vector $v_\star \in \Omega^\circ$ as follows. Consider $\Omega' = \Omega - v_\star$. This is now a polygon with the origin in its interior. We consider the norm $||\cdot||_{\Omega'}$ whose unit ball is Ω' and its dual norm on $(\mathbb{R}^2)^*$

$$||\phi||^*_{\Omega'} := \max\{\phi(v) : v \in \Omega'\}$$

We identify $(\mathbb{R}^2)^*$ with $\mathbb{R}^2$ via the dot product, giving

$$||u||^*_{\Omega'} := \max\{u \cdot v : v \in \Omega'\}$$

Define the length in this pseudonorm of a polygonal path ψ consisting of line segments $v_1, \ldots, v_r$ to be

$$\ell_\Omega^{v_\star}(\psi) := \sum_{i=1}^{r} ||v_i||_{\Omega'}^*$$

Lemma 5.3.18 ([45], Exercise 4.13)**.** *The length of closed polygonal paths measured in $\ell_\Omega^{v_\star}$ is independent of $v_\star$.*

We denote the restriction of $\ell_\Omega^{v_\star}$ to closed polygonal paths by ℓ'_Ω to indicate its independence of $v_\star$.

Theorem 5.3.19 ([44], Theorem 1.11)**.** *Suppose $\Omega \subseteq \mathbb{R}^2$ is a polygon that does not intersect either coordinate axis so that X_Ω is a free convex toric domain. Then*

$$c_k(X_\Omega) = \min\{\ell'_\Omega(\partial\Lambda) : L_\Lambda = k+1\}$$

where the minimum ranges over lattice polygons Λ.

As discussed in [45] Exercise 4.16 it is equivalent to take the minimum over all polygons with edges parallel to edges of Ω and with no constraints on their vertices with the modification that $L_\Lambda \geq k+1$.

Theorem 5.3.20. *Suppose X_Ω is a free rational convex toric domain. Then,*

$$c_k(X_\Omega) = \min\{D \cdot D_\Omega : h^0(Y_{\Sigma(\Omega)}, D) \geq k+1\}$$
$$\mathrm{cap}_{X_\Omega}(r) = \max\{h^0(Y_{\Sigma(\Omega)}, D) : D \cdot D_\Omega \leq r\}$$

where D_Ω is the balance divisor from Definition 5.3.6 and where both extrema range over all nef $\mathbb{Q}$- or $\mathbb{R}$-divisors on $Y_{\Sigma(\Omega)}$.

Proof. As before, the two equalities are equivalent and so we will only show the first. Let $v_\star \in \Omega^\circ$ and set $\Omega' = \Omega - v_\star$. Suppose that u_1, u_2 are outward normals to adjacent faces of Ω. The dual norm $|| \cdot ||_{\Omega'}^*$ is linear on $\mathrm{Cone}(u_1, u_2)$, since the maximum of $v \cdot -$ will be achieved (possibly non-uniquely) at the vertex shared between the two adjacent edges for any $v \in \mathrm{Cone}(u_1, u_2)$. It is hence a support function on the outer normal fan $\Sigma^-(\Omega)$, which is just the negative of the inner normal fan. Notice that for $v \in \mathrm{Cone}(u_1, u_2)$, the dual norm $||v||_{\Omega'}^* = v \cdot p$ where p is the vertex described above, but this is equal to $-v^\perp \times p$ by definition. Note that p is exactly the point of $\partial\Omega'$ at which $-v^\perp$ is tangent to $\partial\Omega'$ so that $p = p_v$ as in the definition of Ω'-length in §5.2. Hence,

$$||v||_{\Omega'}^* = \ell_{\Omega'}(-v^\perp)$$

It follows that

$$\begin{aligned} c_k(X_\Omega) &= \min\{\ell'_\Omega(\partial\Lambda) : L_\Lambda = k+1\} \\ &= \min\{\ell_{\Omega'}(\partial\Xi) : L_\Xi = k+1\} \end{aligned}$$

via the correspondence $\Lambda \mapsto -\Lambda^\perp$, where both minima range over all lattice polygons Λ or Ξ respectively. By a similar (actually simpler) argument to the proof of Theorem 5.3.16, this second minimum can be seen to be equal to

$$\min\{D \cdot D_{\Omega'} : h^0(Y_{\Sigma(\Omega)}, D) \geq k+1\}$$

Now Ω' is just a translate of Ω and so $D \cdot D_\Omega = D \cdot D_{\Omega'}$ for all divisors D, which gives the result. □

We finally observe that all of the machinery developed above works equally well when Ω is an irrational polygon with rational slopes, since rationality is only required on the level of edges to define a fan that will produce a toric variety. The only difference is that D_Ω will no longer be a $\mathbb{Q}$-divisor.

5.4 Eventual expressions for ECH capacities

The aim of this section is to define 'tightly constrained' convex domain and to prove the following theorem giving an explicit eventual description of the cap function for such convex toric domains.

Studying cap functions

Theorem 5.4.1. *Suppose Ω is a tightly constrained convex lattice domain with Ω-perimeter λ. Then, there exists $x_0 \in \mathbb{Z}_{\geq 0}$ such that for all $x \geq x_0$ and for each $r = 0, \ldots, \lambda - 1$,*

$$\mathrm{cap}_{X_\Omega}(r + \lambda x) = \mathrm{ehr}_\Omega(x) + rx + \gamma_r$$

for some constant $\gamma_r \in \mathbb{Z}$ depending only on r. If Ω has a weight equal to 1 then Ω is tightly constrained and moreover one can choose $x_0 = 0$.

In order to do so, we will study the combinatorics of Ω in terms of its weight sequence, and then use this data to compute the cap function recursively. We will discuss the tightly constrained assumption on Ω and how every convex toric lattice domain conjecturally reduces to this case.

We also observe that the periodicity in the expression for cap_{X_Ω} is actually concentrated in the constant coefficient. Indeed, if $y = r + \lambda x$ then

$$\begin{aligned}
\mathrm{cap}_{X_\Omega}(y) &= \mathrm{ehr}_\Omega(\frac{y-r}{\lambda}) + r\frac{y-r}{\lambda} + \gamma_r \\
&= \mathrm{vol}(\Omega)\left(\frac{y-r}{\lambda}\right)^2 + \frac{1}{2}L_{\partial\Omega}\frac{y-r}{\lambda} + 1 + r\frac{y-r}{\lambda} + \gamma_r \\
&= \frac{1}{2\lambda}y^2 + \frac{1}{2\lambda}L_{\partial\Omega}y - \frac{rL_{\partial\Omega}}{2\lambda} - \frac{r^2}{2\lambda} + \gamma_r \\
&= \frac{1}{2\lambda}y^2 + \frac{1}{2\lambda}L_{\partial\Omega}y + O(1)
\end{aligned}$$

Recall that $\lambda = 2\,\mathrm{vol}(\Omega)$, allowing one to express this function in more Ehrhart-theoretic terms.

Examples

We start by presenting some suggestive examples of calculations of capacities and cap functions for some basic convex toric domains.

Example 5.4.2. $\mathrm{cap}_{E(a,b)}(r) = \mathrm{ehr}_Q(r)$, *the Ehrhart quasipolynomial of the rational triangle*

$$Q = \mathrm{Conv}((0,0),(1/a,0),(0,1/b))$$

This is also equal to the Hilbert function of $\mathcal{O}(1)$ for the weighted projective plane $\mathbb{P}(1,a,b)$.

Example 5.4.3. *The cap function for the polydisk $P(a,b)$ has*

$$\mathrm{cap}_{P(a,b)}(2abr) = (ar+1)(br+1) = \mathrm{hilb}_{(\mathbb{P}^1\times\mathbb{P}^1,\mathcal{O}(a,b))}(r)$$

Example 5.4.4. *Let $\Omega(a)$ be the convex hull of the points $(0,0),(0,2a),(a,a),(a,0)$. One has*

$$\mathrm{cap}_{X_{\Omega(a)}}(3ar) = h^0(X, rD)$$

where X is the first Hirzebruch surface, or $\mathbb{P}^2$ blown up in one point, and where $D = 3C+2F$ with C the (-1)-curve and F a fibre in the $\mathbb{P}^1$-bundle structure on X.

Combinatorics of weight sequences

Recall that the weight sequence associated to a convex domain Ω consists of a number and two lists that we will write as $(c; a_i; b_i)$. We will assume that the lists are finite sets of integers, which implies that Ω is rational. From the asymptotics of capacities of convex domains,

$$\mathrm{Vol}(\Omega_2) = \mathrm{Vol}(\amalg_i B(a_i)) = \frac{1}{2}\sum a_i^2$$

and so

$$\ell_\Omega(\partial\Omega) = 2\,\mathrm{Vol}(\Omega) = \mathrm{Vol}(B(c)) - \mathrm{Vol}(\amalg_i B(a_i)) - \mathrm{Vol}(\amalg_i B(b_i)) = c^2 - \sum a_i^2 - \sum b_i^2$$

Consider now the number of lattice points enclosed by a concave domain, excluding those on the upper boundary. Each ball $B(b_i)$ contributes $\frac{1}{2}b_i(b_i+1)$ lattice points; note that the transformation realising the inductive description of the weight sequence is a special affine linear map and so preserves lattice point counts. Hence the number of lower lattice points (i.e. excluding the upper boundary) in Ω_3 is

$$\sum \frac{1}{2}b_i(b_i+1)$$

and thus the number of lattice points enclosed by Ω is

$$\frac{1}{2}(c+1)(c+2) - \sum \frac{1}{2}\alpha_i(\alpha_i+1) - \sum \frac{1}{2}b_j(b_j+1)$$

For future reference will note that this is equal to

$$1 + \frac{1}{2}c(c+3) - \sum \frac{1}{2}\alpha_i(\alpha_i+1) - \sum \frac{1}{2}b_j(b_j+1)$$

Reducing the problem

For a convex domain Ω with weight sequence $w(\Omega) = (c; a_i; b_j)$, Lemma 5.2.6 gives that the ECH capacities of X_Ω are given by

$$c_k(X_\Omega) = \min\{c_{k+k_2+k_3}(B(c)) - c_{k_2}(\amalg_i B(a_i)) - c_{k_3}(\amalg_j B(b_j)) : k_2, k_3 \in \mathbb{Z}_{\geq 0}\}$$

By the disjoint union property of capacities, this is equal to

$$c_k(X_\Omega) = \min\{c_{k+\sum_i k_i + \sum_j m_j}(B(c)) - \sum_i c_{k_i}(B(a_i)) - \sum_j c_{m_j}(B(b_j)) : k_i, m_j \in \mathbb{Z}_{\geq 0}\}$$

It follows that the cap function of X_Ω is given by

$$\begin{aligned}&\mathrm{cap}_{X_\Omega}(r) = \\ &1 + \max\{k : \exists k_i, m_j \text{ with } c_{k+\sum_i k_i + \sum_j m_j}(B(c)) - \sum_i c_{k_i}(B(a_i)) - \sum_j c_{m_j}(B(b_j)) \leq r\}\end{aligned}$$

The capacities of a ball $B(q)$ take the form

$$c_k(B(q)) = dq \text{ when } \frac{1}{2}d(d+1) \leq k \leq \frac{1}{2}\delta(\delta+3)$$

Hence, to maximise k, one may assume that $k_i = \frac{1}{2}\alpha_i(\alpha_i+1)$, $m_j = \frac{1}{2}\beta_j(\beta_j+1)$, and $k + \sum_i k_i + \sum_j m_j = \frac{1}{2}\delta(\delta+3)$ for some $\alpha_i, \beta_j, \delta$. Therefore $\operatorname{cap}_{X_\Omega}(r)$ is 1 plus the maximum of

$$C(\delta, \alpha_i, \beta_j) := \frac{1}{2}\delta(\delta+3) - \sum_i \frac{1}{2}\alpha_i(\alpha_i+1) - \sum_j \frac{1}{2}\beta_j(\beta_j+1)$$

subject to

$$\delta c - \sum_i \alpha_i a_i - \sum_j \beta_i b_i \leq r$$

where $\alpha_i, \beta_j, \delta$ range over nonnegative integers.

Final calculations

Lemma 5.4.5. *Fix a weight sequence $(c; a_1, \ldots, a_s; b_1, \ldots, b_t)$ and let $\lambda = c^2 - \sum a_i^2 - \sum b_i^2$. Suppose $(\delta, \alpha_i, \beta_i)$ maximises $C(\delta, \alpha_i, \beta_j)$ subject to*

$$\delta c - \sum_i \alpha_i a_i - \sum_j \beta_j b_j = r$$

Then the sequence $(\delta + c, \alpha_i + a_i, \beta_i + b_i)$ maximises $C(\delta', \alpha_i', \beta_i')$ subject to

$$\delta' c - \sum_i \alpha_i' b_i - \sum_j \beta_j' b_j = r + \lambda$$

Proof. Suppose there exists $(\delta', \alpha_i', \beta_j')$ with $C(\delta', \alpha_i', \beta_j') > C(\delta + c, \alpha + a_i, \beta_j + b_j)$. We will show that $C(\delta' - c, \alpha_i' - a_i, \beta_j' - b_j) > C(\delta, \alpha_i, \beta_j)$, contradicting maximality since

$$(\delta' - c)c - \sum(\alpha_i' - a_i)a_i - \sum(\beta_j' - b_j)b_j = r$$

For convenience, relabel the b_j as a_{s+j} and β_j as α_{s+j} and write $C(\delta, \alpha_i) = C(\delta, \alpha_i, \beta_j)$. Compute $2C(\delta' - c, \alpha_i' - a_i)$ to be

$$\begin{aligned}
&(\delta' - c)(\delta' - c + 3) - \sum(\alpha_i' - a_i)(\alpha_i' - a_i + 1) \\
&= \delta'(\delta'+3) - \sum \alpha_i'(\alpha_i'+1) - c\delta' + \sum \alpha_i' a_i - c(\delta'+3) + \sum(\alpha_i'+1)a_i + c^2 - \sum a_i^2 \\
&= \delta'(\delta'+3) - \sum \alpha_i'(\alpha_i'+1) - (r+\lambda) - (r+\lambda) - 3c + \sum a_i + \lambda \\
&> (\delta+c)(\delta+c+3) - \sum(\alpha_i + a_i)(\alpha_i + a_i + 1) - 2r - \lambda - 3c + \sum a_i
\end{aligned}$$

This is equal to

$$
\begin{aligned}
&= \delta(\delta+3) - \sum \alpha_i(\alpha_i+1) + c\delta - \sum \alpha_i a_i + c(\delta+3) \\
&\quad - \sum (\alpha_i+1)a_i + c^2 - \sum a_i^2 - 2r - \lambda - 3c + \sum a_i \\
&= C(\delta, \alpha_i, \beta_j) + r + r + 3c - \sum a_i + \lambda - 2r - \lambda - 3c + \sum a_i \\
&= C(\delta, \alpha_i, \beta_j)
\end{aligned}
$$

as desired. □

Definition 5.4.6. *Say that a convex lattice domain Ω (or a convex lattice toric domain X_Ω) with weight sequence $(c; a_i; b_i)$ is* tightly constrained with lower bound r_0 *if for all $r \geq r_0$ the maximum of $C(\delta, \alpha_i, \beta_j)$ subject to*

$$\delta c - \sum_i \alpha_i a_i - \sum_j \beta_i b_i \leq r$$

is attained by some $(\delta, \alpha_i, \beta_j)$ with

$$\delta c - \sum_i \alpha_i a_i - \sum_j \beta_i b_i = r$$

Say that Ω is tightly constrained *if it is tightly constrained with some lower bound r_0.*

Lemma 5.4.7. *Ω being tightly constrained with lower bound r_0 is equivalent to the statement that for every positive integer $r \geq r_0$ there is some $k \in \mathbb{Z}_{\geq 0}$ such that $c_k(X_\Omega) = r$.*

Proof. By definition the cap function of X_Ω is 1 plus the largest value of k such that $c_k(X_\Omega) \leq r$. If there is some k with $c_k(X_\Omega) = r$ then this largest value of k will be achieved by some k with $c_k(X_\Omega) = r$ by monotonicity. The largest value of k corresponds to a value of $C(\delta, \alpha_i, \beta_j)$ from the reasoning above, for which the corresponding capacity takes the value $\delta c - \sum \alpha_i a_i - \sum \beta_j b_j = r$. □

Equivalently, $\mathrm{cap}_{X_\Omega}(r+1) > \mathrm{cap}_{X_\Omega}(r)$ for all $r \geq r_0$, so that cap_{X_Ω} is eventually strictly increasing.

Example 5.4.8. *Suppose $X_\Omega = E(a,b)$ is an ellipsoid with a prime, $a < b$, and $\gcd(a,b) = 1$. Then X_Ω is tightly constrained with lower bound $(a-1)b$ to cover all residues mod a.*

Lemma 5.4.9. *Suppose Ω has at least one weight equal to 1. Then Ω is tightly constrained with lower bound $r_0 = 0$.*

Proof. Suppose $(c; a_i; b_j)$ is a weight sequence with $a_1 = 1$. Let $(\delta, \alpha_i, \beta_j)$ maximise $C(\delta, \alpha_i, \beta_j)$ subject to $\delta c - \sum_i \alpha_i a_i - \sum_j \beta_j b_j \leq r$. Suppose $\delta c - \sum_i \alpha_i a_i - \sum_j \beta_j b_j < r$. Modify $(\delta, \alpha_i, \beta_j)$ to $(\delta, \alpha_i', \beta_j)$ where $\alpha_1' = \alpha_1 - 1$ and $\alpha_i' = \alpha_i$ for $i \geq 2$. This sequence has $C(\delta, \alpha_i', \beta_j) = \frac{1}{2}\delta(\delta+3) - \frac{1}{2}(\alpha_1 - 1)\alpha_1 - \sum \frac{1}{2}\alpha_i(\alpha_i + 1) - \sum \frac{1}{2}\beta_j(\beta_j + 1) > C(\delta, \alpha_i, \beta_j)$ and $\delta c - (\alpha_1 - 1)a_1 - \sum \alpha_i a_i - \sum \beta_j b_j = \delta c - \sum_i \alpha_i a_i - \sum_j \beta_j b_j + b_1 \leq r$. This contradicts the fact that $(\delta, \alpha_i, \beta_j)$ was maximal. $\square$

Conjecture 5.4.10. *Suppose that* $\gcd\{c, a_1, \ldots, a_s, b_1, \ldots, b_t\} = 1$*. Then* Ω *is tightly constrained.*

Notice that the conjecture will certainly fail for weight sequences without the coprimality assumption. For example, the ball $B(2)$ has capacities that are all even numbers and so there can be no odd values of the constraint. We will henceforth make the assumption that Ω is tightly constrained.

Corollary 5.4.11. *For a tightly constrained convex lattice domain* Ω *with lower bound* $r_0 = 0$*,* $\partial\Omega$ *is an optimal path among lattice paths of length at most* $\ell_\Omega(\partial\Omega)$*.*

Proof. Clearly $\mathrm{cap}_{X_\Omega}(0) = 1 + 0$ is attained by $(\delta, \alpha_i, \beta_j) = (0, 0, \ldots, 0)$. Hence,

$$\mathrm{cap}_{X_\Omega}(\lambda) = 1 + C(c, a_i, b_j) = \frac{1}{2}(c+1)(c+2) - \sum a_i(a_i+1) - \sum b_j(b_j+1) = L_\Omega$$

as required. $\square$

Lemma 5.4.12. *Let* Ω *be a tightly constrained convex lattice domain. Denote the* Ω*-perimeter of* Ω *by* λ*. Then there exists* $x_0 \in \mathbb{Z}_{\geq 0}$ *such that for all* $x \geq x_0$ *and for each* $r = 0, \ldots, \lambda - 1$*,*

$$\mathrm{cap}_{X_\Omega}(\lambda x + r) = \mathrm{Vol}(\Omega)x^2 + (\frac{1}{2}L_{\partial\Omega} + r)x + \gamma_r$$

for some $\gamma_r \in \mathbb{Z}$*, where* $L_{\partial\Omega}$ *is the number of lattice points on the boundary of* Ω*.*

Proof. From Lemma 5.4.9 and the assumption that Ω is tightly constrained one has that the maximum value of $C(\delta, \alpha_i, \beta_j)$ subject to $\delta c - \sum_i \alpha_i a_i - \sum_j \beta_j b_j \leq r + \lambda$ is

$$C(\delta', \alpha_i', \beta_j') + r + \frac{1}{2}c(c+3) - \sum \frac{1}{2}b_i(b_i+1) = C(\delta', \alpha_i', \beta_j') + r + L_\Omega - 1$$

when $(\delta', \alpha_i', \beta_j')$ is maximal subject to $\delta' c - \sum_i \alpha_i' a_i - \sum_j \beta_j' b_j \leq r$, at least for large enough r. It follows that, for $r + \lambda x$ large enough,

$$\mathrm{cap}_{X_\Omega}(r + \lambda(x+1)) = \mathrm{cap}_{X_\Omega}(r + \lambda x) + r + \lambda x + L_\Omega - 1 \tag{5.1}$$

This implies that $\mathrm{cap}_{X_\Omega}(r+\lambda x)$ is eventually a quadratic polynomial. Solving the difference equation (5.1) gives the leading term as $\lambda/2$ and gives the linear coefficient as $L_\Omega - \mathrm{Vol}(\Omega) - 1 + r$. By Pick's formula the linear term is equal to $\frac{1}{2}L_{\partial\Omega} + r$, and we have seen that $\lambda/2 = \ell_\Omega(\partial\Lambda)/2 = \mathrm{Vol}(\Omega)$. □

This is the desired quasipolynomial representation of cap_{X_Ω}. However, we would also like this to have an algebro-geometric interpretation. The Ehrhart polynomial of Ω, as a lattice polygon, is

$$\mathrm{ehr}_\Omega(x) = \mathrm{Vol}(\Omega)x^2 + \frac{1}{2}L_{\partial\Omega}x + 1$$

Corollary 5.4.13. *Let Ω be a tightly constrained convex lattice domain of Ω-perimeter λ. Then, for any $r \in \{0,1,\ldots,\lambda-1\}$ and sufficiently large $x \in \mathbb{Z}_{\geq 0}$*

$$\begin{aligned}\mathrm{cap}_{X_\Omega}(r+\lambda x) &= \mathrm{ehr}_\Omega(x) + rx + \gamma_r \\ &= \mathrm{hilb}_{(Y_{\Sigma(\Omega)}, D_\Omega)}(x) + rx + \gamma_r\end{aligned}$$

for some $\gamma_r \in \mathbb{Z}$. In particular, for all sufficiently large $x \in \mathbb{Z}_{\geq 0}$

$$\mathrm{cap}_{X_\Omega}(\lambda x) = \mathrm{ehr}_\Omega(x) + \gamma_0 = \mathrm{hilb}_{(Y_{\Sigma(\Omega)}, D_\Omega)}(x) + \gamma_0$$

We believe that always $\gamma_r = \mathrm{cap}_{X_\Omega}(r) - 1$, which is what one would obtain from the difference equation $(*)$ holding for all $x \in \mathbb{Z}$, not just all sufficiently large x. This would in particular imply that $\gamma_0 = 0$. Suppose X_Ω is not tightly constrained. Assuming Conjecture 5.4.10, one can scale Ω to obtain a convex lattice domain Ω' that is tightly constrained. Let $q\Omega' = \Omega$. Then, using the scaling axiom from §5.2, for any $r = 0,\ldots,q-1$ one has

$$\mathrm{cap}_{X_\Omega}(r+qx) = \mathrm{cap}_{X_\Omega}(qx) = \mathrm{cap}_{X_{\Omega'}}(x)$$

Thus, knowing Theorem 5.4.1 for tightly constrained convex toric lattice domains is sufficient to completely describe the long term behaviour of the cap function for all convex toric lattice domains.

Example 5.4.14. *For $X_\Omega = B(2)$, one has*

$$\mathrm{cap}_{X_\Omega}(r) = \begin{cases} \mathrm{cap}_{B(1)}(\frac{r}{2}) & r \equiv 0 \bmod 2 \\ \mathrm{cap}_{B(1)}(\frac{r-1}{2}) & r \equiv 1 \bmod 2 \end{cases} = \begin{cases} \frac{1}{8}(r+2)(r+4) & r \equiv 0 \bmod 2 \\ \frac{1}{8}(r+1)(r+3) & r \equiv 1 \bmod 2 \end{cases}$$

Conjectures

We conjecture that the word 'eventually' may be dropped in all the above results, and that in fact the cap function of a tightly constrained convex toric lattice domain is given entirely by the quasipolynomial in Theorem 5.4.1. Of course, this is already proven if one of the weights of Ω is equal to 1.

We furthermore believe that the following strengthening of the prior results holds.

Conjecture 5.4.15. *Suppose that X_Ω is a tightly constrained toric domain. Then:*

- *there exist convex lattice domains $\Omega_0, \dots, \Omega_{\lambda-1}$ such that, for any $r = 0, \dots, \lambda - 1$ and any sufficiently large $x \in \mathbb{Z}_{\geq 0}$,*

$$\mathrm{cap}_{X_\Omega}(r + \lambda x) = |(\Omega_r + x\Omega) \cap \mathbb{Z}^2|$$

- *there exist divisors $D_0, \dots, D_{\lambda-1}$ on $Y_{\Sigma(\Omega)}$ such that, for any $r = 0, \dots, \lambda - 1$ and any sufficiently large $x \in \mathbb{Z}_{\geq 0}$,*

$$\mathrm{cap}_{X_\Omega}(r + \lambda x) = h^0(Y_{\Sigma(\Omega)}, D_r + xD_\Omega)$$

Moreover, we conjecture that $\Omega_0 = \{0\}$ so that $\gamma_0 = 0$, and that all of these claims actually hold for all $x \in \mathbb{Z}_{\geq 0}$, not just for all sufficiently large x.

These conjectures state that cap_{X_Ω} is (eventually) given by a 'mixed Ehrhart quasipolynomial' or a 'mixed Hilbert quasipolynomial' as studied in [41]. Cristofaro-Gardiner–Kleinman in [32] have previously approached symplectic embeddings problems for ellipsoids via Ehrhart theory, and one can view some aspects of this chapter as pursuing a related philosophy for convex toric lattice domains.

5.5 Algebraic capacities

In analogy with the quantities appearing on the algebraic side to compute ECH capacities for convex toric domains, we define purely algebraic invariants for polarised surfaces. Denote by $\mathrm{WDiv}(Y)$ the group of Weil $\mathbb{Z}$-divisors on a surface Y. We denote the set of nef divisors, as opposed to nef divisor classes, on Y with coefficients in $\mathbb{K} \in \{\mathbb{Z}, \mathbb{Q}, \mathbb{R}\}$ by $\mathrm{nef}(Y)_{\mathbb{K}} \subseteq \mathrm{WDiv}(Y) \otimes_{\mathbb{Z}} \mathbb{K}$. We also denote the cone of pseudo-effective divisors on Y (or the Mori cone) by $\overline{\mathrm{NE}}(Y)$. By a 'polarised surface' (Y, A) we mean a projective surface Y with an ample $\mathbb{R}$-divisor A.

Construction

We define the 'round-down' $\lfloor D \rfloor$ of a Weil $\mathbb{R}$-divisor $D = \sum a_i D_i$ on Y by

$$\lfloor D \rfloor := \sum \lfloor a_i \rfloor D_i$$

where the D_i are prime Weil divisors on Y. Define also the *index* of a Weil $\mathbb{R}$-divisor D on Y to be

$$I(D) := \lfloor D \rfloor \cdot (\lfloor D \rfloor - K_Y)$$

When Y is smooth Noether's formula gives

$$\chi(D) = \chi(\mathcal{O}_Y) + \frac{1}{2} I(D)$$

where $\chi(D) := \chi(\lfloor D \rfloor) = h^0(\lfloor D \rfloor) - h^1(\lfloor D \rfloor) + h^2(\lfloor D \rfloor)$.

Definition 5.5.1. *Let* (Y, A) *be a polarised* $\mathbb{Q}$*-factorial surface. Define the* kth algebraic capacity *of* (Y, A) *to be*

$$c_k^{\text{alg}}(Y, A) := \inf_{D \in \text{nef}(Y)_{\mathbb{Q}}} \{D \cdot A : \chi(D) \geq k + \chi(\mathcal{O}_Y)\}$$

and the corresponding counting function

$$\text{cap}_{(Y,A)}(r) := \#\{k : c_k^{\text{alg}}(Y, A) \leq r\}$$

Note that we require Y to be $\mathbb{Q}$-factorial so that intersection numbers make sense. When Y is smooth we have

$$c_k^{\text{alg}}(Y, A) = \inf_{D \in \text{nef}(Y)_{\mathbb{Q}}} \{D \cdot A : I(D) \geq 2k\} \tag{5.2}$$

We note that this construction agrees with the algebraic formulation of ECH involving toric surfaces, which are all automatically $\mathbb{Q}$-factorial. Indeed, by Demazure vanishing [23, Theorem 9.3.5] the higher cohomology of toric nef $\mathbb{Q}$-divisors vanishes, and so $\chi(D) = h^0(D)$. As toric varieties are rational $\chi(\mathcal{O}_Y) = 1$ and hence

$$c_k^{\text{alg}}(Y, A) := \inf_{D \in \text{nef}(Y)_{\mathbb{Q}}} \{D \cdot A : h^0(D) \geq k + 1\}$$

as used in Theorem 5.3.16. Of course, taking the infimum over nef divisors with rational coefficients gives the same result as over nef divisors with real coefficients.

The reformulation (5.2) for smooth surfaces in terms of the index is attractive for comparisons with symplectic geometry. For a more general expression of this type, suppose Y has rational singularities. There is a Noether-like formula in this situation: suppose $\varphi\colon Y' \to Y$ is a resolution, then for any Cartier divisor D on Y

$$\chi(D) = \chi(\mathcal{O}_Y) + \frac{1}{2} \varphi^* D \cdot (\varphi^* D - K_{Y'})$$

If further φ is a crepant resolution then we have a true Noether formula on Y

$$\begin{aligned}\chi(D) &= \chi(\mathcal{O}_Y) + \frac{1}{2}\varphi^* D \cdot (\varphi^* D - \varphi^* K_Y) \\ &= \chi(\mathcal{O}_Y) + D \cdot (D - K_Y)\end{aligned}$$

As discussed in Chapter 3, crepant resolutions are available for surfaces only when Y has Du Val singularities. In this case, one can formulate algebraic capacities using (5.2) using Noether's formula just as for the smooth case. Surfaces with only Du Val singularities are called 'canonical surfaces', as the Du Val singularities are exactly the canonical singularities in dimension two [68, §1].

Algebraic Weyl law

Given the analogy between algebraic capacities and ECH capacities in the toric case, one can hope for comparable asymptotic structure to exist in general.

Conjecture 5.5.2 (Algebraic Weyl Law)**.** *For (Y, A) a polarised $\mathbb{Q}$-factorial surface,*

$$\lim_{k\to\infty} \frac{c_k^{\mathrm{alg}}(Y, A)^2}{k} = 2A^2$$

One could hope for this to be true in the generality stated, or restrict to the setting of canonical or smooth surfaces. We know this result is true for any toric surface Y with a smooth torus-fixed point equipped with an ample $\mathbb{R}$-divisor A, since in this case $P(A)$ is affine-equivalent to a convex domain with rational slopes.

We include some speculation as to how a proof in any of these settings might proceed. Heuristics suggest that eventually divisors of the form $D = dA + \delta$ for $d \gg 0$ and for some 'small' divisor δ will be at least approximately optimal for $c_k^{\mathrm{alg}}(Y, A)$ with $k \gg 0$ and so minimising

$$D \cdot A = dA^2 + \delta \cdot A \text{ subject to } D \cdot (D - K_Y) = d^2 A^2 + \text{small error terms}$$

gives

$$c_k^{\mathrm{alg}}(Y, D) \sim \inf_d \{dA^2 : d^2 A^2 \geq 2k\} \sim \sqrt{2A^2 k}$$

and so

$$\lim_{k\to\infty} \frac{c_k^{\mathrm{alg}}(Y, A)^2}{k} = \lim_{k\to\infty} \frac{2A^2 k}{k} = 2A^2$$

Showing that divisors of the form $dA + \delta$ are approximately optimal could be done by proving a recurrence for optimisers for $c_k^{\mathrm{alg}}(Y, A)$. Lastly, we note that Conjecture 5.5.2 is straightforward to verify for canonical surfaces of Picard rank 1 where the existence of such a normal form for optimisers is clear. As a consequence of Theorem 5.3.16 this also reproves the symplectic Weyl law for balls and for dilates of the ellipsoids $E(1, 2)$ and $E(2, 3)$.

5.6 Sub-leading asymptotics for ECH capacities

We return to the world of symplectic embeddings. We know that the ECH capacities asymptotically recover the volume constraint for the existence of symplectic embeddings by

$$\lim_{k\to\infty} \frac{c_k(X,\omega)^2}{k} = 4\,\mathrm{vol}(X)$$

and so we can study the 'sub-leading asymptotics' via the error terms

$$e_k(X,\omega) := c_k(X,\omega) - 2\sqrt{\mathrm{vol}(X,\omega)k}$$

There has been much recent work to understand the asymptotics of $e_k(X)$, which should provide potentially subtler numerical obstructions to the existence of symplectic embeddings; see [46, Cor. 1.13] and Cor. 5.6.3 below. We follow the convention of referring to a compact domain in $\mathbb{R}^4$ whose boundary is smooth and transverse to the radial vector field as a 'nice star-shaped domain'. Sun in [72] showed that when (X,ω) is a nice star-shaped domain

$$e_k(X,\omega) = O(k^{125/252})$$

and Cristofaro-Gardiner–Savale [33] improved this to

$$e_k(X,\omega) = O(k^{2/5})$$

The primary methods used in extracting these asymptotics come from Seiberg–Witten theory. For the case of not-necessarily star-shaped domains in $\mathbb{R}^4$ Hutchings [46] showed by more direct methods that

$$e_k(X,\omega) = O(k^{1/4})$$

The author's understanding is that the expectation for all (X,ω) is

$$e_k(X,\omega) = O(1)$$

which these estimates are approaching. This is the case for all examples that have been computed.

Rational convex toric domains

We use Theorem 5.4.1 to prove a result computing the lim inf and lim sup of the error $e_k(X)$ for many toric domains, in particular showing that $e_k(X,\omega)$ is $O(1)$ in these cases.

Proposition 5.6.1. *Suppose $X = X_\Omega$ is a tightly constrained lattice convex toric domain. Then*

$$\liminf_{k\to\infty} e_k(X_\Omega) = 1 - \frac{1}{2}L_{\partial\Omega} \text{ and } \limsup_{k\to\infty} e_k(X_\Omega) = -\frac{1}{2}L_{\partial\Omega}$$

Proof. Write $V = \mathrm{vol}(X)$ for brevity. Suppose $r \gg 0$ and let k_0 be the largest k such that $c_k(X) = r-1$ and k_1 be the largest k such that $c_k(X) = r$. Note that $k_0(r) = \mathrm{cap}_X(r-1)-1$ and $k_1(r) = \mathrm{cap}_X(r) - 1$. From Theorem 5.4.1 we have

$$\frac{1}{4V}(r-1)^2 + \frac{1}{4V}L_{\partial\Omega}(r-1) + \gamma_i'' = k_0(r) \text{ and } \frac{1}{4V}r^2 + \frac{1}{4V}L_{\partial\Omega}r + \gamma_i'' = k_1(r)$$

where i is the residue of $r \bmod \lambda = 2V$ and $\gamma_i'' = \gamma_i' - 1$. Using the quadratic formula

$$\begin{aligned} r &= 2V\left(-\frac{1}{4V}L_{\partial\Omega} + \sqrt{\left(\frac{1}{4V}\right)^2 - \frac{1}{V}\gamma_i'' + \frac{1}{V}k_1(r)}\right) \\ &= -\frac{1}{2}L_{\partial\Omega} + \sqrt{\frac{1}{4} - 4V\gamma_i'' + 4V\cdot k_1(r)} \end{aligned}$$

and

$$\begin{aligned} r-1 &= 2V\left(-\frac{1}{4V}L_{\partial\Omega} + \sqrt{\left(\frac{1}{4V}\right)^2 - \frac{1}{V}\gamma_i'' + \frac{1}{V}k_0(r)}\right) \\ &= -\frac{1}{2}L_{\partial\Omega} + \sqrt{\frac{1}{4} - 4V\gamma_i'' + 4V\cdot k_0(r)} \end{aligned}$$

Note that for k with $c_k(X) = r$

$$r - 2\sqrt{V\cdot k_1(r)} \le e_k(X) \le r - 2\sqrt{V\cdot(k_0(r)+1)}$$

as $k_1(r)$ is the largest k with $c_k(X) = r$ and $k_0(r)+1$ is the smallest. It follows that

$$\begin{aligned} -\frac{1}{2}L_{\partial\Omega} + \sqrt{\frac{1}{4} - 4V\gamma_i'' + 4V\cdot k_1(r)} - \sqrt{4V\cdot k_1(r)} \le e_k(X) \\ \le 1 - \frac{1}{2}L_{\partial\Omega} + \sqrt{\frac{1}{4} - 4V\gamma_i'' + 4V\cdot(k_0(r)+1)} - \sqrt{4V\cdot(k_0(r)+1)} \end{aligned}$$

with both bounds achieved. Since $\frac{1}{4} - 4V\gamma_i''$ is $O(1)$ and $\sqrt{x+B} - \sqrt{x} \to 0$ as $x \to \infty$ for bounded B we see that

$$\liminf_{k\to\infty} e_k(X) = -\frac{1}{2}L_{\partial\Omega} \text{ and } \limsup_{k\to\infty} e_k(X) = 1 - \frac{1}{2}L_{\partial\Omega}$$

as required. $\square$

Observe that this generalises Hutchings' calculation in [46, Example 1.2] for $B(1)$ with $L_{\partial\Omega} = 3$.

Corollary 5.6.2. *Suppose $\Omega = q\Omega_0$ is a rational or real scaling of a tightly constrained lattice convex domain Ω_0. Then*

$$\liminf_{k\to\infty} e_k(X_\Omega) = q(1 - \frac{1}{2}L_{\partial\Omega}) \text{ and } \limsup_{k\to\infty} e_k(X_\Omega) = -\frac{q}{2}L_{\partial\Omega}$$

If Conjecture 5.4.10 is true then Cor. 5.6.2 applies to any convex domain Ω with rational slopes (in particular, all rational convex domains) since one can scale Ω to a lattice convex domain with coprime weights.

Corollary 5.6.3. *Suppose Ω, Ω' are tightly constrained lattice convex domains of the same volume, and suppose that X_Ω° symplectically embeds in $X_{\Omega'}$. Then $L_{\partial\Omega} \geq L_{\partial\Omega'}$.*

Of course a similar version of this corollary exists for scaling of tightly constrained convex toric domains.

Review of the Ruelle invariant

Hutchings [46] considers the 'Ruelle invariant' $\mathrm{Ru}(X, \omega)$ of a nice star-shaped domain in $\mathbb{R}^4$. We will not recall the fairly involved definition of the Ruelle invariant here, and instead refer the reader to [46, §1.2]. Its relevance to sub-leading asymptotics in ECH comes from the following conjecture and theorem of Hutchings.

Conjecture 5.6.4 ([46, Conjecture 1.5]). *If (X, ω) is a 'generic' nice star-shaped domain in $\mathbb{R}^4$ then*

$$\lim_{k\to\infty} e_k(X, \omega) = -\frac{1}{2}\mathrm{Ru}(X, \omega)$$

Theorem 5.6.5 ([46, Theorem 1.10]). *This conjecture is true whenever (X, ω) is a 'strictly' convex or concave toric domain.*

A strictly convex toric domain is a convex toric domain arising from $\Omega \subseteq \mathbb{R}^2$ where the upper part of the boundary of Ω is the graph of a function f with $f'(0) < 0$ and $f'' < 0$. The definition of 'strictly concave' is similar. Hutchings also computes the Ruelle invariant for many 'generic' toric domains.

Proposition 5.6.6 ([46, Prop. 1.11]). *If X_Ω is a nice toric domain such that the part of $\partial\Omega$ excluding the axes has strictly negative slope everywhere, then*

$$\mathrm{Ru}(X_\Omega) = a + b$$

where $(a, 0)$ and $(0, b)$ are the vertices of Ω lying on the x- and y-axes respectively.

We will consider and prove some versions of these statements for rational or lattice convex toric domains. In some sense these are the opposite of the cases considered by Hutchings since such X_Ω are very non-generic.

Algebraic Ruelle invariant

Just as in the symplectic setting, assuming Conjecture 5.5.2 holds one can define algebraic error terms

$$e_k^{\mathrm{alg}}(Y,A) := c_k^{\mathrm{alg}}(Y,A) - \sqrt{2A^2k}$$

Define the *algebraic Ruelle invariant* $\mathrm{Ru}^{\mathrm{alg}}(Y,A)$ for many polarised toric surfaces (Y,A) by the following. When $P(A)$ is affine-equivalent to a tightly constrained lattice convex domain, define

$$\mathrm{Ru}^{\mathrm{alg}}(Y,A) := -K_Y \cdot A - 1$$

When $P(A)$ is affine-equivalent to a scaling $q\Omega_0$ of a tightly-constrained lattice convex domain Ω_0, define

$$\mathrm{Ru}^{\mathrm{alg}}(Y,A) := q\,\mathrm{Ru}^{\mathrm{alg}}(Y,\frac{1}{q}A)$$

Conjecture 5.4.10 implies that this defines the algebraic Ruelle invariant for all (Y,A) where Y has a smooth torus-fixed point. It follows from the definition that

$$\mathrm{Ru}^{\mathrm{alg}}(Y,qA) = q\,\mathrm{Ru}^{\mathrm{alg}}(Y,A)$$

where defined. From standard toric geometry and Prop. 5.6.1 we have the following result. Note that toric surfaces are automatically $\mathbb{Q}$-factorial.

Proposition 5.6.7. *For (Y,A) a polarised toric surface such that $P(A)$ is (affine-equivalent to) a tightly constrained lattice convex domain*

$$\mathrm{Ru}^{\mathrm{alg}}(Y,A) = L_{\partial P(A)} - 1$$

Furthermore, whenever $P(A)$ is (affine-equivalent to) a real scaling of a tightly constrained lattice convex domain,

$$-\frac{1}{2}\,\mathrm{Ru}^{\mathrm{alg}}(Y,A)$$

is the midpoint between $\liminf_{k\to\infty} e_k^{\mathrm{alg}}(Y,A)$ *and* $\limsup_{k\to\infty} e_k^{\mathrm{alg}}(Y,A)$.

Observe that when Y is a weighted projective space of the form $\mathbb{P}(1,r,s)$ and $A = \mathcal{O}(\mathrm{lcm}(r,s))$, the algebraic Ruelle invariant is the quantity $a+b$ from [46, Prop. 1.11], agreeing with the symplectic Ruelle invariant, in this case for the ellipsoid

$$E\left(\frac{r}{\gcd(r,s)}, \frac{s}{\gcd(r,s)}\right)$$

That said, it is currently unclear how to compare the algebraic and symplectic Ruelle invariants in a meaningful, geometric way.

Weight sequences for algebraic capacities

We suspect that there is a good analog of weight sequences for ample divisors (at least on rational surfaces) that mimics the role of weight sequences from ECH for toric domains; a good starting point for this seems to be the work of Biran [10] on 'nef vectors'. This would provide a reflection of Conjecture 5.4.10 in the algebraic setting, and would refine results such as Prop. 5.6.7 above. The ideal situation that we could conjecture is that $P(A)$ is affine-equivalent to a tightly-constrained lattice convex domain if and only if, in addition to Y having a smooth torus-fixed point, A is a primitive Cartier divisor (that is, A is Cartier and there is no Cartier divisor A_0 with $qA_0 = A$ for $q \in \mathbb{Z}_{\geq 2}$).

This would enable a more general definition of the algebraic Ruelle invariant: for (Y, A) a $\mathbb{Q}$-factorial polarised surface with $A = qA_0$ for a primitive Cartier divisor A_0 and some $q \in \mathbb{R}$, set

$$\mathrm{Ru}^{\mathrm{alg}}(Y, A) := q(-K_Y \cdot A_0 - 1)$$

Using this more general definition we conjecture that the sub-leading asymptotics of the algebraic error terms are centred around the algebraic Ruelle invariant just as in Prop. 5.6.7.

Conjecture 5.6.8. *For any $\mathbb{Q}$-factorial polarised surface* (Y, A)

$$\liminf_{k\to\infty} e_k^{\mathrm{alg}}(Y, A) + \limsup_{k\to\infty} e_k^{\mathrm{alg}}(Y, A) = -\mathrm{Ru}^{\mathrm{alg}}(Y, A)$$

5.7 Connection to minimal hypersurfaces

Lastly, we will outline connections of algebraic capacities to minimal (hyper)surface theory. One of the fundamental tools in sourcing and studying minimal hypersurfaces – for example in Song's recent proof [71] of Yau's conjecture – is *min-max theory*. Some of the principal objects in this theory are min-max widths, which have many striking similarities to capacities in symplectic geometry. To define these widths, we require the notion of a *p-sweepout*. These are continuous maps

$$\Phi\colon X \to \mathcal{Z}^1$$

where X is a finite-dimensional simplicial complex and $\mathcal{Z}^1$ is a certain topological space of codimension one $\mathbb{Z}/2$-chains on M, satisfying a nondegeneracy condition. $\mathcal{Z}^1$ is homotopy equivalent to $\mathbb{RP}^\infty$; denote its $\mathbb{Z}_2$-cohomology ring by $\mathbb{C}[\lambda]$. With this notation, the nondegeneracy condition for p-sweepouts is that $(\Phi^*\lambda)^p \neq 0$. We refer to [71, §2.3] for an actual definition. One should imagine a p-sweepout as being a formal generalisation of a p-dimensional family of hypersurfaces in M. To a p-sweepout Φ one can associate its mass function $\mathbf{M}_\Phi\colon X \to \mathbb{R}$ given at x by taking the g-area of the chain $\Phi(x)$.

The pth min-max width for a compact Riemannian manifold (M, g) is then

$$\omega_p(M, g) := \inf_{\Phi} \sup\{\mathbf{M}_\Phi(x) : x \in \mathrm{dom}(\Phi)\}$$

where the infimum ranges over p-sweepouts Φ with ‘no concentration of mass’ (see [71, §2.3]), and $\mathrm{dom}\,\Phi$ is the domain of Φ. This infimum should essentially be achieved by the g-area of a minimal hypersurface, hence the application to problems such as Yau's conjecture.

One can make a similar construction for p-sweepouts of codimension two, in which one uses a space $\mathcal{Z}^2$ of codimesion two $\mathbb{Z}/2$-chains, which is homotopy equivalent to $\mathbb{CP}^\infty$. This produces codimension two min-max widths

$$\omega_p^2(M, h)$$

defined similarly to the min-max widths above. The crucial example for our purposes is the following.

Example 5.7.1. *Suppose M is a smooth complex projective algebraic variety equipped with an ample divisor A. Let D be a nef (or big) divisor with $h^0(D) = p + 1$. This defines a (real) codimension two p-sweepout for M by pulling back the hyperplane sections of M in the morphism to $\mathbb{P}^p$ by the linear system $|D|$.*

In particular, it follows that the codimension two min-max weights in this situation satisfy

$$\omega_p^2(M, g) \leq c_p^{\mathrm{alg}}(M, A)$$

where g is the metric corresponding to A. It is conceivable that this is actually an equality, and that further ties are present between the theory of minimal hypersurfaces and algebraic capacities.

Chapter 6

Retrospective

This thesis sought to explore and establish connections between the algebraic geometry of orbifolds and other forms of algebra, geometry, and combinatorics. As a concomitant of this, we aimed to find and ask meaningful algebro-geometric questions inspired by phenomena elsewhere in mathematics. The main arenas in which we saw this play out were:

- *McKay correspondence:* uniting orbifold geometry with the representation theory of finite groups and of quivers, as well as with the combinatorics associated to these objects and coming from toric geometry

- *Ehrhart theory:* orbifold Riemann–Roch produces a particular decomposition of the Ehrhart function that enables us to apply deformation and singularity theory to answer subtle questions about lattice point counting

- *Symplectic embeddings:* obstructions to symplectic embeddings prompt various optimisation problems for nef divisors on orbifolds, which are of intrinsic interest and also help to study embeddings for 'special' manifolds that are typically harder to treat with symplectic methods.

We offered evidence new insight supplied by each of these connections, but we firmly believe that in each arena there is much more to uncover. As one example, in the the worlds of embedded contact homology and of minimal hypersurfaces there are powerful consequences to the existence of a Weyl law; what could the presence of a Weyl law for algebraic capacities imply?

Bibliography

[1] Akhtar, M., Coates, T., Corti, A., Heuberger, L., Kasprzyk, A., Oneto, A., Petracci, A., Prince, T., and Tveiten, K. Mirror symmetry and the classification of orbifold del Pezzo surfaces. *Proceedings of the American Mathematical Society 144*, 2 (2016), 513–527.

[2] Akhtar, M., Coates, T., Galkin, S., and Kasprzyk, A. M. Minkowski polynomials and mutations. *Symmetry, Integrability and Geometry: Methods and Applications 8* (2012), 094.

[3] Akhtar, M., and Kasprzyk, A. Singularity content. *arXiv preprint arXiv:1401.5458* (2014).

[4] Akhtar, M. E., and Kasprzyk, A. M. Mutations of fake weighted projective planes. *Proceedings of the Edinburgh Mathematical Society 59*, 2 (2016), 271–285.

[5] Artin, M., and Verdier, J.-L. Reflexive modules over rational double points. *Mathematische Annalen 270*, 1 (1985), 79–82.

[6] Bass, H. On the ubiquity of Gorenstein rings. *Mathematische Zeitschrift 82*, 1 (1963), 8–28.

[7] Batyrev, V. V. On the classification of smooth projective toric varieties. *Tohoku Mathematical Journal Second Series 43(4)* (1991), 569–585.

[8] Batyrev, V. V., and Dais, D. I. Strong McKay correspondence, string-theoretic Hodge numbers and mirror symmetry. *Topology 35*, 4 (1996), 901–929.

[9] Beck, M., and Robins, S. Dedekind sums: a combinatorial-geometric viewpoint. *DIMACS Series in Discrete Mathematics and Theoretical Computer Science 64* (2004), 25–36.

[10] Biran, P. Constructing new ample divisors out of old ones. *Duke Mathematical Journal 98*, 1 (1999), 113–135.

[11] Bocklandt, R., Craw, A., and Vélez, A. Q. Geometric Reid's recipe for dimer models. *Mathematische Annalen 361*, 3-4 (2015), 689–723.

[12] Boole, G. Exposition of a general theory of linear transformations. *Cambridge Math. J. 3* (1841), 106–119.

[13] Bridgeland, T. Stability conditions on triangulated categories. *Annals of Mathematics* (2007), 317–345.

[14] Bridgeland, T., King, A., and Reid, M. The McKay correspondence as an equivalence of derived categories. *Journal of the American Mathematical Society 14*, 3 (2001), 535–554.

[15] Brion, M. Introduction to actions of algebraic groups. *Les cours du CIRM 1(1)* (2010), 1–22.

[16] Buckley, A., Reid, M., and Zhou, S. Ice cream and orbifold Riemann–Roch. *Izvestiya: Mathematics 77*, 3 (2013), 461.

[17] Cautis, S., Craw, A., and Logvinenko, T. Derived Reid's recipe for abelian subgroups of $\mathrm{SL}_3(\mathbb{C})$. *arXiv preprint arXiv:1205.3110* (2012).

[18] Choi, K., Cristofaro-Gardiner, D., Frenkel, D., Hutchings, M., and Ramos, V. G. B. Symplectic embeddings into four-dimensional concave toric domains. *Journal of Topology 7*, 4 (2014), 1054–1076.

[19] Cieliebak, K., Hofer, H., Latschev, J., and Schlenk, F. Quantitative symplectic geometry. *arXiv preprint math/0506191* (2005).

[20] Coates, T., Corti, A., Galkin, S., Golyshev, V., and Kasprzyk, A. Mirror symmetry and Fano manifolds. *arXiv preprint arXiv:1212.1722* (2012).

[21] Cox, D. The homogeneous coordinate ring of a toric variety. *Jour. Algebraic Geom. 4* (1995), 17–50.

[22] Cox, D. A. Recent developments in toric geometry. *arXiv preprint alg-geom/9606016* (1996).

[23] Cox, D. A., Little, J. B., and Schenck, H. K. *Toric varieties.* American Mathematical Soc., 2011.

[24] Craw, A. An explicit construction of the McKay correspondence for A-Hilb $\mathbb{C}^3$. *Journal of Algebra 285*, 2 (2005), 682–705.

[25] Craw, A., Heuberger, L., and Amador, J. T. Combinatorial Reid's recipe for consistent dimer models. *arXiv preprint arXiv:2001.07506* (2020).

[26] Craw, A., and Ishii, A. Flops of G-Hilb and equivalences of derived categories by variation of GIT quotient. *Duke Mathematical Journal 124*, 2 (2004), 259–307.

[27] Craw, A., Ito, Y., and Karmazyn, J. Multigraded linear series and recollement. *Mathematische Zeitschrift 289*, 1-2 (2018), 535–565.

[28] Craw, A., and Reid, M. How to calculate A-Hilb $\mathbb{C}^3$. In *Séminaires & Congrés* (2002), Citeseer.

[29] Crawley-Boevey, W. Geometry of the moment map for representations of quivers. *Compositio Mathematica 126(3)* (2001), 257–293.

[30] Cristofaro-Gardiner, D. Symplectic embeddings from concave toric domains into convex ones. *arXiv preprint arXiv:1409.4378* (2014).

[31] Cristofaro-Gardiner, D., Hutchings, M., and Ramos, V. G. B. The asymptotics of ECH capacities. *Inventiones mathematicae 199*, 1 (2015), 187–214.

[32] Cristofaro-Gardiner, D., and Kleinman, A. Ehrhart polynomials and symplectic embeddings of ellipsoids. *arXiv preprint arXiv:1307.5493* (2013).

[33] Cristofaro-Gardiner, D., and Savale, N. Sub-leading asymptotics of ECH capacities. *arXiv preprint arXiv:1811.00485* (2018).

[34] Dais, Dimitrios, I. Geometric combinatorics in the study of compact toric surfaces. *Contemp. Math 423*, 71–123.

[35] De Loera, J. A., and McAllister, T. B. Vertices of Gelfand–Tsetlin polytopes. *Discrete & Computational Geometry 32*, 4 (2004), 459–470.

[36] De Loera, J. A., and McAllister, T. B. On the computation of Clebsch–Gordan coefficients and the dilation effect. *Experimental Mathematics 15*, 1 (2006), 7–19.

[37] Denef, J., and Loeser, F. Motivic integration, quotient singularities and the McKay correspondence. *Compositio mathematica 131*, 3 (2002), 267–290.

[38] Ehrhart, E. Sur un probleme de géométrie diophantienne linéaire ii. *J. reine angew. Math 227*, 25 (1967), C49.

[39] Ginzburg, V. Lectures on Nakajima's quiver varieties. *arXiv preprint arXiv:0905.0686* (2009).

[40] Gromov, M. Pseudo holomorphic curves in symplectic manifolds. *Inventiones mathematicae 82*, 2 (1985), 307–347.

[41] Haase, C., Juhnke-Kubitzke, M., Sanyal, R., and Theobald, T. Mixed Ehrhart polynomials. *arXiv preprint arXiv:1509.02254* (2015).

[42] Haase, C., and McAllister, T. B. Quasi-period collapse and $\mathrm{GL}_n(\mathbb{Z})$-scissors congruence in rational polytopes. *Contemporary Mathematics 452* (2008), 115–122.

[43] Hacking, P., and Prokhorov, Y. Smoothable del Pezzo surfaces with quotient singularities. *Compositio Mathematica 146*, 1 (2010), 169–192.

[44] Hutchings, M. Quantitative embedded contact homology. *Journal of Differential Geometry 88*, 2 (2011), 231–266.

[45] Hutchings, M. Lecture notes on embedded contact homology. In *Contact and symplectic topology*. Springer, 2014, pp. 389–484.

[46] Hutchings, M. ECH capacities and the ruelle invariant. *arXiv preprint arXiv:1910.08260* (2019).

[47] Ishii, A., and Ueda, K. Dimer models and the special McKay correspondence. *Geometry & Topology 19*, 6 (2016), 3405–3466.

[48] Ito, Y., and Nakajima, H. McKay correspondence and Hilbert schemes in dimension three. *Topology 6*, 39 (2000), 1155–1191.

[49] Ito, Y., and Nakamura, I. McKay correspondence and Hilbert schemes. *Proceedings of the Japan Academy, Series A, Mathematical Sciences 72*, 7 (1996), 135–138.

[50] Kasprzyk, A. M., and Wormleighton, B. Quasi-period collapse for duals to Fano polygons: an explanation arising from algebraic geometry. *arXiv preprint arXiv:1810.12472* (2018).

[51] King, A. D. Moduli of representations of finite dimensional algebras. *The Quarterly Journal of Mathematics 45(4)* (1994), 515–530.

[52] Kollár, J., and Shepherd-Barron, N. I. Threefolds and deformations of surface singularities. *Inventiones mathematicae 91*, 2 (1988), 299–338.

[53] Kronheimer, P. B. The construction of ALE spaces as hyper-Kähler quotients. *J. Diff. Geom 29*, 3 (1989), 665–683.

[54] Macrì, E., and Schmidt, B. Lectures on Bridgeland stability. In *Moduli of curves*. Springer, 2017, pp. 139–211.

[55] McAllister, T. B., and Woods, K. M. The minimum period of the Ehrhart quasi-polynomial of a rational polytope. *arXiv preprint math/0310255* (2003).

[56] McDuff, D. The Hofer conjecture on embedding symplectic ellipsoids. *Journal of Differential Geometry 88*, 3 (2011), 519–532.

[57] McKay, J. Cartan matrices, finite groups of quaternions, and Kleinian singularities. *Proceedings of the American Mathematical Society 81*, 1 (1981), 153–154.

[58] McMullen, P. The polytope algebra. *Advances in Mathematics 78*, 1 (1989), 76–130.

[59] McMullen, P. On simple polytopes. *Inventiones mathematicae 113*, 1 (1993), 419–444.

[60] Migliore, J. C., Nagel, U., and Zanello, F. A characterization of Gorenstein Hilbert functions in codimension four with small initial degree. *arXiv preprint math/0703901* (2007).

[61] Nakajima, H. Instantons on ALE spaces, quiver varieties, and Kac–Moody algebras. *Duke Mathematical Journal 76(2)* (1994), 365–416.

[62] Nakamura, I. Hilbert schemes of abelian group orbits. *Journal of Algebraic Geometry 10*, 4 (2001), 757–780.

[63] Namikawa, Y. Poisson deformations and birational geometry. *J. Math. Sci. Univ. Tokyo 22* (2015), 339–359.

[64] Pick, G. Geometrisches zur Zahlenlehre. *Sitzenber. Lotos (Prague) 19* (1899), 311–319.

[65] Pommersheim, J. E. Toric varieties, lattice points and Dedekind sums. *Mathematische Annalen 295*, 1 (1993), 1–24.

[66] Proudfoot, N. Geometric invariant theory and projective toric varieties. *Contemporary Mathematics 388* (2005), 161.

[67] Proudfoot, N. A survey of hypertoric geometry and topology. In *Proc. Int. Conf. on Toric Topology* (2006), Citeseer.

[68] Reid, M. Young person's guide to canonical singularities. *Algebraic geometry, Bowdoin 46* (1985), 345–414.

[69] Reid, M. La correspondance de McKay. *Asterisque-Societe Mathematique de France 276* (2002), 53–72.

[70] Rosenlicht, M. Some basic theorems on algebraic groups. *American Journal of Mathematics 78(2)* (1956), 401–443.

[71] Song, A. Existence of infinitely many minimal hypersurfaces in closed manifolds. *arXiv preprint arXiv:1806.08816* (2018).

[72] Sun, W. An estimate on energy of min-max Seiberg–Witten Floer generators. *arXiv preprint arXiv:1801.02301* (2018).

[73] Takahashi, K. On essential representations in the McKay correspondence for $\mathrm{SL}_3(\mathbb{C})$. Master's thesis, Nagoya University, 2011.

[74] Wilson, P. M. H. The Kähler cone on Calabi–Yau threefolds. *Inventiones mathematicae 107*, 1 (1992), 561–583.

[75] Wormleighton, B. ECH capacities, Ehrhart theory, and toric varieties. *arXiv preprint arXiv:1906.02237* (2019).

[76] Wormleighton, B. Walls for G-Hilb via Reid's recipe. *arXiv preprint arXiv:1908.05748* (2019).

[77] Wormleighton, B. Reconstruction of singularities on orbifold del Pezzo surfaces from their Hilbert series. *Communications in Algebra 48(1)* (2020), 119–140.

[78] Zagier, D. Higher dimensional Dedekind sums. *Mathematische Annalen 202*, 2 (1973), 149–172.