UC Riverside

UC Riverside Electronic Theses and Dissertations

Title

Representation Theory for Oriented Matroids and Combinatorial Intersection Homology for Gorenstein* Meet-Semilattices

Permalink

https://escholarship.org/uc/item/0k0971v2

Author

Kowalenko, Ethan

Publication Date

2021

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, availalbe at https://creativecommons.org/licenses/by/4.0/

Peer reviewed|Thesis/dissertation

eScholarship.org Powered by the California Digital Library
University of California

UNIVERSITY OF CALIFORNIA
RIVERSIDE

Representation Theory for Oriented Matroids and Combinatorial Intersection Homology for Gorenstein* Meet-Semilattices

A Dissertation submitted in partial satisfaction
of the requirements for the degree of

Doctor of Philosophy

in

Mathematics

by

Ethan Kowalenko

June 2021

Dissertation Committee:

Dr. Carl Mautner, Chairperson
Dr. Wee Liang Gan
Dr. José Gonzalez

Copyright by
Ethan Kowalenko
2021

The Dissertation of Ethan Kowalenko is approved:

Committee Chairperson

University of California, Riverside

Acknowledgments

First and foremost, I would like to thank Carl Mautner for his guidance, support, formal and informal discussions, and time. You helped me to clarify my ideas too many times to count, and I always looked forward to our meetings. I am so happy to have been your student.

For suggestions on rough versions of the two projects in this draft, I would also like to thank Jens Eberhardt, Joe Moeller, and Nick Proudfoot, as well as Jim Lawrence for helping Carl and I with some silly questions.

I am also deeply grateful to Alessandra Constantini, who, in one comment, helped me realize the importance of the word "Cohen-Macaulay," and therefore helped shape a large portion of this thesis.

I also benefited greatly from the various seminars at UCR, including Lie Theory, Algebraic Geometry, and Commutative Algebra.

I would also like to thank Emma, JD, Derek, Mike, Ripley and Katie for helping me be human throughout graduate school. I look forward to sleeping on every one of your couches in the future.

To my parents. No line I write can express how grateful I am.

ABSTRACT OF THE DISSERTATION

Representation Theory for Oriented Matroids and Combinatorial Intersection Homology for Gorenstein* Meet-Semilattices

by

Ethan Kowalenko

Doctor of Philosophy, Graduate Program in Mathematics
University of California, Riverside, June 2021
Dr. Carl Mautner, Chairperson

This dissertation consist of two distinct parts, coming from different parts of algebraic combinatorics.

In the first part, we discuss representation theory for oriented matroids, a joint project with Carl Mautner. In particular, we consider oriented matroid programs, which are a combinatorial abstraction of linear programming problems. To such a program $\mathbf{P}$, together with a *parameter space* U, we associate a finite-dimensional quadratic algebra $A(\mathbf{P}, U)$. These algebras generalize those of Braden-Licata-Proudfoot-Webster in the realizable case. When the program is Euclidean, we prove that our algebras are quasihereditary and Koszul. We also have partial related results in the non-Euclidean case.

In the second part, we develop intersection homology for certain posets. Combinatorial intersection homology for fans (IHF) is a sheaf-based theory modeled on the equivariant intersection cohomology for toric varieties, which was used to prove that the coefficients of toric h-polynomials of polytopes are nonnegative. More generally, Stanley conjectures that the toric h-polynomials of Cohen-Macaulay lower Eulerian meet-semilattices also have

nonnegative coefficients. This paper builds an analogue of IHF for this more general class of posets, and we make a hard Lefschetz conjecture which would solve Stanley's conjecture.

Contents

List of Figures **xi**

I Representation Theory for Oriented Matroids **1**

1 Introduction **2**
1.1 Hypertoric category $\mathcal{O}$ . . . 3
1.2 Matroidal setting . . . 6
1.3 Main results . . . 8
1.4 Derived Morita equivalence . . . 11
1.5 Matroidal Schur algebras . . . 13
1.6 Outline of paper . . . 14

2 Combinatorial Setup **15**
2.1 Matroids and oriented matroids . . . 15
2.1.1 Matroids . . . 16
2.1.2 Oriented matroids . . . 17
2.1.3 Generic oriented matroid programs . . . 18
2.1.4 Bounded feasible topes and sign vectors . . . 23
2.1.5 Cone relation and Euclidean oriented matroid programs . . . 27
2.1.6 On the existence of Euclidean generic programs . . . 30
2.2 Face rings and matroids . . . 32
2.2.1 Linear systems of parameters and duality . . . 34

3 The Algebras **35**
3.1 The algebra A . . . 35
3.1.1 The definition of A . . . 35
3.1.2 Expressions for elements of A . . . 36
3.1.3 Alternative description of A . . . 40
3.2 The quadratic dual of A . . . 42
3.3 The algebra B . . . 48
3.3.1 A topological lemma . . . 48

3.3.2 The algebra $\widetilde{B}$ . . . 53
3.4 The center of B . . . 60
3.4.1 The topology of affine space and feasible Edmonds Mandel face lattices 61
3.4.2 The proof of Theorem 3.4.1 . . . 64

4 Module Categories and their Derived Categories 70
4.1 The module category of A . . . 70
4.1.1 A is a quasi-hereditary algebra when $\mathbf{P}$ is Euclidean . . . 72
4.1.2 The structure of projectives when $\mathbf{P}$ is not Euclidean . . . 75
4.1.3 A is a Koszul algebra when $\mathbf{P}$ is Euclidean . . . 80
4.1.4 Numerical identity for Hilbert polynomials . . . 83
4.1.5 Self-dual projectives . . . 87
4.2 Derived Morita equivalence . . . 91
4.2.1 The definition and properties of the functor . . . 92
4.2.2 Ringel duality and composition of functors . . . 99

II Combinatorial Intersection Homology for Gorenstein* Meet-Semilattices 103

5 Introduction 104
5.1 Historical motivation . . . 104
5.1.1 Toric h-polynomials . . . 106
5.2 Separating the fan from its realization . . . 107
5.3 Outline of the paper . . . 110

6 Posets and Sheaves 112
6.1 Topological properties of posets . . . 114
6.1.1 Dually connected and diamond transitivity . . . 115
6.2 Sheaves on posets . . . 116
6.3 Cellular poset cohomology and transition coefficients . . . 119
6.3.1 Cellular poset homology with coefficients in a sheaf . . . 122
6.3.2 Comparing the chain complex with global sections for subthin posets. 124

7 Sheaf analogues of facewise functions 126
7.1 Linearization sheaves . . . 127
7.1.1 Homogenization . . . 131
7.2 The sheaf of facewise polynomials . . . 133
7.3 Pure sheaves and the minimal extension sheaf . . . 135

8 Equivariant formality and subdivisions 138
8.1 Equivariant formality . . . 138
8.2 Subdivisions and a decomposition theorem . . . 141

9 Parts of the Kähler package **146**
9.1 Duality of pure sheaves . . . 147
9.1.1 The definition of $\mathbb{D}\mathcal{F}$ . . . 148
9.1.2 Properties of the duality . . . 152
9.1.3 Duality for $\mathcal{E}$ and an intersection pairing . . . 159
9.2 Hard Lefschetz conjecture and toric h-vectors . . . 161
9.2.1 Making the intersection pairing canonical . . . 162
9.2.2 Toric h-polynomials . . . 164

Bibliography **166**

List of Figures

1.1 Hyperplane arrangement and corresponding pseudosphere arrangement. . . . 5
1.2 Ringel example 8

3.1 $\|\mathcal{A} \cup \mathcal{A}^{\infty}\|$, $\|\Delta_{ord}(\mathcal{A})\|$, $\|\Xi\|$ 62

6.1 The Hasse diagram of a diamond $[r, t] = \{r, s, s', t\}$. 113

Part I

Representation Theory for Oriented Matroids

Chapter 1

Introduction

In [BLPW10, BLPW12], Braden–Licata–Proudfoot–Webster introduced a class of finite-dimensional algebras related to the combinatorics of hyperplane arrangements, whose representation theory is closely analogous to the integral blocks of Bernstein–Gelfand–Gelfand (BGG) category $\mathcal{O}$. Recall that BGG category $\mathcal{O}$ plays an important role in Lie theory and can be described using the geometry of the Springer resolution. Braden–Licata–Proudfoot–Webster discovered their algebras by analogy, motivated by the geometry of toric hyperkähler (or hypertoric) varieties, but the algebras can be defined from basic linear algebra data. The input for their definition was the data of a *polarized arrangement* $\mathcal{V} = (V, \eta, \xi)$, where $V \subset \mathbb{R}^n$ is a d-dimensional linear subspace, $\eta \in \mathbb{R}^n/V$ is a (generic) vector and $\xi \in V^*$ is a (generic) covector. Braden–Licata–Proudfoot–Webster [BLPW16] and others (e.g., [Los17]) have since introduced and studied other such geometric categories $\mathcal{O}$ associated to conical symplectic resolutions.

In this paper we extend the definition of Braden–Licata–Proudfoot–Webster in a different, more combinatorial direction: from the setting of polarized arrangements to the combinatorics of oriented matroids. More precisely, the role of $V \subset \mathbb{R}^n$ is replaced by a rank d orientable matroid M and parameter space U and the role of η and ξ by an oriented matroid program $(\widetilde{\mathcal{M}}, g, f)$ that extends and lifts an orientation of M. One motivation for our work was a desire to categorify and better understand the matroidal Schur algebras of [BM17, BM19].

To explain our results and motivation, we first recall in more detail the results of Braden–Licata–Proudfoot–Webster.

1.1 Hypertoric category $\mathcal{O}$

In [BLPW10], Braden–Licata–Proudfoot–Webster defined a quadratic algebra $A(\mathcal{V})$. One motivation was a description of a regular block of category $\mathcal{O}$ as arising from a quantization of the structure sheaf of $T^*(G/B)$, the cotangent bundle of a flag variety. When $\mathcal{V}$ is rational (meaning that V, η, and ξ are defined over $\mathbb{Q}$), one may associate to $\mathcal{V}$ a hyperkähler variety $\mathfrak{M}$[1], sometimes called a hypertoric variety, which behaves in various ways like the cotangent bundle of a flag variety. Braden–Licata–Proudfoot–Webster show that in this case the category of representations of $A(\mathcal{V})$ is equivalent to that obtained by applying the same sort of quantization construction to the hypertoric variety for $\mathcal{V}$. Moreover they show:

Theorem 1.1.1 (Braden–Licata–Proudfoot–Webster)**.** *Let $\mathcal{V} = (V, \eta, \xi)$ be a polarized arrangement where η and ξ are generic.*

[1]More precisely, the subspace V and vector η alone determine the variety $\mathfrak{M}$. The covector ξ can be used to endow $\mathfrak{M}$ with a $\mathbb{C}^*$-action.

1. *The algebra $A(\mathcal{V})$ is quadratic with quadratic dual $A(\mathcal{V}^\vee)$, where $\mathcal{V}^\vee = (V^\perp, -\xi, -\eta)$ denotes the Gale dual polarized arrangement.*

2. *The algebra $A(\mathcal{V})$ is quasi-hereditary.*

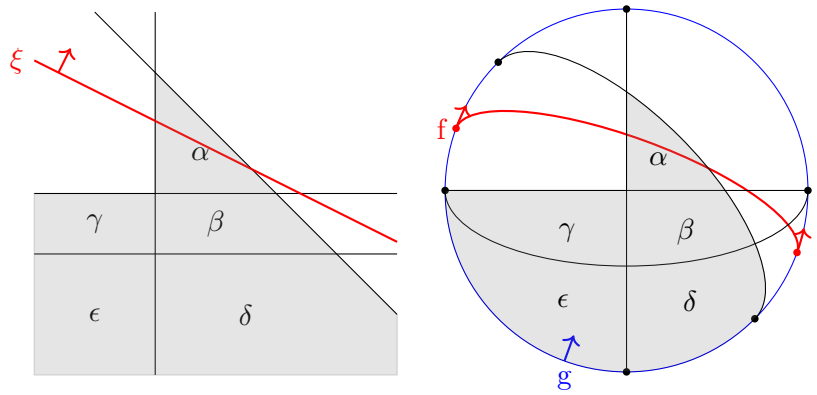
3. *The algebra $A(\mathcal{V})$ is Koszul (and thus by the first result, Koszul dual to $A(\mathcal{V}^\vee)$).*

4. *Up to derived Morita equivalence, the algebra $A(\mathcal{V})$ depends only on $V \subset \mathbb{R}^n$ and not on η or ξ.*

To give a feeling for the representation theory of these algebras, we will describe a labeling of the simple modules for $A(\mathcal{V})$. It is convenient to consider the following hyperplane arrangement defined by $\mathcal{V} = (V, \eta, \xi)$. Note that $\eta \in \mathbb{R}^n/V$ can be viewed as the affine subspace $\eta + V \subset \mathbb{R}^n$, and we consider the arrangement of hyperplanes in $\eta + V$ cut out by the coordinate hyperplanes of $\mathbb{R}^n$. The genericity condition on η is the requirement that the resulting arrangement be *simple*, meaning that the nonempty intersection of m hyperplanes has codimension m. The covector $\xi \in V^*$ lifts to an affine linear functional on $\eta + V$. The genericity condition on ξ is the requirement that ξ be nonconstant on any positive dimensional intersection of V and a coordinate subspace.

Example 1.1.2. The polarized arrangement from Example 2.2 of [BLPW10] consists of a two-dimensional subspace $V \subset \mathbb{R}^4$ and the choices of η and ξ there produces the hyperplane arrangement depicted in Figure 1.1(a).

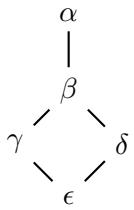
The set $\mathcal{P}$ of chambers of the hyperplane arrangement in $\eta + V$ that are bounded with respect to ξ parametrize the simple modules $\{L_\alpha\}_{\alpha \in \mathcal{P}}$ for $A(\mathcal{V})$. In Example 1.1.2, we can label these chambers $\alpha, \beta, \gamma, \delta, \epsilon$ as in Figure 1.1(a).



(a) (b)

Figure 1.1: Hyperplane arrangement and corresponding pseudosphere arrangement.

For each bounded chamber $\alpha \in \mathcal{P}$, let $\beta \preceq \alpha$ if β is contained in the cone generated by α which originates from the vertex of α which is maximal with respect to ξ. The transitive closure of this relation gives the highest weight partial order on simple objects for the quasihereditary structure in the theorem. In the example above, this produces the poset described by the following Hasse diagram:



More precisely, Braden–Licata–Proudfoot–Webster define standard modules V_α for every $\alpha \in \mathcal{P}$ and prove (see the proof of [BLPW10, Theorem 5.23]):

Theorem 1.1.3 (Braden–Licata–Proudfoot–Webster)**.** *For any $\alpha \in \mathcal{P}$ the indecomposable projective cover P_α of L_α has a filtration with successive subquotients isomorphic to V_β for each $\beta \succeq \alpha$ and each such standard module appears exactly once.*

1.2 Matroidal setting

Fix a field k and a finite index set I. In this paper we will begin with an orientable matroid M of rank d and a choice of parameter space $U \subset k^I$ for M. By parameter space, we mean a subspace $U \subset k^I$ such that the composition $U \hookrightarrow k^n \twoheadrightarrow \mathrm{Span}\{t_i \mid i \in b\}$ is an isomorphism for any basis b of M.

Example 1.2.1. Note that the subspace $V \subset \mathbb{R}^n$ in a polarized arrangement of Braden–Licata–Proudfoot–Webster provides such a pair for $k = \mathbb{R}$: let M be the matroid on the index set $I = \{1, \ldots, n\}$ represented by the coordinate functions of $\mathbb{R}^n$ restricted to V, viewed as vectors $x_1, \ldots, x_n \in V^*$, and let $U = V$.

Let $\mathcal{M}$ be an orientation of M, meaning an oriented matroid $\mathcal{M}$ such that $\underline{\mathcal{M}} = M$, where $\underline{\mathcal{M}}$ denotes the underlying unoriented matroid. (In the polarized arrangement example, there is a natural choice for $\mathcal{M}$, as M is represented by vectors in a real vector space.)

The remaining input data we need is the structure of a oriented matroid program $\mathbf{P} = (\widetilde{\mathcal{M}}, g, f)$, meaning $\widetilde{\mathcal{M}}$ is an oriented matroid on the underlying set $I \sqcup \{g, f\}$ such that g is not a loop, f is not a coloop, and $\mathcal{M}$ is obtained from $\widetilde{\mathcal{M}}$ by deleting f and contracting g. This last condition is written as $(\widetilde{\mathcal{M}} \backslash f)/g = \mathcal{M}$. Like we did for η and ξ, we ask that g and f be sufficiently generic (see Definition 2.1.4).

The matroid M is determined by $\mathbf{P}$, so we can and will omit it from our notation and consider pairs $(\mathbf{P}, U)$ where $\mathbf{P} = (\widetilde{\mathcal{M}}, g, f)$ is a sufficiently generic oriented matroid program and $U \subset k^I$ is a parameter space for the underlying (unoriented) matroid $M = \underline{(\widetilde{\mathcal{M}} \backslash f)/g}$.

Example 1.2.2. Polarized arrangements give a natural class of examples. For a d-dimensional polarized arrangement (V, η, ξ), consider the $d+1$-dimensional subspace $\widetilde{V}$ of $\mathbb{R}^n \times \mathbb{R}_g \times \mathbb{R}_f$ spanned by the graph of $\xi : V \to \mathbb{R}_f$ and the vector $(\overline{\eta}, 1, 0) \in \mathbb{R}^n \times \mathbb{R}_g \times \mathbb{R}_f$, where $\overline{\eta}$ is any representative of the coset $\eta \in \mathbb{R}^n/V$. Let $\widetilde{\mathcal{M}}$ be the oriented matroid on the set $\{1, \ldots, n\} \sqcup \{g, f\}$ defined by the coordinate functions $x_1, \ldots, x_n, x_g, x_f \in \widetilde{V}^*$.

Not every oriented matroid program $\mathbf{P}$ comes from a polarized arrangement, but by the Topological Realization Theorem of Folkman–Lawrence, every loop-free program $\mathbf{P}$ can be expressed a *pseudosphere arrangement* - a topological representation generalizing the notion of a hyperplane arrangement.

Example 1.2.3. Figure 1.1(b) shows the feasible region of the pseudosphere arrangement corresponding to the polarized arrangement from Example 1.1.2.

Example 1.2.4. Figure 1.2 depicts the feasible part of a pseudosphere arrangement, where $|I| = 8$ and M is the uniform rank 2 matroid on 8 points, that defines a non-realizable oriented matroid program $\mathbf{P} = (\widetilde{\mathcal{M}}, g, f)$. Here the oriented submatroid $\widetilde{\mathcal{M}}\backslash f$, a rank 3 oriented matroid on 9 points, is the non-stretchable simple arrangement of 9 pseudolines defined by Ringel [Rin55] as a perturbation the Pappus matroid.

Remark 1.2.5. Every oriented matroid program $\mathbf{P}$ where $d = 2$ and $|I| \leq 7$ is realizable, so the program described in Example 1.2.4 is a minimal non-realizable example.

For a pair $(\mathbf{P} = (\widetilde{\mathcal{M}}, g, f), U)$ as above we define the dual pair $(\mathbf{P}^\vee, U^\perp)$, where $\mathbf{P}^\vee = (\widetilde{\mathcal{M}}^\vee, f, g)$ is the dual oriented matroid program (here the roles of f and g are swapped), and $U^\perp \subset k^I$ is the orthogonal complement.

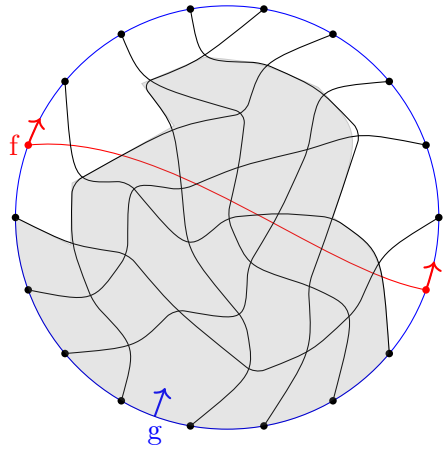


Figure 1.2: Ringel example

Remark 1.2.6. It is an exercise in linear algebra to check that when the oriented matroid program $\mathbf{P} = (\widetilde{\mathcal{M}}, g, f)$ comes from a polarized arrangement $\mathcal{V} = (V, \eta, \xi)$ as in Example 1.2.2, this duality agrees with the standard Gale duality of linear programming. In other words, the dual program $\mathbf{P}^\vee = (\widetilde{\mathcal{M}}^\vee, f, g)$ is the oriented matroid program associated to the Gale dual polarized arrangement $\mathcal{V}^\vee = (V^\perp, -\xi, -\eta)$.

1.3 Main results

As above, let the pair $(\mathbf{P}, U)$ consist of a sufficiently generic oriented matroid program $\mathbf{P} = (\widetilde{\mathcal{M}}, g, f)$ together with a parameter space $U \subset k^I$ for the (unoriented) matroid $M = \underline{(\widetilde{\mathcal{M}} \backslash f)/g}$. Modifying the definition of Braden–Licata–Proudfoot–Webster to this setting, we introduce a finite-dimensional algebra $A(\mathbf{P}, U)$ over k. In particular, in the realizable case of Example 1.2.1 and 1.2.2, one recovers the original algebra $A(\mathbf{P}, V) = A(\mathcal{V})$.

In the more general setting, we observe that part (1) of Theorem 1.1.1 extends without modification:

Theorem 1.3.1. *Let* $(\mathbf{P}, U)$ *be a pair as above. The algebra* $A(\mathbf{P}, U)$ *is quadratic with quadratic dual* $A(\mathbf{P}^\vee, U^\perp)$ *corresponding to the dual pair.*

Similarly to the realizable case, the simple modules for $A(\mathbf{P}, U)$ are labelled by the set $\mathcal{P}$ of bounded, feasible topes. For example, in Example 1.2.4 the bounded, feasible topes correspond to shaded regions in Figure 1.2. Again one can define a cone relation on $\mathcal{P}$ and standard modules V_α for each $\alpha \in \mathcal{P}$.

However, unlike in the realizable case, the transitive closure of the cone relation need not define a poset. An oriented matroid program $\mathbf{P} = (\widetilde{\mathcal{M}}, g, f)$ is said to be *Euclidean* if the transitive closure of the cone relation on bounded, feasible topes of $\mathbf{P}$ is a poset.

Using this condition, we obtain the following analogue of Theorems 1.1.1(2), 1.1.1(3) and 1.1.3.

Theorem 1.3.2. *For a pair* $(\mathbf{P}, U)$ *be as above with the additional assumption that the program* $\mathbf{P}$ *is* ***Euclidean****, the algebra* $A(\mathbf{P}, U)$ *is quasi-hereditary and Koszul.*

Moreover, for any $\alpha \in \mathcal{P}$ *the indecomposable projective cover* P_α *of* L_α *has a filtration with successive subquotients isomorphic to* V_β *for each* $\beta \geq \alpha$ *and each such standard module appears exactly once.*

Remark 1.3.3. While oriented matroid programs are not always Euclidean, every oriented matroid program of rank at most 3 (equivalently d at most 2) is Euclidean. Thus there are plenty of Euclidean, non-realizable programs, such as Example 1.2.4.

We do not know whether or not every non-realizable oriented matroid $\mathcal{M}$ admits a Euclidean program $\mathbf{P} = (\widetilde{\mathcal{M}}, g, f)$ such that $\widetilde{\mathcal{M}}/g\backslash f = \mathcal{M}$. For connections to a well-known conjecture of Las Vergnas, see the discussion surrounding Proposition 2.1.29.

We observe in Example 4.1.10 that in the non-Euclidean case, $A(\mathbf{P}, U)$ need not be quasi-hereditary. In particular, we exhibit a program $\mathbf{P}$ and a projective module for $A(\mathbf{P}, U)$ which does not have a standard filtration. This program $\mathbf{P} = (EFM(8), g, f)$ is a minimal (with respect to rank and number of elements) non-Euclidean program, using a uniform oriented matroid $EFM(8)$ of rank 4 on 8 elements which is hailed as a "counter-example to almost everything" in [BLVS$^+$99].

However, in Theorem 4.1.9 we do prove that for any oriented matroid program $\mathbf{P}$ the following analogue of Theorem 1.1.3 holds on the level of the Grothendieck group of graded $A(\mathbf{P}, U)$-modules.

Theorem 1.3.4. *For any generic oriented matroid program $\mathbf{P}$ and any $\alpha \in \mathcal{P}$, the class of the indecomposable projective P_α in the Grothendieck group can be expressed as the sum:*

$$[P_\alpha] = \sum_{\gamma \geq \alpha} q^{d_{\alpha\gamma}} [V_\gamma],$$

where $d_{\alpha\gamma}$ denotes the distance between the topes α and γ.

While our proof of Koszulity in the Euclidean case relies on $A(\mathbf{P}, U)$ being quasi-hereditary, it is conceivable that $A(\mathbf{P}, U)$ is Koszul more generally. As evidence in this direction, in Theorem 4.1.14 we prove the Hilbert series of $A(\mathbf{P}, U)$ and $A(\mathbf{P}^\vee, U^\perp)$ satisfy the numerical identity discussed in [BGS96, Lemma 2.11.1].

1.4 Derived Morita equivalence

In light of Theorem 1.1.1(4) it seems natural to ask:

Question 1.4.1. Let M be an orientable matroid and U a choice of parameter space for M. For any two orientations $\mathcal{M}_1, \mathcal{M}_2$ of M and generic oriented matroid programs $\mathbf{P}_1 = (\widetilde{\mathcal{M}}_1, g_1, f_1), \mathbf{P}_2 = (\widetilde{\mathcal{M}}_2, g_2, f_2)$ such that $\widetilde{\mathcal{M}}_i/g_i\backslash f_i = \mathcal{M}_i,\ i = 1, 2$, are the algebras $A(\mathbf{P}_1, U)$ and $A(\mathbf{P}_2, U)$ derived Morita equivalent?

If the answer to this question is yes, it would appear to give a rather interesting algebraic invariant of the matroid M. Or weaker, one might still hope for an affirmative answer under the assumption that $\mathcal{M}_1 = \mathcal{M}_2$:

Question 1.4.2. Let $\mathcal{M}$ be an oriented matroid and U a choice of parameter space for $M = \underline{\mathcal{M}}$. For any two generic oriented matroid programs $\mathbf{P}_1 = (\widetilde{\mathcal{M}}_1, g_1, f_1), \mathbf{P}_2 = (\widetilde{\mathcal{M}}_2, g_2, f_2)$ such that $\mathcal{M} = \widetilde{\mathcal{M}}_1/g_1\backslash f_1 = \widetilde{\mathcal{M}}_2/g_2\backslash f_2$, are the algebras $A(\mathbf{P}_1, U)$ and $A(\mathbf{P}_2, U)$ derived Morita equivalent?

If the answer to one or both of these questions is no, the number of derived Morita equivalence classes could also provide a interesting invariant of M or $\mathcal{M}$.

As a partial result in this direction, following the strategy of Braden–Licata–Proudfoot–Webster, we prove the following theorem in Section 4.2.

Theorem 1.4.3. *Fix $\mathcal{M}$ and let $\mathbf{P}_1 = (\widetilde{\mathcal{M}}_1, g_1, f_1)$ and $\mathbf{P}_2 = (\widetilde{\mathcal{M}}_2, g_2, f_2)$ be Euclidean such that $\widetilde{\mathcal{M}}_i/g_i\backslash f_i = \mathcal{M}$ for $i = 1, 2$. Suppose in addition that the oriented matroid program $\mathbf{P}_{mid} = (\widetilde{\mathcal{M}}_{mid}, g_2, f_1)$ is also Euclidean, where $\mathbf{P}_{mid}$ is a generic oriented matroid program*[2]

[2]Note that such an oriented matroid program $(\widetilde{\mathcal{M}}_{\mathrm{mid}}, g_2, f_1)$ always exists and there will typically be many such oriented matroid programs. However the particular choice will not matter for us, because, as mentioned in Remark 2.1.3, all of our constructions depend only on the contraction and restriction oriented matroids $\widetilde{\mathcal{M}}_{\mathrm{mid}}/g_2$ and $\widetilde{\mathcal{M}}_{\mathrm{mid}}\backslash f_1$.

such that $\widetilde{\mathcal{M}}_{mid}/g_2 = \widetilde{\mathcal{M}}_1/g_1, \quad \widetilde{\mathcal{M}}_{mid}\backslash f_1 = \widetilde{\mathcal{M}}_2\backslash f_2$. *Then the bounded derived categories of graded finitely generated* $A(\mathbf{P}_1, U)$*- and* $A(\mathbf{P}_2, U)$*-modules are equivalent.*

This allow us to answer Questions 1.4.1 and 1.4.2 in some simple cases.

Corollary 1.4.4. *Question 1.4.2 has an affirmative answer for any oriented matroid* $\mathcal{M}$ *of rank 2.*

Proof. Recall from Remark 1.3.3, that any oriented matroid program of rank 3 is Euclidean, so for any $\mathbf{P}_1$ and $\mathbf{P}_2$, the three oriented matroid programs $\mathbf{P}_1, \mathbf{P}_2$ and $\mathbf{P}_{\text{mid}}$ are all Euclidean and the result follows from Theorem 1.4.3. □

Corollary 1.4.5. *Question 1.4.1 has an affirmative answer for* $M = U_{2,n}$*, the uniform matroid of rank 2 defined on a set* E_n *of* n *elements.*

Proof. By Corollary 1.4.4, it suffices to show that for any two orientations $\mathcal{M}_1$ and $\mathcal{M}_2$ of $U_{2,n}$, there are generic oriented matroid programs $\mathbf{P}_1 = (\widetilde{\mathcal{M}}_1, g_1, f_1), \mathbf{P}_2 = (\widetilde{\mathcal{M}}_2, g_2, f_2)$ such that $\mathcal{M}_1 = \widetilde{\mathcal{M}}_1/g_1\backslash f_1$ and $\mathcal{M}_2 = \widetilde{\mathcal{M}}_2/g_2\backslash f_2$, for which $A(\mathbf{P}_1, U)$ and $A(\mathbf{P}_2, U)$ are derived equivalent.

But any two orientations of $U_{2,n}$ are related by a relabeling of E_n and reorientation. Note that relabeling and reorientation each induce a canonical isomorphism between the associated algebras. □

The same sort of argument gives a handful of similar examples.

1.5 Matroidal Schur algebras

Motivated in part by [BLPW10], Braden–Mautner defined a *hypertoric Schur algebra* [BM19] - an analogue of the Schur algebra associated to affine hypertoric varieties. Recall that one can construct an affine hypertoric variety $\mathfrak{M}_0$ with the input of a rational subspace $V \subset \mathbb{R}^n$. In this setting the resulting hypertoric Schur algebra $R(V)$ can be interpreted as a convolution algebra for a union of resolutions of stratum slices of $\mathfrak{M}_0$. In particular, for a rational polarized arrangement (V, η, ξ) with the same underlying subspace $V \subset \mathbb{R}^n$, the convolution algebra for the resolution $\mathfrak{M} \to \mathfrak{M}_0$ is a subalgebra of the associated hypertoric Schur algebra. Braden–Proudfoot–Webster showed in [BPW16, Proposition 6.16, Example 6.18] that the convolution algebra of the resolution $\mathfrak{M} \to \mathfrak{M}_0$ is categorified by Harish-Chandra bimodules for hypertoric category $\mathcal{O}$. One expects the entire hypertoric Schur algebra to be similarly categorified by Harish-Chandra bimodules with more general support and similarly to obtain natural a q-deformation of the hypertoric Schur algebra, or q-hypertoric Schur algebra.

In [BM17], Braden–Mautner observed that the hypertoric Schur algebras studied in [BM19] can be defined in terms of the underlying matroid. Following this observation, they defined a *matroidal Schur algebra* $R(M)$ associated to any matroid M.

One motivation for defining the category $\mathcal{O}$ for oriented matroid programs was to provide the foundation to categorify and find natural q-deformations of matroidal Schur algebras for orientable, but non-realizable matroids using an appropriate category of Harish-Chandra bimodules.

1.6 Outline of paper

In Section 2 we describe the combinatorial set-up of oriented matroid programs and parameter spaces, as well as some useful topology. In Section 3.1 we define the algebra $A(\mathbf{P}, U)$ and in Section 3.2 we study its quadratic dual and prove Theorem 3.2.2. In Section 3.3 we define the algebra $B(\mathbf{P}, U)$ and prove Theorem 3.3.13. Section 3.4 develops more topology, resulting in a nice description of the center of $B(\mathbf{P}, U)$.

In Section 4.1, we study the category of finitely-generated (right) $A(\mathbf{P}, U)$-modules. Under the Euclidean assumption, we prove that $A(\mathbf{P}, U)$ is quasihereditary (Theorem 4.1.6) and Koszul (Theorem 4.1.13). In the non-Euclidean setting, we describe the Grothendieck class of the graded indecomposable projective modules (Theorem 4.1.9), prove the Koszulity condition on Hilbert series (Theorem 4.1.14) and show that Theorem 4.1.6 requires the Euclidean assumption (Example 4.1.10). In Section 4.2, we study the derived categories of graded finitely-generated $A(\mathbf{P}, U)$-modules for varying Euclidean $\mathbf{P}$ and a fixed $\mathcal{M}$ and prove Theorem 1.4.3.

Chapter 2

Combinatorial Setup

In this section we briefly introduce the notation we will need to work with oriented matroids, but assume some familiarity with the basic notions. To the uninitiated reader, we recommend [BLVS+99] (particularly the first chapter) for a survey and as a useful reference.

2.1 Matroids and oriented matroids

Matroids are a combinatorial abstraction of notions of arrangements of vectors in a vector space of any field. Oriented matroids are an abstraction of notions coming from hyperplane arrangements over ordered fields. Every oriented matroid $\mathcal{M}$ gives rise to a matroid $\underline{\mathcal{M}}$, called the *underlying matroid* of $\mathcal{M}$. We call a matroid which is the underlying matroid of an oriented matroid *orientable*. There are matroids which are not the underlying matroids of an oriented matroid.

For both matroids and oriented matroids, the various notions one may wish to abstract from each prospective case often lead to different but equivalent (or *cryptomorphic*)

axiom systems. For us, the most important axiom system for a matroid M is given by its set of independent sets $\mathcal{I}(M)$, whereas the most important axiom system for an oriented matroid $\mathcal{M}$ is given by its set of covectors $\mathcal{L}(\mathcal{M})$.

2.1.1 Matroids

Let E be a finite set.

Definition 2.1.1. A set $\mathcal{I} \subset 2^E$ is the *matroid complex* of a matroid M on E if $\varnothing \in \mathcal{I}$, $\mathcal{I}$ is closed under taking subsets, and if $A, B \in \mathcal{I}$ with $|A| > |B|$, then there is an $i \in A \setminus B$ such that $B \cup \{i\} \in \mathcal{I}$.

If $\mathcal{I} = \mathcal{I}(M)$ is the matroid complex of a matroid M on E, then we call the elements of $\mathcal{I}$ the *independent sets* of M, and the maximal elements of $\mathcal{I}$ (with respect to inclusion) are called the *bases* of M. Let $\mathbb{B}$ be the set of bases of a matroid M. The sets $\mathbb{B}$ and $\mathcal{I}$ determine each other for any matroid M.

It is possible to prove that all bases of a matroid M must have the same cardinality. Thus, the matroid complex of a matroid M on E is a pure abstract simplicial complex on the index set E. The *rank function* of a matroid M on E is $r : 2^E \to \mathbb{Z}_{\geq 0}$, given by setting $r(S)$ equal to the maximal $|A|$, where $A \subset E$ ranges over the subsets of S which is also in $\mathcal{I}$. The *rank* of M is $r(E)$, which equals the cardinality of any basis of M.

Every matroid M has a unique dual matroid $M^\vee$ on E, given by declaring

$$\mathcal{I}(M^\vee) := \{S \subset E \mid S \subset I \setminus b \text{ for some } b \in \mathbb{B}\}.$$

The rank of $M^\vee$ is therefore $|E| - d$, where d is the rank of M.

2.1.2 Oriented matroids

For a finite index set I and any function $Z: I \to \{0, +, -\}$, let $\underline{Z} := \{i \mid Z(i) \neq 0\} \subset I$ be the *support* of Z and $z(Z) := I \backslash \underline{Z}$ and be the *zero set* of Z. For any functions $Z, Z': I \to \{0, +, -\}$, their *composition* $Z \circ Z': I \to \{0, +, -\}$ is defined by

$$Z \circ Z'(i) = \begin{cases} Z(i) & \text{if } Z(i) \neq 0 \\ Z'(i) & \text{otherwise} \end{cases}$$

We say that Z is a *face* of Z' if $Z \circ Z' = Z'$.

Definition 2.1.2. A set $\mathcal{L} \subset \{0, +, -\}^I$ is the set of *covectors* of an oriented matroid $\mathcal{M}$ on I if

- $\mathcal{L}$ contains the zero function 0,
- if $Y \in \mathcal{L}$ then $-Y \in \mathcal{L}$,
- if $X, Y \in \mathcal{L}$ then $X \circ Y \in \mathcal{L}$, and
- if $X, Y \in \mathcal{L}$ and $i \in I$ with $X(i) \neq Y(i)$, then there exists $Z \in \mathcal{L}$ such that $Z(i) = 0$ and $Z(j) = X(j)$ whenever $X(j) = Y(j)$.

The set of covectors $\mathcal{L}(\mathcal{M})$ has a natural poset structure defined by $Y \leq X$ if Y is a face of X. The poset $\mathcal{L}(\mathcal{M})$ is pure with minimal element 0. The maximal elements of $\mathcal{L}(\mathcal{M})$ are called *topes* of $\mathcal{M}$, while the minimal elements of $\mathcal{L}(\mathcal{M}) \backslash \{0\}$ are called the *(signed) cocircuits* of $\mathcal{M}$. We will often omit the word "signed" when it will not cause confusion. We denote the set of cocircuits of $\mathcal{M}$ by $C^*(\mathcal{M})$. The set of covectors of $\mathcal{M}$ is in fact determined by $C^*(\mathcal{M})$, since every covector is a composition of cocircuits.

The rank of $Y \in \mathcal{L}(\mathcal{M})$ is given as $\rho(Y) = d - r(z(Y))$, where r is the rank function of the underlying matroid $\underline{\mathcal{M}}$ and d is the size of any (every) basis of $\underline{\mathcal{M}}$. The rank of any cocircuit is 1, while the rank of any tope is d. We also say that the *rank* of $\mathcal{M}$ is the d.

For $Y_1, Y_2 \in \mathcal{L}(\mathcal{M})$, the *join* $Y_1 \vee Y_2$ is the minimal covector that has both Y_1 and Y_2 as faces, which only exists if there is a tope T with both Y_1 and Y_2 as faces. The *meet* $Y_1 \wedge Y_2$ is the maximal covector face of both Y_1 and Y_2. Note that the meet of $Y_1, Y_2 \in \mathcal{L}(\mathcal{M})$ always exists, but is the zero function when Y_1 and Y_2 do not have a common cocircuit face.

Duality for oriented matroids is a bit trickier than duality for matroids. For $X, Y \in \{0, +, -\}^I$, we say that X and Y are *orthogonal* and write $X \perp Y$ if $\underline{X} \cap \underline{Y} = \varnothing$ or there exists $i, j \in \underline{X} \cap \underline{Y}$ such that $X(i) = Y(i)$ and $X(j) = -Y(j)$. Then the dual oriented matroid $\mathcal{M}^\vee$ of $\mathcal{M}$ is given by

$$\mathcal{L}(\mathcal{M}^\vee) := \{X \in \{0, +, -\}^I \mid X \perp Y \text{ for all } Y \in \mathcal{L}(\mathcal{M})\}.$$

We have that $(\mathcal{M}^\vee)^\vee = \mathcal{M}$, and that $\underline{(\mathcal{M}^\vee)} = \underline{\mathcal{M}}^\vee$.

Using duality, we define the *(signed) circuits* of $\mathcal{M}$ to be the cocircuits of $\mathcal{M}^\vee$. The set of signed circuits of $\mathcal{M}$ is $C(\mathcal{M}) := C^*(\mathcal{M}^\vee)$. The set of covectors of $\mathcal{M}$ is in fact determined by $C(\mathcal{M})$.

2.1.3 Generic oriented matroid programs

Let M be an orientable matroid of rank d on the set I. Let $\mathcal{M}$ be an oriented matroid such that $\underline{\mathcal{M}} = M$ is its underlying unoriented matroid.

Let $\mathbb{B}$ denote the set of bases for M. We let $C = C(\mathcal{M})$ denote the set of signed circuits and $C^* = C^*(\mathcal{M})$ the set of signed cocircuits, both regarded as subsets of the set of

functions $I \to \{0, +, -\}$. Note the unoriented matroid $M = \underline{\mathcal{M}}$ has circuits $\{\underline{X} \mid X \in C\}$ and cocircuits $\{\underline{Y} \mid Y \in C^*\}$. The dual oriented matroid $\mathcal{M}^\vee$ is given by switching the roles of the circuits and cocircuits (i.e. $C(\mathcal{M}^\vee) = C^*(\mathcal{M})$ and $C^*(\mathcal{M}^\vee) = C(\mathcal{M})$), while the bases $\mathbb{B}^\vee$ of the underlying matroid $M^\vee = \underline{\mathcal{M}}^\vee$ are the complements in I of the elements of $\mathbb{B}$.

We will often use the following two constructions to get new oriented matroids from old. Let $S \subset I$. Then the set

$$\{X \in C(\mathcal{M}) \mid \underline{X} \subset I \backslash S\}$$

is the set of circuits of an oriented matroid $\mathcal{M}\backslash S$ on $I\backslash S$, called the *deletion* of S from $\mathcal{M}$. The set

$$\{X|_S : S \to \{0, +, -\} \mid X \in C(\mathcal{M}) \text{ and } \underline{X} \cap S \text{ is inclusion minimal and nonempty}\}$$

gives the set of circuits of an oriented matroid $\mathcal{M}/(I\backslash S)$ on S, called the *contraction* of $\mathcal{M}$ to S. Duality exchanges contraction and deletion:

$$(\mathcal{M}/S)^\vee = \mathcal{M}^\vee \backslash S \quad \text{and} \quad (\mathcal{M}\backslash S)^\vee = \mathcal{M}^\vee / S.$$

An element $i \in I$ is a *loop* of $\mathcal{M}$ if $\{i\}$ is the support of a circuit of $\mathcal{M}$. Dually, $i \in I$ is a *coloop* of $\mathcal{M}$ if i is not contained in the support of any circuit of $\mathcal{M}$.

Let I be the set $E_n = \{1, \ldots, n\}$. An *oriented matroid program* $\mathbf{P} = (\widetilde{\mathcal{M}}, g, f)$ is an oriented matroid $\widetilde{\mathcal{M}}$ on the set $E_n \sqcup \{f, g\}$ such that $(\widetilde{\mathcal{M}}\backslash f)/g = \mathcal{M}$, g is not a loop and f is not a coloop. In particular, the rank of $\widetilde{\mathcal{M}}$ is $d + 1$, and this is defined to be the rank of $\mathbf{P}$.

The deletion $\mathcal{N} = \widetilde{\mathcal{M}}\backslash f$ of f from $\widetilde{\mathcal{M}}$ is called the corresponding *affine oriented matroid*.

Remark 2.1.3. Our constructions will only depend on the contraction $\widetilde{\mathcal{M}}/g$, which is a single element extension of $\mathcal{M}$ on $E_n \sqcup \{f\}$, and the deletion $\widetilde{\mathcal{M}}\backslash f$, which is a single element lift of $\mathcal{M}$ on $E_n \sqcup \{g\}$. Thus for our purposes it would be more natural to define an oriented matroid program as a pair, which we have taken to affectionately calling an *eft*, of a single element extension and single element lift of $\mathcal{M}$. We will refrain from doing so in this paper as the original notion appears to be standard in the field.

Definition 2.1.4. We say that the oriented matroid program $\mathbf{P} = (\widetilde{\mathcal{M}}, g, f)$ is *generic* for $\mathcal{M}$ if

1. for any cocircuit Y of $\mathcal{N} = \widetilde{\mathcal{M}}\backslash f$, if $|z(Y)| > d$, then $Y(g) = 0$, and
2. for any circuit X of $\widetilde{\mathcal{M}}/g$, if $|z(X)| > n - d$, then $X(f) = 0$.

Remark 2.1.5. As the rank of the oriented matroid $\widetilde{\mathcal{M}}/g$ on $E_n \sqcup \{f\}$ is d, for any circuit X of $\widetilde{\mathcal{M}}/g$, $|\underline{X}| \leq d+1$ and so the X has *at least* $n-d$ zero entries. In the case of equality, the $\underline{X}$ contains a basis of $\underline{\widetilde{\mathcal{M}}/g}$ and so $z(X)$ is independent in $(\underline{\widetilde{\mathcal{M}}/g})^\vee$.

Dually, for any cocircuit Y of $\mathcal{N}$, $|z(Y)| \geq d$ and if equality holds $z(Y)$ is independent in $\underline{\mathcal{N}}$.

Example 2.1.6. As explained in Example 1.2.2, every polarized arrangement $\mathcal{V} = (V, \eta, \xi)$ naturally gives rise to an oriented matroid program $\mathbf{P}$. If η and ξ are generic in the sense of Theorem 1.1.1, then $\mathbf{P}$ is generic as well.

We now deduce some simple consequences of genericity.

Lemma 2.1.7. *Suppose* $\mathbf{P} = (\widetilde{\mathcal{M}}, g, f)$ *is generic. Then* $\mathcal{N} = \widetilde{\mathcal{M}}\backslash f$ *has no loops and* $\widetilde{\mathcal{M}}/g$ *has no coloops.*

Proof. We prove the first statement and the second follows by duality. By our assumption in the definition of an oriented matroid program, g is not a loop of $\mathcal{N}$, so $g \in \underline{Y}$ for some cocircuit Y of $\mathcal{N}$. By Definition 2.1.4(1), $|z(Y)| = d$ and so Remark 2.1.5 implies $z(Y)$ is independent. If there were a loop i of $\mathcal{N}$, then $i \in z(Y)$ contradicting the fact that $z(Y)$ is independent. □

Lemma 2.1.8. *Assume* $\mathbf{P}$ *is generic. Then there is a natural bijection between the set* $\mathbb{B}$ *of bases for* $\underline{\mathcal{M}}$ *and the set of feasible cocircuits for* $\mathcal{N} = \widetilde{\mathcal{M}}\backslash f$.

Proof. Consider the map that takes a feasible cocircuit Y for $\mathcal{N}$ to its zero set $b := z(Y)$. As Y is feasible, $Y(g) = +$ and so by condition (1), Y must have d zero entries. Then Y has $n+1-d$ non-zero entries and is a circuit of $\mathcal{N}^\vee$ (which has rank $n-d$), so any subset of Y of size $n-d$ is a basis for $\underline{\mathcal{N}}^\vee$. In particular, $\underline{Y}\backslash\{g\}$ is a basis for $\underline{\mathcal{N}}^\vee$, so its complement $b \sqcup \{g\}$ is basis for $\underline{\mathcal{N}}$. Thus b is a basis for $\underline{\mathcal{M}} = \underline{\mathcal{N}}/g$.

To show that this is a bijection, suppose b is a basis for $\underline{\mathcal{M}}$. Then $b \sqcup \{g\}$ is a basis for $\underline{\mathcal{N}}$, its complement $E_n\backslash b$ is a basis for $\underline{\mathcal{N}}^\vee$ and so $(E_n\backslash b) \cup g$ must contain a cocircuit Y for $\mathcal{N}$. By condition (1), either $\underline{Y} = (E_n\backslash b) \cup g$ or $\underline{Y} \subset E_n\backslash b$. But the latter is not possible as $E_n\backslash b$ is a basis for $\underline{\mathcal{N}}^\vee$. We conclude that there is a unique choice feasible cocircuit Y with support $\underline{Y} = (E_n\backslash b) \cup g$. □

For $b \in \mathbb{B}$, we let Y_b be denote the corresponding feasible cocircuit.

We will often use three constructions to obtain new generic oriented matroid programs from a generic oriented matroid program $\mathbf{P}$: duality, deletion, and contraction. Recall that duality for oriented matroid programs takes the program $\mathbf{P} = (\widetilde{\mathcal{M}}, g, f)$ to the program

$\mathbf{P}^\vee = (\widetilde{\mathcal{M}}^\vee, f, g)$ with underlying oriented matroid

$$(\widetilde{\mathcal{M}}^\vee \backslash g)/f = ((\widetilde{\mathcal{M}}/g)\backslash f)^\vee = \mathcal{M}^\vee.$$

For any $S \subset E_n$, we denote the contraction and deletion of S by

$$\mathbf{P}/S := (\widetilde{\mathcal{M}}/S, g, f) \quad \text{and} \quad \mathbf{P}\backslash S := (\widetilde{\mathcal{M}}\backslash S, g, f),$$

respectively.

Note that $\mathbf{P}$ is generic if and only of $\mathbf{P}^\vee$ is generic. If $\mathbf{P}$ is generic and $S \subset b$ for some $b \in \mathbb{B}$, then $\mathbf{P}/S$ is generic and has rank $d + 1 - |S|$. If $\mathbf{P}$ is generic and $S \cap b = \varnothing$ for some $b \in \mathbb{B}$, then $\mathbf{P}\backslash S$ is also generic of rank $d + 1$.

Lemma 2.1.9. *For any oriented matroid $\mathcal{M}$ there exists a generic oriented matroid program $\mathbf{P} = (\widetilde{\mathcal{M}}, g, f)$ such that $(\widetilde{\mathcal{M}}/g)\backslash f = \mathcal{M}$.*

Proof. For example, for any order on E_n, consider the lexicographic extension $\mathcal{M}' = \mathcal{M}[E_n]$ by a point f with respect to this order (Note that this is the same as taking the extension $\mathcal{M}[I_{\min}]$ where $I_{\min}$ is a lexicographically minimal basis of $\mathcal{M}$). By a Lemma of Todd [BLVS+99, Lemma 7.2.6], any circuit X of $\mathcal{M}'$ with more than $n - d$ zero entries satisfies $X(f) = 0$. It then remains to define $\widetilde{\mathcal{M}}$ as a single element of lifting of $\mathcal{M}'$ by a point g, such that $\mathcal{N} = \widetilde{\mathcal{M}}\backslash f$ satisfies property (1) above. This can be done via the dual construction: consider the *colocalization* $\tau : C(\mathcal{M}') \to \{+, -, 0\}$ defined for any $X \in C(\mathcal{M}')$ by $\tau(X) = X_i$, where i is the minimal element of E_n such that $X_i \neq 0$. Let $\widetilde{\mathcal{M}}$ be the lexicographic lifting of $\mathcal{M}'$ defined by τ (in other words the dual of the lexicographic extension of $(\mathcal{M}')^\vee$ associated to τ). □

For the rest of the paper we assume that $\mathbf{P}$ is generic.

2.1.4 Bounded feasible topes and sign vectors

In this section we recall the notions of bounded and feasible topes and show in Corollary 2.1.17 that when $\mathbf{P}$ is generic there is a natural bijection between bases $\mathbb{B}$ of M and bounded feasible topes for $\mathbf{P}$.

Definition 2.1.10. (Feasible covectors and affine space) Let $\mathbf{P} = (\widetilde{\mathcal{M}}, g, f)$ be a generic program and let $\mathcal{L} = \mathcal{L}(\mathcal{N})$ denote the set of covectors for $\mathcal{N} = \widetilde{\mathcal{M}} \backslash f$. The *affine space* of $\mathbf{P}$ is:

$$\mathcal{A} = \{Y \in \mathcal{L} \mid Y(g) = +\}.$$

We call elements of $\mathcal{A}$ *feasible* covectors.

We say that the *boundary of affine space* is:

$$\mathcal{A}^\infty = \{Y \in \mathcal{L} \mid Y(g) = 0\}$$

Notice that $\mathcal{A}^\infty$ defines an oriented matroid on $E_n \cup \{g\}$ which is equal to $\mathcal{M}$ with g adjoined as a loop. Also, notice that the join of covectors in $\mathcal{A}$ is also in $\mathcal{A}$ if it exists, while their meet is in $\mathcal{A}$ if and only if they share a common cocircuit face in $\mathcal{A}$; Otherwise, their meet is in $\mathcal{A}^\infty$.

Definition 2.1.11. (Feasible topes) Let $\mathcal{T}(\mathcal{L})$ denote the set of topes of $\mathcal{N}$. We let

$$\mathcal{F} = \mathcal{A} \cap \mathcal{T}(\mathcal{L})$$

denote the set of feasible topes.

Notice that the definition of feasible topes does not depend on f.

Remark 2.1.12. By Lemma 2.1.7 the topes of $\mathcal{N}$ are the covectors T such that $z(T) = \varnothing$.

A *sign vector* is a function $\alpha: E_n \to \{+,-\}$, usually written as $\alpha \in \{+,-\}^n$. There is an obvious injective map from $\mathcal{F}$ to $\{+,-\}^n$ given by forgetting the value at g (which is always +). We may refer to the sign vectors in the image as *feasible* sign vectors, and in a slight abuse of notation we identify feasible topes with the corresponding sign vectors. When there is a risk of confusion, we will write T_α to denote the feasible tope of $\mathcal{N}$ corresponding to a sign vector α.

Definition 2.1.13. (Directions and optimality)

1. We refer to covectors in the boundary of affine space $Z \in \mathcal{A}^\infty$ as *directions* in $\mathcal{A}$. We say that a direction is *increasing* (resp. decreasing or constant) with respect to f if $Z(f) = +$ (resp. $Z(f) = -$ or $Z(f) = 0$).

2. For a feasible tope $T \in \mathcal{F}$ and a feasible covector face Y of T, we say that the direction $Z \in \mathcal{A}^\infty$ is *feasible* for Y in T if $Z \circ Y$ is a face of T.

3. A feasible covector Y that is a face of $T \in \mathcal{F}$ is an *optimal solution* for T if there is no feasible increasing direction for Y in T.

Definition 2.1.14. (Bounded sign vectors) A sign vector $\alpha \in \{+,-\}^n$ is *unbounded* if there exists a increasing direction $Z \in \mathcal{A}^\infty$ such that $Z|_{E_n} \circ \alpha = \alpha$. If no such Z exists, we say α is *bounded*. Similarly, a tope T is bounded if the sign vector $T|_{E_n}$ is bounded.

Let $\mathcal{B}$ denote the set of bounded sign vectors and $\mathcal{P} = \mathcal{F} \cap \mathcal{B}$ denote the set of bounded and feasible sign vectors.

Remark 2.1.15. Let Y be a bounded, feasible tope. Note that if Y_1, Y_2 are optimal solutions for Y, then so is $Y_1 \circ Y_2$. If Y_1 is an optimal solution and $Y_2 \in \mathcal{A}$ is a face of Y_1, then Y_2 is

also optimal. It follows that if Y has an optimal solution, then it has an optimal cocircuit.

Theorem 2.1.16. *Assume* $\mathbf{P}$ *is generic. Then every feasible, bounded tope* Y *has a unique optimal solution (cocircuit) and the resulting map from* $\mathcal{P}$ *to the set of feasible cocircuits for* $\mathcal{N}$ *is a bijection.*

Proof. Let $Y \in \mathcal{P}$. Recall that an optimal solution exists, without any condition on $\mathbf{P}$, by Bland and Lawrence's Main Theorem of Oriented Matroid Programming (see [BLVS+99, 10.1.13] for a survey).

We first show that if $\mathbf{P}$ is generic, then such a solution is unique.

Suppose that Y has two distinct optimal cocircuits Y_1 and Y_2. Then $Y_1 \circ Y_2 \in \mathcal{A}$ must also be an optimal solution. Replacing Y_2 by another cocircuit face of $Y_1 \circ Y_2$ if necessary, without loss of generality, we may assume that Y_1 and Y_2 are joined by an edge (i.e., the rank of $Y_1 \circ Y_2$ is 2). There then exist two cocircuits at the boundary $\pm Z \in \mathcal{A}^\infty$ on the pseudoline $\overline{Y_1 Y_2}$ such that $Y_1 \circ Y_2 = Y_1 \circ Z$ and $Y_2 \circ Y_1 = Y_2 \circ -Z$. (The cocircuit Z can by obtained via elimination of g from the pair Y_1, Y_2 and this elimination is unique up to sign as Y_1, Y_2 form a modular pair). By optimality of Y_1 and Y_2, we conclude that $\pm Z$ must be constant directions. Note that $z(Y_1 \circ Y_2)$ is an independent set in $\underline{\mathcal{M}}$ of cardinality $d-1$. The contraction $\widetilde{\mathcal{M}}/(\{g\} \cup z(Y_1 \circ Y_2))$ is a rank 1 oriented matroid where f is a loop since the cocircuits are both zero on f. Thus $z(Y_1 \circ Y_2) \cup \{f\}$ contains a circuit X of $\widetilde{\mathcal{M}}/g$ such that $X(f) \neq 0$. This contradicts condition (2) since X is zero on at least $n - d + 1$ entries.

It remains to show that the map from $\mathcal{P}$ to feasible cocircuits is a bijection. Given a cocircuit $Y \in \mathcal{A}$, we would like to show that Y is the optimal solution for a unique tope $T \in \mathcal{P}$. We construct such a T as follows. For any $i \in z(Y)$, we know that $\mathbf{P}_i := \mathbf{P}/(z(Y) \backslash \{i\})$

is a generic program whose affine space is one-dimensional. There is a unique cocircuit $\overline{Z}_i$ in $\mathbf{P}_i$ such that $\overline{Z}_i(f) = -$, and this cocircuit is the restriction of a cocircuit Z_i in $\widetilde{\mathcal{M}}/g$. Then $Z_i(j) = 0$ for $j \in z(Y)\backslash\{i\}$, while $Z_i(i) \neq 0$ and $Z_i(f) = -$. Then T is defined to be the composition $Z \circ Y$, where Z is the composition of all Z_i, $i \in z(Y)$, taken in any order.

This T is feasible since $T(g) = Y(g) = +$, and unique since it agrees on $z(Y)$ with the unique bounded feasible tope of $\mathbf{P}\backslash(E_n\backslash z(Y))$. To show that T is bounded, recall that an equivalent definition for a feasible tope to be bounded is that it must be in the *bounded cone* of some $b \in \mathbb{B}$ (see [BLVS+99, Definition 10.1.8.ii, Corollary 10.1.10.ii]), meaning that it agrees on b with a bounded tope of $\mathbf{P}\backslash(E_n\backslash b)$ (see also Definition 2.1.20). Since T is in the bounded cone of $z(Y) \in \mathbb{B}$, we have that T is bounded. □

Combining the bijections of Lemma 2.1.8 and Theorem 2.1.16 we obtain our desired correspondence.

Corollary 2.1.17. *There is a natural bijection between the set $\mathbb{B}$ of bases for $\underline{\mathcal{M}}$ and the set $\mathcal{P}$ of bounded feasible topes:*

$$\mu : \mathbb{B} \to \mathcal{P},$$

which takes a basis b to the tope whose optimal cocircuit is Y_b. (Recall that Y_b is the feasible cocircuit with $z(Y_b) = b$.)

We conclude this section with a discussion of the effect of duality on the bijection μ. Recall the following result about duality for oriented matroid programs.

Proposition 2.1.18. *[BLVS+99, Corollary 10.1.11] Let $\mathcal{F}^\vee$, $\mathcal{B}^\vee$, and $\mathcal{P}^\vee$ respectively denote the sets of feasible, bounded, and bounded feasible sign vectors for the dual program $\mathbf{P}^\vee$ =*

$(\widetilde{\mathcal{M}}^\vee, f, g)$. *Then*

$$\mathcal{F}^\vee = \mathcal{B} \quad \textit{and} \quad \mathcal{B}^\vee = \mathcal{F}, \quad \textit{and so} \quad \mathcal{P}^\vee = \mathcal{P}.$$

Let $\mathbb{B}^\vee$ denote the set of bases of $\underline{\mathcal{M}}^\vee$. Then $b \mapsto b^c := E_n \backslash b$ defines a bijection $\mathbb{B} \to \mathbb{B}^\vee$. Let $\mu^\vee : \mathbb{B}^\vee \to \mathcal{P}^\vee$ be the bijection for the dual program $\mathbf{P}^\vee$ defined as above. Recall that $\mathcal{P}^\vee = \mathcal{P}$.

Proposition 2.1.19 (Complementary Slackness)**.** *For any $b \in \mathbb{B}$, $\mu(b) = \mu^\vee(b^c)$.*

Proof. Recall that $\mathbf{P}$ is generic if and only if $\mathbf{P}^\vee$ is generic. This is then the "Complementary Slackness" theorem of Bland applied to generic programs (cf. [BLVS+99, Theorem 10.1.12]).

□

2.1.5 Cone relation and Euclidean oriented matroid programs

In this section we consider a binary relation on the set $\mathcal{P}$ of bounded feasible topes (or via μ equivalently on the set of bases $\mathbb{B}$).

Definition 2.1.20. For any basis $b \in \mathbb{B}$ we define the negative cone as

$$\mathcal{B}_b = \{\beta \in \{+,-\}^n \mid \mu(b)(i) = \beta(i) \text{ for all } i \in b\}.$$

Notice that this set of sign vectors depends on $\mathcal{M}$ and f but does not depend on g, in the sense that the signs $\mu(b)(i)$ for $i \in b$ only depend on the cocircuits of $\widetilde{\mathcal{M}}/g$.

Proposition 2.1.21 (Complementary Slackness)**.** *For any $b \in \mathbb{B}$, let X_{b^c} be the feasible cocircuit of $\mathbf{P}^\vee$ with $z(X_{b^c}) = b^c = E_n \backslash b$. Then for any $\alpha \in \mathcal{P} = \mathcal{P}^\vee$, $\alpha \in \mathcal{B}_b$ if and only if X_{b^c} is a face of α.*

Proof. By definition, the tope α is in the cone $\mathcal{B}_b$ if and only if

$$\alpha(i) = \mu(b)(i) \quad \text{for all } i \in b.$$

By Proposition 2.1.19, we can rewrite the above condition as

$$\alpha(i) = \mu^\vee(b^c)(i) \quad \text{for all } i \in b.$$

Under the bijection $\mu^\vee$ from Corollary 2.1.17 for the dual program $\mathbf{P}^\vee$, we have that $\mu^\vee(b^c)(i) = X_{b^c}(i)$ for all $i \in b$. Thus the previous condition becomes

$$\alpha(i) = X_{b^c}(i) \quad \text{for all } i \in b,$$

which in turn is equivalent to X_{b^c} is a face of α in the affine space associated to $\mathbf{P}^\vee$. □

Definition 2.1.22 (Cone relation)**.** For $\alpha, \beta \in \mathcal{P}$, we write $\beta \preceq \alpha$ whenever $\beta \in \mathcal{B}_{\mu^{-1}(\alpha)}$.

The binary relation $\preceq$ on $\mathcal{P}$ is reflexive and anti-symmetric, but not necessarily transitive. Let $\leq$ denote the transitive closure of $\preceq$. In general, the binary relation $\leq$ on $\mathcal{P}$ does not define a poset, as the closure may no longer be anti-symmetric.

In the following sections we will define algebras associated to the program $\mathbf{P}$. To ensure that these algebras are quasi-hereditary, we will need the relation $\leq$ to define a poset. It turns out that this is equivalent to a well-known condition on the oriented matroid program $\mathbf{P}$, namely we ask that $\mathbf{P}$ be *Euclidean*. To recall the definition, we first define the following graph associated to $\mathbf{P}$.

Definition 2.1.23. Let $G_{\mathbf{P}}$ be the graph whose vertices are vertices in $\mathcal{A}$ (i.e., the feasible cocircuits for $\mathcal{N}$) and whose edges are the edges in $\mathcal{A}$ (i.e., the feasible covectors of $\mathcal{N}$ of

rank 2). By our genericity condition, the graph is naturally directed by orienting each edge in increasing direction with respect to f.

Definition 2.1.24. For $b, b' \in \mathbb{B}$, we write $b \leq b'$ if there is a directed path from Y_b to $Y_{b'}$ in the graph $G_{\mathbf{P}}$.

Definition 2.1.25. The program $\mathbf{P}$ is *Euclidean* if the directed graph $G_{\mathbf{P}}$ contains no directed cycles. Equivalently, $\mathbf{P}$ is Euclidean if the binary relation on $\mathbb{B}$ from Definition 2.1.24 is anti-symmetric.

By a result of Edmonds and Fukuda, the Euclidean property is well-behaved under duality:

Proposition 2.1.26. *[BLVS+99, Corollary 10.5.9] An oriented matroid program $\mathbf{P} = (\widetilde{\mathcal{M}}, g, f)$ is Euclidean if and only if its dual program $\mathbf{P}^\vee = (\widetilde{\mathcal{M}}^\vee, f, g)$ is Euclidean.*

Importantly for us, for Euclidean programs, the transitive closure of the cone relations is a poset. In fact these two conditions are equivalent:

Lemma 2.1.27. *The oriented matroid program $\mathbf{P} = (\widetilde{\mathcal{M}}, g, f)$ is Euclidean if and only if the transitive closure of the cone relation is anti-symmetric.*

Proof. By Proposition 2.1.26, it suffices to show that the dual program $\mathbf{P}^\vee$ is Euclidean if and only if the transitive closure on the cone relation is anti-symmetric. By definition, $\mathbf{P}^\vee$ is Euclidean if and only if the binary relation $(\mathbb{B}^\vee, \leq)$ is anti-symmetric.[1]

Thus it suffices to show that the bijection $\mu^\vee : \mathbb{B}^\vee \to \mathcal{P}^\vee = \mathcal{P}$ is order reversing. In other words, we wish to show that for $b_1, b_2 \in \mathbb{B}$,

$$\mu(b_1) \geq \mu(b_2) \quad \text{if and only if} \quad b_1^c \leq b_2^c.$$

[1] Here by $(\mathbb{B}^\vee, \leq)$ we are referring to the binary relation on $\mathbb{B}^\vee$ coming from the directed graph $G_{\mathbf{P}^\vee}$.

As $(\mathcal{P}, \leq)$ is the transitive closure of $\preceq$, without loss of generality we may suppose $\mu(b_1) \geq \mu(b_2)$. Then $\mu(b_2) \in \mathcal{B}_{b_1}$ and by Proposition 2.1.21, $X_{b_1^c}$ is a face of $\mu(b_2) = \mu^\vee(b_2^c)$. As $X_{b_2^c}$ is the (unique) optimal cocircuit of the tope $\mu^\vee(b_2^c)$, there is a directed path from $X_{b_2^c}$ to $X_{b_1^c}$. Thus $b_1^c \leq b_2^c$ as desired.

For the other direction, it suffices to consider the case where there is a directed edge from $X_{b_2^c}$ to $X_{b_1^c}$. Then $X_{b_1^c}$ is a cocircuit face of $\mu^\vee(b_2^c) = \mu(b_2)$. By Proposition 2.1.21 this implies $\mu(b_2) \in \mathcal{B}_{b_1}$ and so $\mu(b_2) \leq \mu(b_1)$, which completes the proof. □

2.1.6 On the existence of Euclidean generic programs

Unlike genericity, it is not clear (at least to the authors) that any oriented matroid $\mathcal{M}$ can be extended to a Euclidean generic oriented matroid program $\widetilde{\mathcal{M}}$.

On the other hand, we have already seen in Example 1.2.4 a non-realizable Euclidean program $(\widetilde{\mathcal{M}}, g, f)$ for which $\mathcal{M} = \widetilde{\mathcal{M}}/g\backslash f$ is realizable. There also appear to be many examples for which the oriented matroid $\mathcal{M}$ is non-realizable, and so our setting significantly generalizes that of [BLPW10].

In this section we first give such an example and then give a criterion on $\mathcal{M}$ for the existence of a Euclidean generic oriented matroid program lifting and extending $\mathcal{M}$.

Example 2.1.28. The Vámos matroid $\mathcal{V}$ (rank 4 on 8 elements) is not representable over any field, but is orientable by [BV78]. We now define a Euclidean generic program $\widetilde{\mathcal{V}}$ by adjoining elements to $\mathcal{V}$ as labelled and orientated in [BV78]. Let $(\widetilde{\mathcal{V}}, g, f)$ be the generic oriented matroid program obtained from $\mathcal{V}$ by adjoining g using the lexicographic one-element lift defined by the cobasis $\{3, 6, 7, 8\}$, and adjoining f by a lexicographic one-

element extension defined by the basis $\{1, 2, 4, 5\}$. One can verify that the graph $G_{\mathbf{P}}$ of this program has no directed cycles, so this program is Euclidean and generic.

We are grateful to Jim Lawrence for suggesting that we consider the following statement.

Proposition 2.1.29. *Let $\mathcal{M}$ be an oriented matroid. If one can adjoin elements g and f to $\mathcal{M}$ to obtain a Euclidean generic oriented matroid program $(\widetilde{\mathcal{M}}, g, f)$, then $\mathcal{M}$ has a simplicial tope.*

Before beginning the proof, we note that it is an open conjecture of Las Vergnas [LV80] (see also [BLVS+99, 7.3.10]) that any oriented matroid $\mathcal{M}$ should have a simplicial tope. Thus, the proposition implies that if for any oriented matroid $\mathcal{M}$ there is a Euclidean generic oriented matroid program $(\widetilde{\mathcal{M}}, g, f)$ for which $\widetilde{\mathcal{M}}/g\backslash f = \mathcal{M}$, then Las Vergnas' conjecture holds.

Proof. Suppose $(\widetilde{\mathcal{M}}, g, f)$ is a Euclidean generic oriented matroid program such that $(\widetilde{\mathcal{M}}/g)\backslash f = \mathcal{M}$.

As $\mathbf{P}$ is Euclidean, there exists a minimal vertex of $G_{\mathbf{P}}$. Let $Y_{\min}$ be the feasible cocircuit of $\mathcal{N}$ corresponding to such a minimal vertex. Let $T_{\min}$ be the tope whose optimal cocircuit is $Y_{\min}$, which exists by Theorem 2.1.16. Then $Y_{\min}$ is the only feasible cocircuit face of $T_{\min}$, so for all other cocircuit faces Y, $Y(g) = 0$. Consider the subtopes covered by $T_{\min}$. There are d such subtopes that have $Y_{\min}$ as a face, all of which are feasible. Any other subtope Z covered by $T_{\min}$ cannot be feasible and as all non-feasible cocircuit faces Y of T satisfy $Y(g) = 0$, then $z(Z) = g$. There can be at most one such subtope covered by $T_{\min}$. By [BLVS+99, Exercise 4.4], any tope of $\widetilde{\mathcal{M}}$ covers at least $d + 1$ subtopes, so there

does exist a subtope Z and $T_{\min}$ covers exactly $d+1$ subtopes. By the same exercise, it follows that $T_{\min}$ is simplicial, hence the subtope Z is as well, which is in turn a tope of $\mathcal{M}$. □

2.2 Face rings and matroids

Recall that $M = \underline{\mathcal{M}}$ denotes the underlying matroid of $\mathcal{M}$. Let $\mathcal{I}$ be the set of independent subsets of M. Then $\mathcal{I}$ is a simplicial complex under inclusion, which we refer to as the *matroid complex* of M.

Fix a field $\mathbf{k}$. Let $\mathbf{k}^n$ be the standard n-dimensional vector space with basis $t_1, \ldots, t_n$.

Definition 2.2.1. For any field $\mathbf{k}$, the *face ring* of the matroid complex Δ of M is defined as

$$\begin{aligned}\mathbf{k}[M] &:= \mathbf{k}[t_1, \ldots, t_n]/(t_S \mid S \notin \Delta) \\ &= \mathbf{k}[t_1, \ldots, t_n]/(t_{\underline{X}} \mid X \in C)\end{aligned}$$

where $t_S := \prod_{i \in S} t_i$ for any $S \subset E_n$. We give $\mathbf{k}[M]$ a $\mathbb{Z}_{\geq 0}$-grading by setting $\deg(t_i) = 2$ for all i.

Definition 2.2.2. Recall d denotes the rank of M. A *linear system of parameters* (l.s.o.p.) for $\mathbf{k}[M]$ is a set

$$\Theta = \{\theta_1, \theta_2, \ldots, \theta_d\} \subset \mathbf{k}^n$$

such that $\mathbf{k}[M]$ is a finitely-generated $\mathbf{k}[\theta_1, \theta_2, \ldots, \theta_d]$-module. Equivalently $\mathbf{k}[M]/(\overline{\Theta})$ is a zero-dimensional ring, where $\overline{\Theta}$ is the image of Θ in $\mathbf{k}[M]_2$.

Remark 2.2.3. Stanley [Sta04] defines an l.s.o.p. as a subset of $\mathbf{k}[M]_2 \subset \mathbf{k}[M]$. The set Θ is a l.s.o.p. in the sense we define above if its image in $\mathbf{k}[M]_2$ is an l.s.o.p. in the sense used by Stanley.

We introduce the following perhaps non-standard definition:

Definition 2.2.4. We call a subspace $U \subset \mathbf{k}^n$ a *parameter space* for $\mathbf{k}[M]$ if the composition $U \hookrightarrow \mathbf{k}^n \twoheadrightarrow \mathrm{Span}\{t_i \mid i \in b\}$ is an isomorphism for any basis $b \in \mathbb{B}$.

Example 2.2.5. If M is realizable as a hyperplane arrangement coming from a k-vector subspace $V \subset \mathbf{k}^n$, we can choose U to be V.

Lemma 2.2.6. *If $U \subset \mathbf{k}^n$ is a parameter space for $\mathbf{k}[M]$, then any basis of U is a l.s.o.p. for $\mathbf{k}[M]$. If Θ is an l.s.o.p. for $\mathbf{k}[M]$, then its span Span$(\Theta) \subset \mathbf{k}^n$ is a parameter space for $\mathbf{k}[M]$.*

Proof. Suppose U is a parameter space. By definition $\dim U = \dim \mathrm{Span}\{t_i \mid i \in b\} = d$. Suppose $\theta_1, \ldots, \theta_d$ is a basis of U. By [Sta04, Lemma 2.4(a)], $\theta_1, \ldots, \theta_d$ is an l.s.o.p. for $\mathbf{k}[M]$ if and only if for every facet of Δ, that is basis $b \in \mathbb{B}$, the list $\theta_1, \ldots, \theta_d \in \mathbf{k}^n$ projects to a spanning set of $\mathrm{Span}\{t_i \mid i \in b\}$. This is true by the definition of U.

Similarly, if $\theta_1, \ldots, \theta_d \in \mathbf{k}^n$ is an l.s.o.p., then for any basis $b \in \mathbb{B}$, the projections to $\mathrm{Span}\{t_i \mid i \in b\}$ of $\theta_1, \ldots, \theta_d$ are a spanning set. Thus the projection from $\mathrm{Span}(\Theta)$ to $\mathrm{Span}\{t_i \mid i \in b\}$ is an isomorphism and $\mathrm{Span}(\Theta)$ is a parameter space. □

By the Noether Normalization Lemma, if $\mathbf{k}$ is an infinite field, then a l.s.o.p. for $\mathbf{k}[M]$ exists. From now on, we assume that an l.s.o.p. exists and fix a choice of l.s.o.p. Θ and its span U.

2.2.1 Linear systems of parameters and duality

Let $u_1, \ldots, u_n$ be the basis of $(\mathbf{k}^n)^*$ dual to $t_1, \ldots, t_n$. It will be convenient for us to view the matroid complex $\mathbf{k}[M^\vee]$ of the dual matroid $M^\vee$ as the appropriate quotient of $\mathbf{k}[u_1, \ldots, u_n]$.

Let $U^\perp$ be the kernel of the natural map $(\mathbf{k}^n)^* \to U^*$.

Lemma 2.2.7. *$U^\perp$ is a parameter space for the face ring of the dual matroid $\mathbf{k}[M^\vee]$.*

Proof. For any $b \in \mathbb{B}$ a basis for M, let $b^c := E_n \backslash b$ be its complement, which is a basis for $M^\vee$. It suffices to show that the projection $U^\perp \to \operatorname{Span}\{u_i \mid i \in b^c\}$ is an isomorphism. As the two vector spaces have the same dimension, it suffices to show the null space is trivial. Suppose x is in the null space, so

$$x \in U^\perp \cap \operatorname{Span}\{u_i \mid i \in b\}.$$

Then for any $u \in U$,

$$0 = \langle x, u \rangle = \langle x, \operatorname{pr}_b u \rangle,$$

where $\operatorname{pr}_b : U \to \operatorname{Span}\{t_i \mid i \in b\}$ is the projection. As U is a parameter space for $\mathbf{k}[M]$, the projection pr_b is an isomorphism and so $\langle x, w \rangle = 0$ for any $w \in \operatorname{Span}\{t_i \mid i \in b\}$. Thus $x = 0$ and we may conclude that the projection $U^\perp \to \operatorname{Span}\{t_i \mid i \in b^c\}$ is an isomorphism. □

We define the dual of the pair $(\mathbf{P}, U)$ to be $(\mathbf{P}, U)^* = (\mathbf{P}^\vee, U^\perp)$.

Proposition 2.2.8. *$\mathbf{k}[M]$ is a free algebra over $\mathbf{k}[\theta_1, \ldots, \theta_d] = \operatorname{Sym} U \subset \mathbf{k}[t_1, \ldots, t_n]$, while $\mathbf{k}[M^\vee]$ is a free algebra over $\operatorname{Sym} U^\perp \subset \mathbf{k}[u_1, \ldots, u_n]$. Both have rank $|\mathbb{B}|$.*

Proof. Recall the face ring of a matroid complex is shellable (e.g., [Sta04, Proposition III.3.1]) and so the result follows [Sta04, Theorem III.2.5]. □

Chapter 3

The Algebras

3.1 The algebra A

3.1.1 The definition of A

Recall that $\mathbf{P} = (\widetilde{\mathcal{M}}, g, f)$ is a generic oriented matroid program and $U \subset \mathbf{k}^n$ is a parameter space for $M = \underline{\mathcal{M}}$.

Let Q denote the quiver with vertex set $\mathcal{F}$, the set of feasible topes, and arrows between topes that differ by exactly one sign. We say that two topes α, β that differ by exactly one sign are *adjacent* and write $\alpha \leftrightarrow \beta$. If $\alpha \leftrightarrow \beta$ and differ in the i-th component, we write $\beta = \alpha^i$.

Let $P(Q)$ denote the path algebra for Q, which is generated by orthogonal idempotents $\{e_\alpha \mid \alpha \in \mathcal{F}\}$ and edge paths $\{p(\alpha, \beta)\}$ where α and β are adjacent and $p(\alpha, \beta)$ is the path from α to β. We write $p(\alpha_1, \ldots, \alpha_k)$ for the element in the quiver algebra obtained as the composition $p(\alpha_1, \alpha_2) \cdot \ldots \cdot p(\alpha_{k-1}, \alpha_k)$.

Definition 3.1.1. Let $\widetilde{A} = \widetilde{A}(\mathbf{P}) = \widetilde{A}(\widetilde{\mathcal{M}}, g, f)$ be the quotient of $P(Q) \otimes_{\mathbf{k}} \mathrm{Sym}(\mathbf{k}^n)^* = P(Q) \otimes_{\mathbf{k}} \mathbf{k}[u_1, \ldots, u_n]$ by the two-sided ideal generated by the following relations:

$(A1)$ $e_\alpha = 0$ for $\alpha \notin \mathcal{P}$,

$(A2)$ $p(\alpha, \gamma, \beta) = p(\alpha, \delta, \beta)$ for any four distinct topes $\alpha, \beta, \gamma, \delta \in \mathcal{F}$ where α and β are each connected to γ and δ by an edge, and

$(A3)$ $p(\alpha, \alpha^i, \alpha) = e_\alpha u_i$ whenever $\alpha, \alpha^i \in \mathcal{F}$ differ only in the sign of $i \in E_n$.

We let $A = A(\mathbf{P}, U) = A(\widetilde{\mathcal{M}}, g, f, U)$ be

$$A := \widetilde{A} \otimes_{\mathrm{Sym}(\mathbf{k}^n)^*} \mathrm{Sym}((\mathbf{k}^n)^*/(U^\perp)) = \widetilde{A} \otimes_{\mathrm{Sym}(\mathbf{k}^n)^*} \mathrm{Sym}\, U^*,$$

or equivalently, the quotient of $\widetilde{A}$ by the additional relations

$(A4)$ $x = 0$ for any $x \in U^\perp \subset \mathrm{Sym}(\mathbf{k}^n)^*$.

Remark 3.1.2. When the pair $(\mathbf{P}, U)$ comes from a polarized arrangement as in Examples 1.2.1 and 1.2.2, there is an equality $A(\mathbf{P}, U) = A(V, \eta, \xi)$.

As in 2.2.1, for bookkeeping, we will use the dual coordinates for the dual matroid program, so we view $\widetilde{A}(\mathbf{P}^\vee)$ and $A(\mathbf{P}^\vee, U^\perp)$ as the analogous quotients of $P(Q) \otimes_{\mathbf{k}} \mathrm{Sym}(\mathbf{k}^n) = P(Q) \otimes_{\mathbf{k}} \mathbf{k}[t_1, \ldots, t_n]$.

3.1.2 Expressions for elements of A

We first introduce some terminology for paths in the quiver Q. To distinguish between paths in Q and their images in A, we will use the notation $\alpha_0 \to \alpha_1 \to \cdots \to \alpha_r$ for a path of length r in the quiver Q.

A path $\alpha_0 \to \cdots \to \alpha_r$ in Q is *taut* if α_0 and α_r differ in exactly r coordinates. To relate paths, we use the following notion.

Definition 3.1.3. Two paths P, P' in Q are related by an *elementary homotopy* (a symmetric relation) if

(i) P is the path $\alpha_0 \to \cdots \to \alpha_j \to \alpha_{j+1} \to \cdots \to \alpha_r$ of length r while P' is the path $\alpha_0 \to \cdots \to \alpha_j \to \beta \to \alpha_j \to \alpha_{j+1} \to \cdots \to \alpha_r$ of length $r+2$ for some β adjacent to α_j, or

(ii) P is the path $\alpha_0 \to \cdots \to \alpha_{j-1} \to \alpha_j \to \alpha_{j+1} \to \cdots \to \alpha_r$ while P' is the path $\alpha_0 \to \cdots \to \alpha_{j-1} \to \alpha_j' \to \alpha_{j+1} \to \cdots \to \alpha_r$ of the same length, where we assume α_{j-1} and α_{j+1} differ in exactly 2 coordinates.

Remark 3.1.4. Under our assumption that $(\widetilde{\mathcal{M}}, g, f)$ is generic, any feasible covector for $\mathcal{N} = \widetilde{\mathcal{M}} \backslash f$ of rank $d-1$ has exactly two zero coordinates. Thus, in this setting, this definition of elementary homotopy coincides with that given in [BLVS+99, Section 4.4, page 184] for paths in the tope graph, because every elementary homotopy of type (ii) in the sense of [BLVS+99] between feasible paths is of the form above.

We will use the following result:

Proposition 3.1.5. *Let P and P' be any two taut paths in Q with the same start and end points. Then P and P' are related by a sequence of elementary homotopies of type (ii) such that every intermediate path is also taut.*

Proof. If we consider instead paths in the entire tope graph and the more general notion of elementary homotopy, this is a result of Cordovil-Moreira [CM93] (see also [BLVS+99, Proposition 4.4.7]). Recall that a subset $\mathcal{R} \subset \mathcal{T}$ is T-convex if it contains every shortest path

between any two of its members and the set of feasible topes is T-convex (see [BLVS+99, Definition 4.2.5] and the discussion that follows). The result then follows for paths in Q. □

Proposition 3.1.6. *Given a path $P = (\alpha_0 \to \cdots \to \alpha_s)$ in Q, let d_i be the number of times the i-th coordinate changes twice. Then for any taut path $P' = (\alpha_0 = \beta_0 \to \cdots \to \beta_r = \alpha_s)$, we have*

$$p(\alpha_0, \ldots, \alpha_s) = p(\beta_0, \cdots, \beta_r) \cdot \prod_{i \in E_n} u_i^{d_i}$$

in $\widetilde{A}$.

Proof. Note that $s - r \geq 0$ with equality if and only if P is a taut path. We prove the proposition by induction on $s - r$. If $s = r$, then both paths are taut, and so by Proposition 3.1.5 they are related by a sequence of elementary homotopies of type (ii). But an elementary homotopy of type (ii) descends to an equality in the path algebra by definition, so

$$p(\alpha_0, \ldots, \alpha_r) = p(\beta_0, \ldots, \beta_r),$$

as desired.

Assume that the statement holds whenever $s - r < k$ for some positive integer k. Suppose that $s - r = k$. There exists a minimal ℓ such that $\alpha_0 \to \cdots \to \alpha_\ell$ is taut, while $\alpha_0 \to \cdots \to \alpha_{\ell+1}$ is not taut. Then α_0 and $\alpha_{\ell+1}$ differ in $\ell - 1$ coordinates and for some i, $\alpha_0(i) = \alpha_{\ell+1}(i) \neq \alpha_\ell(i)$.

Notice that any taut path between α_0 and α_ℓ will have length exactly one more than the length of a taut path from α_0 to $\alpha_{\ell+1}$. Therefore, by Proposition 3.1.5, using a sequence of elementary homotopies of type (ii) we can replace $\alpha_0 \to \cdots \to \alpha_{\ell-1} \to \alpha_\ell$ with a taut path $\alpha_0 \to \alpha'_1 \to \cdots \to \alpha'_{\ell-2} \to \alpha_{\ell+1} \to \alpha_\ell$.

This gives an equality in the algebra A:

$$\begin{aligned} p(\alpha_0, \alpha_1, \ldots, \alpha_{\ell-2}, \alpha_{\ell-1}, \alpha_\ell, \alpha_{\ell+1}) &= p(\alpha_0, \alpha_1', \ldots, \alpha_{\ell-2}', \alpha_{\ell+1}, \alpha_\ell, \alpha_{\ell+1}) \\ &= p(\alpha_0, \alpha_1', \ldots, \alpha_{\ell-2}', \alpha_{\ell+1}) u_i. \end{aligned}$$

We are then reduced to considering the path

$$\alpha_0 \to \alpha_1' \to \cdots \to \alpha_{\ell-2}' \to \alpha_{\ell+1} \to \alpha_{\ell+2} \to \cdots \to \alpha_s$$

of length $s-1$ and the number of times the i-th coordinate changes twice is $d_i - 1$, while the number of times every j-th coordinate changes twice remains d_j for all $j \neq i$. We can then invoke the induction hypothesis to complete the proof. □

The following two corollaries are analogous to [BLPW10, Corollary 3.10].

Corollary 3.1.7. *Consider an element*

$$a = p \cdot \prod_{i \in E_n} u_i^{d_i} \in \widetilde{A},$$

where p is a taut path in Q from α to β. Suppose γ denotes a feasible tope such that if $\alpha(i) = \beta(i)$ and $d_i = 0$, then $\gamma(i) = \alpha(i) = \beta(i)$. Then

$$a = a' \cdot m,$$

where a' is the concatenation of a taut path from α to γ with a taut path from γ to β and m is a product of u_i's.

In particular, if γ is not bounded, then $a = 0$ in both $\widetilde{A}$ and A.

Proof. For all $j \in E_n$, either:

1. $\alpha(i) \neq \beta(j)$ and the j-th coordinate changes in the concatenation a' exactly once.

2. $\gamma(j) = \alpha(j) = \beta(j)$ and the j-th coordinate does not change in the concatenation a'.

3. $\gamma(j) \neq \alpha(j) = \beta(j)$ and so the j-th coordinate changes exactly twice in the concatenation a'.

Proposition 3.1.6 then says:

$$a' = p \cdot \prod_{i \in E_n} u_i^{d_i'},$$

where $d_i' \in \{0, 1\}$ and by our assumption on γ, we have $d_i \geq d_i'$. Thus

$$a = a' \cdot \prod_{i \in E_n} u_i^{d_i - d_i'},$$

as desired. □

Corollary 3.1.8. *Let b be the zero set of any feasible cocircuit face of $\alpha \in \mathcal{P}$. For any $j \in E_n$, $e_\alpha u_j \in A$ can be written as a $\mathbf{k}$-linear combination of paths $\{p(\alpha, \alpha^i, \alpha) \mid i \in b\}$.*

In particular, the image in A of the element $a \in \widetilde{A}$ described in Corollary 3.1.7 can be expressed a linear combination of paths in Q that pass through γ.

Proof. For $j \in b$, the tope α^j is feasible, so $e_\alpha u_j = p(\alpha, \alpha^j, \alpha)$. On the other hand, as U is a parameter space for $\mathbf{k}[M]$, the set $\{u_i | i \in b\}$ restricts to a basis of U^*. Thus for any $j \notin b$, $u_j \in A$ can be expressed as a linear combination of $\{u_i \mid i \in b\}$. □

3.1.3 Alternative description of A

We conclude this section with a slightly different description of our algebra A, which will make it easier to describe how A changes when we modify the choice of generic oriented matroid program.

Let D be the path algebra over $\mathbf{k}$ of the quiver with two vertices labelled by $+$ and $-$ and an arrow in each direction. Let $D_n = D \otimes D \otimes \ldots \otimes D$ denote the n-fold tensor product of D with itself. In particular, D_n is the path algebra on the quiver with vertices labelled by the set $\{+,-\}^n$ of sign vectors, or equivalently vertices of an n-cube, and edges connecting any two sign vectors that differ in exactly one position, modulo the relations that whenever $\alpha, \beta \in \{+,-\}^n$ differ in exactly two positions i and j, we have an equivalence of paths in D_n:

$$p(\alpha, \alpha^i, \beta) = p(\alpha, \alpha^j, \beta).$$

(As before, $\alpha^i \in \{+,-\}^n$ denotes the sign vector that differs from α in exactly the i-th coordinate.)

For any sign vector α, we again let e_α denote the idempotent defined as the trivial path at the vertex labelled by α. Let $e_{\mathcal{P}} = \sum_{\alpha \in \mathcal{P}} e_\alpha$ be the sum of idempotents corresponding to bounded, feasible topes and $e_f = \sum_{\alpha \notin \mathcal{B}} e_\alpha$ be the sum of idempotents corresponding to unbounded sign vectors.

For each $i \in E_n$, we consider the element $\theta_i \in D_n$ defined as the sum

$$\theta_i = \sum_{\alpha \in \{+,-\}^n} p(\alpha, \alpha^i, \alpha).$$

Note that the center $Z(D_n)$ of D_n is a polynomial algebra with generators θ_i. Let $\vartheta : (\mathbf{k}^n)^* \to \mathbf{k}\{\theta_i \mid i \in E_n\}$ be the isomorphism sending u_i to θ_i.

Lemma 3.1.9. *The algebra $\widetilde{A}$ from Definition 3.1.1 is isomorphic to the quotient of $e_{\mathcal{P}} D_n e_{\mathcal{P}}$ by the relation:*

$(A1')$ $e_\alpha = 0$ *for* $\alpha \notin \mathcal{B}$,

and the algebra A is obtained by adding the additional relation:

$(A4')$ $\vartheta(x) = 0$ *for* $x \in U^\perp$.

Equivalently, there are isomorphisms

$$\widetilde{A} \cong e_{\mathcal{P}} D_n e_{\mathcal{P}} / \langle e_f e_{\mathcal{P}} \rangle, \qquad A \cong e_{\mathcal{P}} D_n e_{\mathcal{P}} / \langle e_f e_{\mathcal{P}} \rangle + \langle \vartheta(U^\perp) e_{\mathcal{P}} \rangle.$$

Proof. To distinguish between the two definitions, let $\widetilde{A}_1$ and A_1 denote the original algebras and $\widetilde{A}_2$ and A_2 be the algebras defined as in the lemma. Note that there is an injective homomorphism from $\widetilde{A}_1$ to $\widetilde{A}_2$ uniquely defined by sending a path in Q to the corresponding path in the cube quiver $\{+,-\}^n$ and sending u_i to θ_i. To see that it is surjective, it suffices to observe that for any two topes $\alpha, \beta \in \mathcal{P}$, there exists a taut path in Q from α to β, which follows from [Cor82, Lemme 3.7]. □

3.2 The quadratic dual of A

In this section we observe that $A(\mathbf{P}, U)$ is a quadratic algebra and that its quadratic dual is isomorphic to $A(\mathbf{P}^\vee, U^\perp)$. To ease notational clutter, in this section we will write A for $A(\mathbf{P}, U)$ and $A^\vee$ for $A(\mathbf{P}^\vee, U^\perp)$.

To state and prove the results of this section, we will need some additional notation. Let $Q_{\mathcal{P}} \subset Q$ be the full subquiver with vertices $\mathcal{P} \subset \mathcal{F}$. For $\alpha \in \mathcal{P}$, we let

$$J_\alpha := \{i \mid \alpha^i \in \mathcal{P}\},$$

$$I_\alpha := \{i \mid \alpha^i \in \mathcal{F}\},$$

$$K_\alpha := \{i \mid \alpha^i \in \mathcal{B}\}.$$

Lemma 3.2.1. *The algebra $A = A(\mathbf{P}, U)$ is the quotient of the path algebra $P(Q_{\mathcal{P}})$ by the following relations:*

$(A2'')$ *For distinct topes $\alpha, \beta, \gamma \in \mathcal{P}$, $\delta \in \mathcal{F}$, where α and β are each connected to γ and δ by an edge,*

$$p(\alpha, \gamma, \beta) = \begin{cases} p(\alpha, \delta, \beta) & \text{if } \delta \in \mathcal{P} \\ 0 & \text{otherwise} \end{cases}$$

$(A3/4'')$ *For any $\alpha \in \mathcal{P}$ and $w \in U^{\perp} \cap \mathrm{Span}\{u_i \mid i \in I_\alpha\}$, if $w = \sum_{i \in I_\alpha} w_i u_i$ for some $w_i \in k$, then*

$$\sum_{i \in J_\alpha} w_i p(\alpha, \alpha^i, \alpha) = 0.$$

In particular, it follows that A is a quadratic algebra.

Proof. Note that there is a surjection from $P(Q)$ to $P(Q_{\mathcal{P}})$ (by setting $e_\alpha = 0$ for all $\alpha \notin \mathcal{P}$) and that this map factors through the natural map from $P(Q)$ to A.

We first show that the map from $P(Q)$ to A (and hence from $P(Q_{\mathcal{P}})$ to A) is surjective. To do so, it suffices to show that the image in A of any element of $w \in U^*$ is in the image of the map from $P(Q)$ to A. In A, we can express w as the sum $w = \sum_{\alpha \in \mathcal{F}} we_\alpha$, so it is enough to show that for any $\alpha \in \mathcal{F}$, we_α is in the image.

Because α is feasible, there exists a feasible cocircuit face Y of α and $z(Y)$ is a basis of M. Thus $z(Y) \subset I_\alpha$ and so by our assumption that U is a parameter space, the image of $\{u_i | i \in I_\alpha\} \supset \{u_i | i \in z(Y)\}$ is a spanning set of U^*. For any $w \in U^*$, we can therefore write we_α as a linear combination of elements of the form $p(\alpha, \alpha^i, \alpha)$ where $\alpha^i \in \mathcal{F}$. We conclude that map from $P(Q)$ to A is surjective.

To identify the kernel, note that $(A2'')$ is simply the image of the relation $(A2)$ in $P(Q_{\mathcal{P}})$ and similarly the relations $(A3)$ and $(A4)$ combine to give $(A3/4'')$.

As $P(Q_{\mathcal{P}})$ is generated over its degree zero component by its degree one component and the relations above are quadratic, we conclude that A is a quadratic algebra. $\square$

Recall that A_1 is an A_0-bimodule. Since A is quadratic, A is a quotient of the tensor algebra $T(A_1)$ over A_0 of the form

$$A = T(A_1)/T(A_1) \cdot W \cdot T(A_1)$$

where $W \subseteq N := (T(A_1) \otimes_{A_0} T(A_1))_2$ is the space of quadratic relations in A. The *quadratic dual* of A is defined to be

$$A^! = T(A_1^*)/T(A_1^*) \cdot W^\perp \cdot T(A_1^*)$$

where $W^\perp \subseteq N^* := (T(A_1^*) \otimes_{A_0} T(A_1^*))_2$ is the set of elements orthogonal to W.

Theorem 3.2.2. *There is an isomorphism* $A(\mathbf{P}^\vee, U^\perp) \simeq A(\mathbf{P}, U)^!$.

Proof. Mimicking the proof of [BLPW10, Theorem 3.11], we define an isomorphism between $(A^\vee)_1$ and $(A_1)^*$ and show that space $W^\vee$ of quadratic relations in $A^\vee$ coincides with the space $W^\perp$ of quadratic relations in $A^!$ under this identification.

In degree zero, we have a canonical identification

$$A_0 = \mathbf{k}\{e_\alpha \mid \alpha \in \mathcal{P} = \mathcal{P}^\vee\} = (A^\vee)_0.$$

In degree one,

$$\{p(\alpha, \beta) \mid \alpha, \beta \in \mathcal{P} \text{ such that } \alpha \leftrightarrow \beta\} \subset A_1$$

is a natural basis for A_1.

As $\mathcal{P} = \mathcal{P}^\vee$, to distinguish elements of A and $A^\vee$, we let $p^\vee(\alpha, \beta)$ denote the element of $(A^\vee)_1$ associated to the arrow $\alpha \to \beta$ in $\mathcal{P}^\vee$.

We now identify $(A^\vee)_1$ with $(A_1)^*$ as follows. First we attach a sign $\varepsilon(\alpha \leftrightarrow \beta)$ to each pair $\alpha \leftrightarrow \beta$ of adjacent topes in $\mathcal{P}$ such that for distinct topes $\alpha, \alpha^i, \alpha^j, (\alpha^i)^j$, an odd number of the edges of the square

$$\begin{array}{ccc} \alpha & \longleftrightarrow & \alpha^i \\ \updownarrow & & \updownarrow \\ \alpha^j & \longleftrightarrow & (\alpha^i)^j \end{array}$$

are attached a negative sign.[1]

We then identify $(A^\vee)_1$ with $(A_1)^*$ via the perfect pairing

$$(A^\vee)_1 \times A_1 \to \mathbf{k}$$

$$\langle p^\vee(\alpha, \beta), p(\gamma, \delta) \rangle = \begin{cases} \varepsilon(\alpha \leftrightarrow \beta) & \text{if } \alpha = \delta, \beta = \gamma, \\ 0 & \text{otherwise.} \end{cases}$$

For the remainder of the proof, we let

$$N := A_1 \otimes_{A_0} A_1 = T(A_1)_2,$$

$$N^\vee := (A^\vee)_1 \otimes_{A_0} (A^\vee)_1 = T((A^\vee)_1)_2.$$

We wish to show that our choice of perfect pairing induces an isomorphism between $W^\vee \subseteq N^\vee$ and $W^\perp \subseteq N^*$.

Note that the relations of type $(A2'')$ lie in $e_\alpha W e_\beta$ and the relations of type $(A3/4'')$ lie in $e_\alpha W e_\alpha$. As these relations are homogeneous with respect to the idempotents e_α for

[1] This can be done for the edges of the n-cube $\{+,-\}^n$ by identifying its vertices with monomials in the exterior algebra $\Lambda k\{e_1, \ldots, e_n\}$ and then using the standard differential to attach signs to edges. Restricting to $\mathcal{P}$ then gives a collection of signs as desired.

$\alpha \in \mathcal{P}$, we have a direct sum decomposition

$$W = \bigoplus_{\alpha,\beta \in \mathcal{P}} e_\alpha W e_\beta.$$

Moreover, $e_\alpha W e_\beta$ and $e_\gamma W^\vee e_\delta$ are orthogonal unless $\alpha = \delta$ and $\beta = \gamma$. Thus it is enough to check that $e_\beta W^\vee e_\alpha \subset e_\beta N^\vee e_\alpha$ is the perpendicular complement to $e_\alpha W e_\beta \subset e_\alpha N e_\beta$ for any $\alpha, \beta \in \mathcal{P}$.

Note that $e_\alpha N e_\beta, e_\beta N^\vee e_\alpha$ are zero unless $\alpha = \beta$ or α and β differ in exactly two positions and there is a path from α to β in $\mathcal{P}$.

We first deal with the latter case. Suppose the two elements of E_n where α and β differ are $i \neq j$. We must have that at least one of α^i or α^j is in $\mathcal{P}$ by assumption. We can assume that $a^j \in \mathcal{P}$.

If $a^i \in \mathcal{P}$, then we have that $e_\alpha N e_\beta$ is a two dimensional $\mathbf{k}$-vector space with basis $\{p(\alpha, \alpha^i) \otimes p(\alpha^i, \beta), p(\alpha, \alpha^j) \otimes p(\alpha^j, \beta)\}$ and

$$e_\alpha W e_\beta = k\{p(\alpha, \alpha^i) \otimes p(\alpha^i, \beta) - p(\alpha, \alpha^j) \otimes p(\alpha^j, \beta)\} \subset e_\alpha N e_\beta, , \text{ and}$$

$$e_\beta W^\vee e_\alpha = k\{p^\vee(\beta, \alpha^i) \otimes p^\vee(\alpha^i, \alpha) - p^\vee(\beta, \alpha^j) \otimes p^\vee(\alpha^j, \alpha)\} \subset e_\beta N^\vee e_\alpha.$$

Pairing the two basis vectors together using the form defined above we get

$$\varepsilon(\alpha \leftrightarrow \alpha^i)\varepsilon(\alpha^i \leftrightarrow \beta) + \varepsilon(\alpha \leftrightarrow \alpha^j)\varepsilon(\alpha^j \leftrightarrow \beta).$$

By our choice of signs, the terms cancel and we conclude that $e_\alpha W^\perp e_\beta = e_\beta W^\vee e_\alpha$.

If $\alpha^j \notin \mathcal{P}$, then either $\alpha^j \in \mathcal{F} \backslash \mathcal{P}$ and $\alpha^j \notin \mathcal{B} = \mathcal{F}^\vee$, or $\alpha^j \in \mathcal{B} \backslash \mathcal{P} = \mathcal{F}^\vee \backslash \mathcal{P}^\vee$ and $\alpha^j \notin \mathcal{F} = \mathcal{B}^\vee$. Assume that $\alpha^j \in \mathcal{F} \backslash \mathcal{P}$ and $\alpha^j \notin \mathcal{B} = \mathcal{F}^\vee$. Then we have that

$$e_\alpha N e_\beta = k\{p(\alpha, \alpha^i) \otimes p(\alpha^i, \beta)\} = e_\alpha W e_\beta$$

since $p(\alpha, \alpha^i, \beta) = p(\alpha, \alpha^j, \beta) = 0$ in A. On the other hand, $\alpha^j \notin \mathcal{F}^\vee$ means that $e_\beta N^\vee e_\alpha = k\{p^\vee(\beta, \alpha^i) \otimes p^\vee(\alpha^i, \alpha)\}$ and $(A2)$ does not impose any relations, so $e_\beta W^\vee e_\alpha = 0$. Therefore $e_\beta W^\perp e_\alpha = 0 = e_\beta W^\vee e_\alpha$. The case $\alpha^j \in \mathcal{B}\backslash\mathcal{P} = \mathcal{F}^\vee\backslash\mathcal{P}^\vee$ and $\alpha^j \notin \mathcal{F} = \mathcal{B}^\vee$ follows from the same argument on the dual side.

Finally, consider the case where $\alpha = \beta$. Note that

$$e_\alpha N e_\alpha = \mathbf{k}\{p(\alpha, \alpha^i) \otimes p(\alpha^i, a) \mid i \in J_\alpha\}.$$

We identify $e_\alpha N e_\alpha$ with $\mathbf{k}^{J_\alpha}$ by regarding $p(\alpha, \alpha^i) \otimes p(\alpha^i, a)$ as the standard basis element labelled by $i \in J_\alpha$. We also use the standard pairing on $\mathbf{k}^n$ to view $U^\perp$ as a subspace of $\mathbf{k}^n$. From the relations $(A3/4'')$, we find that $e_\alpha W e_\alpha$ is given by $pr_{J_\alpha}(U^\perp \cap \mathbf{k}^{I_\alpha})$ or equivalently by $(pr_{K_\alpha} U^\perp) \cap \mathbf{k}^{J_\alpha}$, where pr_S denotes the orthogonal projection from $\mathbf{k}^n$ to $\mathbf{k}^S$ for any $S \subset \{1, \ldots, n\}$.

Taking the orthogonal complement of $e_\alpha W e_\alpha \subset e_\alpha N e_\alpha$ using the first description gives:

$$\begin{aligned}(e_\alpha W e_\alpha)^\perp = (pr_{J_\alpha}(U^\perp \cap \mathbf{k}^{I_\alpha}))^\perp &= (pr_{I_\alpha} U) \cap \mathbf{k}^{J_\alpha} \\ &= (pr_{K_\alpha^\vee} U) \cap \mathbf{k}^{J_\alpha^\vee} = e_\alpha W^\vee e_\alpha.\end{aligned}$$

Here the second equality follows from the fact that $(pr_S V)^\perp = V^\perp \cap \mathbf{k}^S$, the third equality uses the fact that $I_\alpha = K_\alpha^\vee$ and $J_\alpha = J_\alpha^\vee$ (see Proposition 2.1.18) and the final equality follows from the second description of W above. □

3.3 The algebra B

Following Braden–Licata–Proudfoot–Webster, we define in this section another algebra $B = B(\mathbf{P}, U)$ associated to the pair $(\mathbf{P}, U)$ and prove that B is isomorphic to the quadratic dual $A^!$ of $A(\mathbf{P}, U)$. We also consider a deformed version $\widetilde{B}(\mathbf{P})$ such that $\widetilde{B}(\mathbf{P}) \cong \widetilde{A}(\mathbf{P}^\vee)$. The algebras B and $\widetilde{B}$ defined in [BLPW10, Section 4.1] coincide with those defined here in the special case when $(\mathbf{P}, U)$ comes from a linear subspace as in Example 2.1.6.

3.3.1 A topological lemma

In this section we assume some familiarity with regular cell complexes, posets and their geometric realizations and refer the reader to [BLVS+99, Section 4.7] for more on these topics.

Recall that $\mathcal{L} = \mathcal{L}(\mathcal{N})$ denotes the poset of covectors for $\mathcal{N} = \widetilde{\mathcal{M}} \backslash f$. This poset is pure with a unique minimal element 0 and a rank function ρ. Every covector is uniquely determined by its cocircuit faces. We will define the algebra $B(\mathbf{P}, U)$ from the poset structure of $\mathcal{L}$, using an affine version of the following notion [BLVS+99, Definition 4.1.2(ii)]:

Definition 3.3.1. For any $Y \in \mathcal{L}(\mathcal{N})$, the *Edmonds-Mandel face lattice of* Y, denoted $\mathfrak{F}_{em}(Y)$, is the set of all faces of Y in $\mathcal{L}(\mathcal{N})$. The opposite poset $\mathfrak{F}_{lv}(Y) := \mathfrak{F}_{em}(Y)^{\mathrm{op}}$ is called the *Las Vergnas face lattice of* Y.

Both $\mathfrak{F}_{em}(Y)$ and $\mathfrak{F}_{lv}(Y)$ are graded posets and by theorems of Folkman–Lawrence and Edmonds–Mandel, they have the following topological interpretation [BLVS+99, Theorem 4.3.5]:

Theorem 3.3.2. *The lattices $\mathfrak{F}_{em}(Y)$ and $\mathfrak{F}_{lv}(Y)$ are each isomorphic to the face lattices (or augmented face posets) of PL regular cell decompositions of the $(\rho(Y)-2)$-sphere.*

Definition 3.3.3. Let $Y \in \mathcal{A}$, and let Y^∞ be the unique maximal face of Y such that $Y^\infty \in \mathcal{A}^\infty$. The *feasible Edmonds-Mandel face lattice of Y* is

$$\mathfrak{F}^{\mathcal{F}}_{em}(Y) := (\mathfrak{F}_{em}(Y) \cap \mathcal{A}) \cup \{0\},$$

while the *feasible Las Vergnas face lattice of Y* is

$$\mathfrak{F}^{\mathcal{F}}_{lv}(Y) := \mathfrak{F}^{\mathcal{F}}_{em}(Y)^{\mathrm{op}}.$$

The *core* of $\mathcal{A}$ is $\mathcal{A}^0 := \{Y \in \mathcal{A} \mid Y^\infty = 0\}$.

Note that $\mathcal{A}$ does not have a cellular interpretation, since the faces of a feasible covector need not be feasible. The same is true of $\mathfrak{F}^{\mathcal{F}}_{em}(Y)$ when $Y^\infty \neq 0$. On the other hand, $\mathfrak{F}^{\mathcal{F}}_{lv}(Y)$ will always have a cellular interpretation, even if the program $\mathbf{P}$ is not generic. The assumption that g is generic implies $\mathfrak{F}^{\mathcal{F}}_{lv}(Y)$ is the face lattice of a pure simplicial complex, whose vertices correspond to the feasible facets of Y and whose maximal simplices correspond to the feasible cocircuits faces of Y. We will use the following lemma on the topology of $\mathfrak{F}^{\mathcal{F}}_{lv}(Y)$ to show that the algebra $B(\mathbf{P}, U)$ we define below is finite-dimensional (which will imply that $A(\mathbf{P}, U)$ is finite dimensional as well).

Lemma 3.3.4. *Let $Y \in \mathcal{A}$.*

- *If $Y^\infty = 0$, then the geometric realization $\|\mathfrak{F}^{\mathcal{F}}_{lv}(Y)\backslash\{0, Y\}\|$ of the proper part of the feasible Las Vergnas face lattice of Y is a PL $(\rho(Y)-2)$-sphere.*
- *If $Y^\infty \neq 0$, then $\|\mathfrak{F}^{\mathcal{F}}_{lv}(Y)\backslash\{0, Y\}\|$ is a PL $(\rho(Y)-2)$-ball.*

Proof. If $Y^\infty = 0$, then $\mathfrak{F}_{lv}^{\mathcal{F}}(Y) = \mathfrak{F}_{lv}(Y)$ and so the statement reduces to Theorem 3.3.2.

Assume that $Y^\infty \neq 0$, and let

$$\Delta := \mathfrak{F}_{lv}^{\mathcal{F}}(Y) \backslash \{0, Y\}$$

denote the proper part of the feasible Las Vergnas face lattice of Y.

The geometric realization $\|\Delta\|$ of Δ is homeomorphic to the geometric realization $\|\Delta_{\mathrm{ord}}(\Delta)\|$ of the order complex $\Delta_{\mathrm{ord}}(\Delta)$ of Δ (which is a subdivision of the former). Further, using the canonical identification $\Delta_{\mathrm{ord}}(\Delta) \cong \Delta_{\mathrm{ord}}(\Delta^{\mathrm{op}})$, we find that $\|\Delta\|$ is homeomorphic to $\|\Delta_{\mathrm{ord}}(\Delta^{\mathrm{op}})\|$.

It will thus suffice to prove that $\|\Delta_{\mathrm{ord}}(\Delta^{\mathrm{op}})\|$ is a PL ball. By restricting to $\mathbf{P}/z(Y)$ if necessary, we may assume that Y is a tope, so $\rho(Y) = d+1$ and $\Delta_{\mathrm{ord}}(\Delta^{\mathrm{op}})$ is a $(d-1)$-dimensional simplicial complex.

We first note that $\Delta_{\mathrm{ord}}(\Delta^{\mathrm{op}}) \cong \Delta_{\mathrm{ord}}(\Delta)$ is pure because Δ is the face poset of a pure simplicial complex. By [BLVS+99, Proposition 4.5.4], the graded poset $\mathfrak{F}_{em}^{\mathcal{F}}(Y)$ admits a recursive coatom ordering, which implies by [BLVS+99, Lemma 4.7.19] that its open interval $\Delta^{\mathrm{op}} = (0, Y)$ is shellable (meaning that the order complex $\Delta_{\mathrm{ord}}(\Delta^{\mathrm{op}})$ is shellable).

To conclude that the shellable $(d-1)$-dimensional simplicial complex $\Delta_{\mathrm{ord}}(\Delta^{\mathrm{op}})$ is a PL $(d-1)$-ball we use the criterion of [BLVS+99, Proposition 4.7.22(ii)]: namely we must show that every $(d-2)$-simplex is the face of one or two $(d-1)$-simplices and at least one $(d-2)$-simplex is the face of exactly one $(d-1)$-simplex.

Note that a $(d-1)$-simplex of $\Delta_{\mathrm{ord}}(\Delta^{\mathrm{op}})$ is a maximal chain $x_0 < x_1 < \cdots < x_{d-1}$ of feasible covector faces of Y, where $x_i < x_{i+1}$ means that x_i is a proper face of x_{i+1}. Now consider any $(d-2)$-simplex of $\Delta_{\mathrm{ord}}(\Delta^{\mathrm{op}})$. It will similarly be a chain of the form

$x_0 < x_1 < \cdots < x_{i-1} < x_{i+1} < \cdots < x_{d-1}$ for some i between 0 and $d-1$. If $i > 0$, then the chain can be completed to a maximal chain in exactly two ways because Δ is a pure simplicial complex. If $i = 0$, then the feasible edge x_1 either has one or two feasible cocircuit faces and so the $(d-2)$-simplex can be completed to either one or two $(d-1)$-simplices. As $Y^\infty \neq 0$, there exists an edge x_1 of Y with only one feasible cocircuit face and thus a $(d-2)$-simplex of $\Delta_{\mathrm{ord}}(\Delta^{\mathrm{op}})$ that is the face of a unique $(d-1)$-simplex.

We conclude that $\|\Delta_{\mathrm{ord}}(\Delta^{\mathrm{op}})\| \cong \|\Delta\|$ is a PL $(d-1)$-ball. □

Although we will not need it in sequel, we can prove the following upgrade (from PL to shellable) for feasible Las Vergnas face posets in Euclidean programs.

Lemma 3.3.5. *If* $\mathbf{P}$ *is Euclidean and* Y *is a feasible covector, then* $\mathfrak{F}_{lv}(Y)$ *is the face lattice of a shellable simplicial complex.*

Proof. Again, it suffices to assume Y is a tope, restricting to $\mathbf{P}/z(Y)$ if necessary. Let $\Delta := \mathfrak{F}_{lv}(Y)\backslash\{0, Y\}$ denote the face poset of the simplicial complex $z(\Delta) := \{z(X) \mid X \in \Delta\}$, and let Δ^c be the set of feasible cocircuit faces of Y. We need to exhibit an ordering $b_1, b_2, \ldots, b_r$ of the maximal faces of $z(\Delta)$ such that for all $1 < \ell \leq r$

$$\left(\bigcup_{m=1}^{\ell-1} b_m\right) \cap b_\ell = \bigcup b_\ell \cap b_m,$$

where the right hand side is the union over all $m < \ell$ such that $|b_\ell \cap b_m| = d-1$. Rephrasing in terms of Δ, we wish to display an ordering $\Delta^c = \{X_1, \ldots, X_r\}$ with the following property: For any $1 < \ell \leq r$ and any $F \in \Delta$ with X_ℓ a face of F, we have X_m is a face of F for some $m < \ell$ if and only if X_t is a face of F for some $t < \ell$ with $\rho(X_\ell \circ X_t) = 2$. We only need to prove the forward implication.

We choose our ordering of Δ^c via the following prescription. Since $\mathbf{P}$ is Euclidean, we know that the directed graph $G_{\mathbf{P}}$ contains no loops, and therefore the full subgraph $G \subset G_{\mathbf{P}}$ supported on only the feasible cocircuit faces of Y again has no loops. Then we take our ordering $X_1, \ldots, X_r$ of Δ^c to be a linear refinement of the transitive closure of the relation

$$Z \leq Z' \text{ if there is an arrow } Z' \to Z \text{ in } G.$$

Thus, X_1 is the optimal cocircuit face of Y, and $1 \leq s < t \leq r$ implies that that if a directed path between X_s and X_t exists in G, then it starts at X_t and ends at X_s.

Suppose now that $1 < \ell \leq r$ and F is a covector face of Y with with cocircuit faces X_ℓ and X_m for some $m < \ell$. Suppose also that no $t < \ell$ exists such that $\rho(X_t \circ X_\ell)$ and X_t is a face of F. In $\mathbf{P}/z(F)$, which is again Euclidean, the respective images $\overline{Y}, \overline{F}, \overline{X}_\ell, \overline{X}_m$ of Y, F, X_ℓ, X_m still have $\overline{X}_\ell$ and $\overline{X}_m$ as cocircuit faces of the tope $\overline{Y} = \overline{F}$. The existence of any edge from $\overline{X}_\ell \to \overline{X}_t$ implies $\overline{X}_t \neq 0$ and $t < \ell$, and any such edge must come from an edge $X_\ell \to X_t$. Thus, we know that if no such edge exists, then $\overline{X}_\ell$ must be the optimal cocircuit face of $\overline{F} = \overline{Y}$. But this implies that there is a directed path from $\overline{X}_m$ to $\overline{X}_\ell$ using only cocircuit faces of $\overline{F}$, which can be found using a simplex algorithm which valid even for non-Euclidean programs (see [BLVS+99, Chapter 10]). We can pull this path back up to a directed path in G from X_m to X_ℓ, but this contradicts our ordering since $m < \ell$. Thus, we are done. □

3.3.2 The algebra $\widetilde{B}$

Recall the face ring of a simplicial complex $\mathbf{\Delta}$ is defined as

$$\mathbf{k}[\mathbf{\Delta}] = \mathbf{k}[t_1, \ldots, t_n]/(t_S \mid S \notin \mathbf{\Delta}),$$

where here and in what follows $t_S = \prod_{i \in S} t_i$ with the convention $t_\varnothing = 1$. As before, we write $\mathbf{k}[M]$ denote the face ring of the matroid complex of M.

For any bounded feasible topes $\alpha_1, \ldots, \alpha_r \in \mathcal{P}$, let $Y = \alpha_1 \wedge \cdots \wedge \alpha_r$ denote the unique maximal common feasible covector face of all $\alpha_1, \ldots, \alpha_r$. Then

$$\mathfrak{F}_{em}^{\mathcal{F}}(Y) = \bigcap_{\ell=1}^{r} \mathfrak{F}_{em}^{\mathcal{F}}(\alpha_\ell).$$

We define

$$\Delta_{\alpha_1 \cdots \alpha_r} := \mathfrak{F}_{lv}^{\mathcal{F}}(Y) \backslash \{0\}$$

and

$$z(\Delta_{\alpha_1 \cdots \alpha_r}) := \{S \subset E_n \mid S \subset z(X) \text{ for some } X \in \Delta_{\alpha_1 \cdots \alpha_r}\}.$$

Note that $z(\Delta_{\alpha_1 \cdots \alpha_r})$ is a simplicial complex and $\Delta_{\alpha_1 \cdots \alpha_r}$ can be realized as the face poset of a simplicial complex.

Definition 3.3.6. For $\alpha_1, \ldots, \alpha_r \in \mathcal{P}$, let

$$\widetilde{R}_{\alpha_1 \cdots \alpha_r} := \mathbf{k}[z(\Delta_{\alpha_1 \cdots \alpha_r})].$$

Remark 3.3.7. When $(\mathbf{P}, U)$ comes from a linear subspace as in Example 2.1.6, then each feasible tope α corresponds to a bounded feasible chamber in the the corresponding hyperplane arrangement and the ring $\widetilde{R}_{\alpha_1 \cdots \alpha_r}$ defined here agrees with the corresponding ones defined in [BLPW10, Definition 4.1].

For any feasible covector Y, the zero set $z(Y) \subset E_n$ is an independent set of M, so there are natural quotient maps $\mathbf{k}[M] \to \widetilde{R}_{\alpha_1 \cdots \alpha_r}$. Notice that for any $\beta \in \mathcal{F}$ there is also a natural quotient map $\widetilde{R}_{\alpha_1 \cdots \alpha_r} \to \widetilde{R}_{\alpha_1 \cdots \alpha_r \beta}$ compatible with the maps from $\mathbf{k}[M]$. Furthermore, the quotient $\mathbf{k}[M] \to \widetilde{R}_{\alpha_1 \cdots \alpha_r}$ makes $\widetilde{R}_{\alpha_1 \cdots \alpha_r}$ a $\operatorname{Sym} U$-module.

Lemma 3.3.8. *For $\alpha_1, \ldots, \alpha_r \in \mathcal{P}$, let $Y = \alpha_1 \wedge \cdots \wedge \alpha_r$. The ring $\widetilde{R}_{\alpha_1 \cdots \alpha_r}$ is a free $\operatorname{Sym} U$-module whose rank is equal to the number of feasible cocircuit faces of Y.*

Proof. Lemma 3.3.4 tells us that $\|\mathfrak{F}_{lv}^{\mathcal{F}}(Y) \backslash \{0, Y\}\| = \|\Delta_{\alpha_1 \ldots \alpha_r}\|$ is a $(\rho(Y) - 2)$-sphere or $(\rho(Y) - 2)$-ball. If $z(Y) = \varnothing$, then the posets $z(\Delta_{\alpha_1 \cdots \alpha_r})$ and $\Delta_{\alpha_1 \cdots \alpha_r}$ are isomorphic. More generally, the geometric realization $\|z(\Delta_{\alpha_1 \cdots \alpha_r})\|$ is the $|z(Y)|$-fold cone over $\|\Delta_{\alpha_1 \cdots \alpha_r}\|$. In any case, $\|z(\Delta_{\alpha_1 \cdots \alpha_r})\|$ is either a $(d-1)$-ball or $(d-1)$-sphere.

By results of Hochster, Reisner, and Munkres (cf [Sta04, Section II.4]) it follows that $\widetilde{R}_{\alpha_1 \cdots \alpha_r} = \mathbf{k}[z(\Delta_{\alpha_1 \cdots \alpha_r})]$ is a Cohen-Macaulay ring. Hence [Sta04, Theorem I.5.10] implies that $\widetilde{R}_{\alpha_1 \cdots \alpha_r}$ is a free module of finite rank over the symmetric algebra of any parameter space. Here U is a parameter space for $\widetilde{R}_{\alpha_1 \cdots \alpha_r}$ by [Sta04, Lemma III.2.4] because the composition $U \hookrightarrow \mathbf{k}^I \twoheadrightarrow \mathbf{k}^{z(X)}$ is an isomorphism for any cocircuit $X \in \mathfrak{F}_{em}^{\mathcal{F}}(Y)$. Also by [Sta04, Lemma III.2.4], the rank is equal to the number of maximal simplices of $z(\Delta_{\alpha_1 \cdots \alpha_r})$, which are in bijection with the cocircuits of $\mathfrak{F}_{em}^{\mathcal{F}}(Y)$. □

Remark 3.3.9. While we will not need it in what follows, we note $z(\Delta_{\alpha_1 \cdots \alpha_r})$ is in fact shellable if $\mathbf{P}$ is Euclidean by 3.3.5. This shelling combined with [Sta04, Theorem III.2.5] gives an explicit basis of $\widetilde{R}_{\alpha_1 \cdots \alpha_r}$ as a free $\operatorname{Sym} U$-module.

For any $\alpha, \beta \in \mathcal{P}$, we let

$$d_{\alpha\beta} := |\{i \mid \alpha(i) \neq \beta(i)\}|,$$

which coincides with the length of any taut path from α to β. For α, β, γ, we let

$$S^{\beta}_{\alpha\gamma} := \{i \mid \alpha(i) = \gamma(i) \neq \beta(i)\},$$

which is the set of $i \in E_n$ such that the concatenation of a taut path α to β and a taut path β to γ changes the i-th coordinate exactly twice.

For a graded vector space (or module) M and integer k, we write $M\langle k \rangle$ to denote the graded vector space shifted down by k, that is $(M\langle k \rangle)_i = M_{i+k}$.

Definition 3.3.10. Let $\widetilde{B} = \widetilde{B}(\mathbf{P})$, as a graded vector space in non-negative degree, be defined as

$$\widetilde{B} := \bigoplus_{(\alpha,\beta)\in\mathcal{P}\times\mathcal{P}} \widetilde{R}_{\alpha\beta}\langle -d_{\alpha\beta}\rangle$$

where the variables t_i are given degree 2.

Following [BLPW10] we define a multiplication $\star\colon \widetilde{B}\otimes\widetilde{B} \to \widetilde{B}$ as zero on $\widetilde{R}_{\alpha\beta}\otimes\widetilde{R}_{\delta\gamma}$ if $\beta \neq \delta$ and for $\alpha, \beta, \gamma \in \mathcal{P}$, by the composition

$$\widetilde{R}_{\alpha\beta}\otimes\widetilde{R}_{\beta\gamma} \to \widetilde{R}_{\alpha\beta\gamma}\otimes\widetilde{R}_{\alpha\beta\gamma} \to \widetilde{R}_{\alpha\beta\gamma} \overset{f_{\alpha\beta\gamma}}{\to} \widetilde{R}_{\alpha\gamma}$$

where the first map is the product of the restrictions, the second is multiplication in $\widetilde{R}_{\alpha\beta\gamma}$, and the third map $f_{\alpha\beta\gamma}$ is induced by multiplication by $t_{S^{\beta}_{\alpha\gamma}}$.

Lemma 3.3.11. *The map $f_{\alpha\beta\gamma}$ is well-defined.*

Proof. It suffices to check that if $t_S = 0$ in $\widetilde{R}_{\alpha\beta\gamma}$, then $t_{S^{\beta}_{\alpha\gamma}} t_S = 0$ in $\widetilde{R}_{\alpha\gamma}$. In other words, we want to show if $S \notin z(Y)$ for all $Y \in \Delta_{\alpha\beta\gamma}$, then $S \cup S^{\beta}_{\alpha\gamma} \notin z(Y')$ for all $Y' \in \Delta_{\alpha\gamma}$. Suppose $Y' \in \Delta_{\alpha\gamma}$ and either: (1) $Y' \notin \Delta_{\beta}$ or (2) $Y' \in \Delta_{\beta}$.

In the first case, $Y' \notin \Delta_\beta$ means that there exists an i such that $Y'(i) = -\beta(i)$, in particular $Y'(i) \neq 0$ and $i \notin z(Y')$. As Y' is a face of α and γ, it follows that $Y'(i) = \alpha(i) = \gamma(i)$ and so $i \in S^\beta_{\alpha\gamma}$. Thus $S \cup S^\beta_{\alpha\gamma} \notin z(Y')$.

In the second case, we have $Y' \in \Delta_{\alpha\beta\gamma}$ hence by assumption $S \notin z(Y')$ and so again $S \cup S^\beta_{\alpha\gamma} \notin z(Y')$. □

Proposition 3.3.12. *The multiplication $\star$ gives $\widetilde{B}$ the structure of a graded ring.*

Proof. We need to check associativity and compatibility with grading. For associativity, the map $\widetilde{R}_{\alpha\beta} \otimes \widetilde{R}_{\beta\gamma} \otimes \widetilde{R}_{\gamma\delta} \to \widetilde{R}_{\alpha\delta}$ given by $x \otimes y \otimes z \mapsto (x \star y) \star z$ is equal to the map given by restricting each of the components to $\widetilde{R}_{\alpha\beta\gamma\delta}$, multiplying in order, and then multiplying by $t_{S^\beta_{\alpha\gamma}} \cdot t_{S^\gamma_{\alpha\delta}}$ to get back into $R_{\alpha\delta}$. For the $x \star (y \star z)$, the only change is that we multiply by $t_{S^\gamma_{\beta\delta}} \cdot t_{S^\beta_{\alpha\delta}}$. To show

$$t_{S^\beta_{\alpha\gamma}} \cdot t_{S^\gamma_{\alpha\delta}} = t_{S^\gamma_{\beta\delta}} \cdot t_{S^\beta_{\alpha\delta}},$$

note that the power of t_i appearing on each side is equal to the number of times a path given by the concatenation of taut paths from α to β, β to γ, and γ to δ changes the i-th coordinate twice.

For the compatibility of gradings, note that

$$d_{\alpha\beta} + d_{\beta\gamma} - d_{\alpha\gamma} = 2|S^\beta_{\alpha\gamma}|.$$

It follows that multiplication $\star$ gives a graded preserving map

$$\widetilde{R}_{\alpha\beta}\langle -d_{\alpha\beta}\rangle \otimes_{\mathbf{k}} \widetilde{R}_{\beta\gamma}\langle -d_{\beta\gamma}\rangle \to \widetilde{R}_{\alpha\gamma}\langle -d_{\alpha\gamma}\rangle.$$

(Recall that $\deg(t_i) = 2$ for all i.) □

Note that the map $\zeta\colon \mathrm{Sym}(\mathbf{k}^n) = \mathbf{k}[t_1, \dots, t_n] \to \widetilde{B}$ given by the composition

$$\mathbf{k}[t_1, \dots, t_n] \hookrightarrow \bigoplus_{\alpha \in \mathcal{P}} \mathbf{k}[t_1, \dots, t_n] \twoheadrightarrow \bigoplus_{\alpha \in \mathcal{P}} \widetilde{R}_{\alpha\alpha} \hookrightarrow \widetilde{B}$$

makes $\widetilde{B}$ into a graded $\mathrm{Sym}(\mathbf{k}^n)$-algebra. Moreover, this map factors through the projection $\mathbf{k}[t_1, \dots, t_n] \to \mathbf{k}[M]$ and so we may also view $\widetilde{B}$ as a graded $\mathbf{k}[M]$-algebra.

Let

$$R_{\alpha_1 \cdots \alpha_r} \coloneqq \widetilde{R}_{\alpha_1 \cdots \alpha_r} \otimes_{\mathrm{Sym}\, U} \mathbf{k} = \widetilde{R}_{\alpha_1 \cdots \alpha_r} \otimes_{\mathbf{k}[M]} \mathbf{k}[M]/(U).$$

We define $B = B(\mathbf{P}, U)$ via

$$B \coloneqq \widetilde{B} \otimes_{\mathrm{Sym}\, U} k = \widetilde{B} \otimes_{\mathbf{k}[M]} \mathbf{k}[M]/(U).$$

Note, B is itself a graded ring whose multiplication we will also denote by $\star$.

Theorem 3.3.13. *There is a natural isomorphism $\widetilde{A}(\mathbf{P}^\vee) \to \widetilde{B}(\mathbf{P})$ as graded $\mathrm{Sym}(\mathbf{k}^n)$-algebras, and this induces an isomorphism $A(\mathbf{P}^\vee, U^\perp) \simeq B(\mathbf{P}, U)$ as graded rings.*

In particular the theorem implies $\widetilde{A}(\mathbf{P}^\vee)$ is in fact a $\mathbf{k}[M]$-module.

Proof. We define a map $\phi\colon \widetilde{A}(\mathbf{P}^\vee) \to \widetilde{B}(\mathbf{P})$ by

1. $e_\alpha \mapsto 1_{\alpha\alpha} \in \widetilde{R}_{\alpha\alpha}$ for all $\alpha \in \mathcal{P}$,
2. $p^\vee(\alpha, \beta) \mapsto 1_{\alpha\beta} \in \widetilde{R}_{\alpha\beta}\langle -1 \rangle$ for adjacent $\alpha, \beta \in \mathcal{P}$, and
3. $f \mapsto \zeta(f)$ for all $f \in \mathrm{Sym}(\mathbf{k}^n) = \mathbf{k}[t_1, \dots, t_n]$.

We will first show that this gives a well-defined homomorphism. We then prove surjectivity and conclude with injectivity.

We must check that the image of the relations $(A2)$ and $(A3)$ for $\widetilde{\mathcal{A}}(\widetilde{\mathcal{M}}^\vee)$ hold in $\widetilde{B}(\widetilde{\mathcal{M}})$.

For $(A2)$, we consider $\alpha \in \mathcal{P}^\vee = \mathcal{P}$ and $i \neq j$ in E_n such that $\gamma = (\alpha^i)^j \in \mathcal{P}^\vee$ and $\alpha^i, \alpha^j \in \mathcal{F}^\vee$. This means that $S^{\alpha^i}_{\alpha\gamma} = \varnothing = S^{\alpha^j}_{\alpha\gamma}$, so that $t_{S^{\alpha^i}_{\alpha\gamma}} = 1$.

If α^i and α^j are both in $\mathcal{P}^\vee = \mathcal{P}$, then we have

$$1_{\alpha\alpha^i} \star 1_{\alpha^i\gamma} = 1_{\alpha\gamma} = 1_{\alpha\alpha^j} \star 1_{\alpha^j\gamma}.$$

Otherwise, by relabeling i and j if necessary, we may assume without loss of generality that $\alpha^i \notin \mathcal{P}^\vee$. This means that in $\widetilde{A}(\mathbf{P}^\vee)$, we have

$$0 = p^\vee(\alpha, \alpha^i, \gamma) = p^\vee(\alpha, \alpha^j, \gamma).$$

On the other hand it also means that $\alpha^i \in \mathcal{F}^\vee \backslash \mathcal{P}^\vee = \mathcal{B} \backslash \mathcal{F}$. But any common face of α and γ must also be a face of α^i, which is infeasible, hence α and γ have no common feasible faces and so $\widetilde{R}_{\alpha\gamma} = 0$, which means both products $1_{\alpha\alpha^i} \star 1_{\alpha^i\gamma}$ and $1_{\alpha\alpha^j} \star 1_{\alpha^j\gamma}$ must be zero.

We now check the relation $(A3)$. Let $\alpha \in \mathcal{P}^\vee$ and $\alpha^i \in \mathcal{F}^\vee$.

If $\alpha^i \in \mathcal{P}^\vee$, we have that

$$\begin{aligned}\phi(e_\alpha t_i) = 1_{\alpha\alpha} \star \zeta(t_i) = t_i \in \widetilde{R}_{\alpha\alpha} \\ = t_{S(\alpha(\alpha^i)\alpha)} \in \widetilde{R}_{\alpha\alpha} \\ = 1_{\alpha(\alpha^i)} \star 1_{(a^i)\alpha} = \phi(p(\alpha, \alpha^i, \alpha)).\end{aligned}$$

If $\alpha^i \notin \mathcal{P}^\vee$, then $\alpha^i \in \mathcal{F}^\vee \backslash \mathcal{P}^\vee$ and so $e_\alpha t_i = e_\alpha p^\vee(\alpha, \alpha^i, \alpha) = 0$ in $\widetilde{A}(\mathbf{P}^\vee)$. On the other hand, $\alpha^i \in \mathcal{B} \backslash \mathcal{P}$ implies that i is not in the zero set of any feasible face of the feasible tope α of $\mathcal{N}$, so $t_i = 0 \in \widetilde{R}_{\alpha\alpha}$. This completes the proof that the relation $(A3)$ is satisfied.

Thus, the homomorphism ϕ is well-defined.

To see that ϕ is surjective, note that $\phi(e_\alpha t_i) = 1_{\alpha\alpha} t_i$ for all $i \in E_n$ and $\alpha \in \mathcal{P}^\vee = \mathcal{P}$. This means that $\bigoplus_\alpha \widetilde{R}_{\alpha\alpha} \subseteq \widetilde{B}$ is contained in the image of ϕ. Since the natural quotient $\widetilde{R}_{\alpha\alpha} \to \widetilde{R}_{\alpha\alpha\beta} = \widetilde{R}_{\alpha\beta}$ is given by multiplication by $1_{\alpha\beta} = \phi(p^\vee_{\alpha,\beta})$ for any $\beta \in \mathcal{P}$ and $p^\vee_{\alpha,\beta}$ representing a taut path from α to β in the quiver $Q^\vee$ associated to $\mathbf{P}^\vee$, we have that ϕ is surjective.

Finally, we must prove that ϕ is injective. It suffices to show that the dimension of $\widetilde{R}_{\alpha\beta}$ in each degree is at least the dimension of the corresponding graded part of $e_\alpha \widetilde{A}(\mathbf{P}^\vee) e_\beta$. To do so, we construct a surjection of graded $\mathrm{Sym}(\mathbf{k}^n)$-modules $\widetilde{R}_{\alpha\beta} \to e_\alpha \widetilde{A}(\mathbf{P}^\vee) e_\beta$.

Let $\chi : \mathrm{Sym}(\mathbf{k}^n) = \mathbf{k}[t_1, \ldots, t_n] \to e_\alpha \widetilde{A}(\mathbf{P}^\vee) e_\beta$ be the map that takes 1 to a taut path p from α to β in the quiver $Q^\vee$. By Proposition 3.1.6 and Proposition 3.1.5 this map is surjective. It remains to show that χ factors through $\widetilde{R}_{\alpha\beta}$, which is equivalent to showing that for any $S \notin z(Y)$ for all $Y \in \Delta_{\alpha\beta}$, we have $\chi(t_S) = p \cdot t_S = 0$. We reduce to the case of $\alpha = \beta$ by replacing S by $S \cup S^\beta_{\alpha\alpha}$.

By Corollary 3.1.7 it suffices to prove the existence of $\gamma \in \mathcal{F}^\vee \backslash \mathcal{P}^\vee = \mathcal{B} \backslash \mathcal{P}$ such that $\gamma(i) = \alpha(i)$ for all $i \notin S$. Since $S \notin z(Y)$ for any $Y \in \Delta_\alpha$, the image $\overline{\alpha}$ of α in $\mathbf{P}/S = (\widetilde{\mathcal{M}}/S, g, f)$ is not feasible. However, $\overline{\alpha}$ is bounded in $\mathbf{P}/S$ since any cocircuit face $\overline{X} \in \mathcal{A}^\infty(\mathbf{P}/S)$ of $\overline{\alpha}$ comes from a cocircuit face $X \in \mathcal{A}^\infty(\mathbf{P})$ of α, and therefore has $\overline{X}(f) = X(f) = -$. Hence $\overline{\alpha}$ is a feasible but unbounded sign vector in $\mathbf{P}^\vee \backslash S$, and therefore allows us to find $\gamma \in \mathcal{F}^\vee \backslash \mathcal{P}^\vee$ with $\gamma(i) = \alpha(i)$ for all $i \notin S$. □

3.4 The center of B

We continue to assume $\mathbf{P} = (\widetilde{\mathcal{M}}, g, f)$ is a generic oriented matroid program and $U \subset \mathbf{k}^n$ is a parameter space for $M = \underline{\mathcal{M}}$. In this section we compute the centers of $\widetilde{B} := \widetilde{B}(\mathbf{P})$ and $B := B(\mathbf{P}, U)$.

Let $\overline{\zeta}$ be the composition

$$\mathbf{k}[M] \hookrightarrow \bigoplus_{\alpha \in \mathcal{P}} \mathbf{k}[M] \to \bigoplus_{\alpha \in \mathcal{P}} \widetilde{R}_{\alpha\alpha} \hookrightarrow \widetilde{B}$$

and recall that this map makes $\widetilde{B}$ a graded $\mathbf{k}[M]$-algebra.

Theorem 3.4.1. *The map $\overline{\zeta}\colon \mathbf{k}[M] \to \widetilde{B}(\mathbf{P})$ is injective, and its image is the center of $\widetilde{B}(\mathbf{P})$. Furthermore, the quotient $\widetilde{B}(\mathbf{P}) \to B(\mathbf{P}, U)$ induces a surjection of centers and the center of $B(\mathbf{P}, U)$ is isomorphic to $\mathbf{k}[M]/(U)$.*

This is the natural generalization of [BLPW10, Theorem 4.16] and we imitate the structure of their proof, which makes use of the *extended algebras*:

$$\widetilde{B}_{\text{ext}} = \widetilde{B}_{\text{ext}}(\mathbf{P}) := \bigoplus_{(\alpha,\beta) \in \mathcal{F} \times \mathcal{F}} \widetilde{R}_{\alpha\beta}\langle -d_{\alpha\beta} \rangle \quad \text{and} \quad B_{\text{ext}} = B_{\text{ext}}(\mathbf{P}, U) := \widetilde{B}_{\text{ext}} \otimes_{\text{Sym}\, U} k,$$

where the product $\star$ is defined as before but we use all feasible topes, not just the bounded feasible topes. Braden–Licata–Proudfoot–Webster first prove their result is true when $\widetilde{B}$ and B are replaced by the extended algebras $\widetilde{B}_{\text{ext}}$ and B_{ext}. This is done by studying a chain complex whose homology homology is the center $Z(\widetilde{B}_{\text{ext}})$. To get the theorem for $\widetilde{B}$ and B, they then use a categorical limit argument.

This section is split into two subsections. In the first subsection we define the necessary notation and lay the topological foundation for the proof. In the second subsection we adapt the arguments of [BLPW10] to our setting.

3.4.1 The topology of affine space and feasible Edmonds Mandel face lattices

When $\mathbf{P}$ realizable by a polarized arrangement (V, η, ξ), it is possible to view $\mathcal{A}$ as the cells of the coordinate hyperplane arrangement in $\eta + V \subset \mathbb{R}^n$. This allows one to find tubular neighborhoods of intersections of hyperplanes, and compute the relative cellular Borel-Moore homology of these tubular neighborhoods using the decomposition by cells. In this section we recall definitions and notation to generalize these notions to our setting.

For any $i \in E_n$ and $S \subseteq E_n$, let

$$H_i^{\mathcal{F}} := \{X \in \mathcal{A} \mid X(i) = 0\} \quad \text{and} \quad H_S^{\mathcal{F}} := \bigcap_{i \in S} H_i^{\mathcal{F}}.$$

Our genericity assumption on g implies that $H_S^{\mathcal{F}} \neq \varnothing$ if and only if S is independent in the underlying matroid M of $\mathcal{M}$, in which case any maximal covector in $H_S^{\mathcal{F}}$ has rank $d+1-|S|$. Note that $H_S^{\mathcal{F}}$ does not have a cellular interpretation[2] unless S is a basis of M, in which case $H_S^{\mathcal{F}}$ consists of a single feasible cocircuit.

It is known [BLVS+99, Theorem 4.5.7.i] that $\|\Delta_{\text{ord}}(\mathcal{A})\|$ is a shellable d-ball. Thus its boundary $\|\Delta_{\text{ord}}(\mathcal{A}\backslash\mathcal{A}^0)\|$ is a PL $(d-1)$-sphere. In $\|\Delta_{\text{ord}}(\mathcal{A})\|$, we have that $\|\Delta_{\text{ord}}(H_S^{\mathcal{F}})\|$ is a $(d-|S|)$-ball (when nonempty), with boundary $\|\Delta_{\text{ord}}(H_S^{\mathcal{F}}\backslash\mathcal{A}^0)\|$.

For any $Y \in \mathcal{A}$, let

$$\sigma_Y := \|\Delta_{\text{ord}}(\mathfrak{F}_{em}^{\mathcal{F}}(Y)\backslash\{0\})\| \subset \|\Delta_{\text{ord}}(\mathcal{A})\|.$$

Lemma 3.4.2. *For any $Y \in \mathcal{A}$, the order complex $\Delta_{\text{ord}}(\mathfrak{F}_{em}^{\mathcal{F}}(Y)\backslash\{0\})$ is a shellable $(\rho(Y)-1)$-ball.*

[2] As was discussed in the case of $H_{\varnothing}^{\mathcal{F}} = \mathcal{A}$ following Definition 3.3.3.

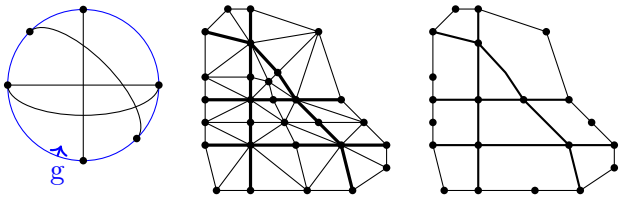


Figure 3.1: $\|\mathcal{A} \cup \mathcal{A}^\infty\|$, $\|\Delta_{ord}(\mathcal{A})\|$, $\|\Xi\|$

Proof. Recalling the notation from the proof of Lemma 3.3.4, note that $\Delta_{\text{ord}}(\mathfrak{F}_{em}^{\mathcal{F}}(Y)\backslash\{0\})$ is the cone on $\Delta_{\text{ord}}(\Delta^{\text{op}})$ with vertex Y. If $Y^\infty = 0$, then $\Delta_{\text{ord}}(\Delta^{\text{op}})$ is a shellable sphere (see [BLVS+99, Theorem 4.3.5(i)]) and if $Y^\infty \neq 0$, then in the proof of Lemma 3.3.4 we showed that $\Delta_{\text{ord}}(\Delta^{\text{op}})$ is a shellable ball. In both cases, the cone is a shellable ball as claimed. □

The boundary of each σ_Y is the union of cells σ_X for the proper faces X of Y and the geometric simplices in $\Delta_{\text{ord}}(\mathfrak{F}_{em}^{\mathcal{F}}(Y)\backslash\{0\})$ corresponding to chains that do not begin with a cocircuit. Let Ξ to be the regular cell complex of cells $\{\sigma_Y\}_{Y\in\mathcal{A}}$ together with the set of (geometric) simplices $\{\sigma\}_{\sigma\in\Delta_{\text{ord}}(\mathcal{A}\backslash\mathcal{A}^0)}$. In particular, $\Delta_{\text{ord}}(\mathcal{A})$ is a subdivision of Ξ. The cells σ_Y for $Y \in \mathcal{A}^0$ define a subcomplex of Ξ, as do those for $Y \in H_S^{\mathcal{F}} \cap \mathcal{A}^0$ when S is independent in M.

Remark 3.4.3. Figure 3.1 shows an example to illustrate some of these definitions. The reason we consider $\Delta_{\text{ord}}(\mathcal{A})$ and Ξ is because $\mathcal{A}$ does not have a cellular interpretation unless we include a boundary. The natural boundary would be $\mathcal{A}^\infty$, but this gives us undesirable topology at the boundary. By introducing $\Delta_{\text{ord}}(\mathcal{A})$ and Ξ we resolve this issue.

Definition 3.4.4. For an independent set $S \subset E_n$ of M, we define

$$\Sigma_S := \{A \in \Delta_{\text{ord}}(\mathcal{A}) | A \subset (x_0 < x_1 < \cdots < x_i) \in \Delta_{\text{ord}}(\mathcal{A}) \text{ and } x_j \in H_S^{\mathcal{F}} \text{ for some } j\}$$

and let $N_S := \|\Sigma_S\|^\circ$ be the interior of $\|\Sigma_S\|$. In particular:

$$N_S := \{A \in \Delta_{\text{ord}}(\mathcal{A}) | A = (x_0 < x_1 < \cdots < x_i) \text{ and } x_j \in H_S^{\mathcal{F}} \cap \mathcal{A}^0 \text{ for some } j\}$$

Proposition 3.4.5. *The Borel-Moore homology of N_S is*

$$H_m^{BM}(N_S; \mathbb{Z}) = \begin{cases} \mathbb{Z} & m = d \\ 0 & m \neq d \end{cases}$$

and can be computed in $\|\Xi\|$ using the relative cellular Borel-Moore homology of N_S via the decomposition by cells $N_S \cap \sigma_Y$ for $Y \in \mathcal{A}$. These $N_S \cap \sigma_Y$ are nonempty exactly when $\mathfrak{F}_{em}^{\mathcal{F}}(Y) \cap H_S^{\mathcal{F}} \neq \varnothing$.

Proof. We first show that N_S is a d-ball, from which the first statement follows.

By [BLVS+99, Theorem 4.5.7(i)], $\|\Delta_{\text{ord}}(H_S^{\mathcal{F}})\|$ is a shellable ball. In particular, $\|\Delta_{\text{ord}}(H_S^{\mathcal{F}})\|$ is collapsible (i.e., it collapses to a point).

Now N_S is a regular neighborhood of $\|\Delta_{\text{ord}}(H_S^{\mathcal{F}})\|$ in the d-ball $\|\Delta_{\text{ord}}(\mathcal{A})\|$ and so by [RS72, Corollary 3.27], N_S of $\|\Delta_{\text{ord}}(H_S^{\mathcal{F}})\|$ is a d-ball.

To see that each intersection $N_S \cap \sigma_Y$ for $Y \in \mathcal{A}$ is a cell, note that $N_S \cap \sigma_Y$ is a regular neighborhood of the shellable ball $H_S \cap \sigma_Y$ in σ_Y. The same collapsibility argument from above then implies $N_S \cap \sigma_Y$ is a ball.

The cells $N_S \cap \sigma_Y$ provide a cellular decomposition of N_S modulo its boundary and the space $N_S \cap \sigma_Y$ is nonempty if and only if a face of Y is contained in $H_S^{\mathcal{F}}$ or equivalently, $\mathfrak{F}_{lv}^{\mathcal{F}}(Y) \cap H_S^{\mathcal{F}} \neq \varnothing$. □

Proposition 3.4.6. *Let $Y \in \mathcal{A}\backslash\mathcal{A}^0$ such that Y has a face in $H_S^{\mathcal{F}} \cap \mathcal{A}^0$. Then*

$$H_m^{BM}(N_S \cap \sigma_Y; \mathbb{Z}) = 0$$

for all m. This can be computed in $\|\Xi\|$ using the decomposition of $N_S \cap \sigma_Y$ by cells $N_S \cap \sigma_X$ for $X \in \mathfrak{F}_{em}^{\mathcal{F}}(Y)$, which are nonempty exactly for X which have a face in $H_S^{\mathcal{F}}$.

Proof. The idea of this proof is to again use relative homology, this time on the pair $(\sigma_Y, \sigma_Y\backslash N_S)$. Recall that σ_Y is a PL $(\rho(Y)-1)$-ball by Lemma 3.4.2. Then σ_Y is contractible, so we are done if we can prove the same for $\sigma_Y\backslash N_S$.

There is a unique maximal element of $H_S^{\mathcal{F}} \cap \mathcal{A}^0 \cap \mathfrak{F}_{em}^{\mathcal{F}}(Y)$, which we call X. The complement of N_S in σ_Y is $\|\Delta\|$, where

$$\Delta := \operatorname{ord}(\mathfrak{F}_{em}^{\mathcal{F}}(Y)\backslash(H_S^{\mathcal{F}} \cap \mathcal{A}^0)) \subset \operatorname{ord}(\mathfrak{F}_{em}^{\mathcal{F}}(Y)).$$

Notice that Δ is nonempty since $Y^\infty \neq 0$, and Δ is equivalent to the cone over $\operatorname{ord}(\mathfrak{F}_{em}^{\mathcal{F}}(Y)\backslash\mathfrak{F}_{em}^{\mathcal{F}}(X))$ with cone point Y. Thus $\|\Delta\|$ is contractible, and we are done. □

3.4.2 The proof of Theorem 3.4.1

We will construct a chain complex with homology isomorphic to the center $Z(\widetilde{B}_{\text{ext}})$.

Let

$$\mathcal{A}_r = \{Y \in \mathcal{A} \mid \rho(Y) - 1 = r\} = \{Y \in \mathcal{A} \mid \dim(\sigma_Y) = r\},$$

where ρ is the poset rank of Y in $\mathcal{L}(\mathcal{N})$, $\mathcal{N} = \widetilde{\mathcal{M}}\backslash f$, and $\sigma_Y = \|\operatorname{ord}(\mathfrak{F}_{em}^{\mathcal{F}}(Y))\| \subset \|\Xi\|$ is a ball by Lemma 3.4.2. For example, $\mathcal{A}_d$ is the set of feasible topes $\mathcal{F}$ and $\mathcal{A}_0$ is the set of feasible cocircuits.

For all $Y \in \mathcal{A}_r$, the space of orientations of σ_Y is a one-dimensional vector space

$$or(Y) \coloneqq H_r^{BM}(\sigma_Y^\circ; \mathbf{k}),$$

where σ_Y° denotes the interior of σ_Y. There is a natural boundary map

$$\partial_Y : or(Y) \to \bigoplus_{X \in \mathcal{A}_{r-1} \cap \mathfrak{F}_{em}^{\mathcal{F}}(Y)} or(X).$$

Assembling all such maps, we obtain a chain complex on $\bigoplus_{Y \in \mathcal{A}} or(Y)$, graded by $\dim(\sigma_Y)$, which computes the cellular Borel-Moore homology of $\|\Xi\|^\circ$. As $\|\Xi\|$ is a closed PL d-ball, this homology is one-dimensional in degree d and zero in all other degrees.

For $Y \in \mathcal{A}_r$, let

$$\widetilde{R}_Y \coloneqq \mathbf{k}[z(\mathfrak{F}_{lv}^{\mathcal{F}}(Y) \backslash \{0\})] = \widetilde{R}_{\alpha_1 \cdots \alpha_\ell},$$

for any choice of $\alpha_1, \ldots, \alpha_\ell \in \mathcal{F}$ such that $Y = \alpha_1 \wedge \cdots \wedge \alpha_\ell$. We define a chain complex $C_\bullet$ such that

$$C_r = \bigoplus_{Y \in \mathcal{A}_r} \widetilde{R}_Y \otimes_{\mathbf{k}} or(Y)$$

with differentials for each $Y \in \mathcal{A}_r$

$$\widetilde{R}_Y \otimes_{\mathbf{k}} or(Y) \to \bigoplus_{X \in \mathcal{A}_{r-1} \cap \mathfrak{F}_{em}^{\mathcal{F}}(Y)} \widetilde{R}_X \otimes_{\mathbf{k}} or(X)$$

induced by the natural boundary maps $or(Y) \to or(X)$ and the quotients $\widetilde{R}_Y \to \widetilde{R}_X$ for each facet X of Y.

Lemma 3.4.7. *Fix an orientation class $\Omega \in H_d^{BM}(\|\Xi\|^\circ; k)$, and let $\Omega_\alpha \in or(\alpha)$ be the restriction of Ω for any $\alpha \in \mathcal{F}$. Let $\psi_\alpha : \mathbf{k}[\underline{\mathcal{M}}] \to \widetilde{R}_\alpha$ denote the natural quotient map. Then the homology of $C_\bullet$ is zero outside of degree d and $\mathbf{k}[M] \simeq H_d(C_\bullet)$ via the map*

$$x \mapsto \sum_{\alpha \in \mathcal{F}} \psi_\alpha(x) \otimes \Omega_\alpha.$$

Proof. Following the proof of [BLPW10, Lemma 4.17], for a monomial $m = \prod_i t_i^{s_i}$, let $S = \{i \mid s_i > 0\}$ and $C_\bullet^m \subset C_\bullet$ be the subcomplex consisting of all images of m, namely

$$C_r^m = \bigoplus_{\substack{Y \in \mathcal{A}_r \\ H_S^{\mathcal{F}} \cap \mathfrak{F}_{em}^{\mathcal{F}}(Y) \neq \varnothing}} or(Y).$$

Note that the complex $C_\bullet$ decomposes as a direct sum $\oplus_m C_\bullet^m$ of subcomplexes because the terms of $C_\bullet$ are direct sums of quotients of $\mathbf{k}[x_1, \ldots, x_n]$ by monomial ideals, while the differentials are induced by the identity map on $\mathbf{k}[x_1, \ldots, x_n]$, up to sign.

If the set $S = \{i \mid s_i > 0\}$ is dependent in M, then $H_S = \varnothing$ and so $C_\bullet^m = 0$. If S is independent in M, then $C_\bullet^m$ is the cellular Borel-Moore complex of the neighborhood $N_S \subset \|\Xi\|$, the homology of which is one-dimensional and concentrated in degree d by Proposition 3.4.5. □

Proposition 3.4.8. *The obvious map $\overline{\zeta}_{\text{ext}}: \mathbf{k}[M] \to \widetilde{B}_{\text{ext}}$ is injective, and its image is the center of $\widetilde{B}_{\text{ext}}$. The quotient homomorphism $\widetilde{B}_{\text{ext}} \to B_{\text{ext}}$ induces a surjection of centers, and yields an isomorphism $Z(B_{\text{ext}}) \cong \mathbf{k}[M]/(U)$.*

Proof. As in the proof of [BLPW10, Proposition 4.18], for an element $z \in \widetilde{B}_{\text{ext}}$ to be in the center, it must commute with each idempotent $1_{\alpha\alpha}$ and thus be of the form $z = \sum_{\alpha \in \mathcal{F}} z_\alpha$ where $z_\alpha \in R_\alpha$. Similarly, using the fact that z must commute with $1_{\alpha\beta}$ for α and β adjacent, we find that

$$\psi_{\alpha\beta}(z_\alpha) = \psi_{\beta\alpha}(z_\beta)$$

where $\psi_{\alpha\beta} : \widetilde{R}_\alpha \to \widetilde{R}_{\alpha\beta}$ and $\psi_{\beta\alpha} : \widetilde{R}_\beta \to \widetilde{R}_{\alpha\beta}$ denote the canonical quotient homomorphisms.

As $\widetilde{B}_{\text{ext}}$ is generated by these element and the image of $\overline{\zeta}_{\text{ext}}$, it follows that:

$$Z(\widetilde{B}_{ext}) \cong \left\{ (z_\alpha) \in \bigoplus_{\alpha \in \mathcal{F}} \widetilde{R}_\alpha \mid \psi_{\alpha\beta}(z_\alpha) = \psi_{\beta\alpha}(z_\beta) \text{ for all } \alpha \leftrightarrow \beta \in \mathcal{F} \right\}. \tag{3.4.1}$$

On the other hand, $H_d(C_\bullet)$ is equal to the set of cycles $y = \sum_{\alpha \in \mathcal{F}} y_\alpha \otimes \Omega_\alpha \in C_d$, which is equivalent to the condition that $\psi_{\alpha\beta}(y_\alpha) = \psi_{\beta\alpha}(y_\beta)$ for all $\alpha \leftrightarrow \beta \in \mathcal{F}$. We conclude that $\overline{\zeta}_{\text{ext}}$ induces an isomorphism

$$\mathbf{k}[M] \cong H_d(C_\bullet) \cong Z(\widetilde{B}_{ext}).$$

Finally, we can define a chain complex $\hat{C}_\bullet$ of free $\operatorname{Sym} U$-modules with $\hat{C}_m = C_m$ for $0 \leq m \leq d$ and $\hat{C}_{d+1} = \ker(\partial_d) \cong \mathbf{k}[M]$. This is acyclic, and thus so is $\hat{C}_\bullet \otimes_{\operatorname{Sym} U} \mathbf{k}$. Arguments analogous to above prove that

$$\mathbf{k}[M]/(U) \cong H_d(C_\bullet \otimes_{\operatorname{Sym} U} \mathbf{k}) \cong Z(B_{\text{ext}})$$

through an isomorphism compatible with the quotient $\widetilde{B}_{\text{ext}} \to B_{\text{ext}}$. □

We now include boundedness into our considerations. Define

$$\mathcal{A}^{\mathcal{P}} := \bigcup_{Y \in \mathcal{P} \subset \mathcal{A}_d} \mathfrak{F}^{\mathcal{F}}_{em}(Y),$$

and notice that we have a chain of proper inclusions $\mathcal{A}^0 \subset \mathcal{A}^{\mathcal{P}} \subset \mathcal{A}$. For any $Y \in \mathcal{A}^{\mathcal{P}}$, let Y^0 be the (unique) maximal face of $\mathfrak{F}^{\mathcal{F}}_{em}(Y) \cap \mathcal{A}^0$.

Note that the description of $Z(\widetilde{B}_{ext})$ in (3.4.1) can be rewritten as asking that $\psi_{\alpha\beta}(z_\alpha) = \psi_{\beta\alpha}(z_\beta)$ for all $\alpha, \beta \in \mathcal{F}$, not necessarily adjacent. This can be rephrased as a limit:

$$Z(\widetilde{B}_{\text{ext}}) \cong \varprojlim_{X \in \mathcal{A}} \widetilde{R}_X. \tag{3.4.2}$$

By the same sort of argument, we find

$$Z(B_{\text{ext}}) \cong \varprojlim_{X \in \mathcal{A}} R_X, \quad Z(\widetilde{B}) \cong \varprojlim_{X \in \mathcal{A}^{\mathcal{P}}} \widetilde{R}_X, \quad \text{and} \quad Z(B) \cong \varprojlim_{X \in \mathcal{A}^{\mathcal{P}}} R_X. \tag{3.4.3}$$

The next lemma allows us to conclude that these centers only depend on $\mathcal{A}^0$.

Lemma 3.4.9. *For any $Y \in \mathcal{A}^{\mathcal{P}}$ and any subcomplex $\mathfrak{D} \subset \mathfrak{F}^{\mathcal{F}}_{em}(Y)$ with $Y^0 \in \mathfrak{D}$, the restrictions*

$$\varprojlim_{X \in \mathfrak{F}^{\mathcal{F}}_{em}(Y)} \widetilde{R}_X \to \varprojlim_{X \in \mathfrak{D}} \widetilde{R}_X \quad \textit{and} \quad \varprojlim_{X \in \mathfrak{F}^{\mathcal{F}}_{em}(Y)} R_X \to \varprojlim_{X \in \mathfrak{D}} R_X \tag{3.4.4}$$

are isomorphisms.

Proof. This proof is simply a rephrasing of [BLPW10, Lemma 4.20] in our setting.

This is trivial if $Y = Y^0$, which includes the case where Y is a feasible cocircuit. So we may assume $Y \neq Y^0$ (equivalently, $Y^\infty \neq 0$), and we may also inductively assume the statement is true for all $X \in \mathfrak{F}^{\mathcal{F}}_{em}(Y) \backslash \{Y\}$.

We will prove the statement first for $\mathfrak{D} = \mathfrak{F}^{\mathcal{F}}_{em}(Y) \backslash \{Y\}$. Let $C^Y_\bullet$ be the subcomplex of $C_\bullet$ consisting only of the summands $\widetilde{R}_X \otimes_{\mathbf{k}} or(X)$ for $X \in \mathfrak{F}^{\mathcal{F}}_{em}(Y)$. As in the proof of Lemma 3.4.7, this complex splits into a direct sum of complexes $C^{Y,m}_\bullet = C^Y_\bullet \cap C^m_\bullet$ for each monomial $m = \prod_i x^{s_i}$. The summand $C^{Y,m}_\bullet$ computes the cellular Borel-Moore homology of $N_S \cap \sigma_Y \subset \sigma_Y$, for $S = \{i \mid s_i > 0\}$. By Proposition 3.4.6, this Borel-Moore homology is trivial, so every $C^{Y,m}_\bullet$ is acyclic and so is $C^Y_\bullet$.

We have that $\varprojlim_{X \in \mathfrak{D}} \widetilde{R}_X$ is isomorphic to the kernel of the boundary map $C^Y_{d-1} \to C^Y_{d-2}$. Since $C^Y_\bullet$ is acyclic, the first map of (3.4.4) is therefore an isomorphism. Similarly, the second map is an isomorphism since $C^Y_\bullet$ is an acyclic complex of free $\operatorname{Sym} U$-modules, which implies $C^Y_\bullet \otimes_{\operatorname{Sym} U} \mathbf{k}$ is an acyclic complex of vector spaces.

For a general $\mathfrak{D}$ containing Y^0, pick an ordering $X_1, \ldots, X_r$ of the elements of $\mathfrak{F}_{em}(Y)\backslash\mathfrak{D}$ such that their ranks are nondecreasing, and let $\mathfrak{D}_\ell = \mathfrak{D} \cup \{X_1, \ldots, X_\ell\}$. Then for $1 \leq \ell \leq r$, we have $\mathfrak{F}_{em}^{\mathcal{F}}(X_\ell)\backslash\{X_\ell\} \subset \mathfrak{D}_{\ell-1}$, so an identical argument shows that

$$\varprojlim_{X \in \mathfrak{D}_\ell} \widetilde{R}_X \to \varprojlim_{X \in \mathfrak{D}_{\ell-1}} \widetilde{R}_X \quad \text{and} \quad \varprojlim_{X \in \mathfrak{D}_\ell} R_X \to \varprojlim_{X \in \mathfrak{D}_{\ell-1}} R_X$$

are isomorphisms. □

Proof of Theorem 3.4.1. Put the equations (3.4.2) and (3.4.3) together with Lemma 3.4.9 to get

$$Z(\widetilde{B}_{\text{ext}}) \cong \varprojlim_{X \in \mathcal{A}} \widetilde{R}_X \cong \varprojlim_{X \in \mathcal{A}^{\mathcal{P}}} \widetilde{R}_X \cong Z(\widetilde{B})$$

and

$$Z(B_{\text{ext}}) \cong \varprojlim_{X \in \mathcal{A}} R_X \cong \varprojlim_{X \in \mathcal{A}^{\mathcal{P}}} R_X \cong Z(B).$$

All of these isomorphisms are compatible with $\overline{\zeta}_{\text{ext}}$, $\overline{\zeta}$ and the natural quotients $\widetilde{B}_{\text{ext}} \to B_{\text{ext}}$ and $\widetilde{B} \to B$, so we are done. □

Chapter 4

Module Categories and their Derived Categories

4.1 The module category of A

In this section, we study the simple modules for $A = A(\mathbf{P}, U)$ and their projective covers using a class of standard modules.

Definition 4.1.1. For any $\alpha \in \mathcal{P}$, let

$$L_\alpha := A/(e_\beta \mid \beta \neq \alpha).$$

Then L_α is the simple one-dimensional A-module supported at α, and let

$$P_\alpha := e_\alpha A$$

be its projective cover. We also define V_α to be P_α/K_α, where

$$K_\alpha := \sum_{i \in b} p(\alpha, \alpha^i) \cdot A \subset P_\alpha$$

and b is the basis of M such that $\mu(b) = \alpha$ under the bijection of Corollary 2.1.17. We refer to V_α as the *standard module* and L_α as the *simple module* associated to α.

Lemma 4.1.2. *Let $\alpha \in \mathcal{P}$. The standard module V_α has a basis consisting of a taut path from α to each $\beta \trianglelefteq \alpha$.*

Proof. We simply copy the argument from [BLPW10, Lemma 5.21]. It is clear that the collection of such taut paths is linearly independent. We now show that the image of any other path is trivial in V_α.

Let $b = \mu^{-1}(\alpha) \in \mathbb{B}$.

Suppose p is a taut path from α to $\gamma \in \mathcal{F}$ and $\gamma \not\trianglelefteq \alpha$. Then $\gamma(i) \neq \alpha(i)$ for some $i \in b$ and $\alpha^i \in \mathcal{F}$. By Corollary 3.1.7 p can be replaced by one of the form $p(\alpha, \alpha^i) \cdot x \in K_\alpha$. Thus $p = 0$ as an element of V_α.

If p is a non-taut path, then by Proposition 3.1.6, we can write

$$p = p' \cdot \prod_i u_i^{a_i} = \left(\prod_i e_\alpha u_i^{a_i} \right) \cdot p'$$

where p' is taut (with the same endpoint as p) and $a_i > 0$ for some i. Corollary 3.1.8 implies that for all $i \in E_n$ and some $c_{i,j} \in \mathbf{k}$, we have

$$e_\alpha u_i = \sum_{j \in b} c_{i,j} p(\alpha, \alpha^j, \alpha) \in K_\alpha.$$

Thus, $p \in K_\alpha$ and we are done. □

Corollary 4.1.3. *The kernel of $V_\alpha \twoheadrightarrow L_\alpha$ has a filtration with subquotients isomorphic to L_β for $\beta \lhd \alpha$, each appearing exactly once.*

4.1.1 A is a quasi-hereditary algebra when P is Euclidean

Recall that a finite-dimensional algebra is *quasi-hereditary* if its category of finitely generated modules is *highest weight* in the following sense.

Definition 4.1.4. Let $\mathcal{C}$ be an abelian, artinian $\mathbf{k}$-linear category and let $\mathcal{I}$ be the set indexing the isomorphism classes of simple objects $\{L_\alpha \mid \alpha \in \mathcal{I}\}$ and indecomposable projective objects $\{P_\alpha \mid \alpha \in \mathcal{I}\}$. Then $\mathcal{C}$ is a *highest weight category* if the set $\mathcal{I}$ can be endowed with a partial order $\leq$ and there exists a collection of objects $\{V_\alpha \mid \alpha \in \mathcal{I}\}$ with surjections

$$P_\alpha \to V_\alpha \to L_\alpha$$

that satisfy:

(i) the kernel of $V_\alpha \to L_\alpha$ has a filtration for which each subquotient is isomorphic to L_γ for some $\gamma < \alpha$, and

(ii) The kernel of $P_\alpha \to V_\alpha$ has a filtration for which each subquotient is isomorphic to V_β for some $\beta > \alpha$.

Now consider the category of finitely generated $A(\mathbf{P}, U)$-modules. As discussed above, the isomorphism classes of simple modules are indexed by the set $\mathcal{P}$ of bounded feasible topes. For $A = A(\mathbf{P}, U)$ to be a quasi-hereditary algebra, we will assume that the oriented matroid program $\mathbf{P}$ is Euclidean. Recall from Section 2.1.5, that this implies there is partial order $\leq$ on $\mathcal{P}$ defined by: $\alpha \leq \beta$ when there exists a directed sequence of edges from $\mu^{-1}(\alpha)$ to $\mu^{-1}(\beta)$ in the graph $G_{\mathbf{P}}$ of the program. By Lemma 2.1.27, this is the same partial order as that defined by the transitive closure of the cone relation $\preceq$.

Suppose $\mathbf{P}$ is Euclidean. Then the category of $A(\mathbf{P}, U)$-modules and the poset $\mathcal{I} = (\mathcal{P}, \leq)$ satisfies condition (i) of Definition 4.1.4 by Corollary 4.1.3. To show that A is quasi-hereditary it remains to show condition (ii).

We will use the following simple lemma.

Lemma 4.1.5. *Suppose $\beta \in \mathcal{P}$ and $i \in \mu^{-1}(\beta)$. Then the feasible sign vector β^i is either unbounded or $\beta^i \geq \beta$.*

Proof. Suppose β^i is bounded. As $\beta^i(j) = \beta(j)$ for all $j \neq i$, if $i \notin \mu^{-1}(\beta^i)$, then for all $j \in \mu^{-1}(\beta^i)$, $\beta^i(j) = \beta(j)$. Thus $\beta \in \mathcal{B}_{\mu^{-1}(\beta^i)}$ or equivalently $\beta \leq \beta^i$.

If β^i is bounded and $i \in \mu^{-1}(\beta^i)$, then the optimal solutions of β and β^i are also optimal solutions of their common subtope $Y = \beta \wedge \beta^i$. But Y is a tope of $\mathbf{P}/\{i\}$ and so Y has a unique optimal solution. Thus $\mu^{-1}(\beta) = \mu^{-1}(\beta^i)$, which contradicts the fact that μ is bijection. □

Theorem 4.1.6. *Assume $\mathbf{P}$ is Euclidean. Then the kernel of the quotient homomorphism $P_\alpha \twoheadrightarrow V_\alpha$ has a filtration with each successive subquotient isomorphic to V_β for $\beta \succ \alpha$. Each of these standard modules appears exactly once. In particular, A is quasihereditary.*

Proof. For any $\gamma \in \mathcal{P}$, we define $P_\alpha^\gamma \subset P_\alpha$ to be the submodule generated by paths which pass through γ. For any $\beta \in \mathcal{P}$, let

$$K_\alpha^\beta := \sum_{\gamma > \beta} P_\alpha^\gamma.$$

After choosing a total order on $\{\beta \in \mathcal{P} \mid \alpha \leq \beta\}$ refining $\leq$, the set of submodules $P_\alpha^\beta + K_\alpha^\beta$ with $\beta \geq \alpha$ forms a filtration of K_α with successive subquotients

$$M_\alpha^\beta := P_\alpha^\beta + K_\alpha^\beta / K_\alpha^\beta.$$

Our goal now is to prove that M_α^β is zero if $\beta \not\geq \alpha$, and is isomorphic to V_β if $\beta \geq \alpha$. Notice that $M_\alpha^\alpha = V_\alpha$.

If $\beta \not\geq \alpha$, then there is an index $i \in \mu^{-1}(\beta)$ such that $\alpha(i) \neq \beta(i)$. By Proposition 3.1.6, any path starting at α and passing through β can be written as $p_{\alpha,\beta} \cdot r$ in $P_\alpha^{\geq\beta}$, where $p_{\alpha,\beta}$ represents a taut path from α to β and r represents a path starting at β. We may then apply Corollary 3.1.7 to the taut path $p_{\alpha,\beta}$ and $\gamma = \beta^i$, to show that $p_{\alpha,\beta}$ can be chosen to pass through $\beta^i \in \mathcal{F}$. By Lemma 4.1.5, $P_\alpha^\beta \subseteq P_\alpha^{\beta^i} \subseteq K_\alpha^\beta$, so $M_\alpha^\beta = 0$.

On the other hand, assume that $\beta \succ \alpha$. There is a natural map $P_\beta \to P_\alpha^\beta$ given by composing any element of P_β with a fixed taut path $p_{\alpha,\beta}$ from α to β. This induces a homomorphism $V_\beta \to M_\alpha^\beta$ that we wish to show is an isomorphism.

By Proposition 3.1.6, any path starting at α and passing through β can be expressed as a product of an element in P_β with some taut path from α to β and by Proposition 3.1.5 the taut path can be chosen to be the one we have fixed. It follows that the map $P_\beta \to P_\alpha^\beta$ is surjective and thus the induced map $V_\beta \to M_\alpha^\beta$ is surjective as well.

Finally, we need to show that $V_\beta \to M_\alpha^\beta$ is injective. We proceed by showing that they have the same dimension. The surjectivity of the map implies

$$\dim_k M_\alpha^\beta \leq \dim_k V_\beta = |\{\gamma \in \mathcal{P} \mid \gamma \leq \beta\}|,$$

so that

$$\dim_k P_\alpha = \sum_{\beta \geq \alpha} \dim_k M_\alpha^\beta \leq |\{(\gamma, \beta) \in \mathcal{P} \times \mathcal{P} \mid \alpha, \gamma \leq \beta\}|.$$

We're done if we can show this is an equality. As $A = \sum_{\alpha \in \mathcal{P}} P_\alpha$, it suffices to prove that

$$\dim_k A = \sum_{\alpha \in \mathcal{P}} \dim_k P_\alpha = \{(\alpha, \gamma, \beta) \in \mathcal{P} \times \mathcal{P} \times \mathcal{P} \mid \alpha, \gamma \leq \beta\}.$$

But recall that $A = A(\mathbf{P}, U) \simeq B(\mathbf{P}^\vee, U^\perp) = \bigoplus_{(\alpha,\gamma)\in\mathcal{P}^\vee\times\mathcal{P}^\vee} R^\vee_{\alpha\gamma}$, and so by Lemma 3.3.8

$$\dim_k B(\mathbf{P}^\vee, U^\perp) = \sum_{(\alpha,\gamma)\in\mathcal{P}^\vee\times\mathcal{P}^\vee} |\{\text{common feasible cocircuit faces of } \alpha \text{ and } \gamma\}|.$$

We are then reduced to showing that the number of common feasible cocircuit faces in $\mathbf{P}^\vee$ of α and γ is equal to the number of bounded feasible topes β of $\mathbf{P}$ such that $\alpha \preceq \beta$ and $\gamma \preceq \beta$. This follows from Complementary Slackness (Proposition 2.1.21). □

4.1.2 The structure of projectives when P is not Euclidean

Note that the definition of the standard modules makes sense for any $\mathbf{P}$ and Lemma 4.1.2 holds even in the non-Euclidean case. However, when $\mathbf{P}$ is not Euclidean, the transitive closure of the cone relation is not a poset and so the standard modules are not part of a highest weight structure.

Nonetheless, one might still hope for a version of Theorem 4.1.6: that the kernel of $P_\alpha \twoheadrightarrow V_\alpha$ has a filtration with successive subquotients isomorphic to V_β for $\beta \succ \alpha$. In this section we observe that this is too optimistic a hope, but that it does hold on the level of graded Grothendieck groups.

Recall from Lemma 3.3.8, that for any $\alpha \in \mathcal{P}$, the dimension $\dim_k R_\alpha$ is equal to the number of feasible cocircuit faces of α. We begin with a graded refinement of this statement.

Lemma 4.1.7. *Let $(h_0, h_1 \ldots, h_{d-1})$ denote the h-vector of $z(\Delta_\alpha)$ or equivalently h_i is equal to the dimension of the graded piece of R_α of degree $2i$. Then h_i is equal to the number of feasible vertices of α with i outgoing edges.*

Proof. We proceed by showing that $z(\Delta_\alpha)$ is partitionable. Recall that a pure simplicial complex Δ is *partitionable*, if it can be expressed as a disjoint union of closed intervals of the form

$$\Delta = [G_1, F_1] \sqcup \ldots \sqcup [G_s, F_s],$$

where each F_i is a facet of Δ. By [Sta04, Proposition III.2.3] the h-polynomial of such a simplicial complex is given by

$$h_i = \#\{j : |G_j| = i\}.$$

Recall that $z(\alpha) = \varnothing$ and $z(\Delta_\alpha)$ is isomorphic as a poset to $\mathfrak{F}_{lv}^{\mathcal{F}}(\alpha)\backslash\{0\}$. Thus we may identify a face of α with the faces of the abstract simplicial complex $z(\Delta_\alpha)$.

The facets $F_1, \ldots, F_s$ of $z(\Delta_\alpha)$ are the zero sets of the feasible vertices (i.e., feasible cocircuit faces) of α. If $F_i = z(X_i)$ for a feasible vertex X_i, let G_i be (the zero set of) the meet in Δ_α of the incoming edges of X_i.

Recall that each face of α has a unique optimal solution (this follows from Theorem 2.1.16). For each feasible face Y of α, the face $z(Y) \in \Delta_\alpha$ is in the interval $[G_j, F_j]$ if and only if X_j is the optimal solution of Y. Thus $[G_1, F_1] \sqcup \ldots \sqcup [G_s, F_s]$ is a partition of $z(\Delta_\alpha)$. Note that

$$|G_j| = d - \#\{\text{incoming edges to } X_j\} = \#\{\text{outgoing edges from } X_j\}.$$

We conclude that

$$h_i = \#\{j : |G_j| = i\} = \#\{\text{feasible vertices of } Y \text{ with } i \text{ outgoing edges}\}.$$

□

Corollary 4.1.8. *Let $(h_0, h_1 \ldots, h_{d-1})$ denote the h-vector of $z(\Delta_{\alpha\beta})$ or equivalently h_i is equal to the dimension of the graded piece of $R_{\alpha\beta}$ of degree $2i$. Then h_i is equal to the number of feasible vertices of $\alpha \wedge \beta$ with i outgoing edges of $\alpha \wedge \beta$.*

Proof. Let γ be the tope in $\mathbf{P}/z(\alpha \wedge \beta)$ given by the restriction of $\alpha \wedge \beta$. Then the simplicial complex $z(\Delta_{\alpha\beta})$ is equal to the simplicial join $z(\Delta_\gamma) * \Gamma$ of $z(\Delta_\gamma)$ with the $(d_{\alpha\beta}-1)$-simplex Γ on the set $z(\alpha \wedge \beta)$. By standard properties of the h-polynomial, we have:

$$h(z(\Delta_{\alpha\beta}), x) = h(z(\Delta_\gamma) * \Gamma, x) = h(z(\Delta_\gamma), x)h(\Gamma, x) = h(z(\Delta_\gamma), x).$$

We conclude that the h-vector of $z(\Delta_{\alpha\beta})$ is equal to that of $z(\Delta_\gamma)$. The result then follows from Lemma 4.1.7. □

For an A-module M, let $[M]$ denote the class of M in the Grothendieck group of A-modules. We will consider the Grothendieck group of the category of graded A-modules as a $\mathbb{Z}[q, q^{-1}]$-module, where $[M\langle -k\rangle] = q^k[M]$. For a graded vector space $V = \oplus_i V_i$, we denote the graded dimension of V by $\mathrm{grdim} V = \sum_i (\dim V_i)\ q^i$.

Theorem 4.1.9. *For any generic oriented matroid program $\mathbf{P}$ and any $\alpha \in \mathcal{P}$, the class of the indecomposable projective P_α in the Grothendieck group can be expressed as the sum:*

$$[P_\alpha] = \sum_{\gamma \geq \alpha} q^{d_{\alpha\gamma}} [V_\gamma].$$

Proof. For any $\beta \in \mathcal{P}$ the graded composition series multiplicity of the simple L_β in the projective P_α is equal to the graded dimension of the space of paths in A that start at α and end at β. In other words we have:

$$[P_\alpha] = \sum_{\beta \in \mathcal{P}} (\mathrm{grdim}\ P_\alpha e_\beta) \cdot [L_\beta] = \sum_{\beta \in \mathcal{P}} (\mathrm{grdim}\ e_\alpha A e_\beta) \cdot [L_\beta].$$

By Theorem 3.3.13, grdim $e_\alpha A e_\beta$ = grdim $R^\vee_{\alpha\beta}\langle -d_{\alpha\beta}\rangle = q^{d_{\alpha\beta}} \cdot$grdim $R^\vee_{\alpha\beta}$. By Corollary 4.1.8, we can express the graded dimension of $R^\vee_{\alpha\beta}$ as

$$\text{grdim } R^\vee_{\alpha\beta} = \sum_{i=0}^{d} \#\{\text{feasible vertices of } \alpha\wedge\beta \text{ in } \mathbf{P}^\vee \text{ with } i \text{ outgoing edges}\} \cdot q^{2i}.$$

Observe that by Proposition 2.1.21 the feasible vertices of $\alpha\wedge\beta$ in $\mathbf{P}^\vee$ (i.e., common feasible vertices of both α and β) are in bijection with the bounded feasible topes δ of $\mathbf{P}$ such that $\alpha \preceq \delta$ and $\beta \preceq \delta$. We claim that the number of outgoing edges of $\alpha\wedge\beta$ of the feasible vertex corresponding to δ is equal to $|S^\delta_{\alpha\beta}|$. As in the proof of Corollary 4.1.8, let γ be the bounded feasible tope $(\alpha\wedge\beta)|_{\underline{\alpha\wedge\beta}}$ in the contraction $\mathbf{P}/z(\alpha\wedge\beta)$. Then the number of outgoing edges of the vertex of γ corresponding to δ is equal to the distance between γ and the restriction $\overline{\delta}$ of δ to $\underline{\alpha\wedge\beta}$. But $\underline{\alpha\wedge\beta} = \{i \in E_n \mid \alpha(i) = \beta(i)\}$ and so the distance between γ and $\overline{\delta}$ is equal to the cardinality of the difference set

$$S(\gamma,\overline{\delta}) = \{i \in E_n \mid \alpha(i) = \beta(i) \neq \delta(i)\} = S^\delta_{\alpha\beta}.$$

Rewriting the sum over topes δ of $\mathbf{P}$ such that $\alpha \preceq \delta$ and $\beta \preceq \delta$ and using the formula $d_{\alpha\delta} + d_{\delta\beta} = d_{\alpha\beta} + 2|S^\delta_{\alpha\beta}|$, we find:

$$\text{grdim } e_\alpha A e_\beta = q^{d_{\alpha\beta}}\text{grdim } R^\vee_{\alpha\beta} = \sum_{\delta\succeq\alpha,\beta} q^{d_{\alpha\beta}+2|S^\delta_{\alpha\beta}|} = \sum_{\delta\succeq\alpha,\beta} q^{d_{\alpha\delta}+d_{\delta\beta}}. \tag{4.1.1}$$

Putting it all together,

$$\begin{aligned}[P_\alpha] &= \sum_\beta \sum_{\delta\succeq\alpha,\beta} q^{d_{\alpha\delta}+d_{\delta\beta}}[L_\beta] \\ &= \sum_{\delta\succeq\alpha} q^{d_{\alpha\delta}} \sum_{\beta\preceq\delta} q^{d_{\beta\delta}}[L_\beta] \\ &= \sum_{\gamma\succeq\alpha} q^{d_{\alpha\delta}}[V_\gamma],\end{aligned}$$

as we wished to show. □

We conclude this section with an example of a generic non-Euclidean program and sign vector α for which the kernel of $P_\alpha \twoheadrightarrow V_\alpha$ does not admit a filtration with successive standard subquotients.

Example 4.1.10. Let $\mathbf{P} = (\mathrm{EFM}(8), g, f)$ be the generic non-Euclidean program defined in [BLVS+99, Section 10.4]. Then $M = \underline{\mathrm{EFM}(8)/g\backslash f}$ is the uniform matroid of rank 3 on E_6. As short hand, we simply write ijk for the basis $\{i, j, k\}$ of M. We denote the sign vector of a bounded feasible tope $\alpha : E_6 \to \{0, +, -\}$ using the string of signs

$$\alpha(1)\alpha(2)\alpha(3)\alpha(4)\alpha(5)\alpha(6).$$

The bijection μ between $\mathbb{B}$ and $\mathcal{P}$ can described as follows, where we have listed the pairs $(b, \mu(b)) \in \mathbb{B} \times \mathcal{P}$ for $\mathbf{P}$:

$$(123, +++++{+}) \ (124, +++-{+}-) \ (123, ++-++{+}) \ (126, +-+++-)$$

$$(134, +-+++{+}) \ (135, -++++-) \ (136, ++++--) \ (145, -+++--)$$

$$(146, +-+-+-) \ (156, +++++-) \ (234, ++--++) \ (235, +++--+)$$

$$(236, -++++{+}) \ (245, ++---+) \ (246, +++-++) \ (256, +-+---)$$

$$(345, ++++-+) \ (346, ++----) \ (356, -+++--) \ (456, +++---).$$

Using this table one can deduce the cone relation $\preceq$ on $\mathcal{P}$ from the fact that $\mu(b) \preceq \mu(b')$ if $\mu(b)(i) = \mu(b')(i)$ for any $i \in b'$. For example, if $\mu(b) \prec \mu(456)$, then $b = 346, 145$ or 256.

Let $\alpha = +++--- \in \mathcal{P}$. Recall the notation α^S denoting the sign vector of a tope which differs from α on exactly the set $S \subset E_6$. Using the above list, we find that

$$\{\beta \in \mathcal{P} \mid \beta \succeq \alpha\} = \{\alpha, \alpha^4, \alpha^5, \alpha^6, \alpha^{\{1,4\}}, \alpha^{\{2,5\}}, \alpha^{\{3,6\}}, \alpha^{\{4,5,6\}}\}$$

Suppose there were a filtration $0 \subset F_1 \subset \ldots \subset F_6 \subset K_\alpha$ of the kernel $K_\alpha := \sum_{i \in b} p(\alpha, \alpha^i) \cdot A$ of $P_\alpha \to V_\alpha$ with nonzero successive standard subquotients $\{V_\gamma \mid \gamma \in \mathcal{C}_\alpha \backslash \{\alpha\}\}$ as in the proof of Theorem 4.1.6. Let $V_\gamma = K_\alpha / F_6$ be the final standard subquotient and suppose that $\gamma = \alpha^S$.

Let p be a taut path from α to α^S. If $p \in F_6$, then $(K_\alpha/F_6)e_{\alpha^S} = 0$, which is a contradiction. Thus we may assume that $p \notin F_6$. For any $i \in S$, $p = p(\alpha, \alpha^i)q$ where q is a taut path from α^i to α^S. As $p \notin F_6$, it follows that $p(\alpha, \alpha^i) \notin F_6$ and L_{α^i} is a quotient of V_γ. This is a contradiction unless $S = \{i\}$. Thus S is either $\{4\}, \{5\}$ or $\{6\}$.

Suppose $S = \{4\}$, so $\gamma = \alpha^4 = + + + + - -$. Note that $\alpha^6 \prec \alpha^4$, so $(V_{\alpha^4})e_{\alpha^6} = (K_\alpha/F_6)e_{\alpha^6} \neq 0$ and thus $p(\alpha, \alpha^6) \notin F_6$. But this would mean that L_{α^6} is a quotient of V_{α^4}, which is a contradiction.

After permuting indices, the same argument shows that neither V_{α^5} nor V_{α^6} is a quotient of K_α. We conclude that P_α doesn't not have the expected filtration. More generally, we will see below in the proof of Theorem 4.1.14 that the change of basis matrix between the standard and simple bases for the Grothendieck group is invertible, so $[P_\alpha]$ cannot be expressed as a different sum of standard classes. Thus P_α does not admit a filtration by standards.

4.1.3 A is a Koszul algebra when P is Euclidean

Definition 4.1.11. Let $M = \bigoplus_{\ell \geq 0} M_\ell$ be a graded $\mathbf{k}$-algebra. A complex

$$\ldots \to P_3 \to P_2 \to P_1 \to P_0$$

of graded projective right M-modules is *linear* if each P_ℓ is generated in degree ℓ. We say that M is *Koszul* if every simple right M-module has a linear projective resolution.

Theorem 4.1.12. *Assume* $\mathbf{P}$ *is Euclidean. Then for all* $\alpha \in \mathcal{P}$*, the standard module* V_α *has a linear projective resolution.*

Proof. We follow the proof of [BLPW10, Theorem 5.24].

Let a be the basis corresponding to the optimal cocircuit for α. We will define the promised resolution as the total complex of the following multicomplex.

For any $S \subseteq a$, let $\alpha^S \in \{+,-\}^n$ be the sign vector which disagrees with α on exactly the entries in S. For example, $\alpha^\varnothing = \alpha$, and $\alpha^{\{i\}} = \alpha^i$ for any $i \in a$. Notice that if $i \in S \subseteq a$ and $\alpha^S, \alpha^{S\setminus i} \in \mathcal{P}$, then there is a degree one map

$$\begin{aligned} \phi_{S,i}\colon \quad P_{\alpha^S} \quad &\longrightarrow \quad P_{\alpha^{S\setminus i}}, \\ q \quad &\mapsto \quad p(\alpha^{S\setminus i}, \alpha^S)\cdot q. \end{aligned}$$

We extend this to all $S \subseteq a$ and $i \in a$ by declaring that $P_{\alpha^S} = 0$ if $\alpha^S \notin \mathcal{P}$ and $\phi_{S,i} = 0$ if $i \notin S$. Consider the module

$$\Pi_\alpha := \bigoplus_{S\subseteq a} P_{\alpha^S},$$

which we view as being graded by the free abelian group $\mathbb{Z}\{\varepsilon_i \mid i \in a\}$ where the summand P_{α^S} is given degree $\varepsilon_S := \sum_{i\in S}\varepsilon_i$. For each $i \in a$, consider the differential $\partial_i\colon \Pi_\alpha \to \Pi_\alpha$ of degree $-\varepsilon_i$ defined as the sum

$$\partial_i := \sum_{S\subseteq a} \phi_{S,i}.$$

Observe that $\partial_i\partial_j = \partial_j\partial_i$ for any $i,j \in a$ by relation $(A2)$ and so we can view Π_α as a multi-complex with differentials ∂_i for each $i \in a$.

Let $\mathbf{\Pi}^\bullet_\alpha$ denote the total complex of Π_α. Then $\mathbf{\Pi}^\bullet_\alpha$ is a linear complex of projective modules and $H_0(\mathbf{\Pi}^\bullet_\alpha) = V_\alpha$. It remains to show that the complex $\mathbf{\Pi}^\bullet_\alpha$ is exact in positive degrees.

To do so, we will filter the multicomplex Π_α. For each $\beta \in \mathcal{P}$, let $(\Pi_\alpha)^\beta \subset \Pi_\alpha$ be the submodule whose ε_S-graded part is defined as

$$\sum_{\gamma \geq \beta, \alpha^S} P^\gamma_{\alpha^S} \subset P_{\alpha^S},$$

that is, the submodule consisting of all paths from α^S passing through some $\gamma \in \mathcal{P}$ where $\gamma \geq \alpha^S$ and $\gamma \geq \beta$. Observe that the differentials ∂_i for $i \in a$ are compatible with the submodules $(\Pi_\alpha)^\beta$ and so we have defined a filtration of Π_α by the poset $\mathcal{P}$.

Computing the associated graded of this filtration yields a multi-complex

$$\widetilde{\Pi}_\alpha = \bigoplus_{\beta \in \mathcal{P}} (\Pi_\alpha)^\beta / (\Pi_\alpha)^{>\beta} = \bigoplus_{\beta \in \mathcal{P}} \left(\bigoplus_S M^\beta_{\alpha^S} \right),$$

where $M^\beta_{\alpha^S}$ is the subquotient of P_{α^S} defined as in the proof of Theorem 4.1.6.

Consider the resulting quotient multi-complexes for each $\beta \in \mathcal{P}$. Let $b = \mu^{-1}(\beta)$. Recall from the proof of Theorem 4.1.6 that $M^\beta_{\alpha^S}$ is non-zero if and only if $\alpha^S \in \mathcal{B}_b$.

If $\beta = \alpha$, then $M^\beta_{\alpha^S} = M^\alpha_{\alpha^S} = 0$ for any non-empty $S \subset a = b$. Thus the only non-zero summand of the α-subquotient is $M^\alpha_\alpha = V_\alpha$ in total degree zero.

If $\beta \neq \alpha$, choose an element $i \in a$ such that $i \notin b$. Consider those subsets $S \subset a$ such that $i \notin S$. Then we have $\alpha^S \in \mathcal{B}_b$ if and only if $\alpha^{S \cup \{i\}} \in \mathcal{B}_b$. If $\alpha^S \in \mathcal{B}_b$, then $M^\beta_{\alpha^{S \cup i}} \simeq V_\beta \simeq M^\beta_{\alpha^S}$ and the differential induced by ∂_i is the isomorphism given by left multiplication with $p(\alpha^S, \alpha^{S \cup \{i\}})$. On the other hand, if $\alpha^S \notin \mathcal{B}_b$, then $M^\beta_{\alpha^{S \cup i}} = 0 = M^\beta_{\alpha^S}$. In particular, the differential induced by ∂_i on the β-component of the associated graded multi-complex is exact.

Recall that if any differential of a multi-complex is exact then the total complex of the multi-complex is also exact. We conclude that the total complex of the associated graded multi-complex is exact in positive degree. It then follows that the total complex of the original multi-complex must also be exact in positive degrees. □

Theorem 4.1.13. *Assume* $\mathbf{P}$ *is Euclidean. Then* A *and* B *are Koszul algebras and* A *is Koszul dual to* B*.*

Proof. By [ÁDL03, Theorem 1] a quasi-hereditary algebra is Koszul if the standard modules have linear projective resolutions and such resolutions exist for A by the previous theorem. Theorem 3.3.13 implies that $B \cong A(\mathbf{P}^\vee, U^\perp)$ must also be Koszul. Finally the Koszul duality follows from the quadratic duality statement of Theorem 3.2.2. □

4.1.4 Numerical identity for Hilbert polynomials

We do not know whether or not the condition that $\mathbf{P}$ be Euclidean is necessary for A to be Koszul. In this section we prove that for any generic oriented matroid program $\mathbf{P}$ the Hilbert polynomial of the algebra $A = A(\mathbf{P}, U)$ satisfies the following numerical identity.

Let $H(A, q)$ denote the Hilbert polynomial of A, which is the $\mathcal{P} \times \mathcal{P}$-matrix with entries

$$H(A, q)_{\alpha,\beta} = \text{grdim } e_\alpha A e_\beta.$$

Recall [BGS96, Lemma 2.11.1] that if A is Koszul, then there is an equality of matrices

$$H(A, q)H(A^!, -q)^T = I.$$

Theorem 4.1.14. *For any generic oriented matroid program* $\mathbf{P}$*, the algebra* $A = A(\mathbf{P}, U)$ *satisfies the numerical identity above, that is, the Hilbert polynomials of* A *and its quadratic dual* $A^!$ *satisfy the matrix equation*

$$H(A,q)H(A^!,-q)^T = I.$$

Remark 4.1.15. This identity does not necessarily imply that A is Koszul. See [Pos95] for an example of a non-Koszul quadratic algebra whose Hilbert series satisfies the numerical identity.

Proof. Using equation (4.2.3) in the proof of Theorem 4.1.9, the (α,β)-entry of $H(A,q)$ is given by

$$H(A,q)_{\alpha,\beta} = \operatorname{grdim} e_\alpha A e_\beta = \sum_{\gamma \geq \alpha,\beta} q^{d_{\alpha\gamma}+d_{\gamma\beta}}.$$

In particular, $H(A,q)$ factors as the product $H(A,q) = XX^T$, where X is the $\mathcal{P}\times\mathcal{P}$-matrix with (α,β)-entry given by $q^{d_{\alpha\beta}}$ if $\beta \geq \alpha$, and 0 otherwise.

Dually, using Proposition 2.1.21 we find that the (α,β)-entry of $H(A^!,-q)$ is given by

$$H(A^!,-q)_{\alpha,\beta} = \sum_i (-q)^i \dim e_\alpha A_i^! e_\beta = \sum_{\gamma \supset Y_{\mu^{-1}(\alpha)}, Y_{\mu^{-1}(\beta)}} (-q)^{d_{\alpha\gamma}+d_{\gamma\beta}},$$

in other words the sum runs over all $\gamma \in \mathcal{P}$ for which the optimal solution (cocircuit) of both α and β are faces of γ. Again this factors as a product $H(A^!,-q) = YY^T$, where Y is the $\mathcal{P}\times\mathcal{P}$-matrix with (α,β)-entry given by $(-q)^{d_{\alpha\beta}}$ if $Y_{\mu^{-1}(\beta)}$ is a face of α, and 0 otherwise.

We wish to show that

$$H(A,q)H(A^!,-q)^T = XX^TYY^T = I.$$

Note that it suffices to show $X^TY = I$.

Computing the product X^TY, we find that its (α,β)-entry is given by

$$(X^TY)_{\alpha,\beta} = \sum_{\gamma\in Q} q^{d_{\alpha\gamma}}(-q)^{d_{\gamma\beta}},$$

where Q is the set of all $\gamma \in \mathcal{P}$ such that $\alpha \geq \gamma$ and $Y_{\mu^{-1}(\beta)}$ is a face of γ. In other words, Q consists of all $\gamma \in \mathcal{P}$ such that

$$\gamma(i) = \alpha(i) \text{ if } i \in \mu^{-1}(\alpha) \quad \text{and} \quad \gamma(i) = \beta(i) \text{ if } i \notin \mu^{-1}(\beta). \tag{4.1.2}$$

We wish to show that $(X^TY)_{\alpha,\beta}$ is 1 if $\alpha = \beta$ and is 0 otherwise. If $\alpha = \beta$, then $Q = \{\alpha\}$ and the sum is equal to $q^{d_{\alpha\alpha}}(-q)^{d_{\alpha\alpha}} = 1$. Now assume that $\alpha \neq \beta$ and let

$$J := \mu^{-1}(\beta)\backslash\mu^{-1}(\alpha) \quad \text{and} \quad J' := \mu^{-1}(\alpha)\backslash\mu^{-1}(\beta)$$

so that

$$J \sqcup J' = (\mu^{-1}(\alpha) \cup \mu^{-1}(\beta))\backslash(\mu^{-1}(\alpha) \cap \mu^{-1}(\beta)).$$

As we have assumed that $\alpha \neq \beta$, J and J' are nonempty.

Note that if $\alpha(i) \neq \beta(i)$ for some $i \in J'$, then by the conditions (4.1.2) Q is empty and $(X^TY)_{\alpha,\beta} = 0$ as desired. Thus we will assume that $\alpha(i) = \beta(i)$ for all $i \in J'$.

Let

$$K := \{i \in \mu^{-1}(\alpha) \cap \mu^{-1}(\beta) \mid \alpha(i) \neq \beta(i)\}.$$

For $\delta \in \mathcal{P}$, $Y_{\mu^{-1}(\beta)}$ is a face of δ if and only if $\delta = \beta^W$ for some subset $W \subset \mu^{-1}(\beta)$. On the other hand, $\delta = \beta^W \preceq \alpha$ if and only if $K \cup J \supset W \supset K$. Thus $Q = \{(\beta^K)^S \mid S \subset J\}$ and

$$(X^TY)_{\alpha,\beta} = \sum_{S\subset J} q^{d_{\alpha,(\beta^K)^S}}(-q)^{d_{(\beta^K)^S,\beta}} = (-1)^{|K|}\sum_{S\subset J}(-1)^{|S|}q^{d_{\alpha,(\beta^K)^S}+d_{(\beta^K)^S,\beta}}$$
$$= (-1)^{|K|}\sum_{S\subset J}(-1)^{|S|}q^{d_{\alpha,\beta}+2|S^{(\beta^K)^S}_{\alpha,\beta}|} = (-1)^{|K|}q^{d_{\alpha,\beta}}\sum_{S\subset J}(-1)^{|S|}q^{2|S^{\beta^S}_{\alpha,\beta}|},$$

where in the last line we have used $S_{\alpha,\beta}^{(\beta^K)^S} = S_{\alpha,\beta}^{\beta^S} = \{i \in S \mid \alpha(i) = \beta(i)\}$. We will need the following lemma to finish this proof.

Lemma 4.1.16. *Assume as above that $\alpha \neq \beta$ and $\alpha(i) = \beta(i)$ for all $i \in J'$. Then there exists an element $t \in J$ such that $\alpha(t) \neq \beta(t)$.*

Proof. Suppose for the sake of contradiction that $\alpha(i) = \beta(i)$ for all $i \in J$, then $\alpha(i) = \beta(i)$ for all $i \in J \sqcup J'$. In the deletion-contraction program

$$\left(\mathbf{P}/(\mu^{-1}(\alpha) \cap \mu^{-1}(\beta))\right) \setminus (\mu^{-1}(\alpha) \cup \mu^{-1}(\beta))^c$$

defined on the set $J \sqcup J'$, the restrictions of the sign vectors of α and β are then equal and so describe the same tope T. Now $Y_{\mu^{-1}(\alpha)}$ is the optimal solution for α and $Y_{\mu^{-1}(\beta)}$ is the optimal solution for β, so the restrictions Y_α and Y_β of $Y_{\mu^{-1}(\alpha)}$ and $Y_{\mu^{-1}(\beta)}$ to $J \sqcup J'$ should both be the unique optimal solution of the tope T. But $z(Y_\alpha) = J' \neq J = z(Y_\beta)$, which is a contradiction. Thus there exists a $t \in J$ such that $\alpha(t) \neq \beta(t)$ as desired. □

In particular if $S \subset J \backslash \{t\}$, we have $S_{\alpha,\beta}^{\beta^{S \sqcup \{t\}}} = S_{\alpha,\beta}^{\beta^S}$.

Using this fact we rewrite the sum above:

$$(X^T Y)_{\alpha,\beta} = (-1)^{|K|} q^{d_{\alpha,\beta}} \sum_{S \subset J} (-1)^{|S|} q^{2|S_{\alpha,\beta}^{\beta^S}|},$$

$$= (-1)^{|K|} q^{d_{\alpha,\beta}} \sum_{S \subset J \backslash \{t\}} \left((-1)^{|S|} q^{2|S_{\alpha,\beta}^{\beta^S}|} + (-1)^{|S \sqcup \{t\}|} q^{2|S_{\alpha,\beta}^{\beta^{S \cup \{t\}}}|} \right)$$

$$= (-1)^{|K|} q^{d_{\alpha,\beta}} \sum_{S \subset J \backslash \{t\}} \left((-1)^{|S|} q^{2|S_{\alpha,\beta}^{\beta^S}|} - (-1)^{|S|} q^{2|S_{\alpha,\beta}^{\beta^S}|} \right) = 0.$$

□

4.1.5 Self-dual projectives

Consider the duality functor $\mathrm{d} : A\text{-mod} \to A\text{-mod}$ defined by composing the equivalence $A^{\mathrm{op}}\text{-mod} \simeq A\text{-mod}$ induced by the isomorphism $A \cong A^{\mathrm{op}}$ given by reversing the arrows of the quiver Q_n in Section 3.1.3 with the induced functor $A^{\mathrm{op}}\text{-mod} \to A\text{-mod}$ coming from vector space duality.

In the following result we will involve bounded feasible topes in the affine spaces of both $\mathbf{P}$ and $\mathbf{P}^\vee$. To distinguish the two topes with a fixed sign vector $\alpha \in \mathcal{P} = \mathcal{P}^\vee$, we write T_α to denote the tope in $\mathbf{P}$ and $T_\alpha^\vee$ to denote the tope in $\mathbf{P}^\vee$.

Theorem 4.1.17. *For any generic oriented matroid program $\mathbf{P}$ and $\alpha \in \mathcal{P}$. The following are equivalent:*

1. *The projective P_α is injective.*
2. *The projective P_α is self-dual.*
3. *The simple L_α is contained in the socle of some standard module V_β.*
4. *The bounded feasible tope T_α covers an infeasible subtope X, meaning $X(g) = 0$.*
5. *The bounded feasible tope $T_\alpha^\vee$ in the dual program $\mathbf{P}^\vee$ is in the core of the affine space for $\mathbf{P}^\vee$. In other words $(T_\alpha^\vee)^\infty = 0$ or equivalently the cocircuit faces of the tope $T_\alpha^\vee$ are all feasible.*

When $\mathbf{P}$ is Euclidean, and so A is quasi-hereditary by Theorem 4.1.6, then the statements above are also equivalent to the following:

(6) The projective P_α is tilting.

Proof. The implications (1) $\Leftrightarrow$ (2) ($\Leftrightarrow$ (6), if A is quasi-hereditary) are standard facts.

(2) $\implies$ (3) : If P_α is self-dual, then the socle of P_α is isomorphic to the cosocle of P_α, which is L_α. Therefore when expressing $[P_\alpha]$ as a sum of simple classes times powers of q in the Grothendieck group, L_α is the only simple to appear in the top degree. On the other hand, by Theorem 4.1.9,

$$[P_\alpha] = \sum_{\beta \in C_\alpha} q^{d_{\alpha\beta}} [V_\beta],$$

so the unique simple class appearing in the highest degree must also appear as the highest degree term of some $[V_\beta]$. We conclude that L_α is the socle of V_β.

(3) $\implies$ (4) : Let $b = \mu^{-1}(\beta)$. Lemma 4.1.2 says that V_β is spanned as a vector space by taut paths p_γ from β to $\gamma \trianglelefteq \beta$. A taut path p_α is in the socle of V_β if there does not exist a longer taut path p_γ that factors through p_α. Note that this is equivalent to the condition: if $i \notin b$ and α has a feasible face Y such that $Y(i) = 0$, then $\alpha(i) \neq \beta(i)$.

Recall that $Y_b \in \mathcal{A}$ denotes the feasible cocircuit of $\mathcal{N} = \widetilde{\mathcal{M}}\backslash f$ that is the optimal solution of the tope T_β. By the covector axioms of an oriented matroid, the composition $T := (-Y_b) \circ T_\alpha$ is also a covector of $\mathcal{N}$ and in particular an infeasible tope such that $T(i) = \alpha(i)$ for all $i \in b$ and for all i which are zero on a feasible face of α. A taut path p from T_α to T exists in the tope graph of $\widetilde{\mathcal{M}}\backslash f$, and this path cannot change the sign corresponding to any feasible facet of T_γ. Thus the first sign change of the path p must be infeasible, which means that T_γ covers an subtope X with $X(g) = 0$.

(4) $\implies$ (5): If the bounded tope T_α in $\widetilde{\mathcal{M}}\backslash f$ covers a subtope X with $X(g) = 0$ then α is a bounded feasible sign vector for both the original program $\mathbf{P}$ as well as the reoriented program ${}_{-g}\mathbf{P} = ({}_{-g}\widetilde{\mathcal{M}}, -g, f)$. Dually, this means that the tope $T_\alpha^\vee$ in $\widetilde{\mathcal{M}}^\vee\backslash g =$

$({}_{-g}\widetilde{\mathcal{M}}^{\vee})\backslash(-g)$ is bounded and feasible in both dual programs $\mathbf{P}^{\vee}$ and ${}_{-g}\mathbf{P}^{\vee} = ({}_{-g}\widetilde{\mathcal{M}}^{\vee}, f, -g)$. In particular, the tope $\alpha \in \mathcal{P}^{\vee}$ does not have any cocircuit face Y with $Y(f) = 0$, since this would imply $T_{\alpha}^{\vee}$ was unbounded in one of these generic programs.

(5) $\implies$ (2): If $(T_{\alpha}^{\vee})^{\infty} = 0$, then $e_{\alpha}Ae_{\alpha} \simeq R_{\alpha}^{\vee} = \mathbf{k}[z(\Delta_{\alpha}^{\vee})]/(U^{\perp})$, where $\Delta_{\alpha}^{\vee} = \mathfrak{F}_{lv}(T_{\alpha}^{\vee})\backslash\{0\}$. By Lemma 3.3.4, $\|\Delta_{\alpha}^{\vee}\| = \|\mathfrak{F}_{lv}(T_{\alpha}^{\vee})\backslash\{0\}\|$ is homeomorphic to a $(n-d-1)$-sphere and so a result of Munkres (see [Sta04, Theorem II.4.3]) implies that $R_{\alpha}^{\vee}$ is Gorenstein, meaning that there is an isomorphism $\int : (e_{\alpha}Ae_{\alpha})_{n-d-1} \to \mathbf{k}$ such that $\langle x, y\rangle = \int xy$ defines a perfect pairing on $e_{\alpha}Ae_{\alpha}$.

We wish to produce an isomorphism of A-modules $\mathrm{d}(P_{\alpha}) \cong P_{\alpha}$. To do so, we will show that the map

$$\langle -, -\rangle : e_{\alpha}A \times Ae_{\alpha} \to k$$

$$(p, q) \mapsto \int pq$$

defines a perfect pairing and so it will follow that $\mathrm{d}(P_{\alpha}) = (e_{\alpha}A)^{*} \cong Ae_{\alpha} = P_{\alpha}$ as right A-modules.

To prove that $\langle -, -\rangle$ is perfect, we first observe that it suffices to show that the map $\cdot p_{\beta,\alpha} : e_{\alpha}Ae_{\beta} \to e_{\alpha}Ae_{\alpha}$ is injective for some taut path $p_{\beta\alpha}$ from β to α. This is because for any nonzero $x \in e_{\beta}Ae_{\alpha}$, if $x \cdot p_{\beta\alpha} \neq 0$, then there exists $y \in e_{\alpha}Ae_{\alpha}$ such that $\int (xp)y = \int x(py) = \langle x, py\rangle \neq 0$. On the B side, this means showing that

$$\cdot u_{S_{\alpha\alpha}^{\beta}} : R_{\alpha\beta}^{\vee} \to R_{\alpha}^{\vee}$$

is injective. We proceed by showing

$$\cdot u_{S_{\alpha\alpha}^{\beta}} : \widetilde{R}_{\alpha\beta}^{\vee} \to \widetilde{R}_{\alpha}^{\vee}$$

is the injective map in a split short exact sequence of $\operatorname{Sym} U^\perp$-modules, which proves the claim by applying the functor $- \otimes_{\operatorname{Sym} U^\perp} \mathbf{k}$.

The claim is obvious if $\alpha = \beta$ or $T_\alpha^\vee \wedge T_\beta^\vee$ is not feasible, so we assume $\alpha \neq \beta$ and $T_\alpha^\vee \wedge T_\beta^\vee$ is a proper non-empty face of $T_\alpha^\vee$. To see the monomial map $\cdot u_{S_{\alpha\alpha}^\beta} : \widetilde{R}_{\alpha\beta}^\vee \to \widetilde{R}_\alpha^\vee$ is injective and to determine its cokernel, consider the image of any non-zero monomial $m = \prod_{i \in S} u_i^{s_i}$ in $\widetilde{R}_{\alpha\beta}^\vee$, where $s_i > 0$ for any $i \in S$. As m is non-zero in $\widetilde{R}_{\alpha\beta}^\vee$, there exists $Y \in \Delta_{\alpha\beta}^\vee$ such that $S \subset z(Y)$. Note that $S_{\alpha\alpha}^\beta = z(T_\alpha^\vee \wedge T_\beta^\vee) \subset z(X)$ for any $X \in \Delta_{\alpha\beta}^\vee$. Thus $S \cup S_{\alpha\alpha}^\beta \subset z(Y)$ and the product $m \cdot u_{S_{\alpha\alpha}^\beta} = \prod_{i \in S \cup S_{\alpha\alpha}^\beta} u_i^{t_i}$, where $t_i > 0$ for $i \in S \cup S_{\alpha\alpha}^\beta$, is non-zero in $\widetilde{R}_\alpha^\vee$.

The computation above also shows that the cokernel of the map $\cdot u_{S_{\alpha\alpha}^\beta} : \widetilde{R}_{\alpha\beta}^\vee \to \widetilde{R}_\alpha^\vee$ is the face ring $\mathbf{k}[\Delta]$ of

$$\Delta = \{S \subset E_n \mid S \subset z(Y) \text{ for some } Y \in \Delta_\alpha \text{ and } S_{\alpha\alpha}^\beta \notin S\}.$$

$$= z(\Delta_\alpha) \backslash \{S \in z(\Delta_\alpha) \mid S_{\alpha\alpha}^\beta \subset S\}.$$

Recall from Lemma 3.3.4 that the geometric realization of the simplicial complex $z(\Delta_\alpha)$ is a PL $(d-1)$-sphere. The subset of simplices $\{S \in z(\Delta_\alpha) \mid S_{\alpha\alpha}^\beta \subset S\}$ is the open star of the simplex on the set $S_{\alpha\alpha}^\beta$ and thus its complement Δ in $z(\Delta_\alpha)$ is a PL $(d-1)$-ball. It follows that $\mathbf{k}[\Delta]$ is Cohen-Macaulay, again with parameter space $U^\perp$. Thus, $\mathbf{k}[\Delta]$ is a free $\operatorname{Sym} U^\perp$-module, and therefore the exact sequence

$$\widetilde{R}_{\alpha\beta}^\vee \hookrightarrow \widetilde{R}_\alpha^\vee \twoheadrightarrow \mathbf{k}[\Delta]$$

splits. □

4.2 Derived Morita equivalence

We conclude with a proof of Theorem 1.4.3. Recall that $\mathcal{M}$ is an oriented matroid, U a parameter space for $M = \underline{\mathcal{M}}$, and $\mathbf{P}_1 = (\widetilde{\mathcal{M}}_1, g_1, f_1), \mathbf{P}_2 = (\widetilde{\mathcal{M}}_2, g_2, f_2)$ and $\mathbf{P}_{\text{mid}} = (\widetilde{\mathcal{M}}_{\text{mid}}, g_2, f_1)$ are Euclidean generic oriented matroid programs such that

$$\mathcal{M} = (\widetilde{\mathcal{M}}_1/g_1)\backslash f_1 = (\widetilde{\mathcal{M}}_2/g_2)\backslash f_2$$

and

$$\widetilde{\mathcal{M}}_{\text{mid}}/g_2 = \widetilde{\mathcal{M}}_1/g_1, \qquad \widetilde{\mathcal{M}}_{\text{mid}}\backslash f_1 = \widetilde{\mathcal{M}}_2\backslash f_2.$$

We wish to show there is an equivalence of categories

$$D(A(\mathbf{P}_1, U)) \cong D(A(\mathbf{P}_2, U)),$$

where $D(A)$ denotes the bounded derived category of graded finitely generated A-modules.

Remark 4.2.1. Note that if $\mathbf{P}_1$ and $\mathbf{P}_2$ are Euclidean, it is not automatic that $\mathbf{P}_{\text{mid}}$ will be Euclidean as well. For example, one could take EFM(8) (see Example 4.1.10) and then change the choice of g and f separately to obtain two realizable (and hence Euclidean) generic oriented matroid programs $\mathbf{P}_1$ and $\mathbf{P}_2$ such that $\mathbf{P}_{\text{mid}}$ is the non-Euclidean program EFM(8).

We will prove Theorem 1.4.3 by reducing it to the following claim.

Proposition 4.2.2. *Suppose* $\mathbf{P}_1 = (\widetilde{\mathcal{M}}_1, g_1, f)$ *and* $\mathbf{P}_2 = (\widetilde{\mathcal{M}}_2, g_2, f)$ *are generic Euclidean programs extending* $\mathcal{M}$ *such that* $\mathcal{M}_1/g_1 = \mathcal{M}_2/g_2$. *Then there is an equivalence of categories*

$$D(A(\mathbf{P}_1, U)) \simeq D(A(\mathbf{P}_2, U)).$$

Before giving a proof of this Proposition, we will use it to deduce Theorem 1.4.3.

Proof that Proposition 4.2.2 implies Theorem 1.4.3. Under the assumptions of Theorem 1.4.3, $\widetilde{\mathcal{M}}_{\text{mid}}/g_2 = \widetilde{\mathcal{M}}_1/g_1$. Then by Proposition 4.2.2 it follows that

$$D(A(\mathbf{P}_1, U)) \simeq D(A(\mathbf{P}_{\text{mid}}, U). \tag{4.2.1}$$

On the other side, duality together with the assumptions of Theorem 1.4.3 give:

$$\widetilde{\mathcal{M}}^{\vee}_{\text{mid}}/f_1 = (\widetilde{\mathcal{M}}_{\text{mid}} \backslash f_1)^{\vee} = (\widetilde{\mathcal{M}}_2 \backslash f_2)^{\vee} = \widetilde{\mathcal{M}}^{\vee}_2/f_2.$$

Viewing f_1 and f_2 as playing the role of g in the Euclidean programs $\widetilde{\mathcal{M}}^{\vee}_{\text{mid}}$ and $\widetilde{\mathcal{M}}^{\vee}_2$ respectively, we can again apply Proposition 4.2.2 to find:

$$D(A(\mathbf{P}^{\vee}_2, U^{\perp})) \simeq D(A(\mathbf{P}^{\vee}_{\text{mid}}, U^{\perp})). \tag{4.2.2}$$

Putting these equivalences (4.2.1) and (4.2.2) together with the equivalences from Koszul duality:

$$D(A(\mathbf{P}_{\text{mid}}, U)) \simeq D(A(\mathbf{P}^{\vee}_{\text{mid}}, U^{\perp})) \quad \text{and} \quad D(A(\mathbf{P}_2, U)) \simeq D(A(\mathbf{P}^{\vee}_2, U^{\perp})),$$

gives the desired result:

$$D(A(\mathbf{P}_1, U)) \simeq D(A(\mathbf{P}_2, U)).$$

□

4.2.1 The definition and properties of the functor

It remains to prove Proposition 4.2.2. For the remainder of the paper we will let

$$\mathbf{P}_1 = (\widetilde{\mathcal{M}}_1, g_1, f) \quad \text{and} \quad \mathbf{P}_2 = (\widetilde{\mathcal{M}}_2, g_2, f)$$

be two Euclidean generic oriented matroid programs such that $\widetilde{\mathcal{M}}_1/g_1 = \widetilde{\mathcal{M}}_2/g_2$.

For $\ell = 1, 2$ let $A_\ell = A(\mathbf{P}_\ell, U)$, $B_\ell = B(\mathbf{P}_\ell, U)$ and $\mathcal{P}_\ell$ be the set of bounded, feasible sign vectors of $\mathbf{P}_\ell$. Note that the set of bounded sign vectors $\mathcal{B}$ is the same for $\mathbf{P}_1$ and $\mathbf{P}_2$.

As in [BLPW10, Section 6], the desired equivalence will come from a derived tensor product with a certain bimodule N.

It is slightly easier to define the bimodule N on the B-side, using the isomorphisms $A_\ell \simeq B_\ell^\vee$ of Theorem 3.3.13 for $\ell = 1, 2$. Namely, let

$$N = \bigoplus_{(\alpha,\beta)\in\mathcal{P}_1\times\mathcal{P}_2} R^\vee_{\alpha\beta}[-d_{\alpha\beta}]$$

with the natural left $B_1^\vee$-action and right $B_2^\vee$-action given by the $\star$ operation.

To translate this to the A-side, recall the following the alternative definition of A from Section 3.1.3,

$$A(\mathbf{P}, U) = e_\mathcal{P} D_n e_\mathcal{P} / \langle e_f e_\mathcal{P} \rangle + \langle \vartheta(U^\perp) e_\mathcal{P} \rangle.$$

We consider an extended version of A that only depends on $\widetilde{\mathcal{M}}/g$ by replacing $e_\mathcal{P}$ by $e_\mathcal{B} = \sum_{\alpha\in\mathcal{B}} e_\alpha$. That is, let

$$A_{\text{ext}}(\mathbf{P}, U) = e_\mathcal{B} D_n e_\mathcal{B} / \langle e_f e_\mathcal{B} \rangle + \langle \vartheta(U^\perp) e_\mathcal{B} \rangle.$$

As $A_{\text{ext}}(\mathbf{P}, U)$ only depends on $\widetilde{\mathcal{M}}/g$ and we have assumed that $\widetilde{\mathcal{M}}_1/g_1 = \widetilde{\mathcal{M}}_2/g_2$, let

$$A_{\text{ext}} := A_{\text{ext}}(\mathbf{P}_1, U) = A_{\text{ext}}(\mathbf{P}_2, U).$$

When viewed as an (A_1, A_2)-bimodule, N can be described as

$$N = e_{g_1} A_{\text{ext}} e_{g_2},$$

where $e_{g_\ell} = \sum_{\alpha\in\mathcal{P}_\ell} e_\alpha$ for $\ell = 1, 2$.

To check that these definitions of N coincide, consider the graded vector space

$$B_{\text{ext}}(\mathbf{P}, U) = \bigoplus_{(\alpha,\beta)\in\mathcal{F}\times\mathcal{F}} R_{\alpha\beta}[-d_{\alpha\beta}],$$

made into an algebra via $\star$, as in the definition of $B(\mathbf{P}, U)$ from Section 3.3.2. Then an easy extension of the proof of Theorem 3.3.13 gives us the following lemma.

Lemma 4.2.3. *There is an isomorphism $A_{\text{ext}}(\mathbf{P}, U) \simeq B_{\text{ext}}(\mathbf{P}^\vee, U^\perp)$. Combining this isomorphism with the isomorphisms $A_\ell \simeq B_\ell^\vee$, we obtain an equivalence between the two definitions of N.*

We define the functor $\Phi : D(A_1) \to D(A_2)$ via

$$\Phi(M) = M \overset{L}{\otimes}_{A_1} N.$$

For $\ell = 1, 2$ and any $\alpha \in \mathcal{P}_\ell$, let P_α^ℓ and V_α^ℓ be the corresponding projective and standard A_ℓ-modules. Define $\nu : \mathcal{P}_1 \to \mathcal{P}_2$ to be the composition

$$\mathcal{P}_1 \xrightarrow{\mu_1^{-1}} \mathbb{B} \xrightarrow{\mu_2} \mathcal{P}_2.$$

Proposition 4.2.4. *If $\alpha \in \mathcal{P}_1 \cap \mathcal{P}_2$, then $\Phi(P_\alpha^1) = P_\alpha^2$.*

Proof. Consider the natural map

$$\Gamma : P_\alpha^2 = e_\alpha A_2 \to e_\alpha A_1 \otimes_{A_1} e_{g_1} A_{\text{ext}} e_{g_2} = P_\alpha^1 \otimes_{A_1} N = \Phi(P_\alpha^1)$$

taking e_α to $e_\alpha \otimes e_{g_1} e_{g_2}$. For paths p, q in the quiver Q with p only passing through nodes in $\mathcal{P}_1$, the equality of the simple tensors

$$e_\alpha p \otimes e_{g_1} q e_{g_2} = e_\alpha \otimes e_\alpha p e_{g_1} q e_{g_2} = e_\alpha \otimes e_{g_1} e_{g_2} \cdot e_\alpha p e_{g_1} q e_{g_2}$$

implies that Γ is an isomorphism. □

Remark 4.2.5. Note that the proposition above and its proof are valid without the assumption that $\mathbf{P}_1$ and $\mathbf{P}_2$ be Euclidean.

Lemma 4.2.6. *For any $\alpha \in \mathcal{P}_1$, the A_2-module $\Phi(P^1_\alpha)$ has a filtration with standard subquotients. For $b \in \mathbb{B}$, if $\alpha \in \mathcal{B}_b$ then the standard module $V^2_{\mu_2(b)}$ appears with multiplicity 1 in the associated graded, and otherwise it does not appear.*[1]

Proof. We have

$$\Phi(P^1_\alpha) = P^1_\alpha \otimes_{A_1} N = e_\alpha A_1 \otimes_{A_1} e_{g_1} A_{\mathrm{ext}} e_{g_2},$$

so elements of $\Phi(P^1_\alpha)$ can be represented as linear combinations of paths in $\mathcal{B}$ which begin at α and end at elements of $\mathcal{P}_2 = \mathcal{B} \cap \mathcal{F}_2$. For $\beta \in \mathcal{P}_2$, let $\Phi(P^1_\alpha)_\beta$ be the submodule generated by paths p such that β is the maximal element of $\mathcal{P}_2$ (with respect to the ordering $\leq_2$ on $\mathcal{P}_2$ coming from our Euclidean assumption on $\mathbf{P}_2$) through which p passes. Then let

$$\Phi(P^1_\alpha)_{>_2 \beta} = \bigcup_{\gamma >_2 \beta} \Phi(P^1_\alpha)_\gamma \quad \text{and} \quad \Phi(P^1_\alpha)_{\geq_2 \beta} = \bigcup_{\gamma \geq_2 \beta} \Phi(P^1_\alpha)_\gamma.$$

We then obtain a filtration

$$\Phi(P^1_\alpha) = \bigcup_{\beta \in \mathcal{P}_2} \Phi(P^1_\alpha)_{\geq_2 \beta}.$$

Suppose $b \in \mathbb{B}$ and let $\beta = \mu_2(b)$. It suffices to show that the quotient $\Phi(P^1_\alpha)_{\geq_2 \beta} / \Phi(P^1_\alpha)_{>_2 \beta}$ is isomorphic to V^2_β if $\alpha \in \mathcal{B}_b$, and is zero otherwise.

Our argument follows the proof of Theorem 4.1.6.

If $\alpha \notin \mathcal{B}_b$, then for some $i \in b$, $\alpha(i) \neq \beta(i)$. Thus any path from α to β can be represented by one passing through β^i. By Lemma 4.1.5, $\beta^i >_2 \beta$ and therefore

$$\Phi(P^1_\alpha)_{\geq_2 \beta} / \Phi(P^1_\alpha)_{>_2 \beta} \cong 0.$$

[1] Recall that the set $\mathcal{B}_b$ defined in Definition 2.1.20 only depends on $\widetilde{\mathcal{M}}/g$.

Otherwise, precomposition with a taut path from α to β gives a surjective map

$$V^2_\beta \twoheadrightarrow \Phi(P^1_\alpha)_{\geq_2 \beta}/\Phi(P^1_\alpha)_{>_2 \beta}.$$

Thus

$$\dim V^2_\beta \geq \dim \Phi(P^1_\alpha)_{\geq_2 \beta}/\Phi(P^1_\alpha)_{>_2 \beta}$$

and it suffices to show equality holds. After summing over all $\beta \in \mathcal{P}_2$:

$$\begin{aligned}\dim(\Phi(P^1_\alpha)) &= \sum_{\beta\in\mathcal{P}_2} \dim \Phi(P^1_\alpha)_{\geq_2 \beta}/\Phi(P^1_\alpha)_{>_2 \beta} \\ &\leq \sum_{\alpha\in\mathcal{B}_b} \dim V^2_{\mu_2(b)} = \#\{(\gamma, b) \in \mathcal{P}_2 \times \mathbb{B} \mid \alpha, \gamma \in \mathcal{B}_b\}.\end{aligned}$$

It suffices to show that equality holds after summing over all $\alpha \in \mathcal{P}_1$. As $N = \oplus_{\alpha\in\mathcal{P}_1}\Phi(P^1_\alpha)$, we have

$$\begin{aligned}\sum_{\alpha\in\mathcal{P}_1} \dim \Phi(P^1_\alpha) = \dim N &= \sum_{(\alpha,\gamma)\in\mathcal{P}_1\times\mathcal{P}_2} \dim R^\vee_{\alpha\gamma} \\ &= \#\{(\alpha, \gamma, b) \in \mathcal{P}_1 \times \mathcal{P}_2 \times \mathbb{B} \mid \alpha, \gamma \in \mathcal{B}_b\}.\end{aligned}$$

Here we are using Lemma 3.3.8 and Proposition 2.1.21 on each $(\alpha, \gamma) \in \mathcal{P}_1 \times \mathcal{P}_2$. □

Remark 4.2.7. Notice that Theorem 4.1.6 can be viewed as the special case $\mathbf{P}_1 = \mathbf{P}_2$, and does not require that $\mathbf{P}_1$ is Euclidean.

For non-Euclidean programs, we can still obtain the following result on the level of Grothendieck groups.

Proposition 4.2.8. *In the graded Grothendieck group of modules for A_2 and for any $\alpha \in \mathcal{P}_1$, we have*

$$[\Phi(P^1_\alpha)] = \sum_{\alpha\in\mathcal{B}_b} q^{d_{\alpha\mu_2(b)}}[V^2_{\mu_2(b)}].$$

Proof. This proof is nearly identical to that of Theorem 4.1.9. Let $\mathcal{A}^\vee$ be the affine space for both $\mathcal{P}_1^\vee$ and $\mathbf{P}_2^\vee$. For any $\beta \in \mathcal{P}_2$, we have isomorphisms of graded vector spaces $\Phi(P_\alpha^1)e_\beta = e_\alpha A_1 \otimes_{A_1} A_{\text{ext}} e_b \simeq R_{\alpha\beta}^\vee\langle -d_{\alpha\beta}\rangle$. Then

$$[\Phi(P_\alpha^1)] = \sum_{\beta\in\mathcal{P}_2} \text{grdim}(R_{\alpha\beta}^\vee)\cdot q^{d_{\alpha\beta}}[L_\beta^2].$$

By Corollary 4.1.8, we can express the graded dimension of $R_{\alpha\beta}^\vee$ as

$$\text{grdim } R_{\alpha\beta}^\vee = \sum_{i=0}^{d} \#\{\text{vertices of } \alpha\wedge\beta \text{ in } \mathcal{A}^\vee \text{ with } i \text{ outgoing edges}\}\cdot q^{2i},$$

where "outgoing" edges are taken with respect to $\mathbf{P}_2^\vee$.

Observe that by Proposition 2.1.21 the vertices of $\alpha\wedge\beta$ in $\mathcal{A}^\vee$ (i.e., common feasible cocircuits of both α and β) are in bijection with the bases $b \in \mathbb{B}$ such that $\alpha, \beta \in \mathcal{B}_b$, where this bijection assigns b to the cocircuit face $Y_{b^c} \in \mathcal{A}^\vee$ of $\alpha\wedge\beta$ such that $\underline{Y_{b^c}} = b$. For a fixed such b, let $\delta = \mu_2(b)$. Let $\gamma := (\alpha\wedge\beta)|_{\underline{\alpha\wedge\beta}}$, which is a bounded feasible tope in $\mathbf{P}_2^\vee/z(\alpha\wedge\beta)$.

Each edge of $\alpha\wedge\beta$ which either starts or ends at Y_b in $\mathbf{P}_2$ is nonzero on exactly one uniquely determined $i \in b^c$. The edge corresponding to i starts at Y_b if and only if $\delta(i) \neq (\alpha\wedge\beta)(i)$. Thus, the number of edges of γ starting at Y_{b^c} is equal to the distance between γ and the restriction $\overline{\delta} := \delta|_{\underline{\alpha\wedge\beta}}$. But $\underline{\alpha\wedge\beta} = \{i \in E_n \mid \alpha(i) = \beta(i)\}$ and so the distance between γ and $\overline{\delta}$ is equal to the cardinality of the difference set

$$S(\gamma,\overline{\delta}) = \{i \in E_n \mid \alpha(i) = \beta(i) \neq \delta(i)\} = S_{\alpha\beta}^\delta.$$

Rewriting the sum over bases $b \in \mathbb{B}$ such that $\alpha, \beta \in \mathcal{B}_b$ and using the formula $d_{\alpha\mu_2(b)} + d_{\beta\mu_2(b)} = d_{\alpha\beta} + 2|S_{\alpha\beta}^{\mu_2(b)}|$, we find:

$$\text{grdim } e_\alpha A e_\beta = q^{d_{\alpha\beta}}\text{grdim } R_{\alpha\beta}^\vee = \sum_{\alpha,\beta\in\mathcal{B}_b} q^{d_{\alpha\beta}+2|S_{\alpha\beta}^{\mu_2(b)}|} = \sum_{\alpha,\beta\in\mathcal{B}_b} q^{d_{\alpha\mu_2(b)}+d_{\beta\mu_2(b)}}. \tag{4.2.3}$$

Putting it all together,

$$[\Phi(P^1_\alpha)] = \sum_{\beta\in\mathcal{P}_2}\sum_{\alpha,\beta\in\mathcal{B}_b} q^{d_{\alpha\mu_2(b)}+d_{\beta\mu_2(b)}}[L^2_\beta]$$

$$= \sum_{\alpha\in\mathcal{B}_b} q^{d_{\alpha\mu_2(b)}} \sum_{\beta\in\mathcal{B}_b\cap\mathcal{P}_2} q^{d_{\beta\mu_2(b)}}[L^2_\beta]$$

$$= \sum_{\alpha\in\mathcal{B}_b} q^{d_{\alpha\mu_2(b)}}[V^2_{\mu_2(b)}],$$

as we wished to show. □

Proposition 4.2.9. *Φ induces an isomorphism of (ungraded) Grothendieck groups, regardless of Euclidean assumptions. If $\mathbf{P}_1$ is Euclidean and $\alpha \in \mathcal{P}_1$, we additionally have $[\Phi(V^1_\alpha)] = [V^2_{\nu(\alpha)}]$ in the Grothendieck group of right A_2-modules.*

Proof. For any $\alpha \in \mathcal{P}_1$, the equalities

$$\sum_{\alpha\in\mathcal{B}_b}[V^2_{\mu_2(b)}] = [\Phi(P^1_\alpha)] = \sum_{\alpha\in\mathcal{B}_b}[\Phi(V^1_{\mu_1(b)})]$$

follow from Proposition 4.2.8 and its special case Theorem 4.1.9.

To show that Φ induces an isomorphism, notice that both Grothendieck groups are free abelian groups of rank $|\mathbb{B}|$, where the Grothendieck group of right A_1-modules has basis $\{[P^1_\alpha]\}_{\alpha\in\mathcal{P}_1}$. We also have that $\{[V^2_\gamma]\}_{\gamma\in\mathcal{P}_2}$ is a basis for the Grothendieck group of right A_2-modules, since the matrix encoding the equalities $[V^2_\gamma] = \sum_{\beta\le\gamma}[L^2_\beta]$ is $X_2\,|_{q=1}$, where X_2 is the $\mathcal{P}_2 \times \mathcal{P}_2$ matrix whose (γ,β)-entry is given by

$$(X_2)_{\gamma,\beta} = \begin{cases} q^{d_{\gamma\beta}} & \text{if } \beta \trianglerighteq \gamma \\ 0 & \text{otherwise.} \end{cases}$$

This matrix was proved to be invertible in Theorem 4.1.14, and remains invertible when evaluated at $q = 1$. The $\mathbb{B} \times \mathcal{P}_1$ matrix whose α-th column has entries 1 if $\alpha \in \mathcal{B}_b$ and 0

otherwise is the matrix expressing the equalities $[\Phi(P^1_\alpha)] = \sum_{\alpha\in\mathcal{B}_b}[V^2_{\mu_2(b)}]$. This matrix is also expresses the equalities $[P^1_\alpha] = \sum_{\alpha\in\mathcal{B}_b}[V^1_{\mu_1(b)}]$, and is therefore given under the bijection $\mu_1 : \mathbb{B} \to \mathcal{P}_1$ by $X_1^T\,|_{q=1}$, for X_1 the $\mathcal{P}_1 \times \mathcal{P}_1$ matrix analogue of X_2 for $\mathbf{P}_1$. Thus, X_1 is invertible, and remains say when evaluated at $q = 1$.

If $\mathbf{P}_1$ is Euclidean, then the second claim follows by induction on the opposite poset of $\mathcal{P}_1$. The base case is for $\alpha \in \mathcal{P}_1$ maximal, so $P^1_\alpha = V^1_\alpha$ and $\Phi(P^1_\alpha) = V^2_{\nu(\alpha)}$. □

Proposition 4.2.10. *Let $\alpha \in \mathcal{P}_1$. Then $\Phi(V^1_\alpha)$ is the quotient of $\Phi(P^1_\alpha)$ by the submodule generated by all paths changing their i-th coordinate for some $i \in \mu_1^{-1}(\alpha)$. In particular, $\mathrm{Tor}^{A_1}_m(V^1_\alpha, N) = 0$ for all $m > 0$.*

Proof. Apply Φ to the linear projective resolution of V^1_α of Theorem 4.1.12. The degree zero homology of the resulting complex is the quotient promised. We wish to show that the resulting complex is a resolution of $V^1_\alpha \otimes_{A_1} N$. This claim follows from argument is analogous to the proof of Theorem 4.1.12, where for each $S \subset \mu_1^{-1}(\alpha)$ we filter each A_2-module $\Phi(P^1_{\alpha^S}) = P^1_\alpha \otimes_{A_1} N$ by standards as in the proof of Lemma 4.2.6. □

Corollary 4.2.11. *If a right A_1-module M admits a filtration by standard modules, we have $\mathrm{Tor}^{A_1}_m(M, N) = 0$ and therefore*

$$\Phi(M) = M \otimes_{A_1} N.$$

4.2.2 Ringel duality and composition of functors

Suppose that $\mathbf{P} = (\widetilde{\mathcal{M}}, g, f)$ is a generic Euclidean extension of $\mathcal{M}$. Consider the program $\overline{\mathbf{P}} = {}_g\mathbf{P} = ({}_g\widetilde{\mathcal{M}}, -g, f)$ obtained from $\mathbf{P}$ by reorientation of g. In other words $\overline{\mathbf{P}}$ is the program on the oriented matroid $\mathcal{M}$ with feasible cocircuits equal to the negative of the

feasible cocircuits of $\mathbf{P}$. This program is also generic and Euclidean. We let $\overline{A} = A(\overline{\mathbf{P}}, U)$ and denote by

$$\Phi^- : D(A) \to D(\overline{A}),$$

the functor Φ for $\mathbf{P}_1 = \mathbf{P}$ and $\mathbf{P}_2 = \overline{\mathbf{P}}$. We will prove that Φ^- is an equivalence and relate it to Ringel duality.

Theorem 4.2.12. *Φ^- is an equivalence, the algebras A and $\overline{A}$ are Ringel dual, and the Ringel duality functor is $d \circ \Phi^- = \Phi^- \circ d$.*

Proof. Notice that for any $\alpha \in \mathcal{P}$, we have that the B-side description of Φ^- gives

$$\Phi^-(P_\alpha) = \bigoplus_{\overline{\beta} \in \overline{\mathcal{P}}} R^\vee_{\alpha\overline{\beta}}[-d_{\alpha\overline{\beta}}].$$

Then we have that the tope corresponding to the restrictions of both α and $\overline{\beta}$ in the oriented matroid $\widetilde{\mathcal{M}}^\vee / S^{\overline{\beta}}_{\alpha\alpha}$ on $E_n \cup \{f\}$ has all cocircuit faces taking the value $+$ on f. As in the proof of Theorem 4.1.17, this implies that $R^\vee_{\alpha\overline{\beta}}$ is Gorenstein and $\Phi^-(P_\alpha)$ is self-dual. By Lemma 4.2.6 it follows that $\Phi^-(P_\alpha)$ is tilting.

It remains to show that Φ^- is an equivalence. With Proposition 4.2.4 and Theorem 4.1.17 in hand, one can repeat the proof of [BLPW10, Theorem 6.10] word for word. □

We now study the composition of functors. Let $\mathbf{P}_1 = (\widetilde{\mathcal{M}}_1, g_1, f)$, $\mathbf{P}_2 = (\widetilde{\mathcal{M}}_2, g_2, f)$, and $\mathbf{P}_3 = (\widetilde{\mathcal{M}}_3, g_3, f)$ be generic Euclidean programs extending $\mathcal{M}$ for which $\widetilde{\mathcal{M}}_1/g_1 = \widetilde{\mathcal{M}}_2/g_2 = \widetilde{\mathcal{M}}_3/g_3$. We can then define the three functors

$$D(A_1) \xrightarrow{\Phi_{12}} D(A_2) \xrightarrow{\Phi_{23}} D(A_3) \quad \text{and} \quad D(A_1) \xrightarrow{\Phi_{13}} D(A_3)$$

as before. We would like to compare Φ_{13} with the composition $\Phi_{23} \circ \Phi_{12}$.

Notice that

$$N_{12} = \Phi_{12}(A_1) = \bigoplus_{\alpha \in \mathcal{P}_1} \Phi_{12}(P_\alpha^1)$$

has a filtration by standard modules as a right A_2-module by Lemma 4.2.6. Then

$$\Phi_{23} \circ \Phi_{12}(M) = (M \overset{L}{\otimes}_{A_1} N_{12}) \overset{L}{\otimes}_{A_2} N_{23} = M \overset{L}{\otimes}_{A_1} (N_{12} \otimes_{A_2} N_{23})$$

by Corollary 4.2.11. The natural map $N_{12} \otimes_{A_2} N_{23} \to N_{13}$ given by concatenation of paths induces a natural transformation $\Phi_{23} \circ \Phi_{12} \to \Phi_{13}$. We also have that $\Phi_{23} \circ \Phi_{12}$ and Φ_{13} induce the same map on Grothendieck groups by Proposition 4.2.9. This implies that

$$\dim_k N_{12} \otimes_{A_2} N_{23} = \dim_k N_{13}$$

since the classes of $N_{12} \otimes_{A_2} N_{23}$ and N_{13} agree in the Grothendieck group of A_3-modules.

We now combine this discussion with the equivalence we have already proved. Let $\mathbf{P}_3 = \overline{\mathbf{P}_1}$, so that

$$\Phi_{13} = \Phi^- : D(A_1) \to D(\overline{A_1}) \quad \text{and} \quad \Phi_{23} = \Phi_{2\bar{1}} : D(A_2) \to D(\overline{A_1}).$$

Lemma 4.2.13. $\Phi^- \cong \Phi_{2\bar{1}} \circ \Phi_{12}$.

Proof. The conclusion follows from the discussion above if we can show that the map $N_{12} \otimes_{A_2} N_{2\bar{1}} \to N_{1\bar{1}}$ is an isomorphism. We have already observed that the source and target have the same dimension, so it will suffice to show that this map is surjective. This means showing that for any $(\alpha, \beta) \in \mathcal{P}_1 \times \overline{\mathcal{P}_1}$, every path from α to β in $e_\alpha A_{\text{ext}} e_\beta$ can be represented as a path that passes through some sign vector γ in $\mathcal{P}_2$. It suffices to do this for a taut path from α to β. Translating this to the B-side, we wish to show that $1^\vee_{\alpha\beta} = 1^\vee_{\alpha\gamma} \star 1^\vee_{\gamma\beta}$ for some $\gamma \in \mathcal{P}_2$.

Let $(\alpha, \beta) \in \mathcal{P}_1 \times \overline{\mathcal{P}_1}$ and suppose that $R^\vee_{\alpha\beta}$ is nonzero. This means the maximal common covector face $T^\vee_\alpha \wedge T^\vee_\beta$ of the topes $T^\vee_\alpha$ and $T^\vee_\beta$ in

$$\widetilde{\mathcal{M}}^\vee := \widetilde{\mathcal{M}}^\vee_1 \backslash g_1 = (\widetilde{\mathcal{M}}_1 / g_1)^\vee = (\widetilde{\mathcal{M}}_{\overline{1}} / g_{\overline{1}})^\vee = \widetilde{\mathcal{M}}^\vee_{\overline{1}} \backslash g_{\overline{1}}$$

is nonzero and all of its nonzero cocircuit faces are feasible, i.e. they take the value + on f. Together with the fact that $z(T^\vee_\alpha \wedge T^\vee_\beta) = S^\beta_{\alpha\alpha}$, this implies that $T^\vee_\alpha \wedge T^\vee_\beta$ restricts to a bounded feasible tope in the oriented matroid program

$$\mathbf{P}_2^\vee / S^\beta_{\alpha\alpha} = (\widetilde{\mathcal{M}}^\vee / S^\beta_{\alpha\alpha}, f, g_2).$$

Let Y be the optimal cocircuit face of $T^\vee_\alpha \wedge T^\vee_\beta$ viewed as a covector in $\mathbf{P}_2^\vee / S^\beta_{\alpha\alpha}$. Then Y lifts to a unique feasible cocircuit of $\mathbf{P}_2^\vee$ and let $\gamma \in \mathcal{P}_2^\vee = \mathcal{P}_2$ be the corresponding sign vector. By construction, $T^\vee_\alpha \wedge T^\vee_\beta$ is a face of $T^\vee_\gamma$.

Thus $R^\vee_{\alpha\gamma\beta} = R^\vee_{\alpha\beta}, S^\gamma_{\alpha\beta} = \varnothing$ and

$$1^\vee_{\alpha\beta} = 1^\vee_{\alpha\gamma} \star 1^\vee_{\gamma\beta}.$$

□

Proof of Proposition 4.2.2. We can set up everything as in Lemma 4.2.13, and we know that Φ^- is an equivalence by Theorem 4.2.12. This gives us that $\Phi_{12} : D(A_1) \to D(A_2)$ is faithful while $\Phi_{2\overline{1}}$ is full and essentially surjective.

Note that $\overline{\mathbf{P}_2}$ is Euclidean if $\mathbf{P}_2$ is Euclidean. Appealing to the same arguments as before, $\Phi_{12} \circ \Phi_{\overline{2}1}$ is an equivalence. We conclude that Φ_{12} is also full and essentially surjective and thus an equivalence. □

Part II

Combinatorial Intersection Homology for Gorenstein* Meet-Semilattices

Chapter 5

Introduction

In this paper, we define an analogue of equivariant intersection cohomology for toric varieties in the setting of certain posets with well-behaved homology. Our analogue is based off of the combinatorial intersection cohomology for fans, hereafter referred to as IHF, developed in [BBFK00],[BL00]. IHF depends on the fact that fans are realized in real vector spaces. The key philosophy of this paper is to define a sheaf on a poset which keeps the key linear algebraic properties of such a realization, even when the poset cannot be realized as the face poset of a fan.

5.1 Historical motivation

Recall that a *(pure) fan* in a real vector space V is a collection Σ of pointed cones $\sigma \subset V$, closed under intersections and taking faces, such that every maximal cone has the same dimension as V. A fan Σ is *simplicial* if every cone $\sigma \in \Sigma$ has $\dim \sigma$ ray faces, where a ray is a 1-dimensional cone in Σ. A *rational* fan is a fan in V with an isomorphism $V \simeq \mathbb{R}^d$

such that the image of every ray of Σ intersects the lattice $\mathbb{Z}^d \backslash \{0\}$. Every rational fan Σ give rise to a toric variety X_Σ.

The *sheaf of conewise polynomials* $\mathcal{A}_\Sigma$ of Σ is the sheaf on the underlying poset $P(\Sigma)$ of the fan which takes every cone to the set of polynomial functions on the real span of that cone, with obvious restrictions. If Σ is simplicial, then $\mathcal{A}_\Sigma$ is a combinatorial invariant and its space of global sections is isomorphic to the face ring of the corresponding simplicial complex. If Σ is simplicial and rational, then $\mathcal{A}_\Sigma(\Sigma)$ is also isomorphic to the equivariant cohomology ring of X_Σ.

On the other hand, if Σ is rational but not simplicial, then $\mathcal{A}_\Sigma$ is not a combinatorial invariant and its global sections no longer compute the equivariant cohomology of X_Σ. In this case, X_Σ is no longer rationally smooth, but it is still stratified by torus orbits. Therefore one may consider the equivariant intersection cohomology sheaf on X_Σ, which can be defined using only the fan Σ. For any Σ, IHF then provides a *minimal extension sheaf* $\mathcal{E}_\Sigma$, coinciding with the equivariant intersection cohomology sheaf in the rational case, using only the data of Σ.

Thus, IHF combinatorially adapts ideas from the (equivariant) cohomology of toric varieties to more general fans, including the central fans of nonrational polytopes. This theory enjoys its own Kähler package, which has a few versions of Poincaré duality developed in [BBFK00], [BL00], [BL03], [BBFK05], and a hard Lefschetz theorem originally proved in [Kar01].

For more details on IHF, we recommend the beautiful survey by Braden [Bra05].

5.1.1 Toric h-polynomials

Inspired by the relationship between cohomology of simplicial toric varieties and the h-polynomials of the corresponding fans, Stanley introduced and studied the *toric h-polynomial* of a lower Eulerian poset in [Sta87]. We briefly recall the recursive definition of these polynomials; poset terminology is recalled in Section 6.

Let $\widehat{P}$ be an Eulerian (hence graded) poset, where removing the maximal element $\hat{1}$ results in P. Then we define polynomials $h(P,x)$ and $g(P,x)$ by the following rules:

1. $h(\varnothing,x) = g(\varnothing,x) = 1$,

2. If $\widehat{P}$ has rank $d+1 \geq 1$ and $h(P,x) = h_0 + h_1 x + \cdots$, then

$$g(P,x) = \sum_{i=0}^{\lfloor d/2 \rfloor} (h_i - h_{i-1}) x^i, \tag{5.1.1}$$

   where we set $h_{i-1} = 0$,

3. If $\widehat{P}$ has rank $d+1 \geq 1$, then

$$x^d h(P,x^{-1}) = \sum_{t \in P} g(\partial t, x)(x-1)^{d-\rho(t)}, \tag{5.1.2}$$

   where $\partial t := \{s \in P \mid s < t\}$.

If P is lower Eulerian with a chain of maximal length d, then define $h(P,x)$ by formula (5.1.2).

When $\widehat{P}$ is the face poset of a rational convex polytope[1], then $h(P,x)$ is the Poincaré polynomial of the intersection cohomology of the corresponding toric variety. If we remove the word rational, then we get a similar result using IHF and Karu's hard

[1]For us, face posets always include the empty face.

Lefschetz theorem: In this case, the central fan Σ of the polytope is *quasiconvex*, meaning that space of global sections $\mathcal{E}_\Sigma(\Sigma)$ is a free graded module over $\operatorname{Sym} V^*$ of finite rank, and $h(P,x)$ is the Poincaré polynomial of $\overline{\mathcal{E}_\Sigma(\Sigma)} = \mathcal{E}_\Sigma(\Sigma)/(\operatorname{Sym} V^*)^+\mathcal{E}_\Sigma(\Sigma)$.

More generally, the following conjecture of Stanley motivates this paper.

Conjecture 5.1.1 (([Sta87], see also [Sta94])). *If P is a Cohen-Macaulay lower Eulerian meet-semilattice, then the toric h-polynomial of P has nonnegative coefficients.*

Here is a quick summary of what is known so far. Let $h(P,x) = h_0 + h_1 x + \cdots + h_d x^d$, for P a Cohen-Macaulay lower Eulerian meet-semilattice. Then Stanley [Sta87] already knew that $h_0 = 1$, $h_1 \geq 0$ and $h_d \geq 0$. Athanasiadis [Ath11] also proved $h_{d-1} \geq 0$ by working with the Möbius function. If P is the face poset of a simplicial complex Δ and is also obtained from a Gorenstein* poset $\widehat{P}$ by removing the maximal element, then Adiprasito-Papadakis-Petrotou [APP21] proved this conjecture for $\widehat{P}$ by showing that there is a hard Lefschetz theorem for the face ring of Δ. Outside of face posets of convex polytopes and these developments, the conjecture is still open.

5.2 Separating the fan from its realization

We now fix an infinite field $\mathbf{k}$. We wish to adapt the ideas of IHF to prove Conjecture 5.1.1 in general. However, we are immediately hindered by the dependence of IHF on the realization of a polytope in a real vector space.

The key philosophy behind this paper is that we can generalize the theory of IHF by separating the fan from its realization. To this end, given a poset[2] P like those appearing

[2]We assume all posets in this paper have a unique minimal element, and are in fact lower graded and lower thin (see Section 6).

in Conjecture 5.1.1, we define a sheaf of vector spaces (over $\mathbf{k}$) $\mathcal{U}$ on P called a *linearization sheaf* (see Definition 7.1.1). In the analogy with the fan case, P plays the role of the face poset of a fan, while $\mathcal{U}$ sends every element of P to the analogue of the linear functions on the span of a cone. In order to have a linearization sheaf, it is necessary for P to be a *meet-semilattice*, meaning the greatest lower bound of any two elements of P exists.

We then associate to such a pair $(P, \mathcal{U})$ the analogue $\mathcal{A} = \mathcal{A}_{(P,\mathcal{U})}$ of the sheaf of conewise polynomials, which we call the *sheaf of facewise polynomials* of $(P, \mathcal{U})$. We use the term "facewise" because we intuitively think of our posets as face posets of regular cell complexes, although no such complex might exist.

We study a certain category of $\mathcal{A}$-modules, called *pure sheaves*, the simple objects of which are in bijection with the elements of P, up to grade shift and isomorphism. The simple pure sheaf corresponding to the minimal element of P is called the *minimal extension sheaf* of $(P, \mathcal{U})$, and denoted by $\mathcal{E} = \mathcal{E}_{(P,\mathcal{U})}$. This $\mathcal{E}$ is the analogue of the equivariant intersection homology sheaf from the toric setting.

To state the main theorems of the paper, we need two more notions. First, we say P is *Gorenstein** (over $\mathbf{k}$) if P is graded and for any interval $[s, t] \subset P$, the order complex of (s, t) has the same homology as a sphere. Intuitively, we think of Gorenstein* meet-semilattices with the maximal element removed as the face posets of regular cellular spheres such that the intersection of any two cells is a cell.

We also make a choice of a *homogenization* U of $\mathcal{U}$ (see Definition 7.1.8). Homogenizations play the role of V^* from the fan example, and are also related to linear systems of parameters for the face rings of simplicial complexes. Our main theorems are as follows.

Theorem 5.2.1. *Let P be a Gorenstein* meet-semilattice of rank d with the maximal element removed. Then fixing a linearization sheaf $\mathcal{U}$ and a homogenization U of $\mathcal{U}$ we have the following:*

1. *The space of global sections $\mathcal{E}(P)$ is a free graded $\operatorname{Sym} U$-module of finite rank (see Theorem 8.1.5).*

2. *There is a decomposition theorem with respect to our notion of subdivision (see Theorem 8.2.3).*

3. *There is a duality functor $\mathbb{D}$ on the category of pure sheaves for $(P, \mathcal{U})$ such that $\mathbb{D}\mathcal{E} \simeq \mathcal{E}$.*

4. *The duality induces an intersection product*

$$\mathcal{E}(P) \times \mathcal{E}(P) \to \operatorname{Sym} U[d]$$

   where $[m]$ denotes the grade shift up by m (Section 9.1).

5. *Combining (1) and (3), we obtain a dual pairing of finite-dimensional graded vector spaces*

$$\overline{\mathcal{E}(P)} \times \overline{\mathcal{E}(P)} \to \mathbf{k}[d],$$

   where $\overline{\mathcal{E}(P)} \coloneqq \mathcal{E}(P)/(\operatorname{Sym} U)^{+}\mathcal{E}(P)$.

To prove this theorem, we rely on an analogue of cellular homology for posets, developed in [ET15]. It is impossible to overstate how important this cellular poset cohomology is for our paper, since our adaptation of it forms the backbone for all of the pieces of Theorem 5.2.1.

More generally, we work with the posets featured in Stanley's conjecture, namely Cohen-Macaulay lower Eulerian meet-semilattices. Examples include the face poset of

- any connected simple graph, and
- any regular cell decomposition of a $\mathbf{k}$-homology sphere with the intersection property.

The latter of these includes the boundary complex of a convex polytope, oriented matroid polytopes, and complete fans which may or may not admit strictly convex linear functions.

If P is the face poset of a simplicial complex, then we define a standard linearization sheaf $\mathcal{S}$ on P in Example 7.1.4. Then $\mathcal{A}(P) = \mathcal{E}(P)$ is the face ring of that simplicial complex and a homogenization U corresponds to a linear system of parameters for that face ring.

On the other hand, if P comes from a fan in a vector space V, then we get a natural choice for a linearization sheaf (with $\mathbf{k} = \mathbb{R}$) by sending every cone to the linear functions on its span. With this choice of linearization sheaf and the corresponding homogenization V^*, our theory reduces to that of IHF.

Finally, we make a hard Lefschetz conjecture for this setting in Section 9.2, and we discuss how Stanley's conjecture would follow as a corollary.

5.3 Outline of the paper

In Section 6, we review the concepts we need for our posets and sheaves, as well as adapt the cellular poset cohomology of [ET15] to our setting. In Section 7, we define linearization sheaves, as well as the associated sheaf of facewise polynomials $\mathcal{A}$, pure sheaves, and the minimal extension sheaf $\mathcal{E}$. In Section 8, we find some sufficient conditions on P

to ensure that $\mathcal{E}(P)$ is a free $\operatorname{Sym} U$-module, for U a homogenization of $\mathfrak{U}$. We then apply these conditions to prove a decomposition theorem with respect to our notion subdivision. We restrict to subthin posets in Section 9, where we define and study a duality of pure sheaves which eventually results in an intersection product for $\mathcal{E}$. We then discuss a hard Lefschetz conjecture in our setting, and relate it to Stanley's conjecture.

Chapter 6

Posets and Sheaves

A poset is a finite set P together with a partial ordering $\leq$ of its elements. If $s < t$ and $s \leq r \leq t$ implies $r = s$ or $r = t$, then we write $s \lessdot t$, and say that t is a *cover* of s and s is a *cocover* of t. We use standard notation for intervals in P. For any $t \in P$, we define its *basic open set* $[t] \coloneqq (-\infty, t]$, its *boundary* $\partial t \coloneqq (-\infty, t)$, and its *star* $\star(t) \coloneqq [t, \infty)$.

We say that a poset P is *lower graded* if P has a unique minimal element $\hat{0}$ and a rank function $\rho : P \to \mathbb{Z}_{\geq 0}$ such that $\rho(\hat{0}) = 0$ and if $s \lessdot t$ then $\rho(t) = \rho(s) + 1$. For any $s < t$ in a lower graded poset P, we define $\rho(s, t) \coloneqq \rho(t) - \rho(s)$. We define $P_\ell = \rho^{-1}(\ell)$ for all $\ell \in \mathbb{Z}$, while $P_{\leq \ell}, P_{< \ell}, P_{\geq \ell}, P_{> \ell}$ all get obvious meanings. We call the elements of P_1 the *atoms* of P. The *rank* of P is the maximal $\ell \in \mathbb{Z}$ such that $P_\ell \neq \varnothing$ and $P_{> \ell} = \varnothing$. We say that P is *pure* if every *facet* – that is, every maximal element – has the same rank.

Suppose P is lower graded. Then we say that P is *lower thin* if whenever $r < t$ and $\rho(r, t) = 2$, then there are exactly two elements in (r, t). If P is lower thin, then a *diamond* in P is an interval $[r, t]$ with $r < t$ and $\rho(r, t) = 2$.

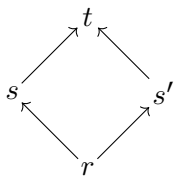


Figure 6.1: The Hasse diagram of a diamond $[r,t] = \{r,s,s',t\}$.

A lower thin P is *subthin* if it is pure and every element of rank 1 less than the rank of P is covered by at most two facets. For a subthin P of rank d, we define ∂P to be the union of all $[s]$ such that $s \in P_{d-1}$ and s has exactly one cover.

We say a lower graded poset is *graded* if it has a unique maximal element. Given a pure lower graded poset P, we define $\widehat{P}$ to be the graded poset on $P \sqcup \{\hat{1}\}$ obtained by adjoining a new element $\hat{1}$ which covers every facet of P. A *thin* poset is a poset which is graded and lower thin. A lower graded pure poset has $\widehat{P}$ thin if and only if P is subthin and $\partial P = \varnothing$.

The *meet* of any two elements $s,t \in P$ is the greatest lower bound of s and t in P, and is denoted $s \wedge t$. Dually, the *join* $s \vee t$ of $s,t \in P$ is the least upper bound of s and t. Neither meets nor joins need to exist between two elements in a poset, but if a meet or join exists then it must be unique. A *meet-semilattice* is a lower graded poset P such that $s \wedge t$ exists for any $s,t \in P$. We say that P is a *lattice* if P is a graded meet-semilattice such that $s \vee t$ also exists for any $s,t \in P$. By [SF99, Proposition 3.3.1], a graded (finite) meet-semilattice is a lattice.

We are concerned with lower thin pure posets throughout. For such P and for any $s \in P$, the basic open, boundary, and star of s are all pure, lower graded, and lower thin as well, where the rank function of $\star(s)$ is $\rho_{\star(s)}(t) \coloneqq \rho(s,t)$.

6.1 Topological properties of posets

We will often be concerned with topological properties of our (pure, lower thin) posets. We fix an infinite field $\mathbf{k}$ throughout the rest of the paper.

For any poset P, we define $\operatorname{ord}(P)$ to be the *order complex* of P, which is a simplicial complex on the set P consisting of all chains in P. We denote the geometric realization of P by $\|P\| := \|\operatorname{ord}(P_{>0})\|$, where we delete the minimal element so that our realizations are not automatically contractible. As a result, any simplicial homology we mention using $\operatorname{ord}(P)$ is naturally the reduced homology of $\|P\|$ with coefficients in $\mathbf{k}$.

We say that a graded poset P is *Eulerian* if $[s,t]$ has as many elements of even rank as odd rank for any $s < t$. Equivalently, by Hall's theorem (see [SF99, Proposition 3.8.5]), we require the Euler characteristic of $\operatorname{ord}([s,t))$ (or equivalently, the reduced Euler characteristic of $\operatorname{ord}((s,t))$) to be $(-1)^{\rho(s,t)}$ for any $s < t$. Notice that any Eulerian poset is thin.

We say that a (pure lower graded) P is *Cohen-Macaulay* (over $\mathbf{k}$) if $\operatorname{ord}([s,t))$ has trivial homology over $\mathbf{k}$ in all degrees less than $\rho(s,t) - 2$ for any $s < t$ in $\widehat{P}$. We will abbreviate Cohen-Macaulay over $\mathbf{k}$ by CM/$\mathbf{k}$, or simply CM, since our field $\mathbf{k}$ is fixed throughout. By a result of Munkres (for instance, see [Bac80, Proposition 3.4]), CM is a topological property, meaning that if P is Cohen-Macaulay and Q is another poset such that $\|Q\|$ is homeomorphic to $\|P\|$, then Q is Cohen-Macaulay.

A poset P is *Gorenstein** (over $\mathbf{k}$) if P is both CM (over $\mathbf{k}$) and Eulerian (and hence graded). Again, the Gorenstein* property depends on $\mathbf{k}$, but we omit it from the notation.

Example 6.1.1. Let P be the face poset of a regular cell complex, where *regular* means that $\|\partial t\|$ is homeomorphic to a $(\rho(t)-2)$-sphere for all $t \in P$. If $\|P\|$ is (homeomorphic to) a sphere, then $\widehat{P}$ is Gorenstein*. If $\|P\|$ is a ball, $\mathbf{k}$-homology sphere, or connected graph, then P is CM.

We say that P is *lower Eulerian* (respectively *lower CM* or *lower Gorenstein**) if for every $t \in P$, the basic open $[t]$ is Eulerian (respectively CM or Gorenstein*). Notice that a poset which is graded and lower Gorenstein* is necessarily Gorenstein*, but a poset which is not graded cannot be Eulerian or Gorenstein*.

Example 6.1.2. The face poset of any regular cell complex is lower Gorenstein*.

6.1.1 Dually connected and diamond transitivity

Let P be graded of rank d and thin. Given a maximal chain $\hat{0} \lessdot s_1 \lessdot \cdots \lessdot s_d \in \mathrm{ord}(P)$, we say we can reach another maximal chain $\hat{0} \lessdot s_1' \lessdot \cdots \lessdot s_d'$ with a *diamond move* if $s_i' = s_i$ for all but one index $1 \leq \ell < d$. If we can reach any maximal chain in P from any other using diamond moves, we say P is *diamond transitive*. As usual, if $[t]$ is diamond transitive for all $t \in P$ then P is *lower diamond transitive*.

Remark 6.1.3. This definition of diamond transitive comes from [Cha19]. This definition is also related to the *compatible families* as described in [ET15], but our scope will be much less general than theirs since all of our posets are lower thin.

A pure poset P is *dually connected* if for any two facets t, t', we can find elements $s_1, \ldots, s_\ell \in P_{d-1}$ and $t_1, \cdots, t_{\ell-1} \in P_d$ such that $t \gtrdot s_1 \lessdot t_1 \gtrdot s_2 \lessdot \cdots \lessdot t_{\ell-1} \gtrdot s_\ell \lessdot t$. We say P is *star normal* if $\star(s)$ is dually connected for any $s \in P$.

Remark 6.1.4. Recall that a pseudomanifold (with or without boundary) is a topological space together with a triangulation whose face poset is pure, subthin, and dually connected. Then a pseudomanifold is *normal* if the link of any simplex of a certain dimension is also a triangulation of a pseudomanifold. Our definition of star normal related to this notion of normal, but emphasizes stars rather than links, and also applies to posets whose order complexes are not necessarily subthin.

Proposition 6.1.5. *Any lower thin CM poset P is star normal and lower diamond transitive.*

Proof. Dually connected follows from the fact that CM posets remain CM after deleting all elements of a given rank (see [Bac80]), together with the fact that CM and connected are equivalent for graphs. Since $\star(s)$ remains lower thin and CM for any $s \in P$, we have P is star normal. Finally, for any $t \in P_{>0}$, $[t]$ is diamond transitive by induction on rank, since $[\hat{0}, t)$ is also CM and lower thin and P is dually connected. □

6.2 Sheaves on posets

Any poset P has a topology whose basic open sets are unsurprisingly the basic opens $\{[t]\}_{t \in P}$, and the basic closed sets are stars $\{\star(s)\}_{s \in P}$. In particular, the open sets are exactly the order ideals. If we think of P as a face poset of a cell complex, then $[t]$ plays the role of a closed cell, and open subsets play the part of subcomplexes.

We also view P as a category, with objects $s \in P$ and a unique arrow $s \to t$ if $s \leq t$. A *sheaf* on P is a (covariant) functor $P^{\mathrm{op}} \to \mathcal{C}$, where $\mathcal{C}$ is some abelian category, usually the category of $\mathbf{k}$-vector spaces Vect or commutative graded $\mathbf{k}$-algebras CommAlg.

Dually, a *copresheaf* on P is a functor $P \to \mathcal{C}$. We will only consider copresheaves of finite dimensional vector spaces in this paper.

If $\mathcal{F}$ is a sheaf on P, we write $\mathcal{F}_s^t := \mathcal{F}(s \leq t) : \mathcal{F}(t) \to \mathcal{F}(s)$ and call this the *restriction* of $\mathcal{F}$ from t to s. Dually, for a copresheaf G on P, we write $G_s^t := G(s \leq t) := G(s) \to G(t)$ and call this the *inclusion* of G from s to t. Notice that the notation in either case reflects the poset, so that the superscript is above the subscript in P.

Fix a sheaf $\mathcal{F}$ on P. For any open subset $X \subset P$, $\mathcal{F}(X)$ is the categorical limit of the diagram $\mathcal{F}|_X : X^{\mathrm{op}} \to \mathcal{C}$ – this justifies our use of the word sheaf rather than presheaf. For any $X \subset Y \subset P$ open, we write $\mathcal{F}_X^Y : \mathcal{F}(Y) \to \mathcal{F}(X)$ to be the restriction induced by the universal properties of limits. We also let $\mathcal{F}(Y, X)$ be the kernel of $\mathcal{F}_X^Y$. When basic opens are involved, we write $\mathcal{F}_X^t := \mathcal{F}_X^{[t]}$, $\mathcal{F}(t, X) := \mathcal{F}([t], X)$, and $\mathcal{F}(t, s) := \mathcal{F}([t], [s])$.

The following is a useful criterion for computing global sections of sheaves on our posets of study, including CM meet-semilattices.

Lemma 6.2.1. *Let $\mathcal{F}$ be a sheaf on a star normal meet-semilattice P of rank d. Then we may compute $\mathcal{F}(P)$ concretely as the set of elements $(x_t) \in \bigoplus_{t \in P_d} \mathcal{F}(t)$ such that*

$$\mathcal{F}_{t \wedge t'}^t(x_t) = \mathcal{F}_{t \wedge t'}^{t'}(x_{t'}) \tag{6.2.1}$$

whenever $t \wedge t' \in P_{d-1}$.

Proof. Computing $\mathcal{F}(P) \subset \bigoplus_{t \in P_d} \mathcal{F}(t)$ using Čech cohomology with respect to the open cover $\{[t]\}_{t \in P_d}$, then $\mathcal{F}(P)$ consists of the elements $(x_t) \in \bigoplus_{t \in P_d} \mathcal{F}(t)$ such that (6.2.1) holds for any distinct $t, t' \in P_d$, regardless of $\rho(t \wedge t')$. Thus our condition is necessary, and now we check to see that it is sufficient, so assume (6.2.1) only holds for $t, t' \in P_d$ with $t \wedge t' \in P_{d-1}$.

Let $s, r \in P_d$. Then $\bigstar(s \wedge r)$ dually connected, so we may find $s = t_0, t_1, \ldots, t_m = r \in \bigstar(s \wedge r) \cap P_d$ such that for all $1 \leq i \leq m$ we have $t_{i-1} \wedge t_i \in P_{d-1}$. Then

$$\mathcal{F}^{t_i}_{s\wedge r}(x_{t_i}) = \mathcal{F}^{t_i \wedge t_{i-1}}_{s\wedge r}(\mathcal{F}^{t_i}_{t_i \wedge t_{i-1}}(x_{t_i})) = \mathcal{F}^{t_i \wedge t_{i-1}}_{s\wedge r}(\mathcal{F}^{t_{i-1}}_{t_i \wedge t_{i-1}}(x_{t_{i-1}})) = \mathcal{F}^{t_{i-1}}_{s\wedge r}(x_{t_{i-1}}),$$

and therefore $\mathcal{F}^{s}_{s\wedge r}(x_s) = \mathcal{F}^{r}_{s\wedge r}(x_r)$. □

The dual notion of a limit is a colimit. If G is a copresheaf of finite dimensional vector spaces on P, and P is pure of rank d, then the *colimit* of G on P is the quotient of $\bigoplus_{t\in P_d} G(t)$ by the subspace

$$\{G^t_s(x) - G^{t'}_s(x) \mid x \in G(s) \text{ and } s \leq t, t'\}.$$

Corollary 6.2.2. *For P as in Lemma 6.2.1 and G a copresheaf of finite-dimensional vector spaces on P, then we have that the colimit of G on P is the quotient of $\bigoplus_{t\in P_d} G(t)$ by*

$$\{G^t_{t\wedge t'}(x) - G^{t'}_{t\wedge t'}(x) \mid t \wedge t' \in P_{d-1}\}.$$

Proof. If $(\cdot)^\vee$ denotes composing with vector space duality, then $G^\vee$ is a sheaf of finite-dimensional vector spaces on P. Therefore, we can compute the limit $G^\vee(P)$ by only restricting our attention to $Q \coloneqq \{t \wedge t' \mid t, t' \in P_d\} \cap P_{\geq d-1} \subset P$. Applying $(\cdot)^\vee$ again turns the limit of this diagram into a colimit of the $G^{\vee\vee} \simeq G$ applied to Q, which is computed as claimed. □

6.3 Cellular poset cohomology and transition coefficients

Suppose $\mathcal{F}$ is a sheaf on a (finite, pure, lower graded, lower thin) poset P of rank d. Then we wish to find a chain complex

$$0 \to \bigoplus_{t \in P_d} \mathcal{F}(t) \to \bigoplus_{t \in P_{d-1}} \mathcal{F}(t) \to \cdots \to \bigoplus_{t \in P_1} \mathcal{F}(t) \to \mathcal{F}(\hat{0}) \to 0$$

whose differentials are the sums of all possible restrictions with signs sprinkled in. Furthermore, we want a relationship between this complex, the cohomology of $\mathrm{ord}(P)$, and the global sections of $\mathcal{F}$. If P is the face poset of a regular cell complex, such a chain complex is given by the cellular chain complex with coefficients in $\mathcal{F}$.

We adapt the relevant parts of [ET15], to our setting. For any $s \in P_{>0}$, let C_s be the vector space generated by the maximal chains in $(\hat{0}, s]$, subject to the relations that $\sigma + \sigma' = 0$ if σ and σ' differ by a diamond move. We then define a chain complex $C^\bullet(P)$ as follows. We set

$$C^n(P) = \bigoplus_{s \in P_{d-n}} C_s,$$

and we let the matrix entry $\overline{\delta}^s_t : C_s \to C_t$ of the differential $\overline{\delta} : C^n \to C^{n-1}$ be given by

$$\overline{\delta}^s_t(\sigma) = \begin{cases} (-1)^{d-n}(\sigma \lessdot t) & s \lessdot t \\ 0 & \text{otherwise,} \end{cases}$$

where σ is any generator of C_s and $(\sigma < t)$ is the generator of C_t given by appending t to σ.

Remark 6.3.1. We think of $C_{\hat{0}}$ as an augmentation of this complex, which makes sense on the nose if we were to change our indices to $C^{d-1-\bullet}(P)$. However, later it will be more natural to keep these swapped indices.

Remark 6.3.2. This chain complex is much less general than the cellular poset cohomology developed in [ET15], since they consider the version with coefficients in a copresheaf on P (or rather, a sheaf on P^{op}). We will not need that generality for the copresheaves, but we will dualize this construction below and use coefficients in a sheaf (on P).

The following is an amalgamation of the main theorems of [ET15], adapted to our situation.

Theorem 6.3.3 ([ET15])**.** *If P is lower CM, then*

$$H^n(C^\bullet(P)) \simeq H^{d-n-1}(\mathrm{ord}(P)) \simeq \widetilde{H}^{d-n-1}(\|P\|).$$

If P is lower Gorenstein, then C_s is one-dimensional for all $s \in P$.*

Proof. Lower CM implies that $(P_{>0})^{\mathrm{op}}$ is *cellular*, and therefore [ET15, Theorem 1] says that $H^n(C^\bullet(P)) \simeq \widetilde{H}^{d-n}(\|\,\mathrm{ord}(P_{>0})^{\mathrm{op}}\|) = \widetilde{H}^{d-n}(\|\,\mathrm{ord}(P_{>0})\|)$ for $n \leq d-1$, where we are considering $C^d = \mathbf{k}\{\hat{0}\}$ makes $C^{d-1-\bullet}(P)$ an augmented complex.

The second claim is obvious for $s = \hat{0}$, so suppose $s \neq \hat{0}$. Then $[s]$ is Gorenstein*, we have that the cohomology of $\mathrm{ord}(\partial s)$ is one-dimensional and concentrated in degree $\rho(s)-1$. Then $C_s \simeq H^{\rho(s)-1}(\mathrm{ord}(\partial s))$ by [ET15, Proposition 6 and Remark 5]. □

Before we adapt this complex to work with sheaves on P, we will make some choices of generators of the various C_s's to ease the process. Assuming P is lower Gorenstein*, then any fixed choice of generators $\sigma_s \in C_s$ for all $s \in P$ results in *transition coefficients* ε_s^t for all $s \lessdot t$ defined by

$$\overline{\delta}_t^s(\sigma_s) = \varepsilon_s^t \sigma_t.$$

Notice that $\varepsilon_{\hat{0}}^{s}$ is a constant for all $s \in P_1$, since the generators of C_s are unique for $s \in P_{\leq 1}$. Also, we necessarily have

$$\varepsilon_s^t \varepsilon_r^s \varepsilon_{s'}^t \varepsilon_r^{s'} = -1 \quad \text{for any diamond } [r,t] = \{r, s, s', t\}, \tag{6.3.1}$$

since for such a diamond we have

$$r \lessdot s \lessdot \sigma_t = -r \lessdot s' \lessdot \sigma_t \in C_r,$$

and one of these sides must be equivalent to $\sigma_r \in C_r$.

Sometimes we will want to have more control over these generators. The next lemma is very important whenever we try to tie the upcoming chain complex with global sections of sheaves.

Lemma 6.3.4. *Suppose P is a subthin lower Gorenstein* dually connected meet-semilattice of rank $d > 1$. Then we can choose generators $\sigma_s \in C_s$ for all $s \in P$ such that the transition coefficients $\varepsilon_s^t \in \{\pm 1\}$ satisfy*

$$\varepsilon_{t\wedge t'}^{t} \varepsilon_{t\wedge t'}^{t'} = -1 \quad \textit{for any } t, t' \in P_d \textit{ with } t \wedge t' \in P_{d-1}. \tag{6.3.2}$$

Proof. If $d = 2$, then P is the face poset of a simple graph such that $\|P\|$ is homeomorphic to either a circle or a line segment. Fixing any single choice $\sigma_t = s \lessdot t$, where we choose $s \in \partial P$ if $\partial P \neq \varnothing$, determines $\sigma_{t'} \coloneqq s' \lessdot t'$ for all other $t' \in P_2$ by the condition that an alternating sequence

$$s \lessdot t \gtrdot s_1 \lessdot t_1 \gtrdot s_2 \lessdot \cdots \gtrdot s_\ell \lessdot t_\ell \gtrdot s' \lessdot t'$$

has $s, s_1, \ldots, s_\ell, s' \in P_1$ all distinct.

If $d > 2$, then choose any (fixed) generators $\sigma_s \in C_s$ for all $s \in P$. We will modify these choices until we reach the σ_s promised in the claim. To this end, consider the graph $\mathrm{ord}(Q)$, where Q is the poset given by deleting all elements of P with ranks not equal to $0, d-1, d$, which is connected by assumption. Fix $t_1 \in P_d$, and pick a labeling $P_d = \{t_1, t_2, \cdots, t_{|P_d|}\}$ which is a linear refinement of the ordering on P_d induced by distance from t_1 in $\mathrm{ord}(Q)$.

For all $1 < k \leq |P_d|$, let

$$\mathcal{I}_k := \{1 \leq \ell < k \mid \sigma_{t_\ell \wedge t_k} \lessdot t_\ell = \sigma_{t_\ell} \text{ and } \sigma_{t_\ell \wedge t_k} \lessdot t_k = \sigma_{t_k}\}$$

be the set of all indices less than k which force a failure of (6.3.2). If $\mathcal{I}_k = \varnothing$ for all k, then we are done.

Otherwise, there is a minimal $1 < k \leq |P_d|$ with $\mathcal{I}_k \neq \varnothing$. We repair this faulty k by replacing

- $\sigma_{t_k \wedge t_\ell}$ with $\sigma'_{t_k \wedge t_\ell} = -\sigma_{t_k \wedge t_\ell} \in C_{t_k \wedge t_\ell}$ for all $\ell < k$ with $\ell \notin \mathcal{I}_k$, and
- σ_{t_k} with $\sigma'_{t_k} = -\sigma_{t_k} \in C_{t_k}$.

With respect to these new generators, we find that $\mathcal{I}_k$ has now become empty, while $\mathcal{I}_\ell$ for $\ell < k$ remains empty because P is subthin. Proceeding to the next faulty k, we can fix all of the remaining errors in our current choice. □

6.3.1 Cellular poset homology with coefficients in a sheaf

We now adapt the complex above to our setting. Assume P is lower Gorenstein* and pure of rank $d \geq 1$, and make a choice of generators $\{\sigma_s\}_{s \in P}$. If P is subthin and diamond transitive, pick these generators as in Proposition 6.3.4.

Definition 6.3.5. Given a sheaf $\mathcal{F}$ on P, we define the chain complex $\mathcal{C}_\bullet(P;\mathcal{F})$ as follows. Put

$$\mathcal{C}_n(P;\mathcal{F}) = \bigoplus_{s\in P_{d-n}} \mathcal{F}(s),$$

and set the matrix entry $\delta_s^t : \mathcal{F}(t) \to \mathcal{F}(s)$ of the differential $\delta : \mathcal{C}_n(P;\mathcal{F}) \to \mathcal{C}_{n+1}(P;\mathcal{F})$ to be $\delta_s^t := \varepsilon_s^t \mathcal{F}_s^t$ for $s \lessdot t$, and 0 otherwise. We make the convention that $\mathcal{C}_\bullet(\varnothing;\mathcal{F})$ is the zero complex. If $X \subset P$ is open, we also define the relative complex

$$\mathcal{C}_\bullet(P,X;\mathcal{F}) := \mathcal{C}_\bullet(P;\mathcal{F})/\mathcal{C}_\bullet(X;\mathcal{F}).$$

Both $\mathcal{C}(P;\mathcal{F})$ and $\mathcal{C}(P,X;\mathcal{F})$ are chain complexes because of (6.3.1). They are functorial in $\mathcal{F}$, and in particular respect direct sums of sheaves and tensoring sheaves with a fixed vector space.

For all $s \in P$, let $\mathcal{J}^s$ denote the constant sheaf on $\bigstar(s) \subset P$ outputting $\mathbf{k}$, extended by zero to the rest of P.

Proposition 6.3.6. *For any $s \in P$, we have an isomorphism*

$$H^m(\mathcal{C}_\bullet(P;\mathcal{J}^s)) \simeq H^{d-\rho(s)-1-m}(\mathrm{ord}(\bigstar(s));\mathbf{k}). \tag{6.3.3}$$

Proof. Let $\underline{\mathbf{k}}$ be the constant sheaf on the lower Gorenstein* $\bigstar(s)$. We have

$$H^m(\mathcal{C}_\bullet(P;\mathcal{J}^s)) = H^m(\mathcal{C}_\bullet(\bigstar(s);\underline{\mathbf{k}})) \simeq H^m(C^\bullet(\bigstar(s)))$$

where the isomorphism comes from applying vector space duality to $\mathcal{C}_\bullet(\bigstar(s);\underline{\mathbf{k}})$, which is itself the isomorphic to the vector space dual of $C^\bullet(\bigstar(s))$. The first claim is then finished by Theorem 6.3.3. □

We also want a relative version, which is not developed in [ET15]. The following will suffice for our purposes.

Lemma 6.3.7. *If P is subthin, CM, and ∂P is empty or CM, then*

$$H^m(\mathcal{C}_\bullet(P, \partial P; \mathcal{J}^s)) = 0$$

for $m \neq 0$ and $s \in P$.

Proof. We may assume $\star(s) = P$ and $\partial P \neq \varnothing$ by using an argument similar to the proof of Proposition 6.3.6 and the fact that $\star(s)$ and $\partial \star(s) = \star(s) \cap \partial P$ are CM. Notice that we get a short exact sequence of chain complexes

$$0 \to \mathcal{C}_\bullet(\partial P; \underline{\mathbf{k}}) \to \mathcal{C}_\bullet(P; \underline{\mathbf{k}}) \to \mathcal{C}_\bullet(P, \partial P; \underline{\mathbf{k}}) \to 0$$

and applying vector space duality gives the exact sequence:

$$0 \to (\mathcal{C}_\bullet(P, \partial P; \underline{\mathbf{k}}))^* \to C^\bullet(P) \to C^\bullet(\partial P) \to 0.$$

Then we get a long exact sequence in cohomology

$$\cdots \to \widetilde{H}^{d-m-2}(\|\partial P\|) \to H^m((\mathcal{C}_\bullet(P, \partial P; \underline{\mathbf{k}}))^*) \to \widetilde{H}^{d-m-1}(\|P\|) \to \widetilde{H}^{d-m-1}(\|\partial P\|) \to \cdots$$

where we have replaced most of the terms using (6.3.3). CM implies that $\widetilde{H}^m(\|P\|) = \widetilde{H}^{m-1}(\|\partial P\|) = 0$ for all $m \neq 0$, so the same must be true of $H^m((\mathcal{C}_\bullet(P, \partial P; \underline{\mathbf{k}}))^*)$. Finally, we reapply vector space duality to get the claim. □

6.3.2 Comparing the chain complex with global sections for subthin posets.

Now we let P be a subthin lower Gorenstein* poset of rank $d > 1$ and $\mathcal{F}$ a sheaf on P. Define $\mathcal{C}_\bullet(P; \mathcal{F})$ with respect to generators chosen as in Proposition 6.3.4. We will often

be most concerned with the 0-cycles of our chain complexes, so we make the definitions

$$Z(P;\mathcal{F}) \coloneqq H_0(\mathcal{C}_\bullet(P;\mathcal{F})) \quad \text{and} \quad Z(P,X;\mathcal{F}) \coloneqq H_0(\mathcal{C}_\bullet(P,X;\mathcal{F})),$$

both of which we view as subsets of $\mathcal{C}_0(P;\mathcal{F}) = \bigoplus_{t\in P_d} \mathcal{F}(t)$.

Proposition 6.3.8. *Suppose P is a subthin star normal lower Gorenstein* meet-semilattice. Then*

$$Z(P,\partial P;\mathcal{F}) \simeq \mathcal{F}(P) \quad \textit{and} \quad Z(P;\mathcal{F}) \simeq \mathcal{F}(P,\partial P).$$

Proof. To save on space, let X be either ∂P or $\varnothing$, and let $Y = \partial P \backslash X$. An element of $\mathcal{F}(P,X)$ is a section $(x_t) \in \mathcal{F}(P)$, written in view of Lemma 6.2.1, which restricts to zero in $\mathcal{F}(X)$, meaning if $s \lessdot t$ with $s \in X \subset \partial P$, then $\mathcal{F}_s^t(x_t) = 0$. On the other hand, $Z(P,Y;\mathcal{F})$ is a subspace of $Z(P;\mathcal{F})$ with the additional condition that the matrix entry

$$\varepsilon_s^t \mathcal{F}_s^t : \mathcal{F}(t) \to \mathcal{F}(s)$$

is zero if $s \in Y \subset \partial P$. This is implies that if $(x_t) \in Z(P,Y;\mathcal{F})$ and $s \in X \cap P_{d-1}$, then $\mathcal{F}_s^t(x_t) = 0$ for the unique t covering s.

Thus, by Lemma 6.2.1, we need only to show that $(x_t) \in Z(P,Y;\mathcal{F})$ satisfies Equation 6.2.1 whenever $t,t' \in P_d$ and $s = t \wedge t' \in P_{d-1}$. This follows from the assumption that $\varepsilon_s^t \varepsilon_s^{t'} = -1$, since this implies that the differential restricted to $\mathcal{F}(t) \oplus \mathcal{F}(t') \to \mathcal{F}(s)$ takes $(x_t, x_{t'})$ to

$$\varepsilon_s^t \mathcal{F}_s^t(x_t) + \varepsilon_s^{t'} \mathcal{F}_s^{t'}(x_{t'}) = \pm(\mathcal{F}_s^t(x_t) - \mathcal{F}_s^{t'}(x_{t'})) = 0,$$

meaning $\mathcal{F}_s^t(x_t) = \mathcal{F}_s^{t'}(x_{t'})$. □

Chapter 7

Sheaf analogues of facewise functions

Central to the study of IHF is the sheaf of conewise polynomials, which becomes the structure sheaf of the fan. In this section, we will define an analogue of this sheaf for our posets. We will then study certain modules for this sheaf, called *pure sheaves*, which will include the minimal extension sheaf we desire.

The strategy will be to first find the analogue of the sheaf of the dual vector spaces of the spans of the cones of a fan, and then to make this algebraic by functorially taking symmetric algebras of the resulting vector spaces.

7.1 Linearization sheaves

Throughout, our posets are assumed to be pure, lower thin, and lower graded. Recall that the pullback of the diagram in vector spaces

$$\begin{array}{ccc} & & U \\ & & \downarrow f \\ V & \xrightarrow{g} & W \end{array}$$

is its limit, represented by the vector space $\{(x, y) \mid f(x) = g(y)\} \subset U \times V$, whereas the pushout of

$$\begin{array}{ccc} U & \xrightarrow{j} & V \\ \downarrow i & & \\ W & & \end{array}$$

is its colimit, which we can represent with $V \oplus W/\{i(u) - j(u) \mid u \in U\}$.

If $\mathcal{F}$ is a sheaf $\mathcal{F}$ of vector spaces on a lower thin poset P, and if $[r, t] = \{r, s, s', t\}$ is a diamond in P, then we refer to the pullback of

$$\begin{array}{ccc} & & \mathcal{F}(s') \\ & & \downarrow \mathcal{F}_r^{s'} \\ \mathcal{F}(s) & \xrightarrow{\mathcal{F}_r^{s}} & \mathcal{F}(r) \end{array}$$

as the *diamond pullback* of $\mathcal{F}$ on $[r, t]$. The following are the key gadgets introduced in this project.

Definition 7.1.1. A *linearization sheaf* on P is a sheaf $\mathcal{U} : P^{\mathrm{op}} \to \mathrm{Vect}$ such that $\dim \mathcal{U}(s) = \rho(s)$ for all $s \in P$, and if $[r, t]$ is a diamond in P, then the induced map from $\mathcal{U}(t)$ to the diamond pullback of $\mathcal{U}$ on $[r, t]$ is an isomorphism.

Dually, we define a *colinearization copresheaf* on P to be a copresheaf $\mathbf{V} : P \to \mathrm{Vect}$ which satisfies the dual properties, namely that $\dim \mathbf{V}(s) = \rho(s)$ and that for any diamond

$[r,t]$ the induced map $\mathbf{V}(t)$ from the pushout of

$$\begin{array}{ccc} \mathbf{V}(r) & \xrightarrow{\mathbf{V}_r^{s'}} & \mathbf{V}(s') \\ \downarrow{\scriptstyle \mathbf{V}_r^{s}} & & \\ \mathcal{U}(s) & & \end{array}$$

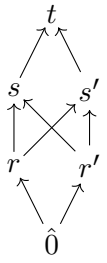
is an isomorphism.

It is clear that linearization sheaves and colinearization copresheaves correspond to each other by composing with vector space duality. The intuition for colinearization copresheaves is that they should behave like a linear realization of the poset, whereas linearization sheaves should act like linear functions on a realization. Most of this paper is written from the perspective of sheaves rather than copresheaves, so we focus on linearization sheaves. A notable exception is that we can directly prove that colinearization copresheaves exist in Lemma 7.1.6, whereas linearization sheaves exist as a corollary.

For a linearization sheaf, we use the shorthand $U_t \coloneqq \mathcal{U}(t)$ whenever no confusion will arise. Notice that for a linearization sheaf, the restrictions $\mathcal{U}_s^t : U_t \to U_s$ must always be surjective by induction. Furthermore, any diamond $[r,t] = \{r,s,s',t\}$ must have the the kernels $\mathcal{U}(t,s)$ and $\mathcal{U}(t,s')$ are distinct. We will use both of these facts throughout.

Remark 7.1.2. Given $t \in P$ and a sheaf $\mathcal{F}$ on P, there is a unique map from $\mathcal{F}(\partial t)$ to the diamond pullback of $\mathcal{F}$ on a diamond $[r,t]$ which commutes with the restrictions by the universal property of limits. For a linearization sheaf, it is then necessary that this map must be surjective for every such diamond. The following example is a poset which cannot admit a linearization sheaf.

Example 7.1.3. Consider the poset with Hasse diagram given by



This is the face poset of a regular cell decomposition of a 2-ball, where the boundary is a circle with two vertices and two edges. Any linearization sheaf $\mathcal{U}$ on $[\hat{0}, t)$ must take $U_s = U_{s'} = U_r \oplus U_{r'}$, with restrictions given by projections, up to isomorphism. But then $\mathcal{U}(\partial t) = U_r \oplus U_{r'}$ is too small to allow us to define U_t. Notice that this poset is not a meet-semilattice.

However, in our two motivating examples of simplicial complexes and fans, we have the following examples.

Example 7.1.4. Suppose P is (the face poset of) a simplicial complex on the vertex set P_1. Then we define the *standard linearization sheaf* $\mathcal{S}$ on P via $\mathcal{S}(s) := S_s = \mathbf{k}^{[s]_1}$, where the restriction maps are given by the standard projections.

Every linearization sheaf $\mathcal{U}$ on P is isomorphic to $\mathcal{S}$: we can identify $\mathcal{U}|_{P_{\leq 1}} = \mathcal{S}|_{P_{\leq 1}}$ by picking generators of U_i for $i \in P_1$. For $t \in P_{\geq 1}$ and $i \in [t]_1$, there is a unique $s \lessdot t$ with $t = s \vee i$, so we choose $v_{i,t} \in U_t$ to be the unique element of $\mathcal{U}(t, s)$ with $\mathcal{U}_i^t(v_i) = 1 \in \mathbf{k}^{\{i\}} = U_i$. Then $\{v_i \mid i \in [t]_1\}$ is a basis of U_t, and our isomorphism follows because if $r \leq t$ we have

$$\mathcal{U}_r^t(v_{i,t}) = \begin{cases} v_{i,r} & i < r \\ 0 & i \not< r. \end{cases}$$

Example 7.1.5. Let $\mathbf{k} = \mathbb{R}$, and suppose P is the face poset of a full-dimensional fan in a real vector space V, meaning $s \leq t$ in P if we have cone containment $\sigma_s \subset \sigma_t$ in the fan. Then we get a colinearization copresheaf $\mathbf{V}$ on P by setting $\mathbf{V}(s)$ to be the linear span of the cone $\sigma_s \subset V$, and setting $\mathbf{V}_s^t : \mathbf{V}(s) \to \mathbf{V}(t)$ to be the inclusion. The corresponding linearization sheaf gives rise to the *sheaf of conewise polynomials* on the fan, which motivates the development of this paper.

Before proceeding any further, we must make sure that our new gadgets exist. The following explains why we deal exclusively with meet-semilattices throughout the rest of the paper.

Lemma 7.1.6. *If P is a lower CM meet-semilattice, then P has a colinearization copresheaf, and therefore a linearization sheaf.*

Proof. We build a colinearization copresheaf by hand. Notice that $P_{\leq 2}$ is the face poset of a simplicial complex, so we may choose $\mathbf{V}(s) \coloneqq \mathbf{k}^{[s]_1}$ for $s \in P_{\leq 2}$, with standard inclusions. We proceed inductively, assuming that we have defined $\mathbf{V}|_{P_{\leq d}}$ on $P_{\leq d}$, where the rank of P is $d+1$. We may also assume $P = [t]$, so that $P_{\leq d} = \partial t$ is CM.

For all $s \in P_{\leq d}$, let $V_s \coloneqq \mathbf{V}(s)$. For any $s, s' \in P_d$, let $W_{s,s'} = V_s \oplus V_{s'}/I_{s,s'}$ be the pushout of

$$\begin{array}{ccc} V_{s\wedge s'} & \xrightarrow{\iota_{s\wedge s'}^{s'}} & V_{s'} \\ \downarrow{\scriptstyle \iota_{s\wedge s'}^{s}} & & \\ V_s & & \end{array}$$

By Corollary 6.2.2, we have that the colimit W of $\mathbf{V}$ on P is the quotient of $\bigoplus_{r\in P_d} V_r$ by

$$I = \sum_{s\wedge s' \in P_{d-1}} I_{s,s'}$$

where we identify $I_{s,s'} \subset V_s \oplus V_{s'}$ with its image under the natural inclusion $V_s \oplus V_{s'} \hookrightarrow \bigoplus_{r \in P_d} V_r$. For any $s, s' \in P_d$ with $s \wedge s' \in P_{d-1}$, we notice that $I \cap (V_s \oplus V_{s'}) = I_{s,s'}$; this is the crucial place where we use the fact that P is a meet-semilattice. Thus, the natural map $W_{s,s'} \to W$ is injective.

To finish the claim, recall that $\mathbf{k}$ is infinite. Then we may generically pick some $I' \subset W$ such that $W = W_{s,s'} \oplus I'$ for any $s, s' \in P_d$ with $s \wedge s' \in P_{d-1}$, and let $\mathbf{V}(t) \coloneqq W/I'$, with the maps $\mathbf{V}_s^t$ determined by following the colimit maps with the quotient by I'. □

Thus, it is imperative that we work with meet-semilattices, as we do for the rest of the paper.

Remark 7.1.7. For a fixed P, (co)linearization (co)sheaves are not unique in general. We refrain from making any preferential choice for linearization sheaves, unless P is the face poset of a simplicial complex as in Remark 7.1.4.

7.1.1 Homogenization

Definition 7.1.8. Suppose P is pure of rank d. A *homogeneous linearization sheaf* on P is a linearization sheaf $\mathcal{U}$ on P together with a vector space U and maps $\eta_s : U \to U_s$ for all $s \in P_d$ such that η_t is an isomorphism for all $t \in P_d$ and the diagram

$$\begin{array}{ccc} U & \xrightarrow{\eta_{s'}} & U_{s'} \\ \downarrow{\scriptstyle \eta_s} & & \downarrow{\scriptstyle \mathcal{U}_r^{s'}} \\ U_s & \xrightarrow{\mathcal{U}_r^{s}} & U_r \end{array}$$

commutes whenever $r \leq s, s'$. We often omit the maps η_s from the notation, and call U a *homogenization* of $\mathcal{U}$. The dual notion for a colinearization copresheaf $\mathbf{V}$ is called a *cohomogenization* V.

Remark 7.1.9. In the case of (face posets of) simplicial complexes, we will see that homogenizations are related to *linear systems of parameters* for the corresponding face rings. On the other hand, in the fan case, cohomogenizations are modeled after the global vector space where the fan lives.

Lemma 7.1.10. *Suppose P is a star normal lower CM lower thin meet-semilattice of rank d. Then any colinearization sheaf $\mathbf{V}$ has a cohomogenization V, and thus any linearization sheaf $\mathcal{U}$ has a homogenization U.*

Proof. A closer inspection of the proof of Lemma 7.1.6 shows that we only restricted to the case $P = [t]$ since it implied that $[r,t) = \star(r) \cap \partial t$ was dually connected for all $r < t$. Our current assumptions allow us to use the same argument as before and 6.2.2 to find a cohomogenization V as a quotient of the colimit of $\mathbf{V}$ on P. Duality of finite-dimensional vector spaces gives the existence of homogenizations. □

Remark 7.1.11. Given a homogeneous linearization sheaf $(\mathcal{U}, U)$ on P, we will also like to know how this relates to sheaves on other posets related to P. If $X \subset P$ is open and pure, then the restriction $\mathcal{U}|_X$ is a linearization sheaf on X. We can make this homogeneous by choosing a subspace $W \subset U$ of the correct dimension generically so to avoid the kernels of any maps $\eta_s : U \to U_s$ for every facet $s \in X$.

On the other hand, let $s \in P$. Then we obtain a homogeneous linearization sheaf on $\star(s)$ by defining $\mathcal{U}' : \star(s)^{\mathrm{op}} \to \mathrm{Vect}$ via $\mathcal{U}'(t) := \mathcal{U}(t,s)$, with homogenization given by the kernel of $\eta_s : U \to U_s$.

For most of the rest of the paper, we omit the colinearization copresheaves from our discussion, with the exception of the existence of subdivisions in Section 8.

7.2 The sheaf of facewise polynomials

From now on, we assume P is a (pure lower thin) star normal lower CM meet-semilattice, so that homogeneous linearization sheaves exist on P. Given a linearization sheaf $\mathcal{U}$ on P, we make the following definition.

Definition 7.2.1. The *sheaf of facewise polynomials* of the pair $(P, \mathcal{U})$ is

$$\mathcal{A} \coloneqq \operatorname{Sym} \circ \mathcal{U} : P^{\mathrm{op}} \to \mathrm{Alg}.$$

We set $A_s \coloneqq \mathcal{A}(s) = \operatorname{Sym} U_s$ for all $s \in P_d$. If U is a homogenization of $\mathcal{U}$, then we set $A = \operatorname{Sym} U$.

As a motivation to study this sheaf, we consider how it relates to fans and simplicial complexes. Let P be the face poset of a simplicial complex Δ with at least one vertex, where we identify P_1 with the set of vertices of Δ. Then the *face ring* of Δ is $\mathbf{k}[\Delta] \coloneqq \mathbf{k}[x_i \mid i \in P_1]/I_\Delta$, defined by the ideal

$$I_\Delta \coloneqq (x_{i_1} \cdots x_{i_l} \mid i_\ell, \ldots, i_\ell \text{ are distinct and } \{i_1, \ldots, i_\ell\} \notin \Delta).$$

Proposition 7.2.2. *If P is the face poset of a simplicial complex Δ and $\mathcal{S}$ is the standard linearization sheaf on P, then there is an isomorphism*

$$\mathbf{k}[\Delta] \xrightarrow{\simeq} \mathcal{A}(P).$$

If P is pure, then this isomorphism defines a bijection between homogenizations U of $\mathcal{S}$ and $\Theta \subset \mathbf{k}[\Delta]_1$ which are spanned by a linear system of parameters.

Proof. Suppose the length of a maximal chain in P is $d > 0$. Define a linear map $\overline{\phi} : \mathbf{k}^{P_1} \to \bigoplus_{t \text{ maximal}} \mathbf{k}^{I(t)}$ with respect to the standard basis $\{x_i \mid i \in P_1\}$ given by sending x_i to the

element which is $e_i \in \mathbf{k}^{I(t)}$ whenever $i \leq t$ and 0 otherwise. Then we extend $\overline{\phi}$ to a map of graded algebras

$$\phi : \operatorname{Sym} \mathbf{k}^{P_1} \to \mathcal{A}(P)$$

by letting $\overline{\phi}$ be the map on generators in degree 1. Through a support argument, we see that $I_\Delta \subset \ker(\phi)$. On the other hand, $A_t = \operatorname{Sym} \mathbf{k}^{I(t)}$ for all $t \in P$ implies ϕ is surjective and that its kernel is no bigger than I_Δ.

Finally, assume P is pure and that $U \subset \mathcal{S}(P)$ is a homogenization of $\mathcal{U}$, which can be viewed as a subspace of $\mathcal{A}(P)_1$. Then we may define $\Theta := \phi^{-1}(U) \subset \mathbf{k}[\Delta]_1$, and we notice that the projection of Θ to the span of $\{x_i \mid i \in I(t)\}$ for any $t \in P_d$ is an isomorphism, so Θ must be spanned by a linear system of parameters as in [Sta04, Lemma III.2.4]. Conversely, such Θ give rise to homogenizations U of $\mathcal{S}$ by defining $U' := \phi(\Theta) \subset \mathcal{A}(P)_1$, and then picking $U \in \mathcal{S}(P)$ such that U restricts to the restriction of U' in $\mathcal{A}(t)_1 = U_t$. □

Remark 7.2.3. When P is the face poset of a simplicial complex Δ, we know that $\mathcal{A}(P)$ is a free graded $A = \operatorname{Sym} U$-module of finite rank if and only if Δ is CM. We will prove something related to this for CM lower Eulerian meet-semilattices in Theorem 8.1.5, which rests on ideas related to cellular poset homology as discussed in Section 6.3.

Remark 7.2.4. For P the face poset of a fan Σ in a vector space V, we also have an interpretation that is central to the study of combinatorial intersection homology for fans. Using the linearization sheaf given by composing the colinearization copresheaf of Example 7.1.5 with vector space duality, we get that $\mathcal{A}$ is the *sheaf of conewise polynomials* on Σ. If P is pure and full-dimensional, then V^* is a homogenization of the underlying linearization sheaf.

7.3 Pure sheaves and the minimal extension sheaf

Let $\mathcal{U}$ be a linearization sheaf on P. The theory of combinatorial intersection homology for fans is based around studying certain modules for the sheaf of conewise polynomials. We will study similar modules for a sheaf of facewise polynomials $\mathcal{A}$.

Recall that a sheaf $\mathcal{F}$ is *flabby* if for any $X \subset Y \subset P$ open, the restriction $\mathcal{F}_X^Y$ is surjective. Notice that $\mathcal{F}$ is flabby if and only if the restrictions $\mathcal{F}(t) \to \mathcal{F}(\partial t)$ are surjective for all $t \in P$.

Definition 7.3.1. A *pure sheaf* is a flabby graded $\mathcal{A}$-module $\mathcal{F}$ which is *locally free*, meaning $\mathcal{F}(t)$ is a free A_t-module of finite rank for all $t \in P$.

Suppose U is a homogenization of $\mathcal{U}$. Then the maps $U \to U_s$ induce maps $A \to A_s$ for all $s \in P$. If $\mathcal{F}$ is an $\mathcal{A}$-module and $X \subset P$ is open, then we get that $\mathcal{F}(X)$ is an A-module, and we define $\overline{\mathcal{F}(X)}$ to be the quotient of $\mathcal{F}(X)$ by the action of the irrelevant ideal $A^+ \subset A$, or equivalently by the action of $A_s^+ \subset A_s$. Then $\mathcal{F}$ is flabby if and only if the induced map of graded vector spaces $\overline{\mathcal{F}(s)} \to \overline{\mathcal{F}(\partial s)}$ is surjective for all $s \in P$.

We now turn our attention to simple pure sheaves. Let $s \in P$. We will define a pure sheaf $\mathcal{L}^s$ for any $s \in P$. Define $\mathcal{L}^s$ to be zero on $P \smallsetminus \star(s)$, and define $\mathcal{L}^s(s) \coloneqq A_s$. Assuming we have defined $\mathcal{L}^s(\partial t)$ for some $t \in \star(s)$, we can define $\mathcal{L}^s(t)$ to be a minimal free $\mathcal{A}_t$-module which surjects onto $\mathcal{L}^s(\partial t)$, where $\mathcal{L}^s(\partial t)$ is an $\mathcal{A}_t$-module via the restriction $\mathcal{A}_t \to \mathcal{A}(\partial t)$. Alternatively, we may define $\mathcal{L}^s(t)$ to be the free A_t-module generated by a vector space basis of $\overline{\mathcal{L}^s(\partial t)}$, and then letting the restriction $\mathcal{L}^s(t) \to \mathcal{L}(\partial t)$ be given by choosing representatives for the basis elements.

Theorem 7.3.2. *The set $\{\mathcal{L}^s\}_{s\in P}$ is a complete list of isomorphism classes of simple pure sheaves up to grade shift. Every pure sheaf is a direct sum of simple pure sheaves.*

Proof. That $\mathcal{L}^s$ is simple is true by construction since it is built minimally from the assumption that $\overline{\mathcal{L}^s(s)} = \mathbf{k}$. We then are reduced to showing that any pure sheaves can be written as a sum of these $\mathcal{L}^s$'s. Suppose we have a pure sheaf $\mathcal{F}$ and $s \in P$ is a minimal element in the support of $\mathcal{F}$. We wish to show that we can write

$$\mathcal{F} \simeq \mathcal{F}' \oplus \mathcal{F}''$$

where $\mathcal{F}' = \mathcal{L}^s \otimes_{\mathbf{k}} \overline{\mathcal{F}(s)}$ and $\mathcal{F}''$ is some other pure sheaf. We can construct this decomposition inductively as follows. First, set $\mathcal{F}''(s) = 0$ and $\mathcal{F}''(t) = \mathcal{F}(t)$ for all $t \in P_{\leq\rho(s)}\backslash\{s\}$. Then suppose we have constructed the decomposition on ∂t for some $t \in \star(s)$. Then we know that $\mathcal{F}(t)$ is a finitely generated free A_t-module which surjects onto $\mathcal{F}(\partial t) \simeq \mathcal{F}'(\partial t) \oplus \mathcal{F}''(\partial t)$. Pick respective bases of $\overline{\mathcal{F}'(\partial t)}, \overline{\mathcal{F}''(\partial t)}$, and let $\mathcal{F}'(t), \mathcal{F}''(t)$ respectively be the free A_t-modules generated by these bases, with respective restrictions given by picking representatives of the basis elements. Since $\mathcal{F}(t), \mathcal{F}'(t), \mathcal{F}''(t)$ are all free with flabby restrictions to the respective values on ∂t, then we get an isomorphism

$$\mathcal{F}(t) \simeq \mathcal{F}'(t) \oplus \mathcal{F}''(t)$$

compatible with the restriction maps. Induction finishes the claim. □

We now reach the most important pure sheaf.

Definition 7.3.3. Given a linearization sheaf $\mathcal{U}$ on P, we let $\mathcal{E} \coloneqq \mathcal{L}^{\hat{0}}$ be the simple pure sheaf such that $\mathcal{E}(\hat{0}) = \mathbf{k}$ is concentrated in degree 0. We call $\mathcal{E}$ the *minimal extension sheaf* of $(P, \mathcal{U})$.

Remark 7.3.4. For the face posets of rational fans with the linearization sheaf given by composing the colinearization copresheaf described in Example 7.1.5 with vector space duality, the sheaf $\mathcal{E}$ is the equivariant cohomology sheaf of the corresponding toric variety. Reducing its global sections by the action of $\operatorname{Sym} V^*$ results in the intersection cohomology of the toric variety. In [BBFK00],[BL00] this was extended to non rational fans. We are similarly extending these ideas to our posets of study.

Notice that $\mathcal{E}|_{P_{\leq 2}} = \mathcal{A}|_{P_{\leq 2}}$. Again, we have more information in the simplicial case.

Proposition 7.3.5. *For the standard linearization sheaf $\mathcal{S}$ on the face poset P of a simplicial complex, we have $\mathcal{A} = \mathcal{E}$.*

Proof. We need only prove that $\mathcal{A}$ is flabby. Both [BBFK00] and [BL00] prove this, but here we follow the latter. Let $t \in P$, and let x_i denote the standard basis vector in $\mathcal{S}(t) = \mathbf{k}^{[t]_1}$ for $i \in [t_1]$. Notice that an element of $\mathcal{A}(\partial s)$ is an $\rho(s)$-tuple $(f_i)_{i\in[t]_1}$ with $f_i \in \operatorname{Sym}(k^{[t]_1\setminus\{i\}})$ such that any monomial $m = \prod_{j\in[t]_1} x_j^{a_j}$ has the same coefficient in all f_i for which it can appear (i.e. for all $i \in [t]_1$ not in the support of m). Thus, we can define $f \in A_t$ by setting the coefficient of any monomial m with support $\neq [t]_1$ to agree with the coefficient of m in all f_i with i not in the support of m, with no condition on the monomials with full support. □

Remark 7.3.6. Suppose U is a homogenization of $\mathcal{U}$. For $s \in P$, we can write $\mathcal{L}^s$ in terms of the structure sheaf $\mathcal{E}^s$ on $\bigstar(s)$ defined using the linearization sheaf $\mathcal{S}$ of Remark 7.1.11. Using a choice of decomposition $U = U_s \oplus \mathcal{U}(U, U_s)$, we then have that $A = A_s \otimes_{\mathbf{k}} \operatorname{Sym} \mathcal{U}(U, U_s)$, and on $\bigstar(s) \subset P$ we have $\mathcal{L}^s \cong A_s \otimes_{\mathbf{k}} \mathcal{E}^s$, and of course $\mathcal{L}^s|_{P\setminus\bigstar(s)} = 0$.

Chapter 8

Equivariant formality and subdivisions

In this chapter, we consider some freeness properties of the minimal extension sheaf for a homogeneous linearization sheaf on a poset. After proving our main results, we define subdivisions in our context and prove a decomposition theorem.

8.1 Equivariant formality

For this section, we fix a lower Gorenstein* meet-semilattice P of rank d, together with a homogeneous linearization sheaf $(\mathcal{U}, U)$.

Definition 8.1.1. If P is additionally subthin, we say $(P, \mathcal{U}, U)$ is *equivariantly formal* if both $\mathcal{E}(P)$ and $\mathcal{E}(P, \partial P)$ are free A-modules.

Remark 8.1.2. In the fan case, our equivariantly formal is called *quasiconvex*. The beautiful characterization of quasiconvex fans [BBFK00, Theorem 3.8] includes all fans Δ such that

$\|\Delta\|$ is a (necessarily PL) ball or sphere of dimension $\dim V - 1$. These quasiconvex fan posets are subthin CM lower Eulerian meet-semilattices whose boundaries are either empty or CM.

We have two main theorems in this section, both of which rely on the following lemma. Recall the sheaves $\mathcal{J}^s$ for $s \in P$, which is the extension by zero of the constant sheaf on $\star(s)$ outputting $\mathbf{k}$.

Lemma 8.1.3. *For any flabby sheaf $\mathcal{F}$ of vector spaces on P, there is a decomposition*

$$\mathcal{F} \cong \bigoplus_{s \in P} \mathcal{J}^s \otimes \mathcal{F}(s, \partial s). \tag{8.1.1}$$

Proof. The decomposition 8.1.1 follows as in Lemma 7.3.2, where we can build the decomposition inductively using the flabbiness of $\mathcal{F}$ and freeness of vector spaces (see also [BBFK00, Lemma 3.6]). □

The next result is the first main theorem of this section. It will play a key role in the discussion of our hard Lefschetz conjecture and Stanley's conjecture in Section 9.2, as well as in Theorem 8.1.5. Recall the chain complex $\mathcal{C}_\bullet(P; \mathcal{F})$ with its 0-cycles $Z(P; \mathcal{F})$, as well as the relative versions.

Proposition 8.1.4. *Suppose P is CM, and let $\mathcal{F}$ be a pure sheaf. Then $Z(P; \mathcal{F}) = H_0(\mathcal{C}_\bullet(P; \mathcal{F}))$ is a free A-module. If P is subthin and ∂P is empty or CM, then the same is true for $Z(P, \partial P; \mathcal{F})$.*

Before getting to the proof, we recall a bit of commutative algebra, where for the following we view $\mathbf{k} = A/A^+$. Let $t \in P$ and $\eta_t : U \to U_t$ be the map guaranteed by

homogenization. Then consider the Koszul complex $\mathcal{K}(t)$ for the ideal $\ker(\eta_t)\cdot A \subset A$, which is a free resolution of $A_t \simeq A/\ker(\eta_t)\cdot A$. Then we have

$$\operatorname{Tor}_m^A(A_t, \mathbf{k}) \simeq H_m(\mathcal{K}(t)\otimes_A \mathbf{k})$$

which must be 0 whenever $m > \dim\ker(\eta_t) = d - \rho(t)$.

Proof. Applying (8.1.1), and remembering that $\bigstar(s)$ is CM for all $s \in P$, we have

$$\mathcal{C}_\bullet(P;\mathcal{E}) \simeq \mathcal{C}_\bullet(P;\bigoplus_{s\in P}\mathcal{F}(s,\partial s)\otimes\mathcal{J}^s)$$

is acyclic by Proposition 6.3.6. Letting $Z_k(P;\mathcal{F}) \subset \mathcal{C}_k(P;\mathcal{F})$ be the set of k-cycles, then fact that $\mathcal{C}_\bullet(P;\mathcal{F})$ is acyclic implies that we obtain short exact sequences

$$0 \to Z_k(P;\mathcal{F}) \to \mathcal{C}_k(P;\mathcal{F}) \to Z_{r+1}(P;\mathcal{F}) \to 0$$

for all k, which lead to exact sequences

$$\operatorname{Tor}_{k+1}^A(Z_{r+1}(P;\mathcal{F}),\mathbf{k}) \to \operatorname{Tor}_k^A(Z_r(P;\mathcal{F}),\mathbf{k}) \to \operatorname{Tor}_k^A(\mathcal{C}_r(P;\mathcal{F}),\mathbf{k}), \tag{8.1.2}$$

from the long exact sequence in homology, where $\mathbf{k} = A/U\cdot A$.

We will prove that $\operatorname{Tor}_k^A(Z_r(P;\mathcal{F})) = 0$ for $k > r$ by downward induction on r, starting at $r = d+1$. The base case is easy since $Z_{d+1}(P;\mathcal{F}) = 0$. Proceeding to $r < d+1$, we assume the first term of 8.1.2 is zero for $k > r$. On the other hand, $\mathcal{C}_r(P;\mathcal{F})$ is isomorphic to a direct sum of shifts of elements of $\{A_t\}_{t\in P_{d-r}}$. By the discussion before the proof, we have $\operatorname{Tor}_k^A(A_t,\mathbf{k}) = 0$ for $k > d-\rho(t) = d-(d-r) = r$. By exactness, the middle term is also zero.

In particular, $\operatorname{Tor}_1^A(Z(P;\mathcal{F}),\mathbf{k}) = 0$, which is equivalent to saying $Z(P;\mathcal{F})$ is free by Nakayama's Lemma (see, for instance, [BBFK00, Section 0.B]). This finishes the first claim, as well as the second if $\partial P = \varnothing$.

Assuming P is subthin and $\partial P \neq \varnothing$ is CM, we may follow a similar argument to the first claim. The only major change is that

$$\mathcal{C}_\bullet(P, \partial P; \mathcal{F}) \simeq \bigoplus_{s \in P} \mathcal{C}_\bullet(P, \partial P; \mathcal{J}^s) \otimes \mathcal{F}(s, \partial s)$$

is acyclic by Lemma 6.3.7. The rest of the above paragraphs go through with obvious changes. □

Finally, we reach the second main theorem of the section.

Theorem 8.1.5. *If P is a subthin CM lower Eulerian meet-semilattice and ∂P is empty or CM, then P is equivariantly formal.*

Proof. We choose signs for the complex $\mathcal{C}_\bullet(P; \mathcal{E})$ as guaranteed by Proposition 6.3.4, so that $\mathcal{E}(P, \partial P) \simeq Z(P; \mathcal{E})$ and $\mathcal{E}(P) \simeq Z(P, \partial P; \mathcal{E})$ by Proposition 6.3.8. Then Proposition 8.1.4 finishes the claim. □

8.2 Subdivisions and a decomposition theorem

In this section, we discuss what we mean by "subdivision", and we prove a decomposition theorem for subdivisions using local equivariant formality. If $\mathcal{V}$ is a linearization sheaf on Q, then we let $\mathcal{A}_{(Q,\mathcal{V})}, \mathcal{E}_{(Q,\mathcal{V})}$ respectively denote the corresponding sheaf of facewise polynomials and minimal extension sheaf.

Let Q, P be (pure) lower Gorenstein* posets of the same rank d, and let $f : Q \to P$ be a *poset map*, meaning a function between the underlying sets which respects the orderings. We always assume that $f(\hat{0}_Q) = \hat{0}_P$.

Now, suppose $(\mathcal{U}, U)$ is a homogeneous linearization sheaf on P.

Definition 8.2.1. A *subdivision* of $(P, \mathcal{U})$ is a tuple $(Q, \mathcal{V}, f, \pi)$ such that

1. Q is lower Gorenstein* meet-semilattice with the same rank as P,
2. $\mathcal{V}$ is a linearization sheaf on Q,
3. $f : Q \to P$ is a poset map which induces an isomorphism $\|f^{-1}(X)\| \simeq \|X\|$ for any open subset $X \subset P$,
4. for all $t \in P_{>0}$, $f^{-1}([t])$ is subthin with boundary $\partial f^{-1}([t]) = f^{-1}(\partial t)$, and
5. $\pi : \mathcal{U} \to f_*\mathcal{V}$ is a surjective morphism of sheaves of vector spaces on P.

If U, V are respective homogenizations of $\mathcal{U}, \mathcal{V}$, then we say a subdivision is *homogeneous* with respect to $(\mathcal{V}, V), (\mathcal{U}, U)$ if the map $\mathcal{U}(P) \to f_*\mathcal{V}(P) = \mathcal{V}(Q)$ induced by π restricts to an isomorphism $U \simeq V$, where we identify U and V with their respective images under the respective unique maps $U \to \mathcal{U}(P)$ and $V \to \mathcal{V}(Q)$.

Notice that for a subdivision, we obtain a morphism of ringed spaces $(Q, \mathcal{A}_{(Q,\mathcal{V})}) \to (P, \mathcal{A}_{(P,\mathcal{U})})$, where the induced map $\overline{\pi} : \mathcal{A}_P \to f_*\mathcal{A}_Q$ is a surjective map of sheaves of graded algebras on P. In particular, $f_*\mathcal{E}_Q$ becomes an $\mathcal{A}_P$-module.

Lemma 8.2.2. *If $(Q, \mathcal{V}, f, \pi)$ is a subdivision of $(P, \mathcal{U})$, and $\mathcal{G}$ is a pure sheaf on $(Q, \mathcal{V})$, then $f_*\mathcal{G}$ is a pure sheaf on $(P, \mathcal{U})$.*

Proof. We need only to verify that $f_*\mathcal{E}_Q$ is locally free. Let $\mathcal{A}_P, \mathcal{A}_Q, \mathcal{E}_P, \mathcal{E}_Q$ denote the respective sheaves of facewise polynomials and minimal extension sheaves of P, Q with respect to $\mathcal{U}, \mathcal{V}$. Notice that $f^{-1}([t])$ is a subthin lower Gorenstein* meet-semilattice with boundary $f^{-1}(\partial t)$, and that both $f^{-1}([t]), \partial f^{-1}([t])$ are CM since CMness is a topological

property. Since π is surjective, we also obtain an isomorphism $\mathcal{U}(t) \simeq \mathcal{V}(q)$ for all $q \in f^{-1}(t) \cap Q_{\rho_P(t)}$, so that $\mathcal{U}(t)$ is a homogenization of the linearization sheaf $\mathcal{V}|_{f^{-1}([t])}$ on $f^{-1}([t])$. Then we have that

$$(f^{-1}([t]), \mathcal{V}|_{f^{-1}([t])}, \mathcal{U}(t))$$

is equivariantly formal by Theorem 8.1.5. Thus,

$$f_*\mathcal{E}_Q(t) = f_*\mathcal{E}_Q|_{f^{-1}([t])}(t) = \mathcal{E}_Q(f^{-1}([t]))$$

is a free $\mathcal{A}_P(t)$-module. □

Theorem 8.2.3. *If $(Q, \mathcal{V}, f, \pi)$ is a subdivision of $(P, \mathcal{U})$, then every decomposition of $f_*\mathcal{E}_{(Q,\mathcal{V})}$ into simple pure sheaves on $(P, \mathcal{U})$ has exactly one direct summand isomorphic to $\mathcal{E}_{(P,\mathcal{U})}$.*

Proof. By Lemma 8.2.2, we already know that $f_*\mathcal{E}_{(Q,\mathcal{V})}$ is a pure sheaf, and therefore has a decomposition into simples by Theorem 7.3.2. The fact that $\mathcal{E}_{(P,\mathcal{U})}$ must show up exactly once is because the restriction of f to $Q_{\leq 1} \to P_{\leq 1}$ must be a poset isomorphism, so that $f_*\mathcal{E}_Q(\hat{0}_P) \simeq \mathbf{k}$. □

The following subdivision will be important to us in Theorem 9.1.9 in the next section.

Example 8.2.4. Let $P = [t]$ be Gorenstein* of rank d and let $\mathcal{U}$ be a linearization sheaf on $[t]$, which automatically comes with homogenization $U = U_t$. Then consider the CM lower Eulerian poset Q of rank d given by the pyramid over ∂t with pyramid point q, where we write the elements of Q as either $s \in \partial t \subset Q$ or $q * s$ for $s \in \partial t$, and we identify $q = q * \hat{0}$.

Choose a one-dimensional vector space L. We define a linearization sheaf $\mathcal{V}$ on Q such that $\mathcal{V}|_{\partial t} = \mathcal{U}|_{\partial t}$, while $\mathcal{V}(q * s) = U_s \oplus L$ and the restrictions are defined to be

$$\mathcal{V}^{q*s}_{q*r} : U_s \oplus L \xrightarrow{(\mathcal{U}^s_r, id)} U_r \oplus L \quad \text{and} \quad \mathcal{V}^{q*s}_{r} : U_s \oplus L \xrightarrow{(\mathcal{U}^s_r, 0)} U_r$$

for $r \leq s$ in ∂t. Define $f : Q \to P$ to be the identity on $\partial t \subset Q$, while $f(q * s) = t$ for all $s \in \partial P$.

Given a choice of $W \subset U$ such that the map $U \to U_s$ restricts to an isomorphism $W \simeq U_s$ for all $s \lessdot t$, and choosing $T \in U \backslash W$, we can obtain an isomorphism $U \simeq W \oplus L$ which makes U a homogenization of $\mathcal{V}$.

Finally, we can define $\pi : \mathcal{U} \to f_* \mathcal{V}$ using the homogenization U of $\mathcal{V}$. Then $(Q, \mathcal{V}, f, \pi)$ is a subdivision of $([t], \mathcal{U})$, which is homogeneous with respect to U.

Remark 8.2.5. In the example of a cone σ realized in a vector space, the subdivision of Example 8.2.4 corresponds to choosing a ray passing through the interior of σ and subdividing σ with respect to this new ray. This simple idea depends both on the cone and the realization, which is why we have to go through defining the pieces of our subdivision from scratch in our case.

We now verify that we can obtain a subdivision from the order complex of P, which we give the standard linearization sheaf.

Lemma 8.2.6. *Let P be a lower Gorenstein* meet-semilattice with linearization sheaf $\mathcal{U}$, and let $Q = \{\hat{0}\} \cup \mathrm{ord}(P_{>0})$. Then we can find a subdivision $(Q, \mathcal{S}, f, \pi)$ of $(P, \mathcal{U})$, for $\mathcal{S}$ the standard linearization sheaf on Q and $f : \mathrm{ord}(P) \to P$ the map sending every chain to its last element.*

Proof. Let $Q^t = f^{-1}(t)$ and $Q^{\partial t} = f^{-1}(\partial t)$ for all $t \in P$. We only need to build $\pi : \mathcal{U} \to f_*\mathcal{S}$, which we do inductively. The base case is easy, since f restricts to a poset isomorphism $f^{-1}(P_{\leq 1}) \simeq P_{\leq 1}$. Now suppose $\pi|_{\partial t} : \mathcal{U}|_{\partial t} \to \mathcal{S}|_{Q^{\partial t}}$ is defined for some $t \in P_{>1}$. We have surjections $U_t \to U_s \to \mathcal{S}(\sigma)$ for $\sigma \in Q^t$ ending at $s < t$, which together induce an injective map

$$\lambda : U_t \to \mathcal{S}(Q^{\partial t}).$$

We also have an isomorphism

$$\mathcal{S}(Q^t) \simeq \mathbf{k} \oplus \mathcal{S}(Q^{\partial t})$$

as in Example 8.2.4. Choosing a generic linear function $\gamma : U_t \to \mathbf{k}$, the map $U_t \to \mathbf{k} \oplus \mathcal{S}(Q^{\partial t})$ defined by $x \mapsto (\gamma(x), \lambda(x))$ then gives a surjection $U_t \to \mathcal{S}(\sigma)$ for all $\sigma \in Q^t$. Thus, the claim is proved. □

Chapter 9

Parts of the Kähler package

The Kähler package from algebraic geometry consists of a trio of theorems which drastically restrict the class of rings that can occur as the cohomology rings of complex smooth projective varieties. The Kähler package for nonsmooth projective varieties rests on the theory of intersection cohomology.

The three theorems which comprise the Kähler package are Poincaré-Verdier duality, hard Lefschetz (HL), and the Hodge-Riemann bilinear relations. For projective toric varieties, these theorems prove Stanley's conjecture for the toric h-polynomials of boundary complex of a convex rational polytope. For nonrational polytopes, there exist combinatorial substitutes for the Kähler package in the theory of IHF. For example, two versions of Poincaré duality for IHF were developed in [BBFK00] and [BL00], and Karu proved the rest of the Kähler package in [Kar01].

We develop a duality in this paper for our setting, following another duality for IHF developed in [BBFK05]. We also conjecture an HL theorem in our setting, and dis-

cuss how it relates to Stanley's conjecture. We do not conjecture an analogue of Hodge-Riemann bilinear relations however, since it seems that these relations depend heavily on convexity/positivity, which we do not have in our setting. Instead, we look to [Adi18] and [APP21] for hope, where Lefschetz elements are shown to exist generically for the face rings of certain simplicial complexes.

9.1 Duality of pure sheaves

Let P be a subthin star normal lower Gorenstein* meet-semilattice of rank d with a homogeneous linearization sheaf $(\mathcal{U}, U)$. We have the associated sheaf of facewise polynomials $\mathcal{A}$ and minimal extension sheaf $\mathcal{E}$. In this section, we define a duality for pure sheaves on P, which eventually gives rise to an intersection pairing for $\mathcal{E}$.

To avoid the heavy machinery of derived categories as used for the duality defined in [BL00],[BL03], we instead follow the ideas of [BBFK05] to define our duality. This choice results in a definition of the dual $\mathbb{D}\mathcal{F}$ of any pure sheaf $\mathcal{F}$ which is elementary but not easy. We obtain a duality by definition, and it takes some work to show that the dual of a pure sheaf is pure.

Recall the determinant of an n-dimensional vector space V is the one-dimensional vector space $\det V := \bigwedge^n V$ concentrated in degree n.

Definition 9.1.1. An *orientation* of $(\mathcal{U}, U)$ is a choice of generators $\omega \in \det U$ and $\omega_s \in \det U_s$ for all $s \in P$ such that ω_t is the image of ω under the isomorphism $\det U \to \det U_t$ induced by the isomorphism $U \to U_t$ for all $t \in P_d$.

Remark 9.1.2. Picking an orientation in this sense is analogous in the fan case to picking an

orientation of the ambient vector space and orientations of all of the spans of cones. There is also an analogy with choosing a generator of the top cohomology group with coefficients in the orientation sheaf on a smooth orientable connected manifold, as discussed in [BL03].

We will also fix the following. Using Proposition 6.3.4, for any pair $s, t \in P$ with $s \lessdot t$, we may define *transition coefficients* $\varepsilon_s^t \in \{\pm 1\}$ such that for any diamond $[r, t] = \{r, s, s', t\}$ we have

$$\varepsilon_s^t \varepsilon_r^s \varepsilon_{s'}^t \varepsilon_r^{s'} = -1,$$

while for any $t, t' \in P_d$ with $t \wedge t \in P_{d-1}$ we have

$$\varepsilon_{t\wedge t'}^t \varepsilon_{t\wedge t'}^{t'} = -1.$$

Fix such a choice, and notice that this does not depend on any sheaves or orientations.

9.1.1 The definition of $\mathbb{D}\mathcal{F}$

Let $\mathcal{F}$ be a pure sheaf. We will first define the values $\mathbb{D}\mathcal{F}(t)$, and then define the restriction homomorphisms in two steps.

The value of $\mathbb{D}\mathcal{F}$ on basic opens:

For $t \in P$, we define

$$\mathbb{D}\mathcal{F}(t) := \mathrm{Hom}_{A_t}(\mathcal{F}(t, \partial t), A_t) \otimes \det U_t,$$

where an unadorned $\otimes$ always means $\otimes_{\mathbf{k}}$.

Restrictions maps for $s \lessdot t$

We will define the restriction maps first for $s \lessdot t$, and then extend this definition to any $r < t$ via diamond transitivity.

Assume $s \lessdot t$, so $\mathcal{U}(t,s)$ is one-dimensional. We will define the restriction

$$(\mathbb{D}\mathcal{F})^t_s := \varepsilon^t_s \phi_k \otimes \psi_k : \mathrm{Hom}(\mathcal{F}(t,\partial t), A_t) \otimes \det U_t \to \mathrm{Hom}(\mathcal{F}(s,\partial s), A_s) \otimes \det U_s$$

from maps

$$\psi_k : \det U_t \to \det U_s \quad \text{and} \quad \phi_k : \mathrm{Hom}(\mathcal{F}(t,\partial t), A_t) \to \mathrm{Hom}(\mathcal{F}(s,\partial s), A_s)$$

which separately depend on a choice of generator $k \in \mathcal{U}(t,s)$. While ϕ_k and ψ_k will depend on k, the definition ϱ^t_s will be independent of k because of equalities

$$\phi_{\lambda k} = \lambda \phi_k \quad \text{and} \quad \psi_{\lambda k} = \lambda^{-1} \psi_k$$

for any $\lambda \in \mathbf{k} \backslash \{0\}$.

Choose a generator $k \in \mathcal{U}(t,s)$, which defines an isomorphism $U_t \simeq \mathcal{U}(t,s) \oplus U_s$ and determines a map

$$\psi_k : \det A_t \simeq \mathcal{U}(t,s) \bigwedge \det U_s \to \det U_s$$

via $\psi_k(k \wedge \alpha) = \kappa(\alpha)$, where

$$\kappa : \bigwedge^{\rho(t)-1} U_t \to \det U_s$$

is the natural map induced by $\mathcal{U}^t_s$. Using κ as an anchor to compare choices of k, we have that $\psi_{\lambda k} = \lambda^{-1} \psi_k$ for any $\lambda \in \mathbf{k}^\times$.

Given any $f : \mathcal{F}(t, \partial t) \to A_t$, we define $\phi_k(f) \in \mathrm{Hom}_{A_s}(\mathcal{F}(s, \partial s), A_s)$ by

$$\phi_k(f)(x) \coloneqq \mathcal{A}_s^t(f(k \cdot \hat{x})) \in A_s,$$

where $\hat{x} \in \mathcal{F}(t)$ is a choice such that $\mathcal{F}_{\partial t}^t(\hat{x})$ is the extension of $x \in \mathcal{F}(s, \partial s)$ by zero to ∂t, where we notice that $k \cdot \hat{x} \in \mathcal{F}(t, \partial t)$. We automatically obtain the equality $\phi_{\lambda k} = \lambda \phi_k$ for $\lambda \in \mathbf{k}^\times$ by linearity. This description of $(\phi_k(f))(x)$ *a priori* may depend on the choice of $\hat{x}$, but we have the following.

Lemma 9.1.3. *The definition of $(\phi_k(f))(x)$ does not depend on the choice of $\hat{x} \in \mathcal{F}(t)$.*

Proof. We may assume $P = [t]$ and that $\mathcal{U}$ is a proper linearization sheaf for ease of notation, and we may also assume that $\mathcal{F} = \mathcal{L}^r$ for some $r \in [t]$ by Theorem 7.3.2. Even better, the case $r \not\leq s$ leaves us nothing to prove, so we can assume $r < s$.

Let $W' \subset U_t$ be a codimension 1 subspace such that $\mathcal{U}(t, r) \not\subset W$ for all $r \lessdot t$. Then the image $W \subset \mathcal{U}(\partial t)$ of $W' \subset \mathcal{U}(\partial t)$ is a homogenization of $\mathcal{W} \coloneqq \mathcal{U}|_{\partial t}$. We therefore have that $\mathcal{F}(\partial t)$ is a free module over $\mathrm{Sym}\, W$ by Theorem 8.1.5. Using the decomposition $U_t = W \oplus \mathcal{U}(t, s)$ and the definition of $\mathcal{F} = \mathcal{L}^r$, we know that any two lifts $\hat{x}, \hat{y} \in \mathcal{F}(t)$ of the extension of $x \in \mathcal{F}(s, \partial s)$ by zero to ∂t must have $\hat{x} = \hat{y} + k \cdot z$ for some $z \in A_t$. Then we have

$$f(k \cdot \hat{x}) = f(k \cdot \hat{y}) + k \cdot f(k \cdot z),$$

which in turn gives $\mathcal{A}_s^t(f(k \cdot \hat{x})) = \mathcal{A}_s^t(f(k \cdot \hat{y}))$ because A_s is annihilated by k. □

Restriction maps for general $s < t$

For any $s < t$ we let $(\mathbb{D}\mathcal{F})_s^t$ be the composition $(\mathbb{D}\mathcal{F})_s^{s_1} \circ \cdots \circ (\mathbb{D}\mathcal{F})_{s_{m-1}}^{s_m} \circ (\mathbb{D}\mathcal{F})_{s_m}^t$ for some maximal chain $s \lessdot s_1 \lessdot \cdots \lessdot s_m \lessdot t$.

Lemma 9.1.4. *The restriction map $\mathbb{D}\mathcal{F}_s^t$ is well-defined.*

Proof. To prove that this is independent of the choice of maximal chain, we use the diamond transitivity of Gorenstein* $[t]$ reduce to checking

$$(\mathbb{D}\mathcal{F})_r^s \circ (\mathbb{D}\mathcal{F})_s^t = (\mathbb{D}\mathcal{F})_r^{s'} \circ (\mathbb{D}\mathcal{F})_{s'}^t$$

for some diamond $[r,t] = \{r,s,s',t\}$. Let $k \in \mathcal{U}(t,s)$ and $h \in \mathcal{U}(t,s')$ be generators, which together form a basis of $\mathcal{U}(t,r) \subset U_t$. In particular, $h_0 := \mathcal{U}_s^t(h)$ generates $\mathcal{U}(s,r)$ and $k_0 := \mathcal{U}_{s'}^t(k)$ generates $\mathcal{U}(s',r)$. Then we have

$$\psi_{k_0} \circ \psi_h = -\psi_{h_0} \circ \psi_k$$

since $\omega_t = a \cdot k \wedge h \wedge \alpha = -a \cdot h \wedge k \wedge \alpha \in \det U_t$ for some fixed $a \in \mathbf{k}^\times$ and $\alpha \in \det(U_t/\mathcal{U}(U,U_r))$, while the transition coefficients correct this sign issue.

Finally, we notice that

$$\phi_{h_0} \circ \phi_k = \phi_{k_0} \circ \phi_h$$

since both sides take $f : \mathcal{F}(t,\partial t) \to A_t$ to the function $F : \mathcal{F}(r,\partial r) \to A_r$ defined by

$$F(x) = \mathcal{A}_r^t(f(kh\hat{\hat{x}}))$$

for some $\hat{\hat{x}} \in \mathcal{F}(t)$ such that $\hat{\hat{x}}$ restricts to the extension of $x \in \mathcal{F}(r,\partial r)$ by zero to the rest of ∂t, which follows from the fact that

$$\mathcal{F}_s^t(h\hat{\hat{x}}) = h_0\mathcal{F}_s^t(\hat{\hat{x}}) \in \mathcal{F}(s,\partial s) \quad \text{and} \quad \mathcal{F}_{s'}^t(k\hat{\hat{x}}) = k_0\mathcal{F}_{s'}^t(\hat{\hat{x}}) \in \mathcal{F}(s',\partial s'),$$

finishing the claim. □

Thus, we have a well-defined sheaf $\mathbb{D}\mathcal{F}$ for any pure sheaf $\mathcal{F}$ on P. For the word "duality" to make sense, we must show that our duality is functorial and takes pure sheaves to pure sheaves.

Lemma 9.1.5. *$\mathbb{D}$ defines a functor from the category of pure sheaves on $(P, \mathcal{U})$ to the category of $\mathcal{A}$-modules.*

Proof. For a morphism $\Phi : \mathcal{F} \to \mathcal{G}$ of pure sheaves, we get A_t-module morphisms $\Phi_t : \mathcal{F}(t) \to \mathcal{G}(t)$ which restrict to morphisms $\Phi_t : \mathcal{F}(t, \partial t) \to \mathcal{G}(t, \partial t)$, and therefore give stalkwise morphisms $\mathbb{D}\mathcal{G}(t) \to \mathbb{D}\mathcal{F}(t)$. To prove that these commute with restrictions, suppose $s \lessdot t$ and $k \in \mathcal{U}(t, s)$ is a generator. We only need to prove that the diagram

$$\begin{array}{ccc} \mathrm{Hom}_{A_t}(\mathcal{G}(t, \partial t), A_t) & \xrightarrow{\Phi_t^*} & \mathrm{Hom}_{A_t}(\mathcal{F}(t, \partial t), A_t) \\ \downarrow{\scriptstyle \phi_k} & & \downarrow{\scriptstyle \phi_k} \\ \mathrm{Hom}_{A_s}(\mathcal{G}(s, \partial s), A_s) & \xrightarrow{\Phi_s^*} & \mathrm{Hom}_{A_s}(\mathcal{F}(s, \partial s), A_s) \end{array}$$

commutes. To this end, let $g : \mathcal{G}(t, \partial t) \to A_t$ and $x \in \mathcal{F}(s, \partial s)$, and let $y \in \mathcal{G}(s, \partial s)$ be the image of x under $\Phi_s : \mathcal{F}(s) \to \mathcal{G}(s)$. Then if $\hat{x} \in \mathcal{F}(t)$ is a lift of the extension of $x \in \mathcal{F}(s, \partial s)$ by zero on the rest of ∂t, we have

$$\phi_k(\Phi_t^*(g))(x) = \mathcal{A}_s^t(\Phi_t^*(g)(k\hat{x})) = \mathcal{A}_s^t(g(k \cdot \Phi_t(\hat{x}))) = \phi_k(g)(\Phi_s(x)) = (\Phi_s^*(\phi_k(g)))(x)$$

since $\Phi_t(\hat{x}) \in \mathcal{G}(t)$ is a lift of the extension of $\Phi_t(x) \in \mathcal{G}(s, \partial s)$ to zero on ∂t. □

9.1.2 Properties of the duality

To prove that $\mathbb{D}$ takes pure sheaves to pure sheaves, we will first show that we can compute the global sections of $\mathbb{D}\mathcal{F}$ in one of our favourite special cases.

Notice there is a natural isomorphism

$$\Theta : \bigoplus_{t \in P_d} \mathbb{D}\mathcal{F}(t) \simeq \mathrm{Hom}(\mathcal{F}(P, P_{<d}), A) \otimes \det U,$$

which is a consequence of the homogenization U and the orientation of $(\mathcal{U}, U)$, together with the application of the additive functor $\mathrm{Hom}_A(\,\cdot\,, A) \otimes \det U$ to the obvious direct sum decomposition

$$\mathcal{F}(P, P_{<d}) \simeq \bigoplus_{t \in P_d} \mathcal{F}(t, \partial t).$$

There are also inclusions

$$\mathrm{Hom}(\mathcal{F}(P), A) \subset \mathrm{Hom}_A(\mathcal{F}(P, \partial P), A) \subset \mathrm{Hom}(\mathcal{F}(P, P_{<d}), A)$$

coming from applying $\mathrm{Hom}_A(\,\cdot\,, A)$ to the inclusions of torsion-free A-modules of the same rank $\mathcal{F}(P, P_{<d}) \subset \mathcal{F}(P, \partial P) \subset \mathcal{F}(P)$. Finally, recall the chain complex $\mathcal{C}_\bullet(P; \mathbb{D}\mathcal{F})$ for subthin P with

$$\mathbb{D}\mathcal{F}(P) \simeq Z(P, \partial P; \mathbb{D}\mathcal{F}) \quad \text{and} \quad \mathbb{D}\mathcal{F}(P, \partial P) \simeq Z(P; \mathbb{D}\mathcal{F})$$

by Proposition 6.3.8.

Theorem 9.1.6. *Fix an orientation of* $(\mathcal{U}, U)$. *Then* Θ *induces isomorphisms*

$$\mathbb{D}\mathcal{F}(P) \simeq \mathrm{Hom}_A(\mathcal{F}(P, \partial P), A) \otimes \det U \quad \textit{and} \quad \mathbb{D}\mathcal{F}(P, \partial P) \simeq \mathrm{Hom}_A(\mathcal{F}(P), A) \otimes \det U.$$

The goal is to prove that Θ provides isomorphisms across each level of the diagram

$$\begin{array}{ccc}
\oplus_{t \in P_d} \mathbb{D}\mathcal{F}(t) & \xrightarrow{\Theta} & \mathrm{Hom}_A(\mathcal{F}(P, P_{<d}), A) \otimes \det U \\
\cup & & \cup \\
Z(P, \partial P; \mathbb{D}\mathcal{F}) & & \mathrm{Hom}_A(\mathcal{F}(P, \partial P), A) \otimes \det U \\
\cup & & \cup \\
Z(P; \mathbb{D}\mathcal{F}) & & \mathrm{Hom}_A(\mathcal{F}(P), A) \otimes \det U
\end{array}$$

The proof of Theorem 9.1.6 will follow from the next two lemmas, which depend on the following discussion and reductions.

Choose $h \coloneqq \prod_{i=1}^{r} h_i \in A$ to be a minimal degree product of some nonzero $h_i \in U$ such that for every $s \in P_{d-1}$, $h_i \in \mathcal{U}(U, U_s)$ for some (uniquely determined) i. We let $K_i \coloneqq \ker(U \to \operatorname{span}(h_i))$, so that if $h_i \in \mathcal{U}(U, U_s)$ then we identify $U_s \simeq K_i$. The orientation of $(\mathcal{U}, U)$ gives isomorphisms $\det U \simeq \det U_t$, which also determine isomorphisms $\det U \to \det K_i \simeq \det U_s$ using ψ_{h_i} for all i.

Thus, we may drop all determinant factors in the diagram above, and for each $s \in P_d \cup P_{d-1}$ replace $\mathbb{D}\mathcal{F}(s)$ with $\operatorname{Hom}_{A_s}(\mathcal{F}(s, \partial s), A_s)$. The restrictions $\mathbb{D}\mathcal{F}(t) \to \mathbb{D}\mathcal{F}(s)$ become $\pm\phi_{h_i}$, with both signs occurring for the distinct covers t, t' of any $s \in P_{d-1} \backslash \partial P$ by our choice of transition coefficients.

Using the inclusions of torsion-free A-modules

$$h \cdot \mathcal{F}(P, \partial P) \subset h \cdot \mathcal{F}(P) \subset \mathcal{F}(P, P_{<d}),$$

we see that given $f : \mathcal{F}(P, P_{<d}) \to A$ and $X \in \{\mathcal{F}(P), \mathcal{F}(P, \partial P)\}$, we have f is the restriction of some A-module map $F : X \to A$ if and only if

$$f(h \cdot x) \in h \cdot A$$

for any $x \in X$, since if this holds we may define $F(x) = f(h \cdot x)/h \in A$, which is well-defined since multiplication by $h \in A$ is injective.

Thus, we can prove Theorem 9.1.6 by proving the following two lemmas using $h \in A$ and our reductions. For both, let

$$\zeta = (\zeta_t) \in \bigoplus_{t \in P_d} \mathbb{D}\mathcal{F}(t).$$

Lemma 9.1.7. *We have*

$$\zeta \in Z(P, \partial P; \mathbb{D}\mathcal{F}) \implies \Theta(\zeta)(h \cdot \mathcal{F}(P, \partial P)) \subset h \cdot A$$

and

$$\zeta \in Z(P; \mathbb{D}\mathcal{F}) \implies \Theta(\zeta)(h \cdot \mathcal{F}(P)) \subset h \cdot A.$$

Proof. Consider the first implication. We want to show that h_i divides

$$\Theta(\zeta)(h \cdot x) = \sum_{t \in P_d} \zeta_t(h \cdot x_t) \in A$$

for all indices i and all $x = (x_t) \in \mathcal{F}(P, \partial P) \subset Z(P, \partial P; \mathcal{F})$. Define $g_i := h/h_i$ for all i. For a fixed index i and $t \in P_d$, if the statement

$$U_s = K_i \text{ for some } s \lessdot t \quad \implies \quad s \in \partial P \tag{9.1.1}$$

is true, then $g_i x_t \in \mathcal{F}(t, \partial t)$ because $h_j x_t$ divides $g_i x_t$ for all $j \neq i$, so $\mathcal{F}_s^t(g_i x_t) \neq 0$ implies $s \in \partial P$, but $\mathcal{F}_s^t(x_t) = 0$ for $s \in \partial P$ because $x \in \mathcal{F}(P, \partial P)$. Thus, $\zeta_t(h x_t) = h_i \cdot \zeta_t(g_i x) \in h \cdot A$.

On the other hand, if (9.1.1) is false, meaning that for a fixed i and t there is an $s \lessdot t$ with $s \notin \partial P$ such that $U_s = K_i$, then there exists a unique $t' \in P_d$ such that $s \lessdot t' \neq t$ because P is subthin. We claim that h_i divides the sum $\zeta_t(h x_t) + \zeta_{t'}(h x_{t'})$, which is equivalent to $\mathcal{A}_s^t(\zeta_t(h x_t)) = -\mathcal{A}_s^{t'}(\zeta_{t'}(h x_{t'}))$ since $h_i \cdot A$ is the annihilator of A_s as an A-module. For this, notice $g_i x_t \in \mathcal{F}(t)$ is a lift of the extension of $\mathcal{F}_s^t(g_i x_t) \in \mathcal{F}(s, \partial s)$ by zero to the rest of ∂t, and as in the definition of ϕ_{h_i} we obtain $\mathcal{A}_s^t(\zeta_t(h x_t)) = (\phi_{h_i}(\zeta_t))(g_i x_t)$. Since we assumed $\zeta \in Z(P, \partial P; \mathbb{D}\mathcal{F})$, we know $(\mathbb{D}\mathcal{F})_s^t(\zeta_t) = (\mathbb{D}\mathcal{F})_s^{t'}(\zeta_{t'})$ in $\mathbb{D}\mathcal{F}(s)$. Then the transition coefficients force $\phi_{h_i}(\zeta_t) = -\phi_{h_i}(\zeta_{t'})$, finishing the first claim.

In the absolute case, $\zeta \in Z(P;\mathbb{D}\mathcal{F}) \implies \Theta(\zeta)(h\cdot\mathcal{F}(P)) \subset h\cdot A$, the argument follows similarly to above, but also introduces the following case. If an index i and $t \in P_d$ have $U_s = K_i$ for $s \lessdot t$ with $s \in \partial P$, then $(\mathbb{D}\mathcal{F})^t_s(\zeta_t) = 0$, so that $(\mathbb{D}\mathcal{F})^t_s(h\zeta_t) = 0$ is automatically in $h_i \cdot A$. □

Lemma 9.1.8. *We have*

$$\Theta(\zeta)(h\cdot\mathcal{F}(P,\partial P)) \subset h\cdot A \implies \zeta \in Z(P,\partial P;\mathbb{D}\mathcal{F})$$

and

$$\Theta(\zeta)(h\cdot\mathcal{F}(P)) \subset h\cdot A \implies \zeta \in Z(P;\mathbb{D}\mathcal{F}).$$

Proof. Consider the first implication. We wish to show that for any $t, t' \in P_d$ with $t \wedge t' = s \in P_{d-1}$, we have $(\mathbb{D}\mathcal{F})^t_s(\zeta_t) = (\mathbb{D}\mathcal{F})^{t'}_s(\zeta_{t'})$. Let i be the index such that $U_s = K_i$. Then for any $x_0 \in \mathcal{F}(s,\partial s)$, extend it by zero to the rest of ∂t and pick a lift $x \in \mathcal{F}(t)$, and do the same for $\partial t'$ and obtain $x' \in \mathcal{F}(t')$. Then there is an $x_1 \in \mathcal{F}(P,\partial P)$ such that $(x_1)_t = x$, $(x_1)_{t'} = x'$, and $(x_1)_{t''} = 0$ for all $t'' \in P_d\backslash\{t,t'\}$. Then the equalities

$$h_i\Theta(\zeta)(x_1) = \Theta(\zeta)(h_i x_1) = \zeta_t(h_i x) + \zeta_{t'}(h_i x'),$$

when restricted via $A \to A_s$ yield

$$0 = \phi_{h_i}(\zeta_t)(x_0) + \phi_{h_i}(\zeta_{t'})(x_0) = ((\mathbb{D}\mathcal{F})^t_s(\zeta_t) - (\mathbb{D}\mathcal{F})^{t'}_s(\zeta_{t'}))(x_0).$$

For the final implication, the only case left to take care of is when $s \in P_{d-1}\cap\partial P$ and $s \lessdot t$. Extend some $x_0 \in \mathcal{F}(s,\partial s)$ by zero to the rest of ∂t and then lift this to $x \in \mathcal{F}(t)$, and further extend x to $x_1 \in \mathcal{F}(P)$. Then similarly to above, we get that $h_i\Theta(\zeta)(x_1) = \zeta_t(h_i x)$ implies $(\mathbb{D}\mathcal{F})^t_s(\zeta_t)(x_0) = 0$, finishing the claim. □

Thus, Theorem 9.1.6 is proved, and we may apply it to obtain the following.

Theorem 9.1.9. *For any pure sheaf $\mathcal{F}$ on any CM lower Eulerian meet-semilattice P, $\mathbb{D}\mathcal{F}$ is also pure.*

Proof. We need only prove that $\mathbb{D}\mathcal{F}(t) \to \mathbb{D}\mathcal{F}(\partial t)$ is surjective for any $t \in P$. Since everything is handled locally, we may assume $P = [t]$ is Gorenstein*, and we assume $U_t = U$. Thus, we can import $Q, L, \mathcal{V}, W, T$ as in Example 8.2.4. Then we set $B = \operatorname{Sym} W$, so $B[T] = A$. There is a natural isomorphism

$$\mathbb{D}\mathcal{F}(\partial t) \simeq \operatorname{Hom}_B(\mathcal{F}(\partial t), B) \otimes \det W$$

by Theorem 9.1.6.

Let $\mathcal{A}_Q$ be the sheaf of facewise polynomials for $(Q, \mathcal{V})$, and let $p : Q \to \partial t$ be the obvious two-to-one poset map. Define the $\mathcal{A}_Q$-module $\mathcal{H} = p^*\mathcal{F}|_{\partial t}$. There are isomorphisms

$$\mathcal{H}(Q) \simeq A \otimes_B \mathcal{F}(\partial t) \quad \text{and} \quad \mathcal{H}(\partial t) \simeq \mathcal{F}(\partial t),$$

as well as a "Thom" isomorphism

$$\mu : \mathcal{H}(Q) \to \tau \cdot \mathcal{H}(Q) = \mathcal{H}(Q, \partial t)$$

given by multiplication with the element $\tau \in \mathcal{A}_Q(Q)_1$ which restricts to $(0, T) \in \mathcal{A}_Q(q * s)_1 = U_s \oplus L$ and $0 \in \mathcal{A}(s)_1 = U_s$ for all $s \in \partial t$. The composition

$$\det U \simeq L \otimes \det W \simeq L \otimes \det U_s \to \det U_s \simeq \det W$$

given by $T \wedge \alpha \mapsto \alpha$ induces isomorphisms $\det U \simeq \det U_s \simeq \det W$ which are independent of $s \lessdot t$, so we may forget the determinant factors in what follows.

We wish to show that we can find a preimage in $\mathbb{D}\mathcal{F}(t) \simeq \mathrm{Hom}_B(\mathcal{F}(t, \partial t), B)$ of any $\beta \in \mathbb{D}\mathcal{F}(\partial t)$. Our candidate γ is the composition

$$\mathcal{F}(t, \partial t) \xrightarrow{i} \mathcal{H}(Q, \partial t) \xrightarrow{(\mu)^{-1}} \mathcal{H}(Q) \simeq A \otimes_B \mathcal{F}(\partial t) \xrightarrow{id \otimes \beta} A \otimes_B B = A,$$

where the map $i = j\,|_{\mathcal{F}(t,\partial t)}$, and $j : \mathcal{F}(t) \to \mathcal{H}(Q)$ is defined to be the map such that the composition

$$\mathcal{F}(t) \xrightarrow{j} \mathcal{H}(Q) \to \mathcal{H}(\partial t) = \mathcal{F}(\partial t)$$

agrees with the restriction $\mathcal{F}(t) \to \mathcal{F}(\partial t)$. Such j exists because $\mathcal{F}(t)$ is a free (hence, projective) A-module and $\mathcal{H}(Q) \to \mathcal{H}(\partial t)$ is surjective. Furthermore, this j necessarily has $j(\mathcal{F}(t, \partial t)) \subset \mathcal{H}(Q, \partial t)$.

We now need to show that our candidate γ restricts to β. We will prove that they agree when restricted to each maximal $s \in \partial t$. Here, we identify the naturally isomorphic algebras B and A_s.

For any $x \in \mathcal{F}(s, \partial s) \subset \mathcal{F}(s)$, every preimage $\overline{x} \in \mathcal{F}(t)$ of x has

$$j(\overline{x}) = 1 \otimes x + \tau \cdot y \in \mathcal{H}(Q) \simeq A \otimes_B \mathcal{F}(\partial t),$$

where $y \in \mathcal{H}(Q)$ is some correction term. Thus, following the element $T \cdot \overline{x}\mathcal{F}(t, \partial t)$ through γ, we get that

$$\gamma(T \cdot \overline{x}) = (id \otimes \beta)(1 \otimes x + \tau \cdot y) = \beta(x) + T \cdot y',$$

for some $y' \in A$, so that $\gamma \in ((\mathbb{D}\mathcal{F})^t_s)^{-1}(\beta)$. □

9.1.3 Duality for $\mathcal{E}$ and an intersection pairing

Finally, we consider how our duality behaves with respect to the minimal extension sheaf $\mathcal{E}$. This will result in an intersection pairing between $\mathcal{E}(P)$ and $\mathcal{E}(P, \partial P)$.

Theorem 9.1.10 (Biduality)**.** *Every pure sheaf is reflexive, meaning that there exists a natural isomorphism*

$$\mathcal{F} \xrightarrow{\cong} \mathbb{D}\mathbb{D}\mathcal{F}.$$

Proof. This follows exactly as in [BBFK05], [BL03]. For any $t \in P$, we compose the following natural isomorphisms:

$$\begin{aligned}
\mathbb{D}\mathbb{D}\mathcal{F}(t) &= \mathrm{Hom}_{A_t}(\mathbb{D}\mathcal{F}(t, \partial t), A_t)) \otimes \det U_t \\
&\cong \mathrm{Hom}_{A_t}(\mathrm{Hom}_{A_t}(\mathcal{F}(t), A_t) \otimes \det U_t, A_t) \otimes \det U_t \\
&\cong \mathrm{Hom}_{A_t}(\mathrm{Hom}_{A_t}(\mathcal{F}(t), A_t) \otimes \det U_t, A_t \otimes \det U_t) \\
&\cong \mathrm{Hom}_{A_t}(\mathrm{Hom}_{A_t}(\mathcal{F}(t), A_t), A_t) \\
&\cong \mathcal{F}(t),
\end{aligned}$$

where the first isomorphism is Theorem 9.1.6 applied to $([t], \mathcal{U}|_{[t]}, U_t)$ and the last is because the free A_t-module $\mathcal{F}(t)$ is reflexive. □

Corollary 9.1.11. *For any $\mathcal{L}^s$, we have that $\mathbb{D}\mathcal{L}^s \simeq \mathcal{L}^s \otimes \det U_s$. In particular, we have a degree zero isomorphism*

$$\mathcal{E} \cong \mathbb{D}\mathcal{E}.$$

Proof. Duality respects direct sum decompositions of pure sheaves, and therefore $\mathbb{D}\mathcal{L}^s$ is simple. Furthermore, the support of $\mathbb{D}\mathcal{F}$ must agree with the support of a pure sheaf $\mathcal{F}$,

since $\mathcal{F}(\partial t)$ is a torsion A_t-module and hence $\mathcal{F}(t) = 0$ if and only if $\mathcal{F}(t, \partial t) = 0$. Since a simple sheaf is uniquely determined up to isomorphism by its support and a grade shift, we are done because $\mathbb{D}\mathcal{L}^s(s) \simeq \mathrm{Hom}_{A_s}(A_s, A_s) \otimes \det U_s \simeq A_s \otimes \det U_s$ by Theorem 9.1.6. □

We can use the duality $\mathbb{D}$ to define an intersection pairing as follows. Given any degree zero sheaf homomorphism $\mathcal{E} \to \mathbb{D}\mathcal{E}$ extending the identification

$$\mathcal{E}(\hat{0}) = \mathbf{k} \xrightarrow{1 \mapsto 1^\vee} \mathbf{k}^* = \mathbb{D}\mathcal{E}(\hat{0}),$$

we can define an *intersection product*

$$\mathcal{E}(P) \times \mathcal{E}(P, \partial P) \to A[d]$$

given equivalently by the composition

$$\mathcal{E}(P) \longrightarrow \mathbb{D}\mathcal{E}(P) \overset{\Theta}{\cong} \mathrm{Hom}_A(\mathcal{E}(P, \partial P), A) \otimes \det U \cong \mathrm{Hom}_A(\mathcal{E}(P, \partial P), A[d]).$$

Theorem 9.1.12. *The isomorphism $\mathcal{E} \cong \mathbb{D}\mathcal{E}$ of Corollary 9.1.11 gives such a map $\mathcal{E} \to \mathbb{D}\mathcal{E}$, and the induced intersection product is a dual pairing of reflexive A-modules. If ∂P is empty or CM, then the intersection product reduces to a dual pairing*

$$\overline{\mathcal{E}(P)} \times \overline{\mathcal{E}(P, \partial P)} \longrightarrow \mathbf{k}[d]$$

of finite-dimensional graded vector spaces.

We do not know if this pairing is canonical, but we will conjecture that it is for certain $(\mathcal{U}, U)$, which we will discuss in the next section.

9.2 Hard Lefschetz conjecture and toric h-vectors

In this section, we state a hard Lefschetz conjecture for our theory, and then prove a few interesting statements assuming it is true. Among these statements is Stanley's conjecture about the positivity of toric h-polynomials and a uniqueness result of an intersection pairing on global sections of $\mathcal{E}$ in certain cases.

We assume that P is a CM lower Eulerian meet-semilattice. Given a homogeneous linearization sheaf $(\mathcal{U}, U)$ on P, we define

$$\mathrm{IH}(P, \mathcal{U}, U) := \overline{\mathcal{E}(P)} = \mathcal{E}(P)/(\mathrm{Sym}\, U)^+ \mathcal{E}(P)$$

to be the *intersection homology* of $(P, \mathcal{U}, U)$. Since P is equivariantly formal, $\mathrm{IH}(P, \mathcal{U}, U)$ is a finite-dimensional graded vector space. The following is the main conjecture of this paper.

Conjecture 9.2.1 (Hard Lefschetz)**.** *Let P be a CM lower Eulerian meet-semilattice. Then there exists*

- *a homogeneous linearization sheaf $(\mathcal{U}, U)$,*
- *an element $\ell \in \mathcal{A}(P)_1$, and*
- *a choice of homogenization $W_s \subset U_s$ of $\mathcal{U}|_{\partial s}$ for all $s \in P$,*

such that for each $s \in P$, acting by $\ell^{\rho(s)-2k}$ induces an isomorphism

$$\mathrm{IH}(\partial s, \mathcal{U}|_{\partial s}, W_s)_k \simeq \mathrm{IH}(\partial s, \mathcal{U}|_{\partial s}, W_s)_{\rho(s)-k}$$

for all $0 \leq k \leq \lfloor \frac{\rho(s)}{2} \rfloor$.

Such $\ell \in \mathcal{A}(P)_1$ are called *Lefschetz elements*. In the fan case, [Kar01] proved that a strictly convex function on projective (hence complete) fan is a Lefschetz element for combinatorial intersection homology of that fan. In our case, we do not have such a specific conjecture about what a Lefschetz element might look like. However, all is not lost: following [Adi18] and [APP21], we simply conjecture that there is at least one Lefschetz element which we can find by avoiding some (yet to be determined) pitfalls.

For the remainder of the paper, we write (HL) to refer to Conjecture 9.2.1.

9.2.1 Making the intersection pairing canonical

With the information we currently have, the pairing of Theorem 9.1.12 may not be canonical. In this direction, we have the following conjecture.

Conjecture 9.2.2. *There is at most one sheaf homomorphism $\mathcal{E} \to \mathbb{D}\mathcal{E}$ of degree 0 extending the identification*

$$\mathcal{E}(\hat{0}) = \mathbf{k} \xrightarrow{1 \mapsto 1^\vee} \mathbf{k}^* = \mathbb{D}\mathcal{E}(\hat{0}).$$

To see how (HL) proves Conjecture 9.2.2, we have the following.

Lemma 9.2.3. *Assume (HL) holds, and gives $(\mathcal{U}, U, \ell, \{W_t\}_{t \in P})$. Then for all $t \in P_{>0}$ the following equivalent conditions hold:*

1. *$\mathcal{E}(t)$ is generated as an A-module in degrees $\leq \rho(s)/2$,*
2. *$\mathcal{E}(t, \partial t)$ is generated in grades $\geq \rho(t)/2$.*

Furthermore, the only endomorphisms of $\mathcal{E}$ as a graded $\mathcal{A}$-module are given by multiplication by $\mathbf{k}$.

Proof. We may assume $t \in P_d$ for ease of notation, and identify $U \simeq U_t$ and $A \simeq A_t$. Then (HL) provides $\ell \in \mathcal{A}(P)_1$ and $W := W_t \subset U$. Let $B := \operatorname{Sym} W \subset A$. Then $\mathcal{E}|_{\partial t}(\partial t) = \mathcal{E}(\partial t)$ is a free B-module by Theorem 8.1.5, and is additionally an $\mathcal{A}(\partial t)$-module by taking global sections of the sheaf map $\mathcal{A}|_{\partial t} \to \mathcal{E}|_{\partial t}$. Choosing a splitting $U \to W$, or equivalently, choosing an element $T \in U \backslash W$, induces an $A \simeq B[T]$-module structure on $\mathcal{E}|_{\partial t}(\partial t)$, where T acts as the Lefschetz element $\ell \in \mathcal{A}(\partial t)_1$.

By the definition of $\mathcal{E}$, there is an isomorphism

$$\mathcal{E}(t)/B^+ \cdot \mathcal{E}(t) \simeq \mathcal{E}(\partial t)/B^+ \cdot \mathcal{E}(\partial t).$$

As A-modules, the action of T is trivial on $\overline{\mathcal{E}(t)}$, which forces $\overline{\mathcal{E}(t)} \simeq \overline{\mathcal{E}|_{\partial t}(\partial t)}/L \cdot \overline{\mathcal{E}|_{\partial t}(\partial t)}$. Thus, by (HL) we have that $\overline{\mathcal{E}(t)}$ can only be nonzero in nonnegative degrees $\leq \lfloor d/2 \rfloor$.

Thus, we have proved (1), and this is equivalent to (2) via the isomorphism $\mathcal{E} \simeq \mathbb{D}\mathcal{E}$ of Corollary 9.1.11. Finally, we must show that an homomorphism $\mathcal{E}(\partial s) \to \mathcal{E}(\partial s)$ extends uniquely to $\mathcal{E}(s) \to \mathcal{E}(s)$. By (1) the A-module $\mathcal{E}(s)$ is generated by homogeneous elements of degrees $< \lfloor \rho(s)/2 \rfloor$, while (2) implies $E(s, \partial s) = 0$ in degrees $\leq \lfloor \rho(s)/2 \rfloor$. Thus, the restrictions $\mathcal{E}(s) \to \mathcal{E}(\partial s)$ must be isomorphisms in degrees $< \lfloor \rho(s)/2 \rfloor$, so that our promised uniqueness follows. The existence of such an extension is due to the fact that $\mathcal{E}(s)$ as a free A_s-module. □

Conjecture 9.2.2 then follows from the uniqueness of the isomorphism $\mathcal{E} \cong \mathbb{D}\mathcal{E}$ which sending $1 \mapsto 1^\vee$.

9.2.2 Toric h-polynomials

The initial motivation for considering this project was to prove Conjecture 5.1.1, that the coefficients of toric h-polynomials are nonnegative for CM lower Eulerian meet-semilattices. To this end, we discuss how our minimal extension sheaves are conjecturally related to these toric h-polynomials and g-polynomials.

Definition 9.2.4. Suppose $(\mathcal{U}, U)$ is a homogeneous linearization sheaf on a subthin CM lower Eulerian meet-semilattice P. Then we let $ih_{(P,\mathcal{U},U)}(x)$ be the (intersection) Poincaré polynomial of $\mathrm{IH}(P, \mathcal{U}, U)$, and let $e_{(P,\mathcal{U})}(x)$ be the (equivariant) Poincaré series of $\mathcal{E}(P)$.

Remark 9.2.5. Suppose P or $\widehat{P}$ is Gorenstein*. Since P is equivariantly formal by Theorem 8.1.5, we have

$$e_{(P,\mathcal{U})}(x) = \frac{ih_{(P,\mathcal{U},U)}(x)}{(1-x)^d}.$$

This is true because $\mathcal{E}(P)$ is a free graded A-module, so we may use the Künneth formula $\mathcal{E}(P) \simeq A \otimes_{\mathbf{k}} \mathrm{IH}(P)$. Then taking Poincaré series respects this formula, and the Poincaré series of $A = \operatorname{Sym} U$ is $(1-x)^{-d}$.

For the minimal extension sheaf $\mathcal{E}$ of $(P, \mathcal{U})$, and for any $X \subset P$ open, notice that $\mathcal{E}|_X$ is the minimal extension sheaf for $(X, \mathcal{U}|_X)$. Fixing $(P, \mathcal{U})$, we can the define, for any $t \in P$, the polynomials $ih_t(x) := ih_{([t],\mathcal{U}|_{[t]},U_t)(x)}$ and $ih_{(\partial t,W_t)}(x) := ih_{(\partial t,\mathcal{U}|_{\partial t},W_t)}(x)$, where W_t is some homogenization of $\mathcal{U}|_{\partial t}$.

For the following, recall the notation $h(P, x)$ and $g(P, x)$ for the toric h-polynomial and toric g-polynomial, respectively. Also recall the set of 0-cycles $Z(P; \mathcal{E})$ of the chain complex $\mathcal{C}_\bullet(P, \mathcal{E})$.

Conjecture 9.2.6. *Let P be a CM lower Eulerian meet-semilattice, and assume Conjecture 9.2.1 is true and gives $(\mathfrak{U}, U, \ell, \{W_t\}_{t\in P})$. Then $x^d h(P, x^{-1})$ is the Poincaré polynomial of $\overline{Z(P;\mathcal{E})}$. In particular, Conjecture 5.1.1 is true.*

Proof. We have $Z(P;\mathcal{E})$ is a free graded A-module by Proposition 8.1.4, and thus the Poincaré polynomial of $\overline{Z(P;\mathbb{E})}$ is

$$f(P;\mathcal{E})(x) = \sum_{t\in P} ih_t(x)(x-1)^{d-\rho(t)}$$

since $\mathcal{C}_\bullet(P,\mathcal{E})$ is acyclic. By induction on rank for Gorenstein* meet-semilattices, we can use the argument given in the proof of Lemma 9.2.3 (with ℓ and W_t), together with the fact that our perfect pairing $\overline{\mathcal{E}(\partial t)} \times \overline{\mathcal{E}(\partial t)} \to k[\rho(t)]$ implies

$$ih_{(\partial t, W_t)}(x) = x^{\rho(t)} ih_{(\partial t, W_t)}(x^{-1}),$$

to show that

$$ih_{(\partial t, W_t)}(x) = h(\partial t, x) \quad \text{and} \quad ih_t(x) = g(\partial t, x).$$

Thus, we have our claim. □

Bibliography

[Adi18] K. Adiprasito. Combinatorial Lefschetz theorems beyond positivity. *arXiv preprint arXiv:1812.10454*, 2018. 147, 162

[ÁDL03] I. Ágoston, V. Dlab, and E. Lukács. Quasi-hereditary extension algebras. *Algebras and Representation Theory*, 6(1):97–117, Mar 2003. 83

[APP21] K. Adiprasito, S. A. Papadakis, and V. Petrotou. Anisotropy, biased pairings, and the Lefschetz property for pseudomanifolds and cycles. *arXiv preprint arXiv:2101.07245*, 2021. 107, 147, 162

[Ath11] C. Athanasiadis. On the Möbius function of a lower Eulerian Cohen–Macaulay poset. *Journal of Algebraic Combinatorics*, 35:373–388, 05 2011. 107

[Bac80] K. Baclawski. Cohen-Macaulay ordered sets. *Journal of Algebra*, 63(1):226–258, 1980. 114, 116

[BBFK00] G. Barthel, J.-P. Brasselet, K.-H. Fieseler, and L. Kaup. Combinatorial Intersection Cohomology for Fans. *Tohoku Mathematical Journal*, 54, 03 2000. 104, 105, 137, 138, 139, 140, 146

[BBFK05] G. Barthel, J.-P. Brasselet, K.-H. Fieseler, and L. Kaup. Combinatorial duality and intersection product: A direct approach. *Tohoku Mathematical Journal*, 57(2):273 – 292, 2005. 105, 146, 147, 159

[BGS96] A. Beilinson, V. Ginzburg, and W. Soergel. Koszul duality patterns in representation theory. *J. Amer. Math. Soc.*, 9(2):473–527, 1996. 10, 83

[BL00] P. Bressler and V. Lunts. Intersection Cohomology on Nonrational Polytopes. *Compositio Mathematica*, 135, 03 2000. 104, 105, 137, 146, 147

[BL03] P. Bressler and V. Lunts. Hard Lefschetz Theorem and Hodge-Riemann Relations for Intersection Cohomology of Nonrational Polytopes. *Indiana University Mathematics Journal*, 54, 03 2003. 105, 147, 148, 159

[BLPW10] T. Braden, A. Licata, N. Proudfoot, and B. Webster. Gale duality and Koszul duality. *Adv. Math.*, 225(4):2002–2049, 2010. 2, 3, 4, 5, 13, 30, 39, 44, 48, 53, 55, 60, 66, 68, 71, 81, 93, 100

[BLPW12] T. Braden, A. Licata, N. Proudfoot, and B. Webster. Hypertoric category $\mathcal{O}$. *Adv. Math.*, 231(3-4):1487–1545, 2012. 2

[BLPW16] T. Braden, A. Licata, N. Proudfoot, and B. Webster. Quantizations of conical symplectic resolutions II: category $\mathcal{O}$ and symplectic duality. *Astérisque*, 384:75–179, 2016. 2

[BLVS+99] A. Björner, M. Las Vergnas, B. Sturmfels, N. White, and G. M. Ziegler. *Oriented matroids*, volume 46 of *Encyclopedia of Mathematics and its Applications*. Cambridge University Press, Cambridge, second edition, 1999. 10, 15, 22, 25, 26, 27, 29, 31, 37, 38, 48, 50, 52, 61, 62, 63, 79

[BM17] T. Braden and C. Mautner. Matroidal Schur algebras. *J. Algebraic Combin.*, 46(1):51–75, 2017. 3, 13

[BM19] T. Braden and C. Mautner. Ringel duality for perverse sheaves on hypertoric varieties. *Adv. Math.*, 344:35–98, 2019. 3, 13

[BPW16] T. Braden, N. Proudfoot, and B. Webster. Quantizations of conical symplectic resolutions I: local and global structure. *Astérisque*, (384):1–73, 2016. 13

[Bra05] T. Braden. Remarks on the Combinatorial Intersection Cohomology of Fans. *Pure and Applied Mathematics Quarterly*, 2, 12 2005. 105

[BV78] R. G. Bland and M. L. Vergnas. Orientability of matroids. *J. Comb. Theory, Ser. B*, 24:94–123, 1978. 30

[Cha19] A. Chandler. Thin posets, CW posets, and categorification. *arXiv preprint arXiv:1911.05600*, 2019. 115

[CM93] R. Cordovil and M. L. Moreira. A homotopy theorem on oriented matroids. volume 111, pages 131–136. 1993. Graph theory and combinatorics (Marseille-Luminy, 1990). 37

[Cor82] R. Cordovil. Sur les matroïdes orientés de rang 3 et les arrangements de pseudodroites dans le plan projectif réel. *European J. Combin.*, 3(4):307–318, 1982. 42

[ET15] B. Everitt and P. Turner. Cellular cohomology of posets with local coefficients. *Journal of Algebra*, 439:134–158, 2015. 109, 110, 115, 119, 120, 124

[Kar01] K. Karu. Hard Lefschetz theorem for nonrational polytopes. *Inventiones mathematicae*, 157:419–447, 2001. 105, 146, 162

[Los17] I. Losev. On categories $\mathcal{O}$ for quantized symplectic resolutions. *Compos. Math.*, 153(12):2445–2481, 2017. 2

[LV80] M. Las Vergnas. Convexity in oriented matroids. *J. Combin. Theory Ser. B*, 29(2):231–243, 1980. 31

[Pos95] L. E. Positselski. The correspondence between Hilbert series of quadratically dual algebras does not imply their having the Koszul property. *Funktsional. Anal. i Prilozhen.*, 29(3):83–87, 1995. 84

[Rin55] G. Ringel. Teilungen der Ebene durch Geraden oder topologische Geraden. *Math. Z.*, 64:79–102 (1956), 1955. 7

[RS72] C. P. Rourke and B. J. Sanderson. *Introduction to piecewise-linear topology*. Springer-Verlag, New York-Heidelberg, 1972. Ergebnisse der Mathematik und ihrer Grenzgebiete, Band 69. 63

[SF99] R. P. Stanley and S. Fomin. *Enumerative Combinatorics*, volume 2 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, 1999. 113, 114

[Sta87] R. Stanley. Generalized H-Vectors, Intersection Cohomology of Toric Varieties, and Related Results. In *Commutative Algebra and Combinatorics*, pages 187–213, Tokyo, Japan, 1987. Mathematical Society of Japan. 106, 107

[Sta94] R. P. Stanley. *A Survey of Eulerian Posets*, pages 301–333. Springer Netherlands, Dordrecht, 1994. 107

[Sta04] R. Stanley. *Combinatorics and Commutative Algebra*. Combinatorics and Commutative Algebra. Birkhäuser Boston, 2004. 33, 34, 54, 76, 89, 134