Toric Quiver Varieties

by

Dániel Joó

Submitted to

Central European University

Department of Mathematics and its Applications

In partial fulfillment of the requirements for the degree of Doctor of Philosophy

Supervisor: Professor Mátyás Domokos

Budapest, Hungary

2015

Declaration

I hereby declare that the disserartion contains no materials accepted for any other degrees in any other institutions.

Budapest, 14 April 2015

Dániel Joó

Declaration

I hereby declare that the dissertation contains no materials previously written and/or published by another person, except where appropriate acknowledgment is made in the form of bibliographical reference, etc.

Budapest, 14 April 2015

Dániel Joó

Declaration

I hereby declare that Dániel Joó is the primary author of the following joint paper which is based on material contained in this thesis.

M. Domokos and D. Joó, On the equations and classification of toric quiver varieties (Proceedings A of The Royal Society of Edinburgh)

Budapest, 14 April 2015

Mátyás Domokos

Acknowledgments

First and foremost I would like to express my sincere gratitude to my supervisor, Mátyás Domokos, for his wisdom and patience in guiding me through this work. He taught me invaluable lessons about turning half-baked ideas into solid proofs, and showed me how deep mathematical truth can be discovered by opening up seemingly trivial details.

I would also like to thank Bernd Sturmfels for bringing [46] to my attention, which initiated an important step forward in my research progress.

Abstract

Toric quiver varieties arise as GIT moduli spaces of quiver representations when the dimension vector is fixed to have value 1 on every vertex, and come with a canonical embedding into projective space associated to a quiver polyhedron. We outline a procedure for their classification and show that up to isomorphism there are only finitely many d-dimensional toric quiver varieties in each fixed dimension d. We study the homogeneous toric ideals of projective toric quiver varieties in the canonical embedding associated to the GIT construction. It is shown that these toric ideals are always generated by elements of degree at most 3. We demonstrate a method of subdividing quiver polytopes into 0-1 polytopes to obtain an estimate on the minimal degree of the generators in their toric ideals. As an application of this method it is then shown that up to dimension 4 the toric ideal of every quiver polytope can be generated in degree 2, with the single exception of the Birkhoff polytope B_3. We then investigate 0-1 polytopes arising from general toric GIT constructions and prove that under certain assumptions on the arrangement of singular points their toric ideals are generated in degree 2. Finally, departing from the toric case, we prove a characterization of triples consisting of a quiver, a dimension vector, and a weight vector that yield smooth GIT moduli spaces in terms of forbidden descendants, which is in the flavour of characterizing classes of graphs by forbidden minors.

Table of Contents

Chapter 1

Introduction

A quiver representation assigns finite dimensional vector spaces to the vertices of a quiver (a finite directed graph) and linear maps to the arrows that go between the vector spaces assigned to their endpoints. The dimension vector α of a representation is the map that assigns to each vertex the dimension of the corresponding vector space. Quiver representations with a fixed dimension vector form a vector space which is endowed with the base change action of the product of general linear groups taken at each vertex, denoted by $GL(\alpha)$. The orbits of $GL(\alpha)$ correspond bijectively to the isomorphism classes of α dimensional representations. The affine quotient of this action, which was first studied by Le Bruyn and Procesi in [33], arises as the spectrum of the algebra of $GL(\alpha)$-invariant polynomials, and its points are in bijection with the isomorphism classes of semisimple representations. This quotient space consists of a single point for quivers without oriented cycles, however it can yield a complicated structure when there are cycles in the quiver. More general quotient constructions were introduced by King in [31], where geometric invariant theory (GIT) was applied to construct quasiprojective (projective when the quiver has no oriented cycles) moduli spaces whose points correspond to S-equivalence classes of representations that are semistable with respect to a fixed character of $GL(\alpha)$.

Most of this thesis studies these moduli spaces in the case when the dimension vector takes value 1 on every vertex. In this special case the group $GL(\alpha)$ is just just an algebraic torus (i.e. a product of copies of the multiplicative group of the base field $\mathbb{C}$ of complex numbers) and the resulting moduli spaces are toric varieties, which we will call *toric quiver varieties*. These varieties come with a canonical embedding into projective space, which can be associated to a lattice polyhedron under a standard construction of toric geometry.

We will refer to the polyhedra arising this way as *quiver polyhedra*.

Toric quiver varieties were studied by Hille [25], [26], [27], Altmann and Hille [2], Altmann and van Straten [4]. Further motivation is provided by Craw and Smith [13], who showed that every projective toric variety is the fine moduli space for stable representations with dimension vector $\alpha = (1, \ldots, 1)$ of an appropriate quiver with relations. Toric quiver varieties also play a role in the study of dimer models (see for ex. [28]). Another application was introduced recently by Carroll, Chindris and Lin [11].

Following the expository part of the thesis, we will begin our investigation in Chapter 3 by describing a classification procedure for toric quiver varieties. One of the important consequences we draw from this procedure is that for each positive d there are only finitely many d-dimensional toric quiver varieties up to isomorphism. The finiteness result itself also follows from the work of Altmann and Straten [4] and Altmann, Nill, Schwentner and Wiercinska [3] (although it was not made explicit). However we provide a self-contained combinatorial derivation and obtain some new results along the way. Theorem 3.13 asserts that if a toric quiver variety is a product of lower dimensional varieties, then these factors need to be toric quiver varieties themselves, and that this decomposition can be read off easily from the combinatorial structure of the quiver. Moreover by Theorem 3.22 any prime d-dimensional ($d > 1$) projective toric quiver variety can be obtained from a bipartite quiver with $5(d-1)$ vertices and $6(d-1)$ arrows, whose skeleton (cf. Definition 3.14) is 3-regular. As an application of the classification procedure we compiled a full list of reflexive quiver polytopes up to dimension 3. While similar work has been done in [3], we obtained a different result. The details of our method can be found in Section 3.4.

The toric variety defined by a polyhedron can be covered by principal affine open sets that correspond to the vertices of the polyhedra. In Theorem 3.28 we show that any toric quiver variety can be embedded as the union of some of the sets in this particular open cover of a projective toric quiver variety.

Vanishing ideals of embeddings of affine or projective toric varieties are often referred to as toric ideals. The study of toric ideals of toric quiver varieties in their canonical embeddings fits into the line of several recent works which were concerned with various toric ideals arising from combinatorial constructions. Some examples are toric ideals of matroids (see for ex. [32]) or edge polytopes of graphs (see for ex. [24]). In the projective case quiver polytopes coincide with the class of flow polytopes whose toric ideals were the focus of the work of Lenz [34], while some special cases also appear in [14] and [22]. One of our key results here is Theorem 4.10 where we prove that the toric ideal of a quiver polytope is always generated by elements of degree at most 3. This is deduced from a recent result of Yamaguchi, Ogawa and Takemura from [46], for which we give a simplified proof. We look further into this topic in Section 4.4 where, applying an idea of Haase and Paffenholz from [22], we break a quiver polytope into "cells" by a hyperplane subdivision. The key feature of this method is that while there are infinitely many quiver polyhedra (associated to finitely many varieties) in any dimension, there are only finitely many cells that can occur, hence it is possible to list them by computer exploration. We carried this task out up to dimension 4 and applied it to prove Theorem 4.21 which asserts that up to dimension 4 the only quiver polytope whose ideal is not generated in degree 2 is the Birkhoff polytope B_3.

The cells appearing in Section 4.4 are polytopes with 0-1 vertices, which we call *binary polytopes*. Toric ideals of binary polytopes have already received some interest, for example in the work of Ohsugi [37]. In Section 4.5 we turn our attention to binary polytopes arising from toric GIT moduli space constructions, and show that under suitable restrictions on the arrangement of singular points their toric ideals can be generated in degree 2, and that - under even stronger assumptions - they possess a quadratic Gröbner basis. (Theorem 4.24 and Theorem 4.26). While our primary motivation for this work was to study the quiver polytopes in this class, our results generalized to this wider setting as well. From a different perspective, one can relate these statements to the Bøgvad Conjecture which asserts that the toric ideal of a projectively normal embedding of a projective toric variety can always be generated in degree 2 (see for ex. [9]).

Turning to the affine case in Theorem 4.6 we prove that the vanishing ideal of an affine toric quiver variety can always be generated by binomials that are the difference of a monomial of degree 2 and a monomial of degree at most $d-1$, where d is the dimension of the variety (and the bound is sharp).

While the main focus of the thesis remains the toric case we recover some results that apply to more general dimension vectors. In Proposition 3.27 we generalize a "toric phenonemon" to certain special dimension vectors to embed the affine moduli space of their local quiver settings (see Section 2.3) as an open subset of the moduli space of the original quiver. This is a refinement of the original result on local quiver settings from [33] which guaranteed the existence of an étale isomorphism between an open subset of the moduli space of the original quiver and a neighborhood of the zero representation in the affine moduli space of the local quiver. Finally in Theorem 3.43 we characterize smooth quiver moduli in terms of not possessing certain forbidden descendants, which is a result in the flavour of characterizing classes of graphs with forbidden minors. The key tool for the proof is a result of Bocklandt [7], which provides an algorithmic method for deciding whether a given affine quiver moduli space is smooth.

Chapter 2 of the thesis is expository. Sections 3.1, 3.2, 3.3, 4.1 and 4.3 closely follow [16], which is joint work of M. Domokos and the author of this thesis. Sections 3.5.1 and 4.2 contain some of the theorems of [29] which is work of the author. The rest of the results in this thesis, to the best of our knowledge, has not appeared elsewhere unless explicitly stated otherwise.

Chapter 2

Preliminaries

2.1 Toric Varieties

Recall that a *toric variety* is a normal algebraic variety X containing the torus $T = (\mathbb{C}^*)^d$ as a Zariski open subset, together with an algebraic action $T \times X \to X$ that extends the natural action of T on itself. Note that in some of the literature toric varieties are not required to be normal, however throughout this thesis we work with this assumption. One of the nicest features of toric varieties is that they can be encoded by a combinatorial construction. In this section we recall the main components of this construction and refer to [12] for the details.

Fix a lattice $N \cong \mathbb{Z}^d$ and set $N_{\mathbb{R}} = N \otimes \mathbb{R}$. Denote the dual lattice of N by $M = Hom_{\mathbb{Z}}(N, \mathbb{Z})$ and similarily set $M_{\mathbb{R}} = M \otimes \mathbb{R}$.

Definition 2.1 A *rational polyhedral cone* in $N_{\mathbb{R}}$ is a set of the form

$$\sigma = \text{Cone(S)} := \{\sum_{u \in S} \lambda_u u \mid \lambda_u \geq 0\} \subseteq N_{\mathbb{R}},$$

where $S \subseteq N_{\mathbb{R}}$ is finite. The *dual cone* of σ is

$$\sigma^{\vee} = \{m \in M_{\mathbb{R}} \mid \forall u \in \sigma : \langle m, u \rangle \geq 0\}.$$

A rational polyhedral cone is called strongly convex if it does not contain a line through the origin. Note that for a rational polyhedral cone σ in $N_{\mathbb{R}}$, the dual cone $\sigma^{\vee}$ is also a rational polyhedral cone in $M_{\mathbb{R}}$, moreover if σ is strongly convex then $\sigma^{\vee}$ is full dimensional (i.e. its affine span is the entire $M_{\mathbb{R}}$). A *face* of a cone is just its intersection with the boundary of some closed half-spaces containing the cone. A *ray* of a cone is a one dimensional face, and in the case of rational polyhedral cones, the first lattice point along each ray is called a *ray generator*. The set $\sigma^{\vee} \cap M$ is a finitely generated submonoid of M and we will denote by $\mathbb{C}[\sigma^{\vee} \cap M]$ the corresponding semigroup algebra. To a strongly convex rational polyhedral cone σ we associate the affine toric variety

$$U_{\sigma} = \mathrm{Spec}(\mathbb{C}[\sigma^{\vee} \cap M]).$$

The torus of U_{σ} is $\mathrm{Spec}(\mathbb{C}[M]) \cong (\mathbb{C}^*)^d$, which is embedded via the dual of the map $\mathbb{C}[\sigma^{\vee} \cap M] \to \mathbb{C}[M]$. The lattice M can be identified with the character lattice of the torus via $(m_1, \ldots, m_k)(t_1, \ldots, t_k) = \prod_{i=1}^{k} t_i^{m_i}$, and under this idenfication the lattice points $\sigma^{\vee} \cap M$ are precisely the characters which extend to coordinate functions of U_{σ}. On the other hand the lattice can N can be identified with the one parameter subgroups of T by assigning to $u \in N$ the one parameter subgroup $\lambda^u(t) = (t^{u_1}, \ldots, t^{u_k})$. Now for a lattice point $u \in N$ we have that $u \in \sigma$ if and only if $\lim_{t \to 0} \lambda^u(t)$ exists in U_{σ}.

General toric varities can be constructed via fans:

Definition 2.2 A *fan* Σ in $N_{\mathbb{R}}$ is a finite collection of strongly convex rational polyhedral cones satisfying:

(i) For all $\sigma \in \Sigma$, each face of σ is also in Σ.

(ii) The intersection of any two cones in Σ is a face of each.

From a fan Σ in $N_{\mathbb{R}}$ one can construct a toric variety X_{Σ} by taking the affine toric varieties U_{σ} for each $\sigma \in \Sigma$ and glueing U_{σ_1} and U_{σ_2} via the maps

$$\mathbb{C}[\sigma_1^{\vee} \cap M] \hookrightarrow \mathbb{C}[(\sigma_1 \cap \sigma_2)^{\vee} \cap M] \hookleftarrow \mathbb{C}[\sigma_2^{\vee} \cap M],$$

for every $\sigma_1, \sigma_2 \in \Sigma$.

Next we define embeddings of toric varieties via lattice polyhedra. A *polyhedron* is the intersection of finitely many closed half-spaces in $\mathbb{R}^d$, and a *polytope* is a bounded polyhedron or equivalently the convex hull of finitely many points. A *facet* of a polyhedron is just a face of maximal dimension. A *facet presentation* of a polyhedron ∇ is given by choosing inward pointing normal vectors u_F for each facet F, and real numbers a_F satisfying

$$\nabla = \{x \in \mathbb{R}^d \mid \langle x, u_F \rangle \geq -a_F \quad \text{for all facets F}\}.$$

A polytope in $M_{\mathbb{R}}$ is called a *lattice polytope* if its vertices belong to the lattice M, and a polyhedron is called a *lattice polyhedron* if it is the Minkowski sum of a lattice polytope and a strongly convex rational polyhedral cone. Note that the assumption that the cone is strongly convex implies that the set of vertices of a lattice polyhedron is always non-empty.

Definition 2.3 The lattice polyhedra $\nabla_i \subset M^i_{\mathbb{R}}$ with lattice $M^i \subset M^i_{\mathbb{R}}$ $(i = 1, 2)$ are *integral-affinely equivalent* if there exists an affine linear isomorphism $\varphi : \mathrm{AffSpan}(\nabla_1) \to \mathrm{AffSpan}(\nabla_2)$ of affine subspaces with the following properties:

(i) φ maps $\mathrm{AffSpan}(\nabla_1) \cap M^1$ onto $\mathrm{AffSpan}(\nabla_2) \cap M^2$;

(ii) φ maps ∇_1 onto ∇_2.

The phrase 'integral-affinely equivalent' was chosen in accordance with [10] (though in [10] full dimensional lattice polytopes are considered).

By the *product of two polyhedra* $\nabla_1 \subseteq \mathbb{R}^{d_1}$ and $\nabla_2 \subseteq \mathbb{R}^{d_2}$ we mean the polyhedron

$$\nabla_1 \times \nabla_2 := \{(x_1, x_2) \in \mathbb{R}^{d_1+d_2} \mid x_1 \in \nabla_1,\ x_2 \in \nabla_2\}.$$

Note that the product of lattice polyhedra (resp. strongly convex rational polyhedral cones) is always a lattice polyhedron (resp. a strongly convex rational polyhedral cone). A lattice polyhedron ∇ is called *normal* if for every integer $k \geq 1$ the lattice points of $k\nabla$ (the Minkowski sum of k copies of ∇) can be written as a sum of k lattice points from ∇.

For every vertex v of a lattice polyhedron ∇ we define the cone $\sigma_v \subseteq N_{\mathbb{R}}$ to be dual to the cone generated by $(\nabla \cap M - v) \subseteq M$. Together these cones form a fan in $N_{\mathbb{R}}$, denoted by Σ_∇. Note that whenever the cone τ is a face of the cone σ, the dual cone $\sigma^\vee$ is a face of $\tau^\vee$, hence we have an inclusion reversing bijection between the faces of ∇ and the cones in Σ_∇. In particular the rays of Σ_∇ are in bijective correspondence with the facets of ∇, and a facet prensentation of ∇ can be given by choosing the u_F to be the ray generators in Σ_∇. When we speak of the facet presentation of a lattice polyhedron, we will always mean this unique presentation. To simplify the notation we shall write U_v for the affine variety U_{σ_v} and note that these varieties give a principal affine cover of the toric variety X_{Σ_∇}. To generalize this take a lattice point $m \in \nabla \cap M$, and let $\mathcal{V}_m$ denote the set of vertices on the minimal face of ∇ containing m. We define the affine open subset $U_m \subseteq X_{\Sigma_\nabla}$, as $U_m = \bigcap_{v \in \mathcal{V}_m} U_v$. By construction U_m is the affine toric variety of the cone $\bigcap_{v \in \mathcal{V}_m} \sigma_v$.

When ∇ is full dimensional in $M_{\mathbb{R}}$ we will define the *toric variety associated to* ∇ to be X_{Σ_∇} and denote it by X_∇. When ∇ is not full dimensional this construction would yield a fan that contains both ρ and $-\rho$ for some ray ρ and hence the resulting toric variety would have torus factors. Throughout this thesis we will only work with toric varieties that contain no torus factors, hence it is more convenient for us to eliminate this possibilty by defining X_∇ to be the variety we obtain from the above construction when we restrict the lattice to the affine span of ∇. To make this more precise pick a lattice point $m \in \mathrm{AffSpan}(\nabla) \cap M$, and let Σ^*_∇ be the fan associated to ∇ when considered in the lattice $(\mathrm{AffSpan}(\nabla) \cap M - m) \subseteq (\mathrm{AffSpan}(\nabla) - m)$. We then define X_∇ to be $X_{\Sigma^*_\nabla}$ and note that it does not depend on the choice of m. We note that this convention is not consistent with [12], but it simplifies the description of toric moduli spaces of quivers which are the central objects of this thesis.

By the product of two fans Σ_1 and Σ_2 we mean the fan

$$\Sigma_1 \times \Sigma_2 = \{\sigma_1 \times \sigma_2 \mid \sigma_1 \in \Sigma_1,\ \sigma_2 \in \Sigma_2\}.$$

It can be easily derived from the definitions that $X_{\Sigma_1 \times \Sigma_2} \cong X_{\Sigma_1} \times X_{\Sigma_2}$ and that $\Sigma_{\nabla_1 \times \nabla_2} = \Sigma_{\nabla_1} \times \Sigma_{\nabla_2}$.

The toric variety X_∇ is quasi-projective, moreover it is projective if ∇ is a polytope. To realize the variety X_∇ more explicitly let u_F and a_F be the inward pointing normal vectors and constants in its facet presentation and define the cone $C(\nabla) \subseteq M_{\mathbb{R}} \times \mathbb{R}$ as

$$C(P) = \{(m, \lambda) \in M_{\mathbb{R}} \times \mathbb{R} \ \mid\ \langle m, u_F \rangle \geq -\lambda a_F \text{ for all F},\ \lambda \geq 0\}.$$

Let S_∇ be the graded semigroup $C(P) \cap M \times \mathbb{Z}$ with the grading defined by setting the grade of (m, λ) to be λ, and $\mathbb{C}[S_\nabla]$ be the corresponding graded semigroup algebra. The projective spectrum $\mathrm{Proj}(\mathbb{C}[S_\nabla])$ is isomorphic to the variety X_∇ (see Theorem 7.1.13 in [12]). We note that it follows from the Proj construction that the open set $U_m \subseteq X_{\Sigma_\nabla}$ is the affine spectrum of the homogeneous localization $\mathbb{C}[S_\nabla]_{(t^{(m,1)})}$, where $t^{(m,1)}$ denotes the element of $\mathbb{C}[S_\nabla]$ corresponding to $(m, 1) \in S_\nabla$.

An affine toric variety U_σ is smooth if and only if the ray generators of σ form a $\mathbb{Z}$-basis of the lattice N (in fact the only smooth affine toric varieties are the affine spaces). When σ is full dimensional in $N_{\mathbb{R}}$ this is equivalent to saying that the ray generators of the dual cone $\sigma^\vee$ are a $\mathbb{Z}$-basis of the character lattice M. It follows that for a polytope ∇ and a vertex $v \in \nabla$ the principal affine open set U_v is smooth if the ray generators of

Cone($\nabla \cap M - v$) are a $\mathbb{Z}$-basis of the lattice (AffSpan(∇) $\cap M - v$), and the variety X_∇ is smooth if U_v is smooth for every vertex v. We will say that v is a *smooth vertex* if U_v is smooth, and that ∇ is a *smooth polyhedron* when X_∇ is smooth. Vertices (resp. polyhedra) that are not smooth will be called singular.

Note that $S(\nabla)$ and consequently $\mathbb{C}[S_\nabla]$ are generated by their elements of degree 0 and 1 when ∇ is normal. In particular when ∇ is a normal polytope, with lattice points $(m_1, \ldots, m_k)$ the abstract variety X_∇ can be identified with the Zariski-closure of the image of the map

$$(\mathbb{C}^*)^d \to \mathbb{P}^{k-1}, \quad t \mapsto (t^{m_1} : \cdots : t^{m_k}) \tag{2.1}$$

where for t in the torus $(\mathbb{C}^*)^d$ and $m \in M$ we write $t^m := \prod_{i=1}^d t(i)^{m(i)}$. From now on for a normal polytope ∇, X_∇ will stand for this particular embedding in projective space of our variety. Normality of ∇ implies that X_∇ is *projectively normal*, that is, its affine cone in $\mathbb{C}^d$ is normal. We point out that in general the closure of the image of the map in (2.1) is isomorphic to X_∇ whenever ∇ is a so-called very ample polytope, however since the polytopes occuring in our work are normal by construction, we decided to omit the discussion of this case.

The homogeneous vanishing ideal of the embedding given in (2.1) can be realized as follows: Let us denote the element of $\mathbb{C}[S(\nabla)]$ corresponding to $(m, \lambda) \in S(\nabla)$ by $t^m z^\lambda$. Consider the morphism of graded rings: $\varphi : \mathbb{C}[x_1, \ldots, x_d] \twoheadrightarrow \mathbb{C}[S(\nabla)]$ defined by $\varphi(x_i) = t^{m_i} z$. The homogeneous ideal $\ker(\varphi)$ is called the *toric ideal* of the polytope ∇. It is well known that $\ker(\varphi)$ is generated by binomials (see for ex. Lemma 4.1 in [44]).

Similarily in the affine case if $(m_1, \ldots, m_k)$ is a set of generators for the semigroup $\sigma^\vee \cap M$ we can identify the affine variety U_σ with the Zariski-closure of the image of the map

$$(\mathbb{C}^*)^d \to \mathbb{C}^k, \quad t \mapsto (t^{m_1}, \ldots, t^{m_k}).$$

Now one has a morphism of algebras $\varphi : \mathbb{C}[x_1, \ldots, x_d] \twoheadrightarrow \mathbb{C}[\sigma^\vee \cap M]$ defined by $\varphi(x_i) = t^{m_i}$ and $\ker(\varphi)$ is called the toric ideal of $\sigma^\vee \cap M$ (corresponding to the set of generators chosen).

It is easy to see that if ∇_1 and ∇_2 are integral-affinely equivalent lattice polyhedra then the graded rings $\mathbb{C}[S_{\nabla_1}]$ and $\mathbb{C}[S_{\nabla_2}]$, and hence the varieties X_{∇_1} and X_{∇_2}, are isomorphic. Moreover if ∇_1 and ∇_2 are integral-affinely equivalent normal lattice polytopes then they can be identified via their embeddings into projective space given in the previous paragraph.

2.2 Quiver representations

A *quiver* is a finite directed graph Q with vertex set Q_0 and arrow set Q_1. Multiple arrows, oriented cycles, loops are all allowed. For an arrow $a \in Q_1$ denote by a^- its starting vertex and by a^+ its terminating vertex. By the *valency* of a vertex $v \in Q_0$ we mean $|\{a \in Q_1 \mid a^- = v\}| + |\{a \in Q_1 \mid a^+ = v\}|$. The in-degree of a vertex v is $|\{a \in Q_1 \mid a^+ = v\}|$ and the out-degree is $|\{a \in Q_1 \mid a^- = v\}|$. By a *primitive cycle* in a quiver, we will mean a minimal oriented cycle, and by an *(primitive) undirected cycle* we will mean a set of arrows of the quiver that form a (minimal) cycle in the underlying undirected graph of the quiver. By a *forest* in a quiver we mean a set of arrows not containing undirected cycles, and a *tree* is just a connected forest. By a *spanning forest* we mean a maximal forest, i.e. one that is a maximal tree for each connected component of the quiver. For an undirected graph Γ we set $\chi(\Gamma) := |\Gamma_1| - |\Gamma_0| + \chi_0(\Gamma)$, where Γ_0 is the set of vertices, Γ_1 is the set of edges in Γ, and $\chi_0(\Gamma)$ is the number of connected components of Γ. Define $\chi(Q) := \chi(\Gamma)$ and $\chi_0(Q) := \chi_0(\Gamma)$ where Γ is the underlying graph of Q, and we say that Q is *connected* if Γ is connected, i.e. if $\chi_0(Q) = 1$. The *Ringel form* of the quiver is the bilinear form on $\mathbb{R}^{Q_0}$ defined as

$$\langle \alpha, \beta \rangle_Q = \sum_{v \in Q_0} \alpha(v)\beta(v) - \sum_{a \in Q_1} \alpha(a^-)\beta(a^+).$$

A *representation* R of Q is given by assigning a vector space $R(v)$ to each $v \in Q_0$, and a linear map $R(a) : R(a^-) \to R(a^+)$ for each $a \in Q_1$. Throughout this thesis we will only consider representations over the complex number field $\mathbb{C}$. A morphism between representations R and R' consists of a collection of linear maps $L(v) : R(v) \mapsto R'(v)$ satisfying $R'(a) \circ L(a^-) = L(a^+) \circ R(a)$ for all $a \in Q_1$. Accordingly the representation R' is a *subrepresentation* of R if R'_v is a subspace of R_v for each $v \in Q_0$ and the linear maps R'_a are restrictions of the maps R_a. A representation R is called *simple* if it has no non-zero proper subrepresentations, and *semisimple* if it is a direct sum of simples. The dimension vector $\alpha : Q_0 \to \mathbb{N}$ of a representation R is given by $\alpha(v) = \dim(X_v)$. The pair (Q, α) is called a *quiver setting*. We say that a representation or a quiver setting is *genuine* when its dimension vector is positive on every vertex. For a fixed dimension vector $\alpha \in \mathbb{N}^{Q_0}$,

$$\mathrm{Rep}(Q, \alpha) := \bigoplus_{a \in Q_1} \hom_{\mathbb{C}}(\mathbb{C}^{\alpha(a^-)}, \mathbb{C}^{\alpha(a^+)})$$

is the *space of α-dimensional representations* of Q. The product of general linear groups $GL(\alpha) := \prod_{v\in Q_0} GL_{\alpha(v)}(\mathbb{C})$ acts linearly on $\mathrm{Rep}(Q,\alpha)$ via

$$g \cdot R := (g(a^+)R(a)g(a^-)^{-1} \mid a \in Q_1) \quad (g \in GL(\alpha), R \in \mathrm{Rep}(Q,\alpha)).$$

The $GL(\alpha)$-orbits in $\mathrm{Rep}(Q,\alpha)$ are in a natural bijection with the isomorphism classes of α-dimensional representations of Q.

In [31] geometric invariant theory (GIT) was applied to constuct moduli spaces parametrizing isomorphism classes of quiver representations with a fixed dimension vector. In a somewhat more general setting of GIT, one has a reductive group G acting on an affine variety X, and a fixed character $\chi : G \to \mathbb{C}^*$. An element f of the coordinate ring $\mathcal{O}(X)$ is called a *relative invariant of weight* χ if for every $g \in G$ we have $f(g\cdot x) = \chi(g)f(x)$. The relative invariants of weight χ constitute a subspace of the coordinate ring denoted by $\mathcal{O}(X)_\chi$. Note that when χ is the trivial character $\mathcal{O}(X)_\chi$ is the ring of G invariant polynomials, which we will denote by $\mathcal{O}(X)^G$. Moreover $\bigoplus_{n=0}^\infty \mathcal{O}(X)_{\chi^n}$ can be turned into a graded ring by setting $\mathcal{O}(X)_{\chi^k}$ to be its degree k part. A point $x \in X$ is called χ-*semistable* if there is an $f \in \mathcal{O}(X)_{\chi^n}$ for $n \geq 1$ such that $f(x) \neq 0$ and it is called χ-*stable* if it is χ-semistable, the orbit $G \cdot x$ is closed in the open subset of semistable points and the stabilizer of x in G is finite. The set of semistable points is denoted by $X^{\chi-ss}$ and the set of stable points is denoted by $X^{\chi-s}$. The *GIT quotient*

$$X /\!\!/_\chi G = \mathrm{Proj}(\bigoplus_{n=0}^{\infty} \mathcal{O}(X)_{\chi^n})$$

is a good categorical quotient for the set $X^{\chi-ss}$ under the action of G, moreover it has an (possibly empty) open subset that is a geometric quotient for $X^{\chi-s}$ (see [36] for details of this construction and an explanation of some of the terminology used here). Note that the embedding $\mathcal{O}(X)^G \hookrightarrow \mathrm{Proj}(\bigoplus_{n=0}^\infty \mathcal{O}(X)_{\chi^n}$ induces a surjective projective morphism $X /\!\!/_\chi G \twoheadrightarrow \mathrm{Spec}(\mathcal{O}(X)^G))$, in particular when $\mathcal{O}(X)^G = \mathbb{C}$ the variety $X /\!\!/_\chi G$ is projective.

Returning to our setting of $GL(\alpha)$ acting on $\mathrm{Rep}(Q,\alpha)$ first note that the characters of $GL(\alpha)$ can be given by functions $\theta : Q_0 \to \mathbb{Z}$, by defining the character χ_θ as $\chi_\theta(g) = \prod_{v\in Q_0} \det(g(v))^{\theta(v)}$. To simplify the notation, we will write θ instead of χ_θ when talking about relative invariants, semistable or stable points in this setting. In particular a relative invariant of weight θ is a polynomial function $f \in \mathcal{O}(\operatorname{Rep}(Q,\alpha))$ satisfying

$$f(g \cdot R) = (\prod_{v \in Q_0} \det(g(v))^{\theta(v)}) f(R),$$

for all $g \in GL(\alpha)$ and $R \in \operatorname{Rep}(Q,\alpha)$. We point the reader to [17] for an explicit method of constructing relative invariant polynomials in this setting. Note that $GL(\alpha)$ always contains the one dimensional algebraic torus, whose elements g_c are given by $g_c(v) = c \cdot \operatorname{Id}(R(v))$ for $c \in \mathbb{C}^*$, which acts trivially on $\operatorname{Rep}(Q,\alpha)$. We modify our notion of stability accordingly and instead of the stabilizer being finite we require that $\dim(GL(\alpha) \cdot R) = \dim(GL(\alpha)) - 1$ for a θ-stable representation R. We point out that from this definition it follows that for a θ-stable representation R with dimension vector α the set of vertices $\{v \in Q_0 \mid \alpha(v) \neq 0\}$ is connected in Q, otherwise R would be stabilized by a higher dimensional subtorus of $GL(\alpha)$. Next we recall from Proposition 3.1 from [31] the description of the semistable and stable points of $\operatorname{Rep}(Q,\alpha)$.

Proposition 2.4 *For a quiver Q, a dimension vector α, an integer weight $\theta : Q_0 \to \mathbb{Z}$ and a representation $R \in \operatorname{Rep}(Q,\alpha)$ we have,*

(i) R is θ-semistable if and only if $\sum_{v \in Q_0} \alpha(v)\theta(v) = 0$ and for all subrepresentations R' of R, we have $\sum_{v \in Q_0} \alpha'(v)\theta(v) \geq 0$, where α' is the dimension vector of R'.

(i) R is θ-stable if and only if $\sum_{v \in Q_0} \alpha(v)\theta(v) = 0$ and for all non-zero proper subrepresentations R' of R, we have $\sum_{v \in Q_0} \alpha'(v)\theta(v) > 0$, where α' is the dimension vector of R'. In particular the θ-stable representations are the simple objects in the category of θ-semistable representations of Q.

The GIT-quotient $\operatorname{Rep}(Q,\alpha) /\!\!/_\theta GL(\alpha)$, which we will denote by $\mathcal{M}(Q,\alpha,\theta)$, is a coarse moduli space for families of θ-semistable α-dimensional representations of Q up to S-equivalence (cf. [36] for the terminology and Theorem 4.1 of [31] for the proof of the statement). When a θ stable representation exists, we will say that the quiver setting (Q,α) is θ-stable. Note that in this case

$$\dim(\mathcal{M}(Q,\alpha,\theta)) = 1 - \langle \alpha, \alpha \rangle_Q = \sum_{a \in Q_1} \alpha(a^-)\alpha(a^+) - \sum_{v \in Q_0} \alpha^2(v) + 1.$$

A notable special case is that of the zero weight. Then the moduli space $\mathcal{M}(Q,\alpha,0)$ is the affine variety whose coordinate ring is the subalgebra of $GL(\alpha)$-invariants in $\mathcal{O}(\operatorname{Rep}(Q,\alpha))$.

This was studied in [33] before the introduction of the case of general weights in [31]. Its points are in a natural bijection with the isomorphism classes of semisimple representations of Q with dimension vector α. For a quiver with no oriented cycles, $\mathcal{M}(Q, \alpha, 0)$ is just a point, and consequently $\mathcal{M}(Q, \alpha, \theta)$ is a projective variety.

2.3 Local structure of quiver moduli spaces

In Section 3.5.1 we will be concerned with local properties of $\mathcal{M}(Q, \alpha, \theta)$. For this purpose we will recall some results that provide information on the local structure of $\mathcal{M}(Q, \alpha, \theta)$ in terms of a (typically smaller) quiver. The first method described here was developed by Le Bruyn and Procesi for the affine case in [33], and then extended to the general case by Adriaenssens and Le Bruyn in [1].

Throughout this section we will denote by ϵ_v the dimension vector with $\epsilon_v(v) = 1$ and $\epsilon_v(u) = 0$ for $v \neq u$. We begin by recalling from Proposition 3.2 from [31] that for a point $\xi \in \mathcal{M}(Q, \alpha, \theta)$ there is a unique orbit $GL(\alpha) \cdot R \subseteq \pi^{-1}(\xi)$, such that the representations in $GL(\alpha) \cdot R$ decompose as a direct sum of θ-stable representations (π denotes the quotient map $\mathrm{Rep}(Q, \alpha)^{\theta-ss} \to \mathcal{M}(Q, \alpha, \theta)$). Let us write $R = R_1^{n_1} \oplus \cdots \oplus R_k^{n_k}$, where the R_i are distinct θ-stable representations. Denote by β_i the dimension vector of R_i and note that the pairs (n_i, β_i) only depend on the choice of ξ and not on the choice of R.

We construct a new quiver setting (Q_ξ, α_ξ) as follows: Q_ξ has vertices $v_1, \ldots, v_k$ corresponding to the θ-stable summands in the decomposition of R, the dimension vector is defined as $\alpha_\xi(v_i) = n_i$ and the number of arrows from v_i to v_j is equal $\delta_{ij} - \langle \beta_i, \beta_j \rangle_Q$, where $\langle , \rangle_Q$ is the Ringel-form defined in Section 2.2. We recall Theorem 4.1 from [1]:

Theorem 2.5 *There is an étale isomorphism between an affine neighborhood of ξ in the moduli space $\mathcal{M}(Q, \alpha, \theta)$ and an affine neighborhood of the image of the 0 representation in $\mathcal{M}(Q_\xi, \alpha_\xi, 0)$.*

Theorem 2.5 was proven in the special case $\theta = 0$ in [33]. Observe that in this case every representation is 0-semistable and the 0-stable representations coincide with the simple representations. Making this result effective a characterization of the quiver settings (Q, α) for which Q possesses a simple representation with dimension vector α was given in Theorem 4 of [33], which we recall below:

Theorem 2.6 *Let (Q, α) be a genuine quiver setting. There exists a simple representation with dimension vector α if and only if:*

(i) Q is a single vertex with no loops, or a single vertex with a single loop, or a directed cycle of length k for $k \geq 2$, and α is 1 on every vertex.

(ii) Q is none of the quivers in (i) but is strongly connected and for all $v \in Q_0$ we have $\langle \alpha, \epsilon_v \rangle_Q \leq 0$ and $\langle \epsilon_v, \alpha \rangle_Q \leq 0$.

Moreover in each case except for the one vertex quiver with no loops there are infinitely many isomorphism classes of simples with dimension vector α.

We note that a similar characterization for triples (Q, α, θ) for which Q possesses an α-dimensional θ-stable representation was also given in Theorem 5.1 of [1] . We opted to recall Theorem 2.6 here, since it is the one we need for the proof of Theorem 3.43. We also point out that in [18] it was shown that Theorem 2.6 holds for representations over an algebrically closed field of arbitrary characteristic, moreover the local quiver technique of Theorem 2.5 was also extended to this more general setting.

Remark 2.7 One of the nice features of Theorem 2.5 is that étale morphisms preserve several important properties of $\mathcal{M}(Q, \alpha, \theta)$. In particular $\mathcal{M}(Q, \alpha, \theta)$ is smooth at the point ξ if and only if $\mathcal{M}(Q_\xi, \alpha_\xi, 0)$ is smooth at 0, which in turn implies that $\mathcal{M}(Q_\xi, \alpha_\xi, 0)$ is an affine space (cf. Theorem 2.1 in [7]). Moreover $\mathcal{M}(Q, \alpha, \theta)$ is locally a complete intersection at ξ if and only if $\mathcal{M}(Q_\xi, \alpha_\xi, 0)$ is locally a complete intersection at 0, which in this case implies that $\mathcal{M}(Q_\xi, \alpha_\xi, 0)$ is globally a complete intersection (see [21]).

An important consequence of Theorem 2.5 is that to understand the local structure of quiver varieties we only need to study the case $\theta = 0$. We will now recall a method from [7] which characterizes the quiver settings (Q, α) for which $\mathcal{M}(Q, \alpha, 0)$ is smooth via certain reduction steps.

Lemma 2.8 *(RI) Let (Q, α) be a quiver setting and v a vertex without loops. Assume that at least one of $\langle \alpha, \epsilon_v \rangle_Q \geq 0$ and $\langle \epsilon_v, \alpha \rangle_Q \geq 0$ holds. Let $a_1, \dots, a_k$ denote the arrows pointing to v and $b_1, \dots, b_l$ denote the arrows leaving from v ($k, l \geq 1$). Let Q' be the quiver we obtain by removing v from Q and the arrows incident to it, and for each pair (i, j) with $i \in \{1, \dots, k\}$ and $j \in \{1, \dots, l\}$ adding a new arrow c_{ij} with $c_{ij}^- = a_i^-$ and $c_{ij}^+ = b_j^+$. Set α' to be the restriction of α to the vertices of Q'. Now we have $\mathcal{M}(Q, \alpha, 0) = \mathcal{M}(Q', \alpha', 0)$.*

Lemma 2.9 *(RII) Let (Q, α) be a quiver setting, v a vertex with $\alpha(v) = 1$ and b a loop incident to v. Let Q' be the quiver we obtain from Q by removing b. Now we have*

$\mathcal{M}(Q,\alpha,0) = \mathcal{M}(Q',\alpha',0) \times \mathbb{C}$.

Lemma 2.10 *(RIII) Let (Q,α) be a quiver setting, v a vertex with $\alpha(v) > 1$ and let b be the only loop in Q incident to v. Further assume that (other than b) there is only one arrow leaving from v which points to a vertex w with $\alpha(w) = 1$ (resp. there is only one arrow pointing to v which leaves from a vertex w with $\alpha(w) = 1$). Let Q' be the quiver we obtain from Q by removing the loop b and adding $\alpha(v) - 1$ new arrows from v to w (resp. from w to v). Now we have $\mathcal{M}(Q,\alpha,0) = \mathcal{M}(Q',\alpha',0) \times \mathbb{C}^k$.*

We will refer to the reduction steps in Lemmas 2.8, 2.9 and 2.10 as RI, RII and RIII. Note that all of these reduction steps preserve the properties of being smooth or a complete intersection. The main result of [7] is the following theorem:

Theorem 2.11 *If (Q,α) is a strongly connected quiver setting with $\alpha(v) > 0$ for each $v \in Q_0$ on which the reduction steps RI, RII and RIII can not be applied, then Q is either a single vertex with no loops, a single vertex with one loop, or a single vertex v with $\alpha(v) = 2$ and two loops.*

2.4 Toric quiver varieties

Most of this thesis deals with the special case when $\alpha(v) = 1$ for all $v \in Q_0$. In this case we will simply write $\mathrm{Rep}(Q)$ for $\mathrm{Rep}(Q,\alpha)$ and $\mathcal{M}(Q,\theta)$ for $\mathcal{M}(Q,\alpha,\theta)$. The coordinate ring $\mathcal{O}(\mathrm{Rep}(Q))$ is the polynomial ring $\mathbb{C}[t_a \mid a \in Q_1]$, where t_a is the function $\mathrm{Rep}(Q) \to \mathbb{C}$ defined by $R \to R(a)$. The group $GL((1,\dots,1))$ is just the algebraic torus $(\mathbb{C}^*)^{Q_0}$. Note that it follows immediately from Theorem 14.2.13 in [12] that $\mathcal{M}(Q,\theta)$ is a quasi-projective toric variety. Moreover when Q is θ-stable we have that $\dim(\mathcal{M}(Q,\theta)) = \chi(Q)$. We will refer to the toric varieties arising as $\mathcal{M}(Q,\theta)$ as *toric quiver varieties*.

For $m \in \mathbb{N}^{Q_1}$ we will denote by t^m the monomial $\prod_{a \in Q_1} t_a^{m(a)} \in \mathcal{O}(\mathrm{Rep}(Q))$. It is easy to check that the space of relative invariants $\mathcal{O}(\mathrm{Rep}(Q))_\theta$ is spanned by monomials t^m such that $\theta(v) = \sum_{a^+=v} m(a) - \sum_{a^-=v} m(a)$ for all $v \in Q_1$. We define the *quiver polyhedron* $\nabla(Q,\theta)$ as

$$\nabla(Q,\theta) = \{x \in \mathbb{R}^{Q_1} \mid \mathbf{0} \le x, \quad \forall v \in Q_0 : \quad \theta(v) = \sum_{a^+=v} x(a) - \sum_{a^-=v} x(a)\}.$$

A special case which we will study in Section 3.4 is that of the *canonical weight* $\delta_Q := \sum_{a\in Q_1}(\varepsilon_{a^+} - \varepsilon_{a^-})$ (here ε_v stands for the characteristic function of $v \in Q_0$). Note that $\nabla(Q,\delta_Q)$ always contains $(1,\ldots,1)$ as an interior point.

Denote by $\mathcal{F} : \mathbb{R}^{Q_1} \to \mathbb{R}^{Q_0}$ the map given by

$$\mathcal{F}(x)(v) = \sum_{a^+=v} x(a) - \sum_{a^-=v} x(a) \qquad (v \in Q_0). \tag{2.2}$$

By definition we have $\nabla(Q,\theta) = \mathcal{F}^{-1}(\theta) \cap \mathbb{R}^{Q_1}_{\geq 0}$. Set $M^Q_{\mathbb{R}} = \mathcal{F}^{-1}(0)$ and $M^Q = M^Q_{\mathbb{R}} \cap \mathbb{Z}^{Q_1}$. Next let F be a spanning forest of Q and note $|F| = |Q_0| - \chi_0(Q)$. For an arrow $a \in Q_1 \setminus F$ let c^a_F denote the unique undirected primitive cycle in $F \cup \{a\}$ and set $e^a_F \in \mathbb{Z}^{Q_1}$ to be 1 on the arrows of c^a_F that are oriented the same way as a along c^a_F, -1 on the arrows of c^a_F that are oriented reversely and 0 on the rest of the arrows. One can check without difficulty that e^a_F is the unique element in M^Q that takes value 1 on a and 0 on $Q_1 \setminus (F \cup \{a\})$, and hence derive the following (well-known) result:

Proposition 2.12 *For any spanning forest F of Q, the set $\{e^a_F \mid a \in Q_1 \setminus F\}$ is a $\mathbb{Z}$-basis of M^Q. In particular $\dim(M^Q_{\mathbb{R}}) = \chi(Q)$ and hence $\dim(\nabla(Q,\theta)) \leq \chi(Q)$.*

Set $S(Q,\theta) = \bigoplus_{n=0}^{\infty} \nabla(Q,n\theta) \cap \mathbb{Z}^{Q_1}$ to be the graded semigroup with degree k part $\nabla(Q,k\theta) \cap \mathbb{Z}^{Q_1}$ and $\mathbb{C}[S(Q,\theta)]$ the corresponding graded semigroup algebra. It follows that $\mathcal{M}(Q,\theta) \cong \mathrm{Proj}(\mathbb{C}[S(Q,\theta)])$. To relate this to the construction of toric varieties from polyhedra, we will need to show that $\nabla(Q,\theta)$ is a lattice polyhedron. By the *support* of $x \in \mathbb{R}^{Q_1}$, denoted by $\mathrm{supp}(x)$, we mean the set $\{a \in Q_1 \mid x(a) \neq 0\} \subseteq Q_1$.

Proposition 2.13 *(i) Denote by $Q^1,\ldots,Q^t$ the maximal subquivers of Q that contain no oriented cycles. Then $\nabla(Q,\theta) \cap \mathbb{Z}^{Q_1}$ has a Minkowski sum decomposition*

$$\nabla(Q,\theta) \cap \mathbb{Z}^{Q_1} = \nabla(Q,0) \cap \mathbb{Z}^{Q_1} + \bigcup_{i=1}^{t} \nabla(Q^i,\theta) \cap \mathbb{Z}^{Q_1}. \tag{2.3}$$

(ii) Let $C_1,\ldots,C_r$ be the primitive cycles of Q and let $\varepsilon_{C_1},\ldots,\varepsilon_{C_r}$ denote their characteristic functions. Then

$$\nabla(Q,0) = \{\sum_{i=1}^{r} \lambda_i \varepsilon_{C_i} \mid \lambda_i \geq 0\},$$

and hence $\nabla(Q,0)$ is a strongly convex rational polyhedral cone.
(iii) The quiver polyhedron $\nabla(Q,\theta)$ is a normal lattice polyhedron.

Proof. (i) It is obvious that $\nabla(Q,\theta)\cap\mathbb{Z}^{Q_1}$ contains the set on the right hand side of (2.3). To show the reverse inclusion take an $x\in\nabla(Q,\theta)\cap\mathbb{Z}^{Q_1}$. If its support contains no oriented cycles, then $x\in\nabla(Q^i,\theta)$ for some i. Otherwise take a minimal oriented cycle $C\subseteq Q_1$ in the support of x. Denote by $\varepsilon_C\in\mathbb{R}^{Q_1}$ the characteristic function of C, and denote by λ the minimal coordinate of x along the cycle C. Then $\lambda\varepsilon_C\in\nabla(Q,0)$ and $y:=x-\lambda\varepsilon_C\in\nabla(Q,\theta)$. Moreover, y has strictly smaller support than x. By induction on the size of the support we are done.

(ii) If $x\in\nabla(Q,0)$ then whenever $\mathrm{supp}(x)$ contains an in-arrow of a vertex it also has to contain an out-arrow and hence $\mathrm{supp}(x)$ contains an oriented cycle. Now the statement follows by the same induction as in (i).

(iii) The same argument as in (i) and taking into account the convexity of $\nabla(Q,\theta)$ yields

$$\nabla(Q,\theta)=\nabla(Q,0)+\mathrm{Conv}(\bigcup_{i=1}^{t}\nabla(Q^i,\theta)).$$

The vertices of $\nabla(Q^i,\theta)$ belong to $\mathbb{Z}^{Q_1}$ by Theorem 13.11 in [41]. The set of vertices of $\mathrm{Conv}(\bigcup_{i=1}^{t}\nabla(Q^i,\theta))$ is a subset of the union of the vertices of the $\nabla(Q^i,\theta)$, so it is a lattice polytope. Taking (ii) into account we see that $\nabla(Q,\theta)$ is the Minkowski sum of a lattice polytope and a strongly convex rational polyhedral cone and hence it is a lattice polyhedron. For normality we need to show that for all positive integers k we have $\nabla(Q,k\theta)\cap\mathbb{Z}^{Q_1}=k(\nabla(Q,\theta)\cap\mathbb{Z}^{Q_1})$. The polytopes $\nabla(Q^i,\theta)$ are normal by Theorem 13.14 in [41]. So by (i) we have $\nabla(Q,k\theta)\cap\mathbb{Z}^{Q_1}=\nabla(Q,0)\cap\mathbb{Z}^{Q_1}+\bigcup_{i=1}^{t}(\nabla(Q^i,k\theta)\cap\mathbb{Z}^{Q_1})=\nabla(Q,0)\cap\mathbb{Z}^{Q_1}+\bigcup_{i=1}^{t}k(\nabla(Q^i,\theta)\cap\mathbb{Z}^{Q_1})\subseteq k(\nabla(Q,0)\cap\mathbb{Z}^{Q_1}+\bigcup_{i=1}^{t}\nabla(Q^i,\theta)\cap\mathbb{Z}^{Q_1})$. ☐

Corollary 2.14 *(i) We have the isomorphism $\mathcal{M}(Q,\theta)\cong X_{\nabla(Q,\theta)}$ of toric varieties.*
(ii) For every vertex $m\in\nabla(Q,\theta)$ the arrow set $\mathrm{supp}(m)$ *is a forest.*
(iii) A principal affine open cover of $\mathcal{M}(Q,\theta)$ is given by the Zariski-open sets $\{U_v \mid v$ is a vertex of $\nabla(Q,\theta)\}$, where σ_v is the cone

$$\sigma_v=Cone(\nabla(Q,\theta)-v)=\mathcal{F}^{-1}(0)\cap\{x\in\mathbb{R}^{Q_1}\ \mid\ \forall a\in Q_1\setminus supp(v):\ x(a)\geq 0\},$$

and $U_v=Spec(\mathbb{C}[\sigma_v\cap\mathbb{Z}^{Q_1}])$.

Proof. For (i) we know from Proposition 2.13 that $\nabla(Q,\theta)$ is a normal lattice polyhedron and by construction both $\mathcal{M}(Q,\theta)$ and $X_{\nabla(Q,\theta)}$ are isomorphic to $\mathrm{Proj}(\mathbb{C}[S(Q,\theta)])$. For (ii) assume that $c\subseteq\mathrm{supp}(m)$ is an undirected cycle, and let $a\in Q_1$ be an arrow in c.

Now one can choose a spanning forest F that contains $c \setminus \{a\}$ and see that for the vector $e_F^a \in M^Q$ from Proposition 2.12 we have $\operatorname{supp}(e_F^a) = c$. Since e_F^a takes values $1, -1$ on every arrow of c we have that $m + e_F^a \in \nabla(Q,\theta)$ and $m - e_F^a \in \nabla(Q,\theta)$ and hence m can not be a vertex. Finally (iii) follows from the description of the toric fan of the variety $X_{\nabla(Q,\theta)}$ given in Section 2.1. $\square$

Let $C_1, \ldots, C_r$ be the primitive cycles in Q. Applying the same induction as in the proof of Proposition 2.13 one can deduce that $\varepsilon_{C_1}, \ldots, \varepsilon_{C_r}$ constitute a generating set for the monoid $\nabla(Q,0) \cap \mathbb{Z}^{Q_1}$. Enumerate the elements in $\{m, \varepsilon_{C_j} + m \mid m \in \bigcup_{i=1}^t \nabla(Q^i,\theta) \cap \mathbb{Z}^{Q_1}, j = 1, \ldots, r\}$ as $m_0, m_1, \ldots, m_k$. For a lattice point $m \in \nabla(Q,\theta) \cap \mathbb{Z}^{Q_1}$ denote by $x^m : \operatorname{Rep}(Q) \to \mathbb{C}$ the function $x \mapsto \prod_{a \in Q_1} R(a)^{m(a)}$. Consider the map

$$\rho : \operatorname{Rep}(Q)^{\theta-ss} \to \mathbb{P}^k, \quad x \mapsto (x^{m_0} : \cdots : x^{m_k}). \tag{2.4}$$

Proposition 2.15 *$\mathcal{M}(Q,\theta)$ can be identified with the locally closed subset $\operatorname{Im}(\rho)$ in $\mathbb{P}^d$.*

Proof. The morphism ρ is $(\mathbb{C}^\times)^{Q_0}$-invariant, hence it factors through the quotient morphism $\operatorname{Rep}(Q,\alpha)^{\theta-ss} \to \mathcal{M}(Q,\alpha,\theta)$, so there exists a morphism $\mu : \mathcal{M}(Q,\theta) \to \operatorname{Im}(\rho)$ with $\mu \circ \pi = \rho$. One can deduce from Proposition 2.13 by the Proj construction of $\mathcal{M}(Q,\theta)$ that μ is an isomorphism. $\square$

Next we recall some of the results in [4] and [26] that give us a combinatorial description of the fan of the toric variety $\mathcal{M}(Q,\theta)$. For a set of vertices $V \subseteq Q_0$ we will write $\theta(V) = \sum_{v \in V} \theta(v)$. For a representation $R \in \operatorname{Rep}(Q)$ we will denote by $\operatorname{supp}(R)$ the set $\{a \in Q_0 \mid R(a) \neq 0\}$. Clearly for any $t^m \in \mathcal{O}(\operatorname{Rep}(Q))$ we have $t^m(R) \neq 0$ if and only if $\operatorname{supp}(m) \subseteq \operatorname{supp}(R)$. Note that a subrepresentation of R is given by choosing a set of vertices $V \subseteq Q_0$ such that no arrow in $\operatorname{supp}(R)$ leaves V. Taking these facts and Proposition 2.4 into account one can derive Proposition 2.16 and Proposition 2.17. For $A \subseteq Q_1$ we will call a set $H \subseteq Q_0$ *A-successor closed* if there is no $a \in A$ with $a^- \in H$ and $a^+ \in Q_0 \setminus H$, we will call a Q_1-successor closed set simply *successor closed*.

Proposition 2.16 *The following are equivalent:*

(i) R is θ-semistable.

(ii) $\theta(Q_0) = 0$ and for any $\operatorname{supp}(R)$-successor closed vertex set V we have $\theta(V) \geq 0$.

Proposition 2.17 *The following are equivalent:*

(i) R is θ-stable.

(ii) $\theta(Q_0) = 0$ and for any $\mathrm{supp}(R)$-successor closed, nonempty, proper subset of vertices $V \subset Q_0$, we have $\theta(V) > 0$.

An important consequence of Propositions 2.16 and 2.17 is that the property of being θ-semistable or θ-stable only depends on *supp*(R) and θ. Note that it also follows that whenever $\mathrm{supp}(R) \subseteq \mathrm{supp}(R')$, the semistability (resp. stability) of R implies the semistability (resp. stability) of R'. In particular Q is θ-stable if and only if the representations R with $\mathrm{supp}(R) = Q_1$ are θ-stable.

Following [4] we call a representation R *θ-polystable* if the connected components of the quiver with vertex set Q_0 and arrow set *supp*(R) are θ-stable. We call a set of arrows $A \subseteq Q_1$ polystable (resp. stable) if there is a polystable (resp. stable) representation R with *supp*$(R) = A$. For easier formulation of the upcoming propostions we fix the convention that a quiver without arrows is θ-stable if and only if $\theta = 0$, note that this implies that the zero representation (or equivalently the empty arrow set) is θ-polystable if and only if $\theta = 0$. With this convention in mind we allow the forests in Corollary 2.19 to be empty. We recall (4) from Lemma 7 in [4]:

Proposition 2.18 *The following are equivalent:*

(i) R is θ-polystable.

(ii) There exists an $x \in \nabla(Q, \theta)$ such that supp(x) = supp(R).

Corollary 2.19 *(i) The map that assigns to a set of arrows $A \subseteq Q_1$ the set $\{x \in \nabla(Q, \theta) \mid \forall a \in Q_1 \setminus A : x(a) = 0\}$ establishes an inclusion preserving bijection between the set $\{\mathrm{supp}(R) \mid R \text{ is } \theta\text{-polystable}\}$ and the faces of $\nabla(Q, \theta)$.*

(ii) The cones in the fan of the toric variety $\mathcal{M}(Q, \theta)$ are in inclusion reversing bijection with the set $\{\mathrm{supp}(R) \mid R \text{ is } \theta\text{-polystable}\}$. The cones of maximal dimension correspond to the θ-polystable undirected forests.

Part (i) of Corollary 2.19 is the same as Corollary 8 from [4]. Part (ii) was proven in [26] for the special case when θ is a *generic weight*, meaning that the θ-stable and θ-semistable representations coincide. Since stability of a representations implies polystability, which in turn implies semistability, if we assume that θ is generic one can replace the term "polystable" by "stable" in both Proposition 2.18 and Corollary 2.19.

Remark 2.20 Let us now assume that the arrow sets $Q_1 \setminus \{a\}$ are supports of stable representations for every $a \in Q_1$ (we will see in Section 3.1 that essentially this is the only interesting case). Recall that in this case $\dim(\mathcal{M}(Q,\theta)) = \chi(Q)$. Fix an arbitrary lattice point $v^0 \in \mathrm{AffSpan}(\nabla(Q,\theta)) \cap \mathbb{Z}^{Q_1}$. Now the polytope $\nabla(Q,\theta) - v^0$ is full dimensional in the lattice M^Q, moreover by Corollary 2.19 its maximal faces correspond bijectively to the arrows of Q. Set u_a to be the element in the dual lattice of M^Q, which is the image of the coordinate function corresponding to the arrow a. It follows that the (unique) facet representation of $\nabla(Q,\theta) - v^0$ is:

$$\{x \in M^Q_{\mathbb{R}} \mid \forall a \in Q_1 : \ \langle x, u_a \rangle \geq -v^0(a)\}.$$

We will denote the dual lattice of M^Q by N^Q. One can check without difficulty that the subgroup $\{n \in \mathbb{Z}^{Q_1} \mid \forall m \in M^Q : \ \langle m,n \rangle = 0\}$ is generated by the elements $\sum_{a^+=v} \varepsilon_a - \sum_{a^-=v} \varepsilon_a$. Hence we have

$$N^Q = \mathbb{Z}^{Q_1} / \langle \sum_{a^+=v} \varepsilon_a - \sum_{a^-=v} \varepsilon_a \ \mid \ v \in Q_0 \rangle.$$

We shall conclude this section by noting that some of the results (for example those from [34]) we refer to in this thesis are formulated for *flow polytopes*. Flow polytopes are defined as follows: Given an integral vector $\theta \in \mathbb{Z}^{Q_0}$ and non-negative integral vectors $\mathbf{l}, \mathbf{u} \in \mathbb{N}_0^{Q_1}$ the polytope

$$\nabla = \nabla(Q,\theta,\mathbf{l},\mathbf{u}) = \{x \in \mathbb{R}^{Q_1} \mid \mathbf{l} \leq x \leq \mathbf{u}, \ \forall v \in Q_0 : \ \theta(v) = \sum_{a^+=v} x(a) - \sum_{a^-=v} x(a)\},$$

is called a flow polytope. As we shall point out in Proposition 2.21 below, up to integral-affine equivalence, the class of flow polytopes coincides with the class of quiver polytopes, so the class of quiver polyhedra is the most general among the above classes.

Proposition 2.21 *For any flow polytope $\nabla(Q,\theta,\mathbf{l},\mathbf{u})$ there exists a quiver Q' with no oriented cycles and a weight $\theta' \in \mathbb{Z}^{Q'_1}$ such that the polytopes $\nabla(Q,\theta,\mathbf{l},\mathbf{u})$ and $\nabla(Q',\theta')$ are integral-affinely equivalent.*

Proof. Note that $x \in \mathbb{R}^{Q_1}$ belongs to $\nabla(Q,\theta,\mathbf{l},\mathbf{u})$ if and only if $x - \mathbf{l}$ belongs to $\nabla(Q,\theta',\mathbf{0},\mathbf{u}-\mathbf{l})$ where θ' is the weight given by $\theta'(v) = \theta(v) - \sum_{a^+=v} \mathbf{l}(a) + \sum_{a^-=v} \mathbf{l}(a)$. Consequently $X_{\nabla(Q,\theta,\mathbf{l},\mathbf{u})} = X_{\nabla(Q,\theta',\mathbf{0},\mathbf{u}-\mathbf{l})}$. Therefore it is sufficient to deal with the flow polytopes $\nabla(Q,\theta,\mathbf{0},\mathbf{u})$. Define a new quiver Q' as follows: add to the vertex set of Q two new vertices v_a, w_a for each $a \in Q_1$, and replace the arrow $a \in Q_1$ by three arrows a_1, a_2, a_3, where a_1 goes from a^- to v_a, a_2 goes from w_a to v_a, and a_3 goes from w_a to a^+. Let $\theta' \in \mathbb{Z}^{Q'_0}$ be the weight with $\theta'(v_a) = \mathbf{u}(a) = -\theta'(w_a)$ for all $a \in Q_1$ and $\theta'(v) = \theta(v)$ for all $v \in Q_0$. Consider the linear map $\varphi : \mathbb{R}^{Q_1} \to \mathbb{R}^{Q'_1}$, $x \mapsto y$, where $y(a_1) := x(a)$, $y(a_3) := x(a)$, and $y(a_2) = \mathbf{u}(a) - x(a)$ for all $a \in Q_1$. It is straightforward to check that φ is an affine linear transformation that restricts to an isomorphism $\mathrm{AffSpan}(\nabla(Q, \theta, \mathbf{0}, \mathbf{u})) \to \mathrm{AffSpan}(\nabla(Q', \theta'))$ with the properties (i) and (ii) in Definition 2.3. $\square$

Well-studied examples of flow polytopes are the Birkhoff polytopes B_n, which are usually defined as the set of $n \times n$ real matrices with non-negative entries satisfying that every row and column sum is 1. To realize them as quiver polytopes consider the complete bipartite quiver $K(n, n)$ that has n sources and n sinks, and an arrow pointing from each source to each sink and set θ to be -1 on the sources and 1 on the sinks. Now after identifying $\mathbb{R}^{K(n.n)_1}$ with the vector space of $n \times n$ real matrices we see that $B_n = \nabla(K(n, n), \theta)$.

Chapter 3

Classification results

3.1 Classification of toric quiver varieties

Throughout this section Q stands for a quiver and $\theta \in \mathbb{Z}^{Q_0}$ for a weight such that $\nabla(Q, \theta)$ is non-empty. We say that we *contract an arrow* $a \in Q_1$ which is not a loop when we pass to the pair $(\hat{Q}, \hat{\theta})$, where $\hat{Q}$ is obtained from Q by removing a and glueing its endpoints a^-, a^+ to a single vertex $v \in \hat{Q}_0$, and setting $\hat{\theta}(v) := \theta(a^-) + \theta(a^+)$ whereas $\hat{\theta}(w) = \theta(w)$ for all vertices $w \in \hat{Q}_0 \setminus \{v\} = Q_0 \setminus \{a^-, a^+\}$.

Definition 3.1 Let Q be a quiver, $\theta \in \mathbb{Z}^{Q_0}$ a weight such that $\nabla(Q, \theta)$ is non-empty.

(i) An arrow $a \in Q_1$ is said to be *removable* if $\nabla(Q, \theta)$ is integral-affinely equivalent to $\nabla(Q', \theta)$, where Q' is obtained from Q by removing the arrow a, so $Q'_0 = Q_0$ and $Q'_1 = Q_1 \setminus \{a\}$.

(ii) An arrow $a \in Q_1$ is said to be *contractible* if $\nabla(Q, \theta)$ is integral-affinely equivalent to $\nabla(\hat{Q}, \hat{\theta})$, where $(\hat{Q}, \hat{\theta})$ is obtained from (Q, θ) by contracting the arrow a.

(iii) The pair (Q, θ) is called *tight* if there is no removable or contractible arrow in Q_1.

An immediate corollary of Definition 3.1 is the following statement:

Proposition 3.2 *Any quiver polyhedron $\nabla(Q, \theta)$ is integral-affinely equivalent to some $\nabla(Q', \theta')$, where (Q', θ') is tight. Moreover, (Q', θ') is obtained from (Q, θ) by successively removing or contracting arrows.*

Remark 3.3 A characterization of the notions in Definition 3.1 in terms of stability conditions can be deduced from Sections 3 and 4 of [4]. In particular it follows from Lemma 13 and the elements of Lemma 7 and Corollary 8 in [4], which we recalled in Proposition 2.18 and Corollary 2.19 that an arrow a is contractible if and only if $Q_1 \setminus \{a\}$ is θ-polystable. Moreover an arrow is removable if and only if it is not contained in the unique maximal θ-polystable subquiver of Q. These results along with part (i) of Corollary 2.19 imply Corollary 3.6 below, for which we give a direct derivation from Definition 3.1. We further note that in [4] a quiver is defined to be θ-tight if $Q \setminus \{a\}$ is θ-stable for all a, hence by the above discussion we see that a quiver is θ tight in the sense of Definition 12 of [4] if it is connected and (Q, θ) is tight in the sense used in this thesis.

Lemma 3.4 *(i) Denote by $\hat{Q}, \hat{\theta}$ the quiver and weight obtained by contracting $a \in Q_1$. $\nabla(\hat{Q}, \hat{\theta})$ is integral-affinely equivalent to the polyhedron*

$$\{x \in \mathbb{R}^{Q_1} \mid \forall b \in Q_1 \setminus \{a\} : x(b) \geq 0\} \cap \mathcal{F}^{-1}(\theta).$$

(ii) The arrow a is contractible if and only if in the affine space $\mathcal{F}^{-1}(\theta)$ the halfspace $\{x \in \mathcal{F}^{-1}(\theta) \mid x(a) \geq 0\}$ contains the polyhedron $\{x \in \mathcal{F}^{-1}(\theta) \mid x(b) \geq 0 \quad \forall b \in Q_1 \setminus \{a\}\}$.

Proof. Since the set of arrows of $\hat{Q}$ can be identified with $\hat{Q}_1 = Q_1 \setminus \{a\}$, we have the projection map $\pi : \mathcal{F}^{-1}(\theta) \to \mathcal{F}'^{-1}(\hat{\theta})$ obtained by forgetting the coordinate $x(a)$. The equation

$$x(a) = \theta(a^+) - \sum_{b \in Q_1 \setminus \{a\}, b^+ = a^+} x(b) + \sum_{b \in Q_1 \setminus \{a\}, b^- = a^+} x(b)$$

shows that π is injective, hence it gives an affine linear isomorphism $\mathcal{F}^{-1}(\theta) \cap \mathbb{Z}^{Q_1}$ and $\mathcal{F}'^{-1}(\hat{\theta}) \cap \mathbb{Z}^{\hat{Q}_1}$, and maps the lattice polyhedron

$$\{x \in Q_1 \mid \forall x \in Q_1 \setminus \{a\} : x \geq 0\} \cap \mathcal{F}^{-1}(\theta)$$

onto $\nabla(\hat{Q}, \hat{\theta})$. This proves (i). Now from (i) it follows, that a is contractible if and only if on the affine space $\mathcal{F}^{-1}(\theta)$ the inequality $x(a) \geq 0$ is a consequence of the inequalities $x(b) \geq 0$ $(b \in Q_1 \setminus \{a\})$, proving (ii). □

Lemma 3.5 *(i) Denote by $\hat{Q}, \hat{\theta}$ the quiver and weight obtained by removing $a \in Q_1$. $\nabla(\hat{Q}, \hat{\theta})$ is integral-affinely isomorphic to the polyhedron $\nabla(Q, \theta) \cap \{x \in \mathbb{R}^{Q_1} \mid x(a) = 0\}$.*

(ii) The arrow a is removable if and only if $x(a) = 0$ for all $x \in \nabla(Q, \theta)$.

Proof. (i) follows similarily to (i) of Proposition 3.4 by considering that the projection map $\pi : \mathcal{F}^{-1}(\theta) \to \mathcal{F}'^{-1}(\hat{\theta})$ induces a linear isomorphism of lattices $\{x \in \mathbb{R}^{Q_1} \mid x(a) = 0\} \cap \mathcal{F}^{-1}(\theta) \cap \mathbb{Z}^{Q_1}$ and $\mathcal{F}'^{-1}(\hat{\theta}) \cap \mathbb{Z}^{\hat{Q}_1}$, and (ii) follows trivially from (i). □

For an arrow $a \in Q_1$ set $\nabla(Q, \theta)_{x(a)=0} := \{x \in \nabla(Q, \theta) \mid x(a) = 0\}$.

Corollary 3.6 *(i) The pair (Q, θ) is tight if and only if the assignment $a \mapsto \nabla(Q, \theta)_{x(a)=0}$ gives a bijection between Q_1 and the facets (codimension 1 faces) of $\nabla(Q, \theta)$.*

(ii) If (Q, θ) is tight, then $\dim(\nabla(Q, \theta)) = \chi(Q)$.

Proof. Lemmas 3.4 and 3.5 show that (Q, θ) is tight if and only if $\mathrm{AffSpan}(\nabla(Q, \theta)) = \mathcal{F}^{-1}(\theta)$ and $\{x(a) = 0\} \cap \mathcal{F}^{-1}(\theta)$ $(a \in Q_1)$ are distinct supporting hyperplanes of $\nabla(Q, \theta)$ in its affine span. □

The following simple sufficient condition for contractibility of an arrow turns out to be sufficient for our purposes. For a subset $S \subseteq Q_0$ set $\theta(S) := \sum_{v \in S} \theta(v)$. By (2.2) for $x \in \mathcal{F}^{-1}(\theta)$ we have

$$\theta(S) = \sum_{a \in Q_1, a^+ \in S} x(a) - \sum_{a \in Q_1, a^- \in S} x(a) = \sum_{a^+ \in S, a^- \notin S} x(a) - \sum_{a^- \in S, a^+ \notin S} x(a). \qquad (3.1)$$

Proposition 3.7 *Suppose that $S \subset Q_0$ has the property that there is at most one arrow a with $a^+ \in S$, $a^- \notin S$ and at most one arrow b with $b^+ \notin S$ and $b^- \in S$. Then a (if exists) is contractible when $\theta(S) \geq 0$ and b (if exists) is contractible when $\theta(S) \leq 0$.*

Proof. By (3.1) we have $\theta(S) = x(a) - x(b)$, hence by Lemma 3.4 a or b is contractible, depending on the sign of $\theta(S)$. □

Corollary 3.8 *(i) Suppose that the vertex $v \in Q_0$ has valency 2, and $a, b \in Q_1$ are arrows such that $a^+ = b^- = v$. Then the arrow a is contractible when $\theta(v) \geq 0$ and b is contractible when $\theta(v) \leq 0$.*

(ii) Suppose that for some $c \in Q_1$, c^- and c^+ have valency 2, and $a, b \in Q_1 \setminus \{c\}$ with $a^- = c^-$ and $b^+ = c^+$. Then a is contractible when $\theta(c^-) + \theta(c^+) \leq 0$ and b is contractible when $\theta(c^-) + \theta(c^+) \geq 0$.

Proof. Apply Proposition 3.7 with $S = \{v\}$ to get (i) and with $S = \{c^-, c^+\}$ to get (ii). □

Proposition 3.9 *Suppose that there are exactly two arrows $a, b \in Q_1$ (none is a loop) adjacent to some vertex v, and either $a^+ = b^+ = v$ or $a^- = b^- = v$. Let Q' be the quiver obtained after replacing*

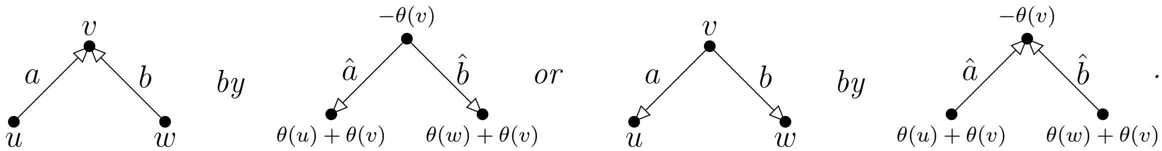

That is, replace the arrows a, b by $\hat{a}$ and $\hat{b}$ obtained by reversing them, and consider the weight $\theta' \in \mathbb{Z}^{Q'_1}$ given by $\theta'(v) = -\theta(v)$, $\theta'(u) = \theta(u) + \theta(v)$ when $u \neq v$ is an endpoint of a or b , and $\theta'(z) = \theta(z)$ for all other $z \in Q'_0 = Q_0$. Then the polyhedra $\nabla(Q, \theta)$ and $\nabla(Q', \theta')$ are integral-affinely equivalent.

Proof. It is straightforward to check that the map $\varphi : \mathbb{R}^{Q_1} \to \mathbb{R}^{Q'_1}$ given by $\varphi(x)(\hat{a}) = x(b)$, $\varphi(x)(\hat{b}) = x(a)$, and $\varphi(x)(c) = x(c)$ for all $c \in Q'_1 \setminus \{\hat{a}, \hat{b}\} = Q_1 \setminus \{a, b\}$ restricts to an isomorphism between $\mathrm{AffSpan}(\nabla(Q, \theta))$ and $\mathrm{AffSpan}(\nabla(Q', \theta'))$ satisfying (i) and (ii) in Definition 2.3. $\square$

Remark 3.10 Proposition 3.9 can be interpreted in terms of *reflection transformations*: it was shown in Sections 2 and 3 in [30] (see also Theorem 23 in [43]) that reflection transformations on representations of quivers induce isomorphisms of algebras of semi-invariants. Now under our assumptions a reflection transformation at vertex v fixes the dimension vector $(1, \ldots, 1)$.

Proposition 3.11 *Suppose that Q is the union of its subquivers Q', Q'' which are either disjoint or have a single common vertex v. Identify $\mathbb{R}^{Q'_1} \oplus \mathbb{R}^{Q''_1} = \mathbb{R}^{Q_1}$ in the obvious way, and let $\theta' \in \mathbb{Z}^{Q'_0} \subset \mathbb{Z}^{Q_0}$, $\theta'' \in \mathbb{Z}^{Q''_0} \subset \mathbb{Z}^{Q_0}$ be the unique weights with $\theta = \theta' + \theta''$ and $\theta'(v) = -\sum_{w \in Q'_0 \setminus \{v\}} \theta(w)$, $\theta''(v) = -\sum_{w \in Q''_0 \setminus \{v\}} \theta(w)$ when $Q'_0 \cap Q''_0 = \{v\}$.*

(i) Then the quiver polyhedron $\nabla(Q, \theta)$ is the product of the polyhedra $\nabla(Q', \theta')$ and $\nabla(Q'', \theta'')$.

(ii) We have $\mathcal{M}(Q, \theta) \cong \mathcal{M}(Q', \theta') \times \mathcal{M}(Q'', \theta'')$.

Proof. (i) A point $x \in \mathbb{R}^{Q_1}$ uniquely decomposes as $x = x' + x''$, where $x'(a) = 0$ for all $a \notin Q'_1$ and $x''(a) = 0$ for all $a \notin Q''_1$. It is obvious by definition of quiver polyhedra that $x \in \nabla(Q, \theta)$ if and only if $x' \in \nabla(Q', \theta')$ and $x'' \in \nabla(Q'', \theta'')$.

(ii) was observed already in [25] and follows from (i) by Corollary 2.14. $\square$

Definition 3.12 (i) We call a connected undirected graph Γ (with at least one edge) *prime* if it is not the union of proper subgraphs Γ', Γ'' having only one common vertex (i.e. it is 2-vertex-connected). A quiver Q will be called *prime* if its underlying graph is prime.

(ii) We call a positive dimensional toric variety (resp. a lattice polyhedron) *prime* if it is not the product of lower dimensional toric varieties (resp. lattice polyhedra).

Obviously any positive dimensional toric variety is the product of prime toric varieties, and this product decomposition is unique up to the order of the factors (see for example Theorem 2.2 in [23]). It is not immediate from the definition, but we shall show in Theorem 3.13 (iii) that the prime factors of a toric quiver variety (resp. a quiver polyhedron) are quiver varieties (resp. quiver polyhedra) as well.

Note that a toric quiver variety associated to a non-prime quiver may well be prime, and conversely, a toric quiver variety associated to a prime quiver can be non-prime, as it is shown by the following example:

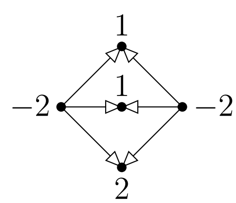

The quiver in the picture is prime but the moduli space corresponding to this weight is $\mathbb{P}^1 \times \mathbb{P}^1$. However, as shown by Theorem 3.13 below, when the tightness of some (Q, θ) is assumed, decomposing Q into its unique maximal prime components gives us the unique decomposition of $\mathcal{M}(Q, \theta)$ as a product of prime toric varieties.

Theorem 3.13 *(i) Let Q^i $(i = 1, \dots, k)$ be the maximal prime subquivers of Q, and denote by $\theta^i \in \mathbb{Z}^{Q_0^i}$ the unique weights satisfying $\sum_{i=1}^k \theta^i(v) = \theta(v)$ for all $v \in Q_0$ and $\sum_{v \in Q_0^i} \theta^i(v) = 0$ for all i. Then $\nabla(Q, \theta)$ is integral-affinely equivalent to $\prod_{i=1}^k \nabla(Q^i, \theta^i)$, and hence $\mathcal{M}(Q, \theta) \cong \prod_{i=1}^k \mathcal{M}(Q^i, \theta^i)$. Moreover, if (Q, θ) is tight, then the (Q^i, θ^i) are all tight.*

(ii) If (Q, θ) is tight then $\mathcal{M}(Q, \theta)$ (resp. $\nabla(Q, \theta)$) is prime if and only if Q is prime.

(iii) Any positive dimensional toric quiver variety (resp. quiver polyhedron) is the product of prime toric quiver varieties (resp. quiver polyhedra).

Proof. The isomorphism $\mathcal{M}(Q,\theta) \cong \prod_{i=1}^{k} \mathcal{M}(Q^i,\theta^i)$ follows from Proposition 3.11 and induction on the number of prime components. The second statement in (i) follows from this isomorphism and Corollary 3.6.

Next we turn to the proof of (ii), so suppose that (Q,θ) is tight. If Q is not prime, then $\chi(Q^i) > 0$ for all i, hence neither of $\mathcal{M}(Q,\theta)$ and $\nabla(Q,\theta)$ are prime by (i). To show the reverse implication for $\mathcal{M}(Q,\theta)$ assume on the contrary that Q is prime, and $\mathcal{M}(Q,\theta) \cong X' \times X''$ where X', X'' are positive dimensional toric varieties. Note that then Q_1 does not contain loops. Let $\{\varepsilon_a \mid a \in Q_1\}$ be a $\mathbb{Z}$-basis of $\mathbb{Z}^{Q_1}$, and for each vertex $v \in Q_0$ let us define $C_v := \sum_{a^+=v} \varepsilon_a - \sum_{a^-=v} \varepsilon_a$. Recall from Remark 2.20 that we can identify the lattice of one-parameter subgroups N^Q of $\mathcal{M}(Q,\theta)$ with $\mathbb{Z}^{Q_1}/\langle C_v \mid v \in Q_0\rangle$, and the ray generators of the fan with the cosets of the ε_a. We will write simply N instead of N^Q for the rest of the proof and denote by N', Σ' and N'', Σ'' the one-parameter subgroups and fans of X' and X'' respectively. Recall from Section 2.1 that $\Sigma = \Sigma' \times \Sigma'' = \{\sigma' \times \sigma'' \mid \sigma' \in \Sigma', \sigma'' \in \Sigma''\}$. Denote by $\pi' : N \to N'$, $\pi'' : N \to N''$ the natural projections to the sets of one-parameter subgroups of the tori in X' and X''. For each ray generator ε_a we have either $\pi'(\varepsilon_a) = 0$ or $\pi''(\varepsilon_a) = 0$. Since (Q,θ) is tight we obtain a partition of Q_1 into two disjoint non-empty sets of arrows: $Q_1' = \{a \in Q_1 \mid \pi''(a) = 0\}$ and $Q_1'' = \{a \in Q_1 \mid \pi'(a) = 0\}$. Since Q is prime, it is connected, hence there exists a vertex w incident to arrows both from Q_1' and Q_1''. Let Π' and Π'' denote the projections from $\mathbb{Z}^{Q_1}$ to $\mathbb{Z}^{Q_1'}$ and $\mathbb{Z}^{Q_1''}$. By choice of w we have $\Pi'(C_w) \neq 0$ and $\Pi''(C_w) \neq 0$. Writing φ for the natural map from $\mathbb{Z}^{Q_1}$ to $N \cong \mathbb{Z}^{Q_1}/\langle C_v \mid v \in Q_0\rangle$ we have $\varphi \circ \Pi' = \pi' \circ \varphi$ and $\varphi \circ \Pi'' = \pi'' \circ \varphi$, so $\ker(\varphi) = \langle C_v \mid v \in Q_0\rangle$ is closed under Π' and Π''. Taking into account that $\sum_{v \in Q_0} C_v = 0$ we deduce that $\Pi'(C_w) = \sum_{v \in Q_0 \setminus \{w\}} \lambda_v C_v$ for some $\lambda_v \in \mathbb{Z}$. Set $S' := \{v \in Q_0 \mid \lambda_v \neq 0\}$. Since each arrow appears in exactly two of the C_v, it follows that S' contains all vertices connected to w by an arrow in Q_1', hence S' is non-empty. Moreover, the set of arrows having exactly one endpoint in S' are exactly those arrows in Q_1' that are adjacent to w. Thus $S'' := Q_0 \setminus (S' \cup \{w\})$ contains all vertices that are connected to w by an arrow from Q_1'', hence S'' is non-empty. Furthermore, there are no arrows in Q_1 connecting a vertex from S' to a vertex in S''. It follows that Q is the union of its full subquivers spanned by the vertex sets $S' \cup \{w\}$ and $S'' \cup \{w\}$, having only one common vertex w and no common arrow. This contradicts the assumption that Q was prime. We have shown that if Q is prime then $\mathcal{M}(Q,\theta)$ is prime, which in turn - by the properties of products we recalled in Section 2.1 - implies that $\nabla(Q,\theta)$ is prime, hence we are done with (ii).

Statement (iii) follows from (i), (ii) and Proposition 3.2. $\square$

Note that if $\chi(\Gamma) \geq 2$ and Γ is prime, then Γ contains no loops (i.e. an edge with identical endpoints), every vertex of Γ has valency at least 2, and Γ has at least two vertices with valency at least 3.

Definition 3.14 For $d = 2, 3, \ldots$ denote by $\mathcal{L}_d$ the set of prime graphs Γ with $\chi(\Gamma) = d$ in which all vertices have valency at least 3. Let $\mathcal{R}_d$ stand for the set of quivers Q obtained from a graph $\Gamma \in \mathcal{L}_d$ by orienting some of the edges somehow and putting a sink on the remaining edges (that is, we replace an edge by a path of length 2 in which both edges are pointing towards the new vertex in the middle). We shall call Γ the *skeleton* $\mathcal{S}(Q)$ of Q; note that $\chi(Q) = \chi(\mathcal{S}(Q))$.

Starting from Q, its skeleton $\Gamma = \mathcal{S}(Q)$ can be recovered as follows: Γ_0 is the subset of Q_0 consisting of the valency 3 vertices. For each path in the underlying graph of Q that connects two vertices in Γ_0 and whose inner vertices have valency 2 we put an edge. Clearly, a quiver Q with $\chi(Q) = d \geq 2$ belongs to $\mathcal{R}_d$ if and only if the following conditions hold:

(i) Q is prime.

(ii) There is no arrow of Q connecting valency 2 vertices.

(iii) Every valency 2 vertex of Q is a sink.

Furthermore, set $\mathcal{R} := \bigsqcup_{d=1}^{\infty} \mathcal{R}_d$ where $\mathcal{R}_1$ is the 2-element set consisting of the 2-Kronecker quiver 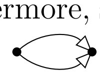 and the quiver with a single vertex and a loop.

Proposition 3.15 *For any $d \geq 2$, $\Gamma \in \mathcal{L}_d$ and $Q \in \mathcal{R}_d$ we have the inequalities*

$$|\Gamma_0| \leq 2d - 2, \quad |\Gamma_1| \leq 3d - 3, \quad |Q_0| \leq 5(d-1), \quad |Q_1| \leq 6(d-1).$$

In particular, $\mathcal{L}_d$ and $\mathcal{R}_d$ are finite for each positive integer d.

Proof. Take $\Gamma \in \mathcal{L}_d$ where $d \geq 2$. Then Γ contains no loops, and denoting by e the number of edges and by v the number of vertices of Γ, we have the inequality $2e \geq 3v$, since each vertex is adjacent to at least three edges. On the other hand $e = v - 1 + d$. We conclude that $v \leq 2d-2$ and hence $e \leq 3d-3$. For $Q \in \mathcal{R}_d$ with $\mathcal{S}(Q) = \Gamma$ we have that $|Q_0| \leq v+e$ and $|Q_1| \leq 2e$. $\square$

Theorem 3.16 *(i) Any d-dimensional prime quiver polyhedron is integral affinely equivalent to $\nabla(Q,\theta)$, where (Q,θ) is a tight-pair and $Q \in \mathcal{R}_d$. Hence any d-dimensional prime toric quiver variety $\mathcal{M}(Q,\theta)$ can be realized by a tight pair (Q,θ) where $Q \in \mathcal{R}_d$ (consequently $|Q_0| \leq 5(d-1)$ and $|Q_1| \leq 6(d-1)$ when $d \geq 2$).*

(ii) For each positive integer d up to isomorphism there are only finitely many d-dimensional toric quiver varieties.

Proof. It follows from Propositions 3.2, Corollary 3.8 and Proposition 3.9 that any d-dimensional prime toric quiver variety (or polyhedron) can be realized by a tight pair (Q,θ) where $Q \in \mathcal{R}_d$; the bounds on vertex and arrow sets of the quiver follow by Proposition 3.15. Statement (ii) follows from (i) and the well-known finiteness of possible GIT-quotients (cf. [45]). More concretely, for a given quiver Q we say that the weights θ and θ' are equivalent if $\mathrm{Rep}(Q)^{\theta-ss} = \mathrm{Rep}(Q)^{\theta'-ss}$; this implies that $\mathcal{M}(Q,\theta) = \mathcal{M}(Q,\theta')$. For a given representation R of Q, the set of weights θ for which R is θ-semistable is determined by the set of dimension vectors of subrepresentations of R (see Proposition 2.4). Since there are finitely many possibilities for the dimension vectors of a subrepresentation of a representation with dimension vector $(1,\dots,1)$, up to equivalence there are only finitely many different weights, hence there are finitely many possible moduli spaces for a fixed Q. □

Remark 3.17 Part (i) of Theorem 3.16 could be deduced from the results in [3] and [4]. From the proof of Theorem 7 in [3] it follows that the bounds on the number of vertices and edges hold whenever the canonical weight is tight for a quiver. While in [3] it is assumed that Q has no oriented cycles, their argument for the bound applies to the general case as well. Moreover Lemma 13 in [4] shows that every toric quiver variety can be realized by a pair (Q,θ) where Q is tight with the canonical weight. These two results imply part (i) of Theorem 3.16.

Remark 3.18 We mention that for a fixed quiver Q it is possible to give an algorithm to produce a representative for each of the finitely many equivalence classes of weights. The change of the moduli spaces of a given quiver when we vary the weight is studied in [25], [26], where the inequalities determining the chamber system were given. To find an explicit weight in each chamber one can use the Fourier-Motzkin algorithm.

Theorem 3.16 is sharp, and the reductions on the quiver are optimal, in the sense that in general one can not hope for reductions that would yield smaller quivers:

Proposition 3.19 *For each natural number $d \geq 2$ there exists a d-dimensional prime toric quiver variety $\mathcal{M}(Q,\theta)$ with $|Q_1| = 6(d-1)$, $|Q_0| = 5(d-1)$, such that for any other quiver and weight Q', θ' with $\mathcal{M}(Q,\theta) \cong \mathcal{M}(Q',\theta')$ (isomorphism of toric varieties) we have that $|Q_1'| \geq |Q_1|$ and $|Q_0'| \geq |Q_0|$.*

Proof. In Example 3.20 for each $d \geq 2$ we show a connected tight pair (Q,θ) with Q prime, $|Q_1| = 6(d-1)$, $|Q_0| = 5(d-1)$, hence $d = \chi(Q) = \dim(\mathcal{M}(Q,\theta))$. Take another quiver and weight Q', θ' with $\mathcal{M}(Q,\theta) \cong \mathcal{M}(Q',\theta')$. Since contracting or removing an arrow does not increase the number of arrows or vertices, there exists a tight pair (Q'', θ'') with $\mathcal{M}(Q',\theta') \cong \mathcal{M}(Q'',\theta'')$, and $|Q_1'| \geq |Q_1''|$, $|Q_0'| \geq |Q_0''|$. By Corollary 3.6 (i) $|Q_1|$ equals the number of facets of the polytope $\nabla(Q,\theta)$, which equals the number of rays in the toric fan of $\mathcal{M}(Q,\theta)$. This is an invariant of the toric variety, implying by $\mathcal{M}(Q'',\theta') \cong \mathcal{M}(Q,\theta)$ that $|Q_1''| = |Q_1|$. Moreover, by Corollary 3.6 (ii) we have $\chi(Q'') = \chi(Q)$, thus $|Q_0''| = |Q_1''| - \chi(Q'') + \chi_0(Q'') = |Q_1| - \chi(Q) + \chi_0(Q'') \geq |Q_1| - \chi(Q) + 1 = |Q_0|$. $\square$

Example 3.20 For $d \geq 3$ consider the graph below with $2(d-1)$ vertices. Removing any two edges from this graph we obtain a connected graph. Now let Q be the quiver obtained by putting a sink on each of the edges (so the graph below is the skeleton of Q). Then (Q, δ_Q) is tight by Corollary 3.33 (δ_Q takes value 2 on each sink and value -3 on each source).

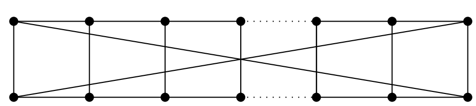

Relaxing the condition on tightness it is possible to come up with a shorter list of quivers whose moduli spaces exhaust all possible projective toric quiver varieties. A key role is played by the following statement:

Proposition 3.21 *Suppose that Q has no oriented cycles and $a \in Q_1$ is an arrow such that contracting it we get a quiver (i.e. the quiver $\hat{Q}$ described in Definition 3.1) that has no oriented cycles. Then for a sufficiently large integer d we have that a is contractible for the pair $(Q, \theta + d(\varepsilon_{a^+} - \varepsilon_{a^-}))$, where $\varepsilon_v \in \mathbb{Z}^{Q_0}$ stands for the characteristic function of $v \in Q_0$.*

Proof. Set $\psi_d = \theta + d(\varepsilon_{a^+} - \varepsilon_{a^-})$, and note that $\hat{\psi}_d = \hat{\theta}$ for all d. Considering the embeddings $\pi : \mathcal{F}^{-1}(\psi_d) \to \mathcal{F}'^{-1}(\hat{\theta})$ described in the proof of Lemma 3.4, we have that for any d, any $y \in \mathcal{F}^{-1}(\psi_d)$ and $b \in Q_1 \setminus \{a\}$,

$$\min\{x(b) \mid x \in \nabla(\hat{Q},\hat{\theta})\} \leq y(b) \leq \max\{x(b) \mid x \in \nabla(\hat{Q},\hat{\theta})\}$$

Since we assumed that $\hat{Q}$ has no oriented cycles, the minimum and the maximum in the inequality above are finite. Now considering the arrows incident to a^- we obtain that for any $x \in \mathcal{F}^{-1}(\psi_d)$ we have $x(a) = d - \theta(a^-) + \sum_{b^+=a^-} x(b) - \sum_{b^-=a^-, b\neq a} x(b)$. Thus for $d \geq \theta(a^-) - \min\{\sum_{b^+=a^-} x(b) - \sum_{b^-=a^-, b\neq a} x(b) \mid x \in \mathcal{F}'^{-1}(\hat{\theta})\}$ the arrow a is contractible for (Q, ψ_d) by Lemma 3.4. □

For $d \geq 2$ introduce a partial ordering $\geq$ on $\mathcal{L}_d$: we set $\Gamma \geq \Gamma'$ if Γ' is obtained from Γ by contracting an edge, and take the transitive closure of this relation. Now for each positive integer $d \geq 2$ denote by $\mathcal{L}'_d \subseteq \mathcal{L}_d$ the set of undirected graphs $\Gamma \in \mathcal{L}_d$ that are maximal with respect to the relation $\geq$, and set $\mathcal{L}'_1 := \mathcal{L}_1$. It is easy to see that for $d \geq 2$, $\mathcal{L}'_d$ consists of 3-regular graphs (i.e. graphs in which all vertices have valency 3). Now denote by $\mathcal{R}'_d$ the quivers which are obtained by putting a sink on each edge from a graph from $\mathcal{L}'_d$.

Theorem 3.22 *For $d \geq 2$ any prime d-dimensional quiver polytope is integral affinely equivalent to $\nabla(Q, \theta)$ where $Q \in \mathcal{R}'_d$, and hence any prime d-dimensional projective toric quiver variety is isomorphic to $\mathcal{M}(Q, \theta)$ where $Q \in \mathcal{R}'_d$.*

Proof. This is an immediate consequence of Theorem 3.16 and Proposition 3.21. □

Example 3.23 $\mathcal{L}'_3$ consists of two graphs:

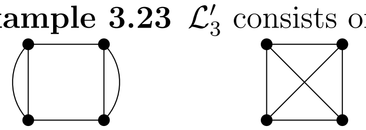

Now put a sink on each edge of the above graphs. The first of the two resulting quivers is not tight for the canonical weight. After tightening we obtain the following two quivers among whose moduli spaces all 3-dimensional prime projective toric quiver varieties occur:

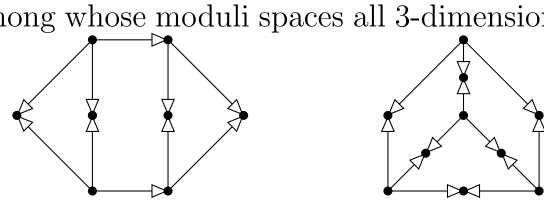

3.2 The 2-dimensional case

As an illustration of the general classification scheme explained in Section 3.1, we quickly reproduce the classification of 2-dimensional toric projective quiver varieties (this result is known, see Theorem 5.2 in [25] and Example 6.14 in [20]):

Proposition 3.24 *(i) A 2-dimensional projective toric quiver variety is isomorphic to one of the following:*
The projective plane $\mathbb{P}^2$, the blow up of $\mathbb{P}^2$ in one, two, or three points in general position, or $\mathbb{P}^1 \times \mathbb{P}^1$.

(ii) The above varieties are realized (in the order of their listing) by the following quiver-weight pairs:

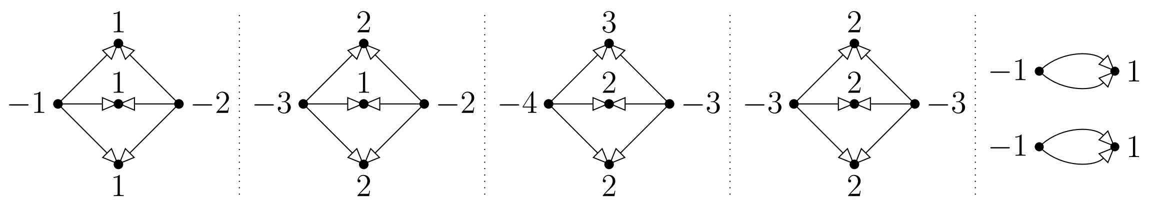

Proof. The only acyclic quiver in $\mathcal{R}_1$ is the Kronecker quiver. The only weights yielding a non-empty moduli space are $(-1, 1)$ and its positive integer multiples, hence the corresponding moduli space is $\mathbb{P}^1$. Thus $\mathbb{P}^1 \times \mathbb{P}^1$, the product of two projective lines occurs as a 2-dimensional toric quiver variety, say for the disjoint union of two copies of 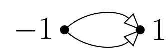

.

$\mathcal{L}_2$ consists of the graph with two vertices and three edges connecting them (say by Proposition 3.15). Thus $\mathcal{R}_2'$ consists of the following quiver:

A:

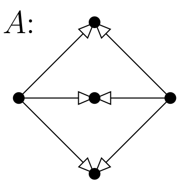

Choosing a spanning tree T in Q, the $x(a)$ with $a \in Q_1 \setminus T_1$ can be used as free coordinates in $\mathrm{AffSpan}(\nabla(Q, \theta))$. For example, take in the quiver A the spanning tree with thick arrows

Figure 3.1: *The polytope* $\nabla(A,\theta)$

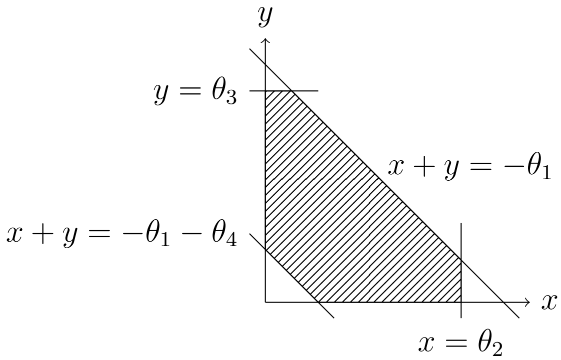

Defining inequalitites:

$$0 \le x \le \theta_2$$
$$0 \le y \le \theta_3$$
$$-\theta_4-\theta_1 \le x+y \le -\theta_1.$$

in the following figure:

AffSpan($\nabla(A,\theta)$):

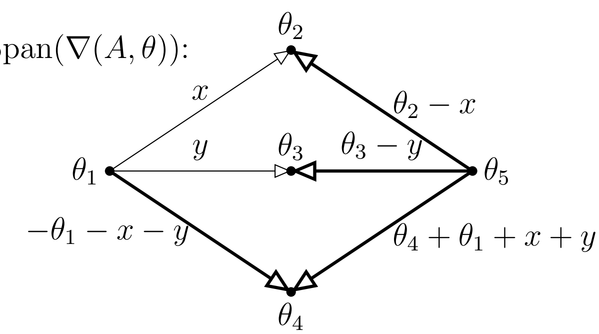

$\sum_{i=1}^{5}\theta_i = 0$

Clearly $\nabla(A,\theta)$ is integral-affinely equivalent to the polytope in $\mathbb{R}^2 = \{(x,y) \mid x,y \in \mathbb{R}\}$ shown on Figure 2.1. Depending on the order of $-\theta_1, \theta_3, -\theta_1-\theta_4, \theta_2$, its normal fan is one of the following:

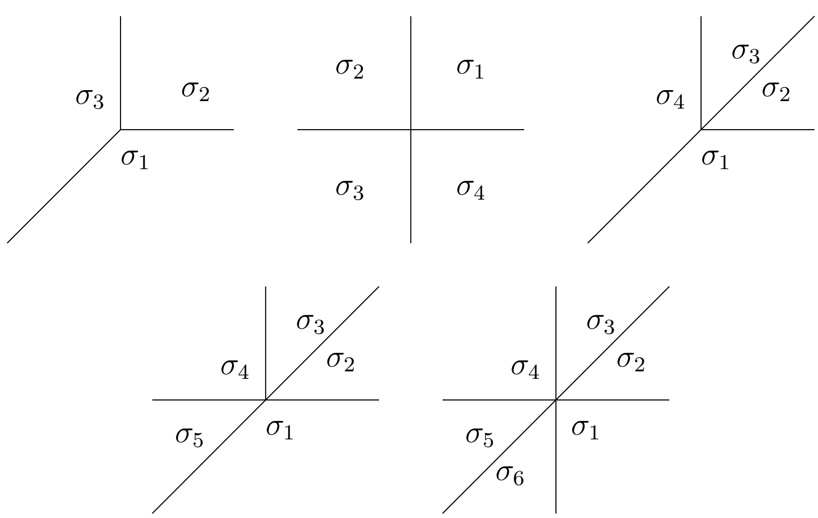

It is well known that the corresponding toric varieties are the projective plane $\mathbb{P}^2$, $\mathbb{P}^1 \times \mathbb{P}^1$ and the projective plane blown up in one, two, or three points in general position, so (i) is proved. Taking into account the explicit inequalities in Figure 2.1, we see that for the pairs (A, θ) given in (ii), the variety $X_{\nabla(A,\theta)} = \mathcal{M}(A, \theta)$ has the desired isomorphism type. $\square$

Remark 3.25 (i) Since the toric fan of the blow up of $\mathbb{P}^2$ in three generic points has 6 rays, to realize it as a toric quiver variety we need a quiver with at least 6 arrows and hence with at least 5 vertices (see Proposition 3.19).

(ii) It is also notable in dimension 2 that each toric quiver variety can be realized as $\mathcal{M}(Q, \delta_Q)$ by precisely one quiver Q from $\mathcal{R}_d$. This does not hold in higher dimensions. For example consider the following quivers:

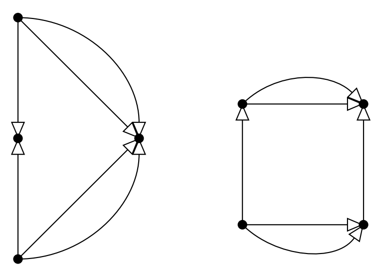

These quivers are both tight with their canonical weights, and give isomorpic moduli, since they are both obtained after tightening:

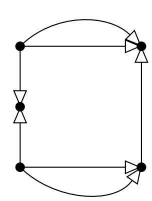

3.3 Affine toric quiver varieties

One of the nice features of toric quiver varieties is that the prinicipal affine open sets corresponding to the vertices of quiver polyhedra are toric quiver varieties themselves. Recall from Corollary 2.14 that for a vertex $v \in \nabla(Q, \theta)$ the set $\operatorname{supp}(v)$ is a forest. Fix a vertex v and let $T^i \quad (i = 1, \dots, r)$ denote the trees in $\operatorname{supp}(v)$. Let (Q^v, θ^v) denote the quiver we get by contracting every arrow in $\operatorname{supp}(v)$. It follows easily (say by Proposition 2.17) that $\theta(T^i) = 0$ for all i, hence $\theta^v = 0$.

Proposition 3.26 *For any vertex v of the quiver polyhedron $\nabla(Q,\theta)$ the affine open toric subvariety U_v in $\mathcal{M}(Q,\theta)$ is isomorphic to $\mathcal{M}(Q^v,0)$.*

Proof. Comparing (iii) from Corollary 2.14 and (i) from Lemma 3.4 we see that the cone of both varieties are integral-affinely equivalent to

$$\{x \in \mathbb{R}^{Q_1} \mid \forall b \in Q_1 \setminus \operatorname{supp}(v) : \ x(b) \geq 0\} \cap \mathcal{F}^{-1}(0).$$

□

Now for a vertex $v \in \nabla(Q,\theta)$ let R be a representation with $\operatorname{supp}(R) = \operatorname{supp}(m)$, then R is polystable and it decomposes as a direct sum of θ-stable representations that are obtained by restricting R to some connected component of $\operatorname{supp}(R)$. Setting $\xi \in U_m$ to be the image of R in $\mathcal{M}(Q,\theta)$ one can check without difficulty that the local quiver Q_ξ is the same as Q^v and α_ξ takes value 1 on every vertex. Hence we obtained that in this particular situation the "étale isomorphism" of Theorem 2.5 can be replaced by a (toric) isomorphism, that maps surjectively onto $\mathcal{M}(Q_\xi,\alpha_\xi,0)$. We will show that this is not entirely specific to the toric case.

Consider the following situation. Let F be a (not necessarily full) subforest of Q which is the disjoint union of trees $F = \coprod_{i=1}^{r} T^i$. Let α be a dimension vector taking the same value d_i on the vertices of each T^i ($i = 1,\dots,r$). Let $\theta \in \mathbb{Z}^{Q_0}$ be a weight such that there exist positive integers n_a ($a \in F_1$) with $\theta(v) = \sum_{a\in F_1:a^+=v} n_a - \sum_{a\in F_1:a^-=v} n_a$. The representation space $\operatorname{Rep}(Q,\alpha)$ contains the Zariski dense open subset

$$U_F := \{R \in \operatorname{Rep}(Q,\alpha) \mid \forall a \in F_1 : \ \det(R(a)) \neq 0\}.$$

Note that U_F is a principal affine open subset in $\operatorname{Rep}(Q,\alpha)$ given by the non-vanishing of the relative invariant $f : R \mapsto \prod_{a\in F_1} \det^{n_a}(R(a))$ of weight θ, hence U_F is contained in $\operatorname{Rep}(Q,\alpha)^{\theta-ss}$. Moreover, U_F is π-saturated with respect to the quotient morphism $\pi : \operatorname{Rep}(Q,\alpha)^{\theta-ss} \to \mathcal{M}(Q,\alpha,\theta)$, hence π maps U_F onto an open subset $\pi(U_F) \cong U_F /\!\!/ GL(\alpha)$ of $\mathcal{M}(Q,\alpha,\theta)$ (here for an affine $GL(\alpha)$-variety X we denote by $X /\!\!/ GL(\alpha)$ the affine quotient, that is, the variety with coordinate ring the ring of invariants $\mathcal{O}(X)^{GL(\alpha)}$), see [36]. Denote by $\hat{Q}$ the quiver obtained from Q by contracting each connected component T^i of F to a single vertex t_i ($i = 1,\dots,r$). So $\hat{Q}_0 = (Q_0 \setminus F_0) \coprod \{t_1,\dots,t_r\}$ and its arrow set can be identified with $Q_1 \setminus F_1$, but if an end vertex of an arrow belongs to T^i in Q then viewed as an arrow in $\hat{Q}$ the correspoding end vertex is t_i (in particular, an arrow in $Q_1 \setminus F_1$

connecting two vertices of T^i becomes a loop at vertex t_i). Denote by $\hat{\alpha}$ the dimension vector obtained by contracting α accordingly, so $\hat{\alpha}(t_i) = d_i$ for $i = 1, \ldots, r$ and $\hat{\alpha}(v) = \alpha(v)$ for $v \in \hat{Q}_0 \setminus \{t_1, \ldots, t_r\}$. Sometimes we shall identify $GL(\hat{\alpha})$ with the subgroup of $GL(\alpha)$ consisting of the elements $g \in GL(\alpha)$ with the property that $g(v) = g(w)$ whenever v, w belong to the same component T^i of F. We have a $GL(\hat{\alpha})$-equivariant embedding

$$\iota : \mathrm{Rep}(\hat{Q}, \hat{\alpha}) \to \mathrm{Rep}(Q, \alpha) \tag{3.2}$$

defined by $\iota(x)(a) = x(a)$ for $a \in \hat{Q}_1$ and $\iota(x)(a)$ the identity matrix for $a \in Q_1 \setminus \hat{Q}_1$. Clearly $\mathrm{Im}(\iota) \subseteq \mathrm{Rep}(Q, \alpha)^{\theta-ss}$.

Proposition 3.27 *(i) $U_F \cong GL(\alpha) \times_{GL(\hat{\alpha})} \mathrm{Rep}(\hat{Q}, \hat{\alpha})$ as affine $GL(\alpha)$-varieties.*

(ii) The map ι induces an isomorphism $\bar{\iota} : \mathcal{M}(\hat{Q}, \hat{\alpha}, 0) \xrightarrow{\cong} \pi(U_F) \subseteq \mathcal{M}(Q, \alpha, \theta)$.

Proof. (i) Set $p := \iota(0) \in \mathrm{Rep}(Q, \alpha)$. Clearly $GL(\hat{\alpha})$ is the stabilizer of p in $GL(\alpha)$ acting on $\mathrm{Rep}(Q, \alpha)$, hence the $GL(\alpha)$-orbit O of p is isomorphic to $GL(\alpha)/GL(\hat{\alpha})$ via the map sending the coset $gGL(\hat{\alpha})$ to $g \cdot p$. On the other hand O is the subset consisting of all those points $R \in \mathrm{Rep}(Q, \alpha)$ for which $\det(R(a)) \neq 0$ for $a \in F_1$ and $R(a) = 0$ for all $a \notin F_1$. This can be shown by induction on the number of arrows of F, using the assumption that F is the disjoint union of trees. Recall also that the arrow set of $\hat{Q}$ is identified with a subset $Q_1 \setminus F_1$. This yields an obvious identification $U_F = \mathrm{Rep}(\hat{Q}, \hat{\alpha}) \times O$. Projection $\varphi : U_F \to O$ onto the second component is $GL(\alpha)$-equivariant by construction. Moreover, the fibre $\varphi^{-1}(p) = \iota(\mathrm{Rep}(\hat{Q}, \hat{\alpha})) \cong \mathrm{Rep}(\hat{Q}, \hat{\alpha})$ as $\mathrm{Stab}_{GL(\alpha)}(p) = GL(\hat{\alpha})$-varieties. It is well known that this implies the isomorphism $U_F \cong GL(\alpha) \times_{GL(\hat{\alpha})} \mathrm{Rep}(\hat{Q}, \hat{\alpha})$, see for example Lemma 5.17 in [8].

(ii) It follows from (i) that $U_F /\!\!/ GL(\alpha) \cong \mathrm{Rep}(\hat{Q}, \hat{\alpha}) /\!\!/ GL(\hat{\alpha}) = \mathcal{M}(\hat{Q}, \hat{\alpha}, 0)$ by standard properties of associated fiber products. Furthermore, taking into account the proof of (i) we see $U_F /\!\!/ GL(\alpha) = \pi(\varphi^{-1}(p)) = \pi(\iota(\mathrm{Rep}(\hat{Q}, \hat{\alpha}))$ where π is the quotient morphism $\mathrm{Rep}(Q, \alpha)^{\theta-ss} \to \mathcal{M}(Q, \alpha, \theta)$. □

Now, still in the situation of Proposition 3.27, consider the representation $R \in \mathrm{Rep}(Q, \alpha)$ which is the identity matrix on the arrows of F and 0 elsewhere, clearly we have $R \in U_F$. Let R_i ($i = 1, \ldots, r$) denote the representation whose dimension vector takes values 1 on the vertices of T^i and 0 on the rest of the vertices, and $R_i(a) = 1$ when $a \in T^i_1$ and $R_i(a) = 0$ otherwise. We have $R = R_1^{d_1} \oplus \cdots \oplus R_i^{d_i}$ and it follows from $n_a > 0$ and

Proposition 2.17 that the representations R_i are θ-stable. Now one easily checks from the definition that the local quiver setting $(Q_{\pi(R)}, \alpha_{\pi(R)})$ is the same as $(\hat{Q}, \hat{\alpha})$, hence the situation in Proposition 3.27 can be viewed as a special case of Theorem 2.5 in which one has an isomorphism between an open subset of the moduli space and the affine moduli space of the local quiver setting. We also point out that Proposition 3.26 can be regarded as a special case of Proposition 3.27.

Conversely, any affine toric quiver variety $\mathcal{M}(Q', 0)$ can be obtained as $U_m \subseteq \mathcal{M}(Q, \theta)$ for some projective toric quiver variety $\mathcal{M}(Q, \theta)$ and a vertex m of the quiver polytope $\nabla(Q, \theta)$. In fact we have a more general result, which is a refinement for the toric case of Theorem 2.2 in [15]:

Theorem 3.28 *For any quiver polyhedron $\nabla(Q, \theta)$ with k vertices there exists a bipartite quiver $\tilde{Q}$, a weight $\theta' \in \mathbb{Z}^{\tilde{Q}_1}$, and a set $m_1, \ldots, m_k$ of vertices of the quiver polytope $\nabla(\tilde{Q}, \theta')$ such that the quasi-projective toric variety $\mathcal{M}(Q, \theta)$ is isomorphic to the open subvariety $\bigcup_{i=1}^k U_{m_i}$ of the projective toric quiver variety $\mathcal{M}(\tilde{Q}, \theta')$.*

Proof. Double the quiver Q to get a bipartite quiver $\tilde{Q}$ as on page 56 in [40]: to each $v \in Q_0$ there corresponds a source v_- and a sink v_+ in $\tilde{Q}$, for each $a \in Q_1$ there is an arrow in $\tilde{Q}$ denoted by the same symbol a, such that if $a \in Q_1$ goes from v to w, then $a \in \tilde{Q}_1$ goes from v_- to w_+, and for each $v \in Q_0$ there is a new arrow $e_v \in \tilde{Q}_1$ from v_- to v_+. Denote by $\tilde{\theta} \in \mathbb{Z}^{\tilde{Q}_0}$ the weight $\tilde{\theta}(v_-) = 0$, $\tilde{\theta}(v_+) = \theta(v)$, and set $\kappa \in \mathbb{Z}^{\tilde{Q}_0}$ with $\kappa(v_-) = -1$ and $\kappa(v_+) = 1$ for all $v \in Q_0$.

Suppose that F is a θ-stable subtree in Q. Denote by $\tilde{F}$ the subquiver of $\tilde{Q}$ consisting of the arrows with the same label as the arrows of F, in addition to the arrows e_v for each $v \in F_0$. It is clear that $\tilde{F}$ is a subtree of $\tilde{Q}$. We claim that $\tilde{F}$ is $(\tilde{\theta}+d\kappa)$-stable for sufficiently large d. Obviously $(\tilde{\theta} + d\kappa)(\tilde{F}_0) = 0$. Let $\tilde{S}$ be a proper successor closed subset of $\tilde{F}_0$ in $\tilde{Q}$. Denote by $S \subset F_0$ the set consisting of $v \in F_0$ with $v_+ \in \tilde{S}$ (note that $v_- \in S$ implies $v_+ \in S$, since $e_v \in \tilde{F}$). We have the equality $(\tilde{\theta} + d\kappa)(\tilde{S}) = \theta(S) + \sum_{v^+ \in \tilde{S}, v_- \notin \tilde{S}} (\theta(v) + d)$. If the second summand is the empty sum (i.e. $v_+ \in \tilde{S}$ implies $v_- \in \tilde{S}$), then S is successor closed, hence $\theta(S) > 0$ by assumption. Otherwise the sum is positive for sufficiently large d. This proves the claim. It follows that if d is sufficiently large, then for any vertex m of $\nabla(Q, \theta)$, setting $F := \mathrm{supp}(m)$, there exists a vertex $\tilde{m}$ of $\nabla(\tilde{Q}, \tilde{\theta} + d\kappa)$ with $\mathrm{supp}(\tilde{m}) = \tilde{F}$.

Denote by $\mu : \mathrm{Rep}(Q) \to \mathrm{Rep}(\tilde{Q})$ the map defined by $\mu(x)(e_v) = 1$ for each $v \in Q_0$, and $\mu(x)(a) = x(a)$ for all $a \in \tilde{Q}_1$. This is equivariant, where we identify $(\mathbb{C}^\times)^{Q_0}$ with the stabilizer of $\mu(0)$ in $(\mathbb{C}^\times)^{\tilde{Q}_0}$. The above considerations show that $\mu(\mathrm{Rep}(Q)^{\theta-ss}) \subseteq$

$\mathrm{Rep}(Q)^{(\tilde{\theta}+d\kappa)-ss}$, whence μ induces a morphism $\bar{\mu} : \mathcal{M}(Q,\theta) \to \mathcal{M}(\tilde{Q}, \tilde{\theta} + d\kappa)$. Restrict $\bar{\mu}$ to the affine open subset $U_m \subseteq \mathcal{M}(Q,\theta)$, and compose $\bar{\mu}|_{U_m}$ with the isomorphism $\bar{\iota} : \mathcal{M}(Q^m, 0) \to U_m \subseteq \mathcal{M}(Q,\theta)$ from Proposition 3.27. By construction we see that $\bar{\mu}|_{U^m} \circ \bar{\iota}$ is the isomorphism $\mathcal{M}(Q^m, 0) \to U_{\tilde{m}}$ of Proposition 3.26. It follows that $\bar{\mu}|_{U_m} : U_m \to U_{\tilde{m}}$ is an isomorphism. As m ranges over the vertices of $\nabla(Q,\theta)$, these isomorphisms glue together to the isomorphism $\bar{\mu} : \mathcal{M}(Q,\theta) \to \bigcup_{\tilde{m}} U_{\tilde{m}} \subseteq \mathcal{M}(\tilde{Q}, \tilde{\theta})$. □

We note that similarly to Theorem 2.2 in [15], it is possible to embed $\mathcal{M}(Q,\theta)$ as an open subvariety into a projective variety $\mathcal{M}(\tilde{Q}, \theta')$, such that for any vertex m' of $\nabla(\tilde{Q}, \theta')$ the affine open subvariety $U_{m'} \subseteq \mathcal{M}(\tilde{Q}, \theta')$ is isomorphic to $U_m \subseteq \mathcal{M}(Q,\theta)$ for some vertex m of $\nabla(Q,\theta)$ (but of course typically $\nabla(\tilde{Q}, \theta')$ has more vertices than $\nabla(Q,\theta)$). In particular, a smooth variety $\mathcal{M}(Q,\theta)$ can be embedded into a smooth projective toric quiver variety $\nabla(\tilde{Q}, \theta')$, where $\tilde{Q}$ is bipartite.

Next we apply our results from Section 3.1 to the special case $\theta = 0$. It follows from Proposition 2.17 that Q is 0-stable if and only if Q is *strongly connected*, that is, for any ordered pair $v, w \in Q_0$ there is an oriented path in Q from v to w.

Proposition 3.29 *Let Q be a prime quiver with $\chi(Q) \geq 2$, such that $(Q,0)$ is tight. Then $|Q_0| \leq \chi(Q) - 1$ and consequently $|Q_1| = |Q_0| + \chi(Q) - 1 \leq 2(\chi(Q) - 1))$.*

Proof. Since Q is prime and is not just a single loop, it contains no loops at all. Suppose $v \in Q_0$ and $a \in Q_1$ is the only arrow with $a^- = v$. The equations (2.2) imply that for any $x \in \nabla(Q,0)$ we have $x(a) = \sum_{b^+=v} x(b)$, so by Lemma 3.4 the arrow a is contractible. The case when a is the only arrow with $a^+ = v$ is similar. Thus for any $v \in Q_0$ we have $|\{a \in Q_1 \mid a^- = v\}| \geq 2$ and $|\{a \in Q_1 \mid a^+ = v\}| \geq 2$ (this is shown also in Lemma 13 (iii) of [4]). In particular, the valency of any vertex is at least 4, hence similar considerations as in the proof of Proposition 3.15 yield the desired bound on $|Q_0|$. □

Denote by $\mathcal{R}''_d$ the set of prime quivers Q with $\chi(Q) = d$ and $(Q,0)$ tight. Then $\mathcal{R}''_1$ consists of the one-vertex quiver with a single loop, $\mathcal{R}''_2$ is empty, $\mathcal{R}''_3$ consists of the quiver with two vertices, and two arrows from each vertex to the other (so four arrows in total). $\mathcal{R}''_4$ consists of the three quivers:

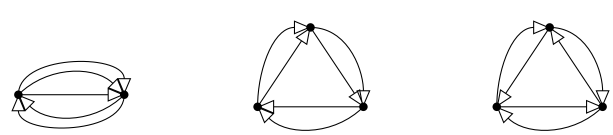

Example 3.30 Consider the quiver Q with d vertices and $2d$ arrows $a_1, \ldots, a_d$, $b_1, \ldots, b_d$, where $a_1 \ldots a_d$ is a primitive cycle and b_i is obtained by reversing a_i for $i = 1, \ldots, d$. Then $\chi(Q) = d + 1$, and after the removal of any of the arrows of Q we are left with a strongly connected quiver. So $(Q, 0)$ is tight, showing that the bound in Proposition 3.29 is sharp. The coordinate ring $\mathcal{O}(\mathcal{M}(Q, 0))$ is the subalgebra of $\mathcal{O}(\mathrm{Rep}(Q))$ generated by $\{x(a_i)x(b_i), x(a_1) \cdots x(a_d), x(b_1) \cdots x(b_d) \mid i = 1, \ldots, d\}$, so it is the factor ring of the $(d+2)$-variable polynomial ring $\mathbb{C}[t_1, \ldots, t_{d+2}]$ modulo the ideal generated by $t_1 \cdots t_d - t_{d+1}t_{d+2}$.

3.4 Reflexive polytopes

A full dimensional lattice polytope in $\mathbb{R}^d$ is called *reflexive* if its facet presentation is

$$P = \{m \in \mathbb{R}^d \mid \langle m, u_F \rangle \geq -1 \text{ for all facets F}\}.$$

From the point of view of toric geometry the study of reflexive polytopes is motivated by the fact that the class of toric varieties associated to them are precisely the Gorenstein Fano toric varieties (see Chapter 8.3 of [12] for an explanation of the terminology and some basic results). Reflexive polytopes also play a role in the study of mirror symmetry (see [5]). As an application of the classification results on quiver polyhedra, we have compiled a full list of quiver polytopes that are integral-affinely equivalent to a reflexive polytope up to dimension 3. Similar work has been done in [3], however the result there is different from ours - they found 39 different reflexive flow polytopes in dimension 3, whereas our list consists of 53 - hence it seemed reasonable to provide the details of our computations.

Note that the property of being reflexive is not invariant under translations, however it is invariant under linear automorphisms of the character lattice. Moreover reflexive polytopes are by definition full dimensional. Consequently we are interested in finding quiver polytopes that can be translated into a lattice polytope ∇ such that ∇ is reflexive when considered in the lattice $\mathrm{AffSpan}(\nabla) \cap \mathbb{Z}^{Q_1}$. The following proposition shows us, that it is enough to consider quivers with their canonical weights.

Proposition 3.31 *If (Q, θ) is tight then $\nabla(Q, \theta)$ is integral-affinely equivalent to a reflexive polytope if and only if θ is the canonical weight.*

Proof. Note that the origin is contained in every reflexive polytope (in fact, it is the unique interior lattice point). Recall from Remark 2.20 that when (Q, θ) is tight the polytope

$\nabla(Q,\theta) - v$ for any $v \in \mathrm{AffSpan}(\nabla(Q,\theta)) \cap \mathbb{Z}^{Q_1}$ has facet representation

$$\{x \in M_{\mathbb{R}}^{Q} \mid \forall a \in Q_1 : \langle x, u_a \rangle \leq -v(a)\}.$$

It follows that $\nabla(Q,\theta)$ is integral-affinely equivalent to a reflexive polytope if and only if $(1,\dots,1) \in \nabla(Q,\theta)$, which happens precisely when θ is the canonical weight. $\square$

By slight abuse of the terminology, we will be calling a quiver polytope reflexive, whenever it is integral-affinely equivalent to a reflexive polytope and note that by the above argument this is equivalent to saying that $\nabla(Q,\theta) - (1,\dots,1)$ is a reflexive polytope in the lattice M^Q. Next we derive a combinatorial characterization for quivers that are tight with their canonical weight δ_Q.

Proposition 3.32 *(i) (Q,δ_Q) is tight if and only if $(Q',\delta_{Q'})$ is tight for every maximal prime subquiver Q' of Q.*

(ii) For a connected quiver Q the pair (Q,δ_Q) is tight if and only if there is no partition $Q_0 = H \coprod H'$, with $\emptyset \neq H \subsetneq Q_0$, such that there is at most one arrow from H to H' and there is at most one arrow from H' to H.

Proof. (i) follows immediately from Theorem 3.13. For (ii) first note that for any $H \subseteq Q_0$ we have that

$$\delta_Q(H) = \#\{a \in Q_1 \mid a^+ \in H,\ a^- \in Q_0 \setminus H\} - \#\{a \in Q_1 \mid a^- \in H,\ a^+ \in Q_0 \setminus H\}.$$

Now assume that there is a partition of Q_0 as in the statement. If there is an arrow a from H' to H and no arrows from H to H' then $x(a) = 1$ for all $x \in \nabla(Q,\theta) \cap \mathbb{Z}^{Q_1}$, hence a is contractible. If there is an arrow a from H' to H and an arrow a' from H to H' then we have $\theta(H) = \theta(H') = 0$ and it follows that $x(a) = x(a')$ for all $x \in \nabla(Q,\theta) \cap \mathbb{Z}^{Q_1}$, which again implies that both a and a' are contractible. For the other direction, if (Q,δ_Q) is not tight then - as it has been explained in Remark 3.3 - it follows from the results in [4] that for some arrow a the set $Q_1 \setminus \{a\}$ is not θ-polystable. If $Q \setminus \{a\}$ is disconnected then its connected components give us a partition of Q_0 as in the Proposition. If $Q \setminus \{a\}$ is connected then it follows from Proposition 2.17 that there is a set of vertices $\emptyset \neq H \subsetneq Q_0$ such that H is $(Q_1 \setminus \{a\})$-successor closed and $\delta_Q(H) \leq 0$. It follows that no arrow other than a goes from H to $Q_0 \setminus H$ and there is at most one arrow going from $Q_0 \setminus H$ to H. $\square$

Corollary 3.33 *Let $Q \in \mathcal{R}_d$ be such that the skeleton $\mathcal{S}(Q)$ is 3-edge-connected (i.e. it remains connected after removing any two of the edges), then (Q, δ_Q) is tight.*

Proof. The case when Q does not contain valency 2 sinks is clear from Proposition 3.32. Now assume that Q contains some valency 2 sinks and there are $H, H' \subseteq Q_0$ as in Proposition 3.32. Clearly every valency 2 sink lies on the same side of the partition $H \coprod H'$ as at least one of its neighbours, otherwise there would be at least 2 arrows going from H' to H or H to H'. Now pick such a neighbour for every valency 2 sink and contract the arrows between them to obtain a new quiver $Q^1 \in \mathcal{R}_d$ that has the same skeleton as Q, and denote by H^1, H'^1 the partition of Q_0^1 we obtain by contracting H and H'. Since the number of arrows running between H and H' is the same as between H^1 and H'^1 we see that the skeleton of Q^1 can not be 3-edge-connected, a contradiction. □

Remark 3.34 (i) Since every quiver polytope can be realized by a tight pair, we see that to obtain a complete (but generally redundant) list of reflexive polytopes up to dimension d it is satisfactory to list prime quivers without oriented cycles that are tight with their canonical weights.

(ii) In dimension 1 the 2-Kronecker quiver is the only quiver that is tight under its canonical weight.

(iii) In dimension 2 it is easy to verify that all the four quivers without oriented cycles in $\mathcal{R}_2$ are tight with their canonical weights, and the resulting reflexive polytopes are pairwise non-isomorphic. Taking into account the product of the single one dimensional reflexive polytope with itself, we see that there are a total of 5 reflexive quiver polytopes in dimension 2.

(iv) Comparing Proposition 3.24 with Section 3.3 in [3] we conclude that for each isomorphism class of 2-dimensional toric quiver varieties there is a quiver Q such that $\mathcal{M}(Q, \delta_Q)$ belongs to the given isomorphism class, in particular in dimension 2 every projective toric quiver variety is Gorenstein Fano. This is explained by the following two facts: (1) in dimension 2, a complete fan is determined by the set of rays; (2) if (Q, θ) is tight, then (Q, δ_Q) is tight. Now (1) and (2) imply that if (Q, θ) is tight and $\chi(Q) = 2$, then $\mathcal{M}(Q, \theta) \cong \mathcal{M}(Q, \delta_Q)$.

(v) The above does not hold in dimension three or higher. Consider for example the quiver-weight pairs:

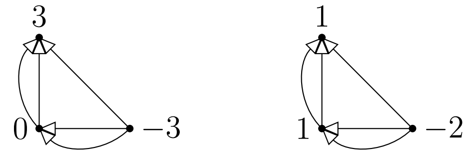

The weight on the left is the canonical weight δ_Q for this quiver, and it is easy to check that (Q, δ_Q) is tight and $\mathcal{M}(Q, \delta_Q)$ is a Gorenstein Fano variety with one singular point. The weight on the right is also tight for this quiver, however it gives a smooth moduli space which can not be isomorphic to $\mathcal{M}(Q, \delta_Q)$, consequently it also can not be Gorenstein Fano since the rays in its fan are the same as those in the fan of $\mathcal{M}(Q, \delta_Q)$.

The following proposition reduces the number of cases that have to be considered for finding all reflexive quiver polytopes in a given dimension:

Proposition 3.35 *(i) For $d \geq 2$ let $Q \in \mathcal{R}_d$ be a quiver without oriented cycles. Then Q can always be obtained from its skeleton $\mathcal{S}(Q)$ by choosing an acyclic orientation on the edges and then adding some valency 2 sinks.*

(ii) For $d \geq 2$ let $Q \in \mathcal{R}_d$ be a quiver such that $\mathcal{S}(Q)$ is not 3-edge-connected. Then there is a quiver $Q' \in \mathcal{R}_d$ such that $\mathcal{S}(Q')$ is 3-edge-connected and $\nabla(Q', \delta_{Q'})$ is integral-affinely equivalent to $\nabla(Q, \delta_Q)$

Proof. For (i) first let us note that the statement is not immediately obvious from the definition, since it can easily happen that a cyclic orientation on the edges of a skeleton becomes acyclic after adding some (for example all) valency 2 sinks to the arrows. The case when Q contains no valency 2 sinks is trivial. Otherwise pick a valency 2 sink v in Q_0 and denote by w_1 and w_2 its neighbours. If there were directed paths going both ways between w_1 and w_2 in Q, then Q would contain an oriented cycle, so without loss of generality, we can assume that there is no directed path going from w_1 to w_2. Now remove v and the arrows incident to it from Q and add an arrow from w_2 to w_1 to obtain a new quiver Q'. Q' contains no oriented cycles, has the same skeleton as Q and one less valency 2 vertex. Now (i) follows by induction on the number of valency 2 vertices.

We prove (ii) by induction on the number of edges in the skeleton of Q. First note that for each $d \geq 2$ the skeleton with the smallest number of edges has 2 vertices and $d+1$ edges running between them and hence it is 3-edge-connected. Now let Q be as in the proposition. If (Q, δ_Q) is not tight then we are done after tightening it and applying the induction hypothesis. Let us now assume that (Q, δ_Q) is tight. Since $\mathcal{S}(Q)$ is not 3-edge-connected, one can easily verify that there is a partition $H \coprod H'$ of the vertices of Q such that there are at most two arrows running between H and H'. Since (Q, δ_Q) is tight, it follows from Proposition 3.32 that there are precisely two arrows between H and H' and that these arrows need to go the same way, say from H to H'. Let us denote these arrows by a_1 and a_2. If one of a_1^+ and a_2^+ is a valency 2 sink then by moving this sink from H' to H and applying Proposition 3.32 we see that (Q, δ_Q) can not be tight, contradicting our assumption. Now let Q' be the quiver we get by removing a_1 and adding a valency 2 sink v with arrows pointing to v from the endpoints of a_1 and Q'' be the quiver we get after contracting the arrow a_2 in Q'. Let us denote by b the arrow incident to v in Q' that is oriented reversely to a_1. Now by applying the same argument as in the proof of Proposition 3.32, we see that for all $x \in \nabla(Q', \delta_{Q'})$ we have that $x(a_2) = x(b)$, hence both a_2 and b are contractible in $\nabla(Q', \delta_{Q'})$. By contracting b in $\nabla(Q', \delta_{Q'})$ we obtain $\nabla(Q, \delta_Q)$ and by contracting a_2 we obtain $\nabla(Q'', \delta_{Q''})$, hence $\nabla(Q, \delta_Q)$ is integral-affinely equivalent to $\nabla(Q'', \delta_{Q''})$. Moreover since a_1^+ was not a valency 2 sink we have that $Q'' \in \mathcal{R}_d$, and $\mathcal{S}(Q'')$ has one less edge than $\mathcal{S}(Q)$ and we are done by induction.

□

Corollary 3.36 *For $d \geq 2$ every reflexive quiver polytope in dimension d is integral-affinely equivalent to $\nabla(Q, \delta_Q)$ for a quiver Q that can be obtained by choosing a 3-edge-connected skeleton from $\mathcal{L}_d$, endowing its arrows with an acyclic orientation and then replacing some (possibly none) of the arrows with valency 2 sinks.*

Corollary 3.36 provides us with a method for obtaining a complete list of reflexive polytopes in a given dimension, however in general it is difficult to eliminate from this list the polytopes that appear multiple times. The following proposition provides a useful invariant for this purpose.

Proposition 3.37 *(i) Let Q be a connected quiver and θ a weight such that (Q, θ) is tight, and let $a_1, a_2 \in Q_1$ be distinct arrows of Q. Then the facets of $\nabla(Q, \theta)$ corresponding to a_1 and a_2 are parallel if and only if $Q \setminus \{a_1, a_2\}$ is disconnected.*

(ii) Let $Q \in \mathcal{R}_d$ be a quiver such that $\mathcal{S}(Q)$ is 3-edge-connected, and let $a_1, a_2 \in Q_1$ be distinct arrows of Q. Then the facets of $\nabla(Q, \theta)$ corresponding to a_1 and a_2 are parallel if and only if a_1 and a_2 are incident to the same valency 2 sink.

(iii) Let $Q^1, Q^2 \in \mathcal{R}_d$ be quivers such that $\mathcal{S}(Q^1)$ and $\mathcal{S}(Q^2)$ are 3-edge-connected. If $\nabla(Q^1, \delta_{Q^1})$ is integral-affinely equivalent to $\nabla(Q^2, \delta_{Q^2}))$, then Q^1 and Q^2 have the same number of valency 2 sinks and $\mathcal{S}(Q^1)$ and $\mathcal{S}(Q^2)$ have the same number of edges.

Proof. For (i) recall from Corollary 3.6 that when Q is tight and connected we have $dim(\nabla(Q,\theta)) = \chi(Q) = |Q_1| - |Q_0| - 1$ and hence

$$\text{AffSpan}(\nabla(Q,\theta)) = \{x \in \mathbb{R}^{Q_1} \mid \forall v \in Q_0 : \theta(v) = \sum_{a^+=v} x(a) - \sum_{a^-=v} x(a)\}.$$

Recall that when (Q,θ) is tight, $M_{\mathbb{R}}^Q$ is the linear subspace parallel to AffSpan$\nabla(Q,\theta)$. Now clearly the facets corresponding to some a_1 and a_2 are parallel if and only if

$$M_{\mathbb{R}}^Q \cap \{x \in \mathbb{R}^{Q_1} \mid x(a_1) = 0\} = M_{\mathbb{R}}^Q \cap \{x \in \mathbb{R}^{Q_1} \mid x(a_2) = 0\}.$$

It follows from Proposition 2.12 that the support of any element in $M_{\mathbb{R}}^Q$ is an undirected cycle of Q. If $Q \setminus \{a_1, a_2\}$ is disconnected then there is no undirected cycle in Q containing precisely one of them, hence from the above it follows that the corresponding facets are parallel. For the other direction if $Q \setminus \{a_1, a_2\}$ is connected then there is a spanning tree $T \subseteq Q_1 \setminus \{a_1, a_2\}$ of Q. Recall the definition of e_F^a for a spanning forest F from the discussion preceding Proposition 2.12. Now $e_T^{a_1} \in M_{\mathbb{R}}^Q \cap \{x \in \mathbb{R}^{Q_1} \mid x(a_2) = 0\}$ but $e_T^{a_1} \notin \{x \in \mathbb{R}^{Q_1} \mid x(a_1) = 0\}$ hence the facets corresponding to a_1 and a_2 are not parallel.

For (ii) we can argue the same way as in the proof of Corollary 3.33 to see that $Q \setminus \{a_1, a_2\}$ is disconnected if and only if a_1 and a_2 are incident to the same valency 2 sink, and then comparing with (i) the statement follows.

For (iii) note that the number of parallel pairs of facets is invariant under integral-affine equivalence, hence the first part of the statement follows immediately from (ii) and the second part from the fact that the number of edges of $\mathcal{S}(Q)$ and the number of valency 2 sinks determine the number of arrows of Q, which in turn equals the number of facets of $\nabla(Q,\theta)$. □

The following proposition is easy to verify:

Proposition 3.38 *Every quiver without oriented cycles in $\mathcal{R}_3$ that has a 3-edge-connected skeleton and contains no valency 2 sinks is isomorphic to one of the following, or the* *opposite quiver of Q^{III} in this list:*

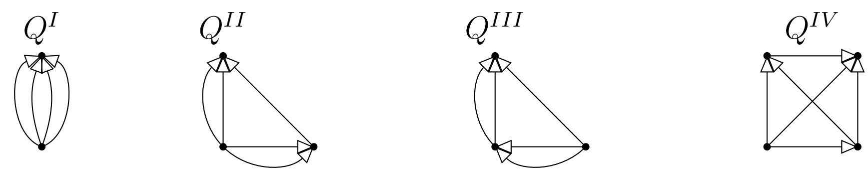

Note that the opposite quiver of Q^{III} need not to be considered, since it yields the same quiver polytopes as Q^{III}. We proceeded to compile a list of pairwise non-isomorphic quivers one can obtain by adding sinks to the quivers in Proposition 3.38, and excluded those cases that can be obtained from ones that we had already listed by a series of anti-isomorphisms (i.e. reversing all arrows), and reflections (i.e. reversing arrows on a valency 2 sink or source) as in Proposition 3.9. We calculated the lattice points of the quiver polytopes we obtain when considering quivers from this list with their canonical weights. This essentially comes down to solving a set of integer inequalities for each of Q^{I}, Q^{II}, Q^{III} and Q^{IV}, and then removing a subset of the solutions depending on the valency 2 sinks added, as illustrated by the following example.

Example 3.39 Label the edges of Q^{II} as in the picture below:

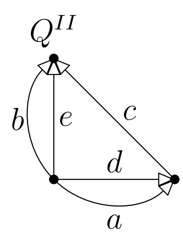

By definition of the quiver polytope $\nabla(Q^{II}, \delta_{Q^{II}})$, we have that $m \in \nabla(Q^{II}, \delta_{Q^{II}}) \cap \mathbb{Z}^{Q_1^{II}}$ if and only if the entries of m are non-negative and satisfy $m(b) + m(c) + m(e) = 3$ and $m(a) + m(d) - m(c) = 1$. Taking a, b, c as free coordinates in $\mathrm{AffSpan}(\nabla(Q^{II}, \delta_{Q^{II}}))$, to list the lattice points of $\nabla(Q^{II}, \delta_{Q^{II}})$, we have to find the integer solutions for the set of inequalities: $m(a), m(b), m(c) \geq 0$, $m(b)+m(c) \leq 3$, $m(a) \leq 1+m(c)$. One can check easily that there are 30 solutions altogether, giving us the number of lattice points in $\nabla(Q^{II}, \delta_{Q^{II}})$. Recall from Corollary 3.33 that Q^{II} is tight with its canonical weight hence each arrow corresponds to a facet and these are given by the equations $x(a) = 0$, $x(b) = 0$, $x(c) = 0$, $x(b) + x(c) = 3$ and $x(a) - x(c) = 1$. One can obtain the list of vertices of $\nabla(Q^{II}, \delta_{Q^{II}})$ by considering that a lattice point is a vertex if and only if it lies on at least 3 facets. Writing m

as $(m(a), m(b), m(c))$ the 6 vertices are $(0,0,0),(0,3,0),(0,0,3),(1,3,0),(1,0,0),(4,0,3)$.

Now consider the quiver Q^{IIa} we obtain by placing a sink on the arrow a of Q^{II}. By the same argument as in the proof of Proposition 2.21 one sees that $\nabla(Q^{IIa}, \delta_{Q^{IIa}})$ is integral affinely equivalent to $\nabla(Q^{II}, \delta_{Q^{II}}) \cap \{x(a) \leq 2\}$. One checks easily that $\nabla(Q^{II}, \delta_{Q^{II}})$ has 4 lattice points satisfying $m(a) \geq 3$ (in particular these are $(3,2,0),(3,2,1),(3,3,0)$ and $(4,3,0)$), hence $\nabla(Q^{IIa}, \delta_{Q^{IIa}})$ has 26 lattice points. The equations defining the facets of $\nabla(Q^{II}, \delta_{Q^{II}})$ define facets of $\nabla(Q^{II}, \delta_{Q^{II}}) \cap \{x(a) \leq 2\}$ as well, and it has one more facet given by the equation $x(a) = 2$. Again considering that a vertex needs to lie on at least 3 facet we obtain that $\nabla(Q^{II}, \delta_{Q^{II}}) \cap \{x(a) \leq 2\}$ has 8 vertices:

$$(0,0,0),(0,3,0),(0,0,3),(1,3,0),(1,0,0),(2,0,1),(2,2,1),(2,0,3).$$

Note that the second task is relatively easy since adding a valency 2 sink is the same as removing the lattice points whose corresponding coordinate is strictly greater than 2 (cf. the argument in the proof of Proposition 2.21). Surprisingly the resulting polytopes turned out to be pairwise non-isomorphic.

In each case we recorded the following invariants: the number of valency 2 sinks of the quiver, the number of lattice points, the number of faces, the number of vertices and smooth vertices of the polytope. The results of these calculations are summarized in the tables of Appendix A. Luckily these invariants are sufficient to separate each but one of the cases. For the remaining one pair of polytopes (with the notation in the Appendix: Q^{IV} with valency 2 sinks on the arrows "a, b" in the first case and "e, f" in the second case) we showed that they are non-isomorphic by counting the lattice points on their facets. We concluded that there are 48 non-isomorphic reflexive polytopes in dimension 3 that can be obtained from prime quivers that are tight with their canonical weight. By Theorem 3.13 none of these 48 are products of lower dimensional polytopes. Moreover there are 5 more reflexive quiver polytopes that are obtained as products of lower dimensional quiver polytopes. Hence we obtained that there are a total of 53 reflexive quiver polytopes in dimension 3.

3.5 Characterizing smooth quiver moduli spaces

3.5.1 Forbidden descendants

The aim of this section is to give a characterization of the triples (Q, α, θ) which yield smooth or locally complete intersection moduli spaces $\mathcal{M}(Q, \alpha, \theta)$ via certain "forbidden descendants", which when specialized to the toric case will yield a theorem in the flavour of characterizing classes of graphs in terms of not containing certain forbidden minors. Note that by a complete intersection we always mean an ideal theoretic complete intersection.

We begin by showing that removing arrows from a quiver preserves the properties of being smooth or a complete intersection (special cases of the statement for the affine variety $\mathcal{M}(Q, \alpha, 0)$ were stated in [6] and [7] without a complete proof). Recall that an injective morphism of algebras $\iota : R \hookrightarrow S$ is called an *algebra retract* if there is a surjective morphism $\varphi : S \twoheadrightarrow R$ such that $\varphi \circ \iota = id_R$. The morphism φ is called the *retraction map*. When R and S are both graded and the morphisms ι, φ are graded morphisms we call $\iota : R \hookrightarrow S$ a *graded algebra retract.*.

We obtain Proposition 3.40 by applying the results from [19] to the graded case. Note that by complete intersection we always mean an ideal theoretic complete intersection. Recall that for the graded algebra S the points of the scheme $\operatorname{Proj} S$ are the homogeneous prime ideals in S that do not contain the irrelevant ideal S_+, and the stalk at the point $p \in \operatorname{Proj} S$ is the homogeneous localization $S_{(p)}$ defined as the subring of degree zero elements in the localized ring $T^{-1}S$, where T consists of the homogeneous elements that are not in p.

Proposition 3.40 *Let $\iota : R \hookrightarrow S$ be a graded algebra retract. Then if the variety* $\operatorname{Proj} S$ *is smooth (resp. locally a complete intersection) then* $\operatorname{Proj} R$ *is also smooth (resp. locally a complete intersection).*

Proof. Let φ denote the retraction morphism $S \twoheadrightarrow R$. We recall from Proposition 2.10 of [19] that for any prime ideal $p \in \operatorname{Spec}(R)$ and $q = \varphi^{-1}(p) \in \operatorname{Spec}(S)$ we have a natural algebra retract of the localized rings $R_p \hookrightarrow S_q$. Since ι, φ preserve the grading if p is a homogeneous prime ideal of R then $q = \varphi^{-1}(p)$ is a homogeneous prime ideal of S and there is a natural algebra retract of the homogeneous localizations $R_{(p)} \hookrightarrow S_{(q)}$, moreover if p does not contain the irrelevant ideal of R then q does not contain the irrelevant ideal of S. Now the proposition follows from Theorem 3.2 of [19] which asserts that every algebra retract of a regular (resp. locally complete intersection) ring is also a regular (resp. locally complete intersection) ring. □

Proposition 3.41 *For any quiver Q' obtained from Q by removing an arrow $a \in Q_1$, we have*

(i) If $\mathcal{M}(Q,\alpha,\theta)$ is smooth, then $\mathcal{M}(Q',\alpha,\theta)$ is also smooth (or empty).

(ii) If $\mathcal{M}(Q,\alpha,\theta)$ is locally a complete intersection, then $\mathcal{M}(Q',\alpha,\theta)$ is also locally a complete intersection (or empty).

(iii) If the affine algebraic variety $\mathcal{M}(Q,\alpha,0)$ is a complete intersection, then $\mathcal{M}(Q',\alpha,0)$ is also a complete intersection.

Proof. View $\mathrm{Rep}(Q',\alpha)$ as a direct summand of $\mathrm{Rep}(Q,\alpha)$ in the obvious way, and denote by ι the algebra retract $\mathcal{O}(\mathrm{Rep}(Q',\alpha)) \hookrightarrow \mathcal{O}(\mathrm{Rep}(Q,\alpha))$ induced by the projection $\mathrm{Rep}(Q,\alpha) \to \mathrm{Rep}(Q',\alpha)$, and by φ the corresponding retraction map induced by the embedding $\mathrm{Rep}(Q',\alpha) \hookrightarrow \mathrm{Rep}(Q,\alpha)$. Since both ι and φ are $GL(\alpha)$ equivariant for any weight θ it follows that that $\iota(\mathcal{O}(\mathrm{Rep}(Q',\alpha))_{n\theta}) \subseteq \mathcal{O}(\mathrm{Rep}(Q,\alpha))_{n\theta}$ and $\varphi(\mathcal{O}(\mathrm{Rep}(Q,\alpha))_{n\theta}) \subseteq \mathcal{O}(\mathrm{Rep}(Q',\alpha))_{n\theta}$, and hence we have the graded algebra retract

$$\bigoplus_{n=0}^{\infty} \mathcal{O}(\mathrm{Rep}(Q',\alpha))_{n\theta} \hookrightarrow \bigoplus_{n=0}^{\infty} \mathcal{O}(\mathrm{Rep}(Q,\alpha))_{n\theta}.$$

Now (i) and (ii) follow from Proposition 3.40, and (iii) follows from Theorem 3.2 of [19] and that as a special case of the above the ring $\mathcal{O}(\mathrm{Rep}(Q',\alpha))_0 = \mathcal{O}(\mathrm{Rep}(Q',\alpha))^{GL(\alpha)}$ is a retraction of the ring $\mathcal{O}(\mathrm{Rep}(Q,\alpha))_0 = \mathcal{O}(\mathrm{Rep}(Q,\alpha))^{GL(\alpha)}$.

□

Taking into account the results we recalled in Section 2.3 we have obtained three tools to reduce the structure of a quiver while preserving the properties of being smooth or a complete intersection: constructing local quiver settings, applying the reduction steps RI-III and taking subquivers (note that vertices without arrows can be removed). We will say that the quiver setting (Q',α') is a *descendant* of the quiver setting (Q,α) if (Q',α') can be obtained from (Q,α) by repeteadly taking subquivers, applying RI-III or constructing local quivers with the 0 weight. We will also call (Q',α') a *descendant* of the triple (Q,α,θ) if it is a descendant of a local quiver setting (Q_ξ,α_ξ) for some $\xi \in \mathcal{M}(Q,\alpha,\theta)$. Clearly being a descendant is a partial ordering on the set of quiver settings, hence local properties that are preserved by all of the reduction methods can be characterized by some list of forbidden descendants (for example take the list that contains every quiver setting without the property). The surprising fact however is that in the case of the property of being smooth taking a single forbidden descendant is enough. We note that certain elements of the proof of Lemma 3.42 could be recovered from the proof of the main theorem in [7].

Lemma 3.42 *Let Q be a strongly connected quiver on at least two vertices and $\alpha \neq (1,\dots,1)$ a genuine dimension vector. Let $w \in Q_0$ be a vertex on which the value of α is maximal. If none of the reduction steps RI-III can be applied to (Q,α) then there exists a simple representation of Q with dimension vector $\alpha - \epsilon_w$, moreover we have $\langle \alpha - \epsilon_w, \epsilon_w \rangle_Q \leq -2$ and $\langle \epsilon_w, \alpha - \epsilon_w \rangle_Q \leq -2$.*

Proof. By the assumption that RII can not be applied to (Q,α), we have that there are no loops on vertices of dimension 1. Moreover the assumption that RI can not be applied to (Q,α) implies that for a vertex v without loops we have $\langle \alpha, \epsilon_v \rangle_Q < 0$ and $\langle \epsilon_v, \alpha \rangle_Q < 0$, and for vertices with loops the same inequalities follow from the assumption that Q is strongly connected and $|Q_0| \geq 2$. In particular Q is not a cycle, otherwise we would have $\langle \alpha, \epsilon_w \rangle_Q \geq 0$ by the maximality of $\alpha(w)$. Thus it follows from Theorem 2.6 that a simple representation of Q with dimension vector $\alpha - \epsilon_w$ exists, if the inequalities $\langle \alpha - \epsilon_w, \epsilon_v \rangle_Q \leq 0$ and $\langle \epsilon_v, \alpha - \epsilon_w \rangle_Q \leq 0$ hold for all $v \in Q_0$

By the maximality of $\alpha(w)$ we have

$$\langle \alpha, \epsilon_v \rangle_Q \leq \alpha(v) + \alpha(w)\langle \epsilon_w, \epsilon_v \rangle_Q \leq \alpha(w)(1 + \langle \epsilon_w, \epsilon_v \rangle_Q).$$

Now the inequality

$$\langle \alpha - \epsilon_w, \epsilon_v \rangle_Q = \langle \alpha, \epsilon_v \rangle_Q - \langle \epsilon_w, \epsilon_v \rangle_Q \leq 0$$

follows from $\langle \alpha, \epsilon_v \rangle_Q < 0$ when $\langle \epsilon_w, \epsilon_v \rangle_Q \geq -1$ and from $\langle \alpha, \epsilon_v \rangle_Q \leq \alpha(w)(1 + \langle \epsilon_w, \epsilon_v \rangle_Q)$ when $\langle \epsilon_w, \epsilon_v \rangle_Q < -1$. The inequality $\langle \epsilon_v, \alpha - \epsilon_w \rangle_Q \leq 0$ can be derived similarly.

Moreover if $\langle \epsilon_w, \epsilon_w \rangle_Q = 1$ then $\langle \alpha - \epsilon_w, \epsilon_w \rangle_Q = \langle \alpha, \epsilon_w \rangle_Q - \langle \epsilon_w, \epsilon_w \rangle_Q \leq -2$. Assume that $\langle \epsilon_w, \epsilon_w \rangle_Q \leq 0$. We have $\langle \alpha - \epsilon_w, \epsilon_w \rangle_Q = \sum_{v \neq w} \alpha(v)\langle \epsilon_v, \epsilon_w \rangle_Q + (\alpha(w) - 1)\langle \epsilon_w, \epsilon_w \rangle_Q$. For $v \neq w$ we have $\langle \epsilon_v, \epsilon_w \rangle_Q \leq 0$ and since Q is strongly connected there is strict inequality for at least one v. Hence $\langle \alpha - \epsilon_w, \epsilon_w \rangle_Q \geq -1$ would imply that $\sum_{v \neq w} \alpha(v)\langle \epsilon_v, \epsilon_w \rangle_Q = -1$ and $\langle \epsilon_w, \epsilon_w \rangle_Q = 0$. It follows that w has a loop and the single arrow pointing to w departs from a vertex of dimension 1. Hence we are in the situation of Lemma 2.10 and RIII can be applied contradicting our assumptions. We obtained that $\langle \alpha - \epsilon_w, \epsilon_w \rangle_Q \leq -2$, and $\langle \epsilon_w, \alpha - \epsilon_w \rangle_Q \leq -2$ follows similarly.

$\square$

Theorem 3.43 *$\mathcal{M}(Q,\alpha,\theta)$ is smooth if and only if $(\hat{Q},\hat{\alpha})$ is not a descendant of (Q,α,θ), where $\hat{Q}$ is the quiver*

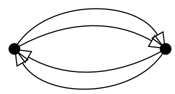

and $\hat{\alpha}$ takes value 1 on both vertices.

Proof. We will say that the quiver setting (Q,α) is smooth (resp. singular) if the moduli space $\mathcal{M}(Q,\alpha,0)$ is smooth (resp. singular), moreover that the quiver setting (Q',α') is a local quiver setting of (Q,α) if it can be obtained as a local quiver setting of $(Q,\alpha,0)$. By a non-trivial descendant of (Q,α) we simply mean one that is not (Q,α) itself. If (Q,α,θ) has a descendant that is singular then by the discussion preceeding the theorem $\mathcal{M}(Q,\alpha,\theta)$ is also singular. Now recall from Remark 2.7 that whenever $\xi\in\mathcal{M}(Q,\alpha,\theta)$ is a singular point, the local quiver setting (Q_ξ,α_ξ) is singular at 0, hence it suffices to prove that $(\hat{Q},\hat{\alpha})$ is a descendant of every singular quiver setting. Next note that for a non trivial local-quiver setting (i.e. not the one at the 0 representation) (Q',α') of (Q,α) we have $\sum_{v\in Q'_0}\alpha'(v)<\sum_{v\in Q_0}\alpha(v)$ and the same holds when (Q',α') is obtained from (Q,α) via the reduction step RI. Furthermore applying the reduction step RII or RIII both decrease the number of loops in the quiver without increasing $\sum_{v\in Q_0}\alpha(v)$. Finally taking non-trivial subquivers reduces the number of arrows or vertices of the quiver. It follows that one can take non-trivial descendants only finitely many times before arriving to a one vertex quiver with no arrows. Hence we are left to prove that every singular quiver setting (Q,α) that is not $(\hat{Q},\hat{\alpha})$ has a non-trivial singular descendant.

It follows from Lemma 2.4 of [7] that if Q is not strongly connected then

$$\mathcal{M}(Q,\alpha,0)\cong\prod_{i=1}^{k}\mathcal{M}(Q^i,\alpha^i,0),$$

where $Q^{1\dots k}$ are the strongly connected components of Q and the α^i are restrictions of α. Hence if Q is not strongly connected and (Q,α) is singular, then (Q,α) has a non-trivial strongly connected descendant. We assume for the rest of the proof that Q is strongly connected.

First we treat the case $\alpha = (1, \ldots, 1)$. Record that the reduction step RI can be applied to a vertex without loops if and only if its in-degree or out-degree is 1, and the reduction step RIII can never be applied in this case. So if Q contains a loop or a vertex with in or out-degree 1 we can apply RI or RII to obtain a non-trivial singular descendant and we are done. Let us assume then that there are no loops in Q and every vertex has in- and out-degrees at least 2, which implies that Q has at least 2 vertices. Note that it follows from Theorem 2.6 (or alternatively from Proposition 2.17) that there is a simple representation with dimension vector $(1, \ldots, 1)$ if and only if Q is strongly connected.

Let us now assume that Q contains a cycle C such that C does not run through every vertex of the quiver, and write $W = \{w_1, \ldots, w_k\}$ for the set of vertices that do not belong to C. Let α_C denote the dimension vector with $\alpha(w) = 0$ for $w \in W$ and $\alpha(v) = 1$ for $v \in Q_1 \setminus W$. Let R_C be the representation of Q with dimension vector α_C that takes values 1 on the arrows of C and 0 on the rest of the arrows. For each $w_i \in W$ set R_{w_i} to be the representation with dimension vector ϵ_{w_i} that is 0 on every arrow. Now let R denote the semisimple representation $R_C \oplus (\bigoplus_{i=1}^k R_{w_i})$, and ξ the image of R in $\mathcal{M}(Q, \alpha, 0)$. We will show that the quiver setting (Q_ξ, α_ξ) is singular. To see this first observe that the vertices of Q_ξ corresponding to R_{w_i} have the same in and out-degrees as w_i in Q and are not incident to loops, hence there is at most one vertex in Q_ξ on which RI and RII can be applied.

Now we claim that any strongly-connected quiver setting on at least 2 vertices with $\alpha = (1, \ldots, 1)$ that has only one vertex on which RI or RII can be applied is singular. Indeed if Q has exactly two vertices then if one has in and out-degrees of at least 2 then so does the other and the claim follows from Theorem 2.11. For a quiver on 3 or more vertices, that satisfies the condition in the claim, it is not difficult to see that if we apply RI or RII the resulting quiver setting still has at most one vertex on which the reduction steps can be applied, and hence the claim follows from induction.

We have proven that when Q has a cycle that is not incident to every vertex then $(Q, (1, \ldots, 1))$ has a non-trivial singular descendant and are left to deal with the case when every cycle in Q runs through every vertex. It follows that Q is a quiver which we can obtain from a single cycle by adding some (at least one) copies of each arrow. Now repeteadly remove arrows from Q and apply RI whenever possible to see that it has a $(\hat{Q}, \hat{\alpha})$ descendant.

Now we turn to the case when $\alpha \neq (1, \ldots, 1)$. Again we can assume that none of RI-III can be applied to the quiver. If Q has only one vertex v and k loops then by

Theorem 2.11 either $\alpha(v) = 2$ and $k \geq 3$ or $\alpha(v) \geq 3$ and $k \geq 2$. By Theorem 2.6 we can choose non-isomorphic simple representations $R_1, \ldots, R_{\alpha(v)}$ such that the dimension vector of any of them is given by $\alpha(v) = 1$. Now let ξ be the image of the semisimple representation $\bigoplus_{i=1}^{\alpha(v)} R_i$ in in $\mathcal{M}(Q, \alpha, 0)$. The local quiver setting (Q_ξ, α_ξ) has dimension vector $\alpha_\xi = (1, \ldots, 1)$. Moreover Q_ξ has $\alpha(v)$ vertices with $k-1$ arrows going both ways between any two distinct vertices. After removing loops (by RII) we see that each vertex has in and out-degrees at least 2 hence by Theorem 2.11 (Q_ξ, α_ξ) is singular.

If Q has more than one vertex then we are in the situation of Lemma 3.42 and see that for a vertex w on which α takes maximum value there is a simple representation R with dimension vector $\alpha - \epsilon_w$. Set R_w to be a simple representation with dimension vector ϵ_w, and ξ to be the image of the semisimple representation $R \oplus R_w$ in $\mathcal{M}(Q, \alpha, 0)$. Now (Q_ξ, α_ξ) has two vertices with dimension 1 and there is $-\langle \alpha - \epsilon_w, \epsilon_w \rangle_Q$ arrows running in one direction and $-\langle \epsilon_w, \alpha - \epsilon_w \rangle_Q$ in the other, both of which are at least 2 by Lemma 3.42. Hence (Q_ξ, α_ξ) is singular by Theorem 2.11. We have shown that (Q, α) has a non-trivial singular descendant completing the proof. $\square$

In Theorem 3.44 we use forbidden descendants to characterize triples (Q, α, θ) that yield a complete intersection moduli space in the toric case $\alpha = (1, \ldots, 1)$. For the affine case the statement was proven by the author in [29], and the general case follows similarly to Theorem 3.43 by considering a local quiver setting at a point where the moduli space is not locally a complete intersection. The key idea for the proof in [29] uses a technique which was developed in the author's masters thesis (without the introduction of forbidden descendants).

Theorem 3.44 *For $\alpha = (1, \ldots, 1)$, $\mathcal{M}(Q, \alpha, \theta)$ is locally a complete intersection if and only if none of $(\hat{Q}_1, \hat{\alpha}_1)$ and $(\hat{Q}_2, \hat{\alpha}_2)$ is a descendant of (Q, α, θ), where $\hat{Q}_1$ and $\hat{Q}_2$ are the quivers:*

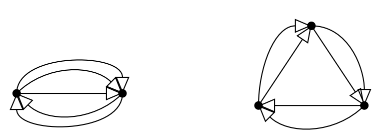

and the dimension vectors $\hat{\alpha}_1$ and $\hat{\alpha}_2$ take value 1 on every vertex.

3.5.2 Generic weights

A weight θ of a quiver setting (Q, α) is called *generic* if the θ-stable and θ-semistable α-dimensional representations coincide. It is well known that in this case the moduli space

$\mathcal{M}(Q, \alpha, \theta)$ is smooth. We will show that in the toric case for tight pairs (Q, θ) the reverse implication also holds.

Lemma 3.45 *If Q is a connected quiver and the pair (Q, θ) is tight then for any $\emptyset \neq V \subset Q_0$ with $\theta(V) = 0$ we have that there are at least two arrows from V to $Q_0 \setminus V$ and at least two arrows from $Q_0 \setminus V$ to V.*

Proof. Set $\text{Out}(V) = \{a \in Q_1 \mid a^- \in V, a^+ \in Q_0 \setminus V\}$ and $\text{In}(V) = \{a \in Q_1 \mid a^+ \in V, a^- \in Q_0 \setminus V\}$. Without loss of generality we may assume that $|\text{Out}(V)| \geq |\text{In}(V)|$ and then since Q is connected $|\text{Out}(V)| \geq 1$. If $\text{In}(V) = \emptyset$ then for any $b \in \text{Out}(V)$ and $x \in \nabla(Q, \theta)$ we have $x(b) = 0$, hence b is removable. If $\text{In}(V) = \{a\}$ and $\text{Out}(V) = \{b_1, \dots, b_k\}$ then for any $x \in \nabla(Q, \theta)$ we have $x(a) = \sum_{i=1}^k x(b_i)$. It follows that the face $\nabla(Q, \theta)_{x(a)=0}$ is contained in $\bigcap_{i=1}^k \nabla(Q, \theta)_{x(b_i)=0}$, hence by (i) of Corollary 3.6 (Q, θ) can not be tight. □

Proposition 3.46 *Let Q be a connected quiver. If (Q, θ) is tight the following are equivalent:*

(i) θ is a generic weight.

(ii) $\mathcal{M}(Q, \theta)$ is smooth.

(iii) For any $m \in \nabla(Q, \theta)$ the quiver with vertices Q_0 and arrows $\text{supp}(m)$ *is connected.*

Moreover if the tightness of (Q, θ) is not assumed (i) and (iii) are still equivalent and they both imply (ii).

Proof. We first show the implications that hold without the assumption that (Q, θ) is tight. If there is $m \in \nabla(Q, \theta)$ such that the quiver with vertices Q_0 and arrows $\text{supp}(m)$ is not connected, then pick a representation R with $\text{supp}(R) = \text{supp}(m)$. R is θ-semistable and so are its non-trivial subrepresentations that correspond to the connected components of $\text{supp}(m)$, hence R can not be θ-stable, proving that (i) implies (iii). If θ is not generic then there is a representation R that is θ-semistable but not stable. By θ-semistability of R there is a lattice point $m \in \nabla(Q, \theta)$ with $\text{supp}(m) \subseteq \text{supp}(R)$. Since R is not θ-stable by Proposition 2.17 there is a $\text{supp}(R)$-successor closed set of vertices $\emptyset \neq V \subset Q_0$ that satisfies $\theta(V) = 0$. Then V is also $\text{supp}(m)$-successor closed, hence by $\theta(V) = 0$ the value of m has to be 0 on every arrow running between $Q_1 \setminus V$ and V, and it follows that V is a connected component of the quiver with vertices Q_0 and arrows $\text{supp}(m)$, proving that (iii) implies (i). Finally from (iii) it follows that the quivers Q^v in Proposition 3.26 have only 1 vertex and hence the open sets U_v in the affine open cover of $\mathcal{M}(Q,\theta)$ are affine spaces, proving that (iii) implies (ii).

Let us turn to the case when (Q,θ) is tight and show that (ii) implies (iii). Assume for a contradiction that (iii) does not hold. Since the supports of the vertices are minimal sets amongst the supports of the lattice points of $m \in \nabla(Q,\theta)$ we can assume that there is a vertex $v \in \nabla(Q,\theta)$ such that the quiver with vertices Q_0 and arrows $\mathrm{supp}(v)$ have connected components $V_1,\dots,V_k$ for $k > 0$. Then we have $\theta(V_i) = 0$ and the vertices of Q^v are obtained by contracting each of the V_i into a single vertex. Since $\mathcal{M}(Q,\theta)$ is smooth $U_v \cong \mathcal{M}(Q^v,0)$ is also smooth. Since Q^v has at least 2 vertices by Theorem 2.11 we see that the reduction step RI can be applied to it after possible loops have been removed. Hence there is a vertex in Q^v, corresponding to some connected component V_i, such that its in or out-degree is at most 1 (not counting loops). Note that the number of arrows leaving (resp. arriving to) the vertex corresponding to V_i in Q^v are the same as the number of arrows running from V_i to $Q_0 \setminus V_i$ (resp. from $Q_0 \setminus V_i$ to V_i) in Q. Hence by Lemma 3.45 we see that (Q,θ) can not be tight.

$\square$

Chapter 4

Toric Ideals of Quivers

4.1 Presentations of semigroup algebras

The aim of this chapter is to study the toric ideals of affine and projective toric quiver varieties. Recall from Section 2.1 that both cases come down to studying relations amongst the generators of certain semigroup algebras, that are obtained from subsemigroups of $\mathbb{Z}^{Q_1}$. We begin by formulating a statement (Lemma 4.2; a version of it was introduced in [29]) in a slightly more general situation than what is needed here. Let S be any finitely generated commutative monoid (written additively) with non-zero generators $s_1, \dots, s_d$, and denote by $\mathbb{Z}[S]$ the corresponding semigroup algebra over $\mathbb{Z}$: its elements are formal integral linear combinations of the symbols $\{x^s \mid s \in S\}$, with multiplication given by $x^s \cdot x^{s'} = x^{s+s'}$. Write $R := \mathbb{Z}[t_1, \dots, t_d]$ for the d-variable polynomial ring over the integers, and $\phi : R \to \mathbb{Z}[S]$ the ring surjection $t_i \mapsto x^{s_i}$. Set $I := \ker(\phi)$. It is well known and easy to see that

$$I = \ker(\phi) = \mathrm{Span}_{\mathbb{Z}}\{t^a - t^b \mid \sum_{i=1}^{d} a_i s_i = \sum_{j=1}^{d} b_j s_j \in S\} \tag{4.1}$$

where for $a = (a_1, \dots, a_d) \in \mathbb{N}_0^d$ we write $t^a = t_1^{a_1} \dots t_d^{a_d}$.

Introduce a binary relation on the set of monomials in R: we write $t^a \sim t^b$ if $t^a - t^b \in R_+ I$, where R_+ is the ideal in R consisting of the polynomials with zero constant term. Obviously $\sim$ is an equivalence relation. Let Λ be a complete set of representatives of the equivalence classes. We have $\Lambda = \coprod_{s \in S} \Lambda_s$, where for $s \in S$ set $\Lambda_s := \{t^a \in \Lambda \mid \sum a_i s_i = s\}$.

For the $s \in S$ with $|\Lambda_s| > 1$, set $\mathcal{G}_s := \{t^{a_1} - t^{a_i} \mid i = 2, \dots, p\}$, where $t^{a_1}, \dots, t^{a_p}$ is an arbitrarily chosen ordering of the elements of Λ_s.

Lemma 4.1 *(i) $\coprod_{s \in S: |\Lambda_s|>1} \mathcal{G}_s$ is a minimal generating system of the ideal I, in particular, I is generated by $\sum_{s \in S}(|\Lambda_s| - 1)$ elements.*
(ii) Suppose that $S = \coprod_{k=0}^{\infty} S_k$ is graded (i.e. $S_k + S_l \subseteq S_{k+l}$) and $S_0 = \{0\}$ (i.e. the generators $s_1, \dots, s_d$ have positive degree). Then $\coprod_{s \in S: |\Lambda_s|>1} \mathcal{G}_s$ is a minimal homogeneous generating system of the ideal I, where the grading on $\mathbb{Z}[t_1, \dots, t_d]$ is defined by setting the degree of t_i to be equal to the degree of s_i.

Proof. It is easy to see that a $\mathbb{Z}$-module direct complement of $R_+ I$ in R is $\sum_{t^a \in \Lambda} \mathbb{Z} t^a$. Thus the statement follows by the graded Nakayama Lemma. $\square$

Next for a cancellative commutative monoid S we give a more explicit description of the relation $\sim$ (a special case occurs in [29]). For some elements $s, v \in S$ we say that s *divides* v and write $s \mid v$ if there exists an element $w \in S$ with $v = s + w$. For any $s \in S$ introduce a binary relation $\sim_s$ on the subset of $\{s_1, \dots, s_d\}$ consisting of the generators s_i with $s_i \mid s$ as follows:

$$\begin{aligned} s_i \sim_s s_j \text{ if } i = j \text{ or there exist } u_1, \dots, u_k \in \{s_1, \dots, s_d\} \\ \text{with } u_1 = s_i,\ u_k = s_j,\ u_l + u_{l+1} \mid s \text{ for } l = 1, \dots, k-1. \end{aligned} \tag{4.2}$$

Obviously $\sim_s$ is an equivalence relation, and

$$s = s_{i_1} + \dots + s_{i_r} \quad \text{implies} \quad s_{i_1} \sim_s s_{i_2} \sim_s \dots \sim_s s_{i_r}. \tag{4.3}$$

Moreover, $s_i \sim_s s_j$ implies $s_i \sim_t s_j$ for any $s \mid t \in S$.

Lemma 4.2 *Let S be a cancellative commutative monoid generated by $s_1, \dots, s_d$. Take $t^a - t^b \in I$, so $s := \sum_{i=1}^d a_i s_i = \sum_{j=1}^d b_j s_j \in S$. Then the following are equivalent:*

(i) $t^a - t^b \in R_+ I$;

(ii) For some $t_i \mid t^a$ and $t_j \mid t^b$ we have $s_i \sim_s s_j$;

(iii) For all $t_i \mid t^a$ and $t_j \mid t^b$ we have $s_i \sim_s s_j$.

Proof. (ii) and (iii) are equivalent by (4.3).

To show that (ii) implies (i) assume that for some $t_i \mid t^a$ and $t_j \mid t^b$ we have $s_i \sim_s s_j$. If $s_i = s_j$, then t^a and t^b have a common variable, say t_1, so $t^a = t_1 t^{a'}$ and $t^b = t_1 t^{b'}$ for some $a', b' \in \mathbb{N}_0^d$. We have

$$x^{s_1}\phi(t^{a'} - t^{b'}) = \phi(t_1(t^{a'} - t^{b'})) = \phi(t^a - t^b) = 0$$

hence $x^{s_1}\phi(t^{a'}) = x^{s_1}\phi(t^{b'})$. Since S is cancellative, we conclude $\phi(t^{a'}) = \phi(t^{b'})$, thus $t^{a'} - t^{b'} \in I$, implying in turn that $t^a - t^b = t_1(t^{a'} - t^{b'}) \in R_+I$. If $s_i \neq s_j$, then there exist $z_1, \ldots, z_k \in \{t_1, \ldots, t_d\}$ such that $u_l \in S$ with $\phi(z_l) = x^{u_l}$ satisfy (4.2). Then there exist monomials (possibly empty) $w_0, \ldots, w_k$ in the variables $t_1, \ldots, t_d$ such that

$$z_1 w_0 = t^a, \quad \phi(z_l z_{l+1} w_l) = x^s \quad (l = 1, \ldots, k-1), \quad z_k w_k = t^b.$$

It follows that

$$t^a - t^b = z_1(w_0 - z_2 w_1) + \sum_{l=2}^{k-1} z_l(z_{l-1}w_{l-1} - z_{l+1}w_l) + z_k(z_{k-1}w_{k-1} - w_k). \tag{4.4}$$

Note that $\phi(z_1 w_0) = x^s = \phi(z_1 z_2 w_1)$, hence $\phi(z_1)\phi(w_0) = \phi(z_1)\phi(z_2 w_1)$. Since S is cancellative, we conclude that $\phi(w_0) = \phi(z_2 w_1)$, so $w_0 - z_2 w_1 \in I$, implying in turn that $z_1(w_0 - z_2 w_1) \in R_+I$. Similarly all the other summands on the right hand side of (4.4) belong to R_+I, hence $t^a - t^b \in R_+I$.

Finally we show that (i) implies (ii). Suppose that $t^a - t^b \in R_+I$. By (4.1) we have

$$t^a - t^b = \sum_{l=1}^{k} t_{i_l}(t^{a_l} - t^{b_l}) \;\text{ where }\; t^{a_l} - t^{b_l} \in I \text{ and } i_l \in \{1, \ldots, d\} \tag{4.5}$$

After a possible renumbering and cancellations we may assume that

$$t_{i_1} t^{a_1} = t^a,\; t_{i_l} t^{b_l} = t_{i_{l+1}} t^{a_{l+1}} \;\text{ for }\; l = 1, \ldots, k-1, \;\text{ and }\; t_{i_k} t^{b_k} = t^b. \tag{4.6}$$

Observe that if $t_{i_l} = t_{i_{l+1}}$ for some $l \in \{1, \ldots, k-1\}$, then necessarily $t^{b_l} = t^{a_{l+1}}$, hence $t_{i_l}(t^{a_l} - t^{b_l}) + t_{i_{l+1}}(t^{a_{l+1}} - t^{b_{l+1}}) = t_{i_l}(t^{a_l} - t^{b_{l+1}})$. Thus in (4.5) we may replace the sum of the lth and $(l+1)$st terms by a single summand $t_{i_l}(t^{a_l} - t^{b_{l+1}})$. In other words, we may achieve that in (4.5) we have $t_{i_l} \neq t_{i_{l+1}}$ for each $l = 1, \ldots, k-1$, in addition to (4.6). If $k = 1$, then t^a and t^b have a common variable and (ii) obviously holds. From now on assume that $k \geq 2$. From $t_{i_l} t^{b_l} = t_{i_{l+1}} t^{a_{l+1}}$ and the fact that t_{i_l} and $t_{i_{l+1}}$ are different variables in $\mathbb{Z}[t_1, \ldots, t_d]$ we deduce that $t^{b_l} = t_{i_{l+1}} t^{c_l}$ for some $c_l \in \mathbb{N}_0^d$, implying that $x^s = \phi(t_{i_l} t^{b_l}) = \phi(t_{i_l} t_{i_{l+1}} t^{c_l}) = \phi(t_{i_l})\phi(t_{i_{l+1}})\phi(t^{c_l})$. Thus $u_l := s_{i_l}$ satisfy (4.2) and hence $s_{i_1} \sim_s s_{i_k}$. □

Corollary 4.3 *Suppose that $S = \coprod_{k=0}^{\infty} S_k$ is a finitely generated graded cancellative commutative monoid generated by $S_1 = \{s_1, \ldots, s_d\}$. The kernel of $\phi : \mathbb{Z}[t_1, \ldots, t_d] \to \mathbb{Z}[S]$, $t_i \mapsto x^{s_i}$ is generated by homogeneous elements of degree at most r (with respect to the standard grading on $\mathbb{Z}[t_1, \ldots, t_d]$) if and only if for all $k > r$ and $s \in S_k$, the elements in S_1 that divide s in the monoid S form a single equivalence class with respect to $\sim_s$.*

Proof. This is an immediate consequence of Lemma 4.1 and Lemma 4.2. □

Remark 4.4 For polytopal semigroup algebras a different proof of Corollary 4.3 can be derived from Theorem 12.12 and Corollary 12.13 in [44].

4.2 Toric ideals in the affine case

Recall from Section 2.4 that when we set $\theta = 0$ the toric quiver variety $\mathcal{M}(Q, 0)$ is an affine variety with coordinate ring $\mathbb{C}[\nabla(Q, 0) \cap \mathbb{Z}^{Q_1}]$. We will denote the semigroup $\nabla(Q, 0) \cap \mathbb{Z}^{Q_1}$ by $S(Q)$. As we pointed out in Section 2.4, denoting the (oriented) primitive cycles of Q by $C_1, \ldots, C_r$ the semigroup $S(Q)$ is generated by the characteristic vectors $\varepsilon_{C_1}, \ldots, \varepsilon_{C_r}$. Throughout this section φ will denote the map $\mathbb{C}[t_1, \ldots, t_r] \to \mathbb{C}[S(Q)]$ defined by $\varphi(t_i) = t^{\varepsilon_{C_i}}$, and we will denote the corresponding toric ideal $\ker(\varphi)$ by $\mathcal{I}_0(Q)$.

We will denote by $\geq$ the usual partial ordering on $\mathbb{R}^{Q_1}$, i.e. for $x_1, x_2 \in \mathbb{R}^{Q_1}$ we define $x_1 \geq x_2$ if and only $x_1(a) \geq x_2(a)$ for all $a \in Q_1$. Note that in the monoid $S(Q)$ we have that $m \mid n$ for some m, n if and only if $m \leq n$. For a non-zero element $s \in \nabla(Q, 0) \cap \mathbb{Z}^{Q_1}$ that is not a primitive cycle and a primitive cycle C with $\varepsilon_C \leq s$ we will say that C is *s-strong* if $s - \varepsilon_C$ is the characteristic vector of a primitive cycle and that C is *s-weak* otherwise. We will denote by $\sim_s$ the equivalence relation from Lemma 4.2.

Lemma 4.5 *For $s \in S(Q)$ and s-weak primitive cycles $C_1, C_2 \leq s$, we have $\varepsilon_{C_1} \sim_s \varepsilon_{C_2}$.*

Proof. We prove by induction on $|s| = \sum_{a \in Q_1} s(a)$. The statement is trivial when $|\mathrm{supp}(s)| = 1$, so we can assume that $|\mathrm{supp}(s)| \geq 2$. If $\mathrm{supp}(s)$ contains a loop L then the characteristic vector ε_L is $\sim_s$ equivalent to every $\varepsilon_C \leq s$, and the statement of the Lemma is clear. We assume for the rest of the proof that $\mathrm{supp}(s)$ does not contain a loop. For distinct arrows $a_1, a_2 \in \mathrm{supp}(s)$ with $a_1^+ = a_2^-$, we will denote by $Q^{(a_1,a_2)}$ the quiver we get from Q by adding a new arrow a_{12} with $a_{12}^- = a_1^-$ and $a_{12}^+ = a_2^+$, and identify the remaining arrow set of $Q(a_1, a_2)$ with that of Q. We will denote by $s_{(a_1,a_2)}$ the element of $S(Q^{(a_1,a_2)})$ defined by $s_{(a_1,a_2)} = s - \varepsilon_{a_1} - \varepsilon_{a_2} + \varepsilon_{a_{12}}$ (where ε_a denotes the characteristic vector of a). Clearly $|s_{(a_1,a_2)}| = |s| - 1$ so we can apply the induction hypothesis on $s_{(a_1,a_2)}$. Let μ denote the map $\{m \in S(Q^{(a_1,a_2)}) \mid m \leq s_{(a_1,a_2)}\} \to \{m \in S(Q) \mid m \leq s\}$ that is defined by setting $\mu(m) = m$ when $m(a_{12}) = 0$ and $\mu(m) = m + \varepsilon_{a_1} + \varepsilon_{a_2} - \varepsilon_{a_{12}}$ when $m(a_{12}) = 1$. Clearly $\mu(s_{(a_1,a_2)}) = s$ and $m_1 + m_2 \leq s_{(a_1,a_2)}$ implies $\mu(m_1) + \mu(m_2) \leq s$. Moreover for an element $m \in S(Q)$, $m \leq s$ we have $m \in \mathrm{Im}(\mu)$ unless $m(a_1) = s(a_1)$ and $m(a_2) = 0$ or $m(a_1) = 0$ and $m(a_2) = s(a_2)$. The preimages of characteristic vectors of primitive cycles in $\mathrm{Im}(\mu)$ are also characteristic vectors of primitive cycles.

Now we claim that if for primitive cycles C_1, C_2 of $Q^{(a_1,a_2)}$ we have $\varepsilon_{C_1} \sim_{s_{(a_1,a_2)}} \varepsilon_{C_2}$ and $\mu(\varepsilon_{C_1}), \mu(\varepsilon_{C_2})$ are characteristic vectors of primitive cycles then $\mu(\varepsilon_{C_1}) \sim_s \mu(\varepsilon_{C_2})$. Indeed $\varepsilon_{C_1} \sim_{s_{(a_1,a_2)}} \varepsilon_{C_2}$ implies that there is series of primitive cycles $C'_1, \dots, C'_k$ such that $C_1 = C'_1$, $C_2 = C'_k$ and $\varepsilon_{C'_i} + \varepsilon_{C'_{i+1}} \leq s_{(a_1,a_2)}$ for all $i = 1, \dots, k-1$. Now by the above we have $\mu(\varepsilon_{C'_i}) + \mu(\varepsilon_{C'_{i+1}}) \leq s$ for all $i = 1, \dots, k-1$. Now we consider the sequence $\mu(\varepsilon_{C'_1}), \dots \mu(\varepsilon_{C'_1})$ and replace each $\mu(\varepsilon_{C'_i})$ by a characteristic vector of a primitive cycle contained in its support to obtain a sequence as in Lemma 4.2 for $\mu(\varepsilon_{C_1})$ and $\mu(\varepsilon_{C_2})$ proving the claim.

We will say that the pair $a_1, a_2 \in \mathrm{supp}(s)$ is *good* for some s-weak primitive cycle C of Q if $\varepsilon_C = \mu(\varepsilon'_C)$ for an $\sim_{s_{(a_1,a_2)}}$-weak primitve cycle C'. Now let $l, k \geq 3$ and $C_1, \dots, C_l, D_1, \dots, D_k$ be primitive cycles in Q satisfying

$$s = \varepsilon_{C_1} + \cdots + \varepsilon_{C_l} = \varepsilon_{D_1} + \cdots + \varepsilon_{D_k}.$$

To prove the lemma we need to show that $\varepsilon_{C_i} \sim_s \varepsilon_{D_j}$ for some $i \in \{1, \dots, l\}$ and $j \in \{1, \dots, k\}$. By the claim proven in the previous paragraph if we can show that there is a pair of arrows $a_1, a_2 \in \mathrm{supp}(s)$ with $a_1^+ = a_2^-$ that is good for at least one of the $C_{1,\dots,l}$ and one of the $D_{1,\dots,k}$ then we are done by induction. Suppose now that there is no such pair of arrows. Since $\mathrm{supp}(s)$ contains no loops all of the primitive cycles listed contain at least two arrows. A pair of arrows a_1, a_2 that are consecutive in some C_i will clearly be good for all of the $C_{1,\dots,l}$. Moreover there is at least one $p \in \{1, \dots, k\}$ with $\varepsilon_{D_p}(a_1) \leq s(a_1)$ and $\varepsilon_{D_p}(a_2) \leq s(a_2)$, implying that $\varepsilon_{D_p} \in \mathrm{Im}(\mu)$ (μ now denotes the map corresponding to this particular a_1, a_2). Let D'_p denote the primitive cycle of $Q^{(a_1,a_2)}$ with $\mu(\varepsilon_{D'_p}) = \varepsilon_{D_p}$

and $\varepsilon_{D'_p}(a_{12}) = 0$. If the pair a_1, a_2 is not good for D_p then D'_p has to be a $s_{(a_1,a_2)}$-strong primitive cycle in $Q^{(a_1,a_2)}$. In other words $s - \varepsilon_{D_p} = \sum_{i \neq p} D_i$ becomes the characteristic vector of a primitive cycle after we replace a copy of a_1 and a_2 with a_{12}. This can only occur when $k = 3$ and $s - \varepsilon_{D_p}$ is the sum of the characteristic vectors of two primitive cycles E_1 and E_2, such that $a_1 \in E_1$, $a_2 \in E_2$ and the only vertex contained in both E_1 and E_2 is $a_1^+ = a_2^-$. This implies that $a_1^- \neq a_2^+$ and consequently that C_i has length at least 3. Also observe that for each of the $D_{1,\dots,k}$ there are at most two pairs of consecutive arrows a_1, a_2 such that $\varepsilon_{D_p}(a_1) \leq s(a_1)$ and $\varepsilon_{D_p}(a_2) \leq s(a_2)$ and D'_p is strong in $s_{(a_1,a_2)}$. So there are at most 6 pairs of such arrows satisfying these conditions for some of the $D_{1,\dots,k}$. We obtained that if one of the $C_{1,\dots,l}$ is of length 2 then its arrows will be good for one of the $D_{1,\dots,k}$, on the other hand if all of the $C_{1,\dots,l}$ are of at least length 3 then together they contain at least 9 pairs of consecutive arrows, and by the above at least 3 of those will be good for some of the $D_{1,\dots,k}$ as well. $\square$

We note that in the next theorem by the degree of a monomial in $\mathbb{C}[t_1, \dots, t_r]$ we mean the one that can be obtained from the standard grading $\deg(t_i) = 1$.

Theorem 4.6 *Let Q be a quiver such that $d := \dim(\mathcal{M}(Q, 0)) > 0$. Then $\mathcal{I}_0(Q)$ is generated by binomials that are the difference of a monomial of degree 2 and a monomial of degree at most $d - 1$.*

Proof. From Lemma 4.2 and Lemma 4.5 we obtain that a generating set of binomials can be given such that at least one monomial of each binomial is of degree 2. For the bound on the degree of the other monomial first consider that up to dimension 2 the only affine toric quiver varieties are the affine spaces. Suppose from now on that $d \geq 3$. Clearly it is sufficient to deal with the case when $(Q, 0)$ is tight and Q is prime. Suppose that $\varepsilon_{D_1} + \dots + \varepsilon_{D_k} = \varepsilon_{C_1} + \varepsilon_{C_2}$. Note that each D_i has an arrow contained in C_1 but not in C_2, and has an arrow contained in C_2 but not in C_1. It follows that $\text{length}(C_1) + \text{length}(C_2) \geq 2k$, implying that Q has at least k vertices. By Proposition 3.29 we conclude that $d - 1 = \chi(Q) - 1 \geq |Q_0| \geq k$. $\square$

Remark 4.7 The quiver in Example 3.30 shows that the bound $d - 1$ on the degree in Theorem 4.6 is sharp.

4.3 Toric ideals in the projective case

Corollary 4.3 applies for the monoid $S(Q,\theta)$, where the grading is given by $S(Q,\theta)_k = \nabla(Q,k\theta)\cap\mathbb{Z}^{Q_1}$. Recall that when Q contains no oriented cycles $\nabla(Q,\theta)$ is a normal polytope. We set the notation

$$F := \mathbb{C}[t_m \mid m \in \nabla(Q,\theta)\cap\mathbb{Z}^{Q_1}]$$

and $\mathcal{A}(Q,\theta) = \mathbb{C}[S(Q,\theta)]$ and recall from Section 2.1 that toric ideal corresponding to $\nabla(Q,\theta)$ is given by the kernel of the map:

$$\varphi : F \to \mathcal{A}(Q,\theta), \qquad t_m \mapsto x^m. \tag{4.7}$$

The ideal $\ker(\varphi)$ is a homogeneous ideal in the polynomial ring F (endowed with the standard grading) which we will denote by $\mathcal{I}(Q,\theta)$.

The following statement is a special case of the main result (Theorem 2.1) of [46]:

Proposition 4.8 *Let $Q = K(n,n)$ be the complete bipartite quiver with n sources and n sinks, with a single arrow from each source to each sink. Let θ be the weight with $\theta(v) = -1$ for each source and $\theta(v) = 1$ for each sink. Then the ideal $\mathcal{I}(Q,\theta)$ is generated by elements of degree at most 3.*

For sake of completeness we present a proof. The argument below is based on the key idea of [46], but we use a different language and obtain a very short derivation of the result. For this quiver and weight generators of $\mathcal{A}(Q,\theta)$ correspond to perfect matchings of the underlying graph of $K(n,n)$. Recall that a *perfect matching* of $K(n,n)$ is a set of arrows $\{a_1,\dots,a_n\}$ such that for each source v there is a unique i such that $a_i^- = v$ and for each sink w there is a unique j such that $a_j^+ = w$. Now $\nabla(Q,\theta)\cap\mathbb{Z}^{Q_1}$ in this case consists of the characteristic functions of perfect matchings of $K(n,n)$. By a *near perfect matching* we mean an incomplete matching that covers all but 2 vertices (1 sink and 1 source). Abusing language we shall freely identify a (near) perfect matching and its characteristic function (an element of $\mathbb{N}_0^{Q_1}$). First we show the following lemma:

Lemma 4.9 *Let θ be the weight for $Q = K(n,n)$ as above, and $m_1+\dots+m_k = q_1+\dots+q_k$ for some $k \geq 4$ and $m_i, q_j \in \nabla(Q,\theta)\cap\mathbb{Z}^{Q_1}$. Furthermore let us assume that for some $0 \leq l \leq n-2$ there is a near perfect matching p such that $p \leq m_1 + m_2$ and p contains l arrows from q_1. Then there is a $j \geq 3$ and $m_1', m_2', m_j' \in \nabla(Q,\theta)\cap\mathbb{Z}^{Q_1}$ and a near perfect* *matching p' such that $m_1 + m_2 + m_j = m_1' + m_2' + m_j'$, $p' \leq m_1' + m_2'$ and p' contains $l+1$ arrows from q_1.*

Proof. Let $v_1, \ldots, v_n$ be the sources and $w_1, \ldots, w_n$ the sinks of Q, and let us assume that p covers all vertices but v_1 and w_1. Let a be the arrow incident to v_1 in q_1. If a is contained in $m_1 + m_2$ then pick an arbitrary $j \geq 3$, otherwise take j to be such that m_j contains a. We can obtain a near perfect matching $p' < m_1 + m_2 + m_j$ that intersects q_1 in $l+1$ arrows in the following way: if a connects v_1 and w_1 we add a to p and remove one arrow from it that was not contained in q_1 (this is possible due to $l \leq n-2$); if a connects v_1 and w_i for some $i \neq 1$ then we add a to p and remove the arrow from p which was incident to w_i (this arrow is not contained in q_1). Set $r := m_1 + m_2 + m_j - p' \in \mathbb{N}_0^{Q_1}$, and denote by S the subquiver of $K(n,n)$ with $S_0 = Q_0$ and $S_1 = \{c \in Q_1 \mid r(c) \neq 0\}$. We have $S_0 = S_0^- \coprod S_0^+$ where S_0^- denotes the set of sources and S_0^+ denotes the set of sinks. For a vertex $v \in S_0$ set $\deg_r(v) := \sum_{v \in \{c^-, c^+\}} r(c)$. We have that $\deg_r(v) = 3$ for exactly one source and for exactly one sink, and $\deg_r(v) = 2$ for all the remaining vertices of S. Now let A be an arbitrary subset of S_0^-, and denote by B the subset of S_0^+ consisting of the sinks that are connected by an arrow in S to a vertex in A. We have the inequality $\sum_{v \in A} \deg_r(v) \leq \sum_{w \in B} \deg_r(w)$. Since on both sides of this inequality the summands are 2 or 3, and 3 can occur at most once on each side, we conclude that $|B| \geq |A|$. Applying the König-Hall Theorem (cf. Theorem 16.7 in [41]) to S we conclude that it contains a perfect matching. Denote the characteristic vector of this perfect matching by m_j'. Take perfect matchings m_1' and m_2' of S with $m_1 + m_2 + m_j - m_j' = m_1' + m_2'$ (note that m_1', m_2' exist by normality of the polytope $\nabla(Q, \theta)$, which in this case can be seen as an imediate consequence of the König-Hall Theorem). By construction we have $m_1 + m_2 + m_j = m_1' + m_2' + m_j'$, $p' \leq m_1' + m_2'$, and p' has $l+1$ common arrows with q_1. $\square$

Proof of Proposition 4.8 By Corollary 4.3 it is sufficient to show that if $s = m_1 + \cdots + m_k = q_1 + \cdots + q_k$ where $m_i, q_j \in \nabla(Q, \theta) \cap \mathbb{Z}^{Q_1}$ and $k \geq 4$, then the m_i, q_j all belong to the same equivalence class with respect to $\sim_s$. Note that since $k \geq 4$, by (4.3) the elements m_1', m_2', m_j' from the statement of Lemma 4.9 belong to the same equivalence class with respect to $\sim_s$ as $m_1, \ldots, m_k$. Hence repeatedly applying Lemma 4.9 we may assume that there is a near perfect matching $p \leq m_1 + m_2$ such that p and q_1 have $n-1$ common arrows. The only arrow of q_1 not belonging to p belongs to some m_j, hence after a possible renumbering of $m_3, \ldots, m_k$ we may assume that $q_1 \leq m_1 + m_2 + m_3$. It follows that $q_1 \sim_s m_4$, implying in turn by (4.3) that the m_i, q_j all belong to the same quivalence class with respect to $\sim_s$. $\square$

Now we are in position to state and prove the main result of this section (this was stated in [34] as well, but was withdrawn later, see [35]):

Theorem 4.10 *Let Q be a quiver with no oriented cycles, $\theta \in \mathbb{Z}^{Q_1}$ a weight such that $\nabla(Q,\theta)$ is non-empty. Then the ideal $\mathcal{I}(Q,\theta)$ is generated by elements of degree at most* 3.

Proof. By Proposition 3.21 and the double quiver construction (cf. the proof of Theorem 3.28) it is sufficient to deal with the case when Q is bipartite and $\nabla(Q,\theta)$ is non-empty. This implies that $\theta(v) \leq 0$ for each source vertex v and $\theta(w) \geq 0$ for each sink vertex w. Note that if $\theta(v) = 0$ for some vertex $v \in Q_0$, then omitting v and the arrows adjacent to v we get a quiver Q' such that the lattice polytope $\nabla(Q,\theta)$ is integral-affinely equivalent to $\nabla(Q', \theta|_{Q'_0})$, hence we may assume that $\theta(v) \neq 0$ for each $v \in Q_0$. We shall apply induction on $\sum_{v \in Q_0}(|\theta(v)| - 1)$.

The induction starts with the case when $\sum_{v \in Q_0}(|\theta(v)| - 1) = 0$, in other words, $\theta(v) = -1$ for each source v and $\theta(w) = 1$ for each sink w. This forces that the number of sources equals to the number of sinks in Q. The case when Q is the complete bipartite quiver $K(n,n)$ having n sinks and n sources, and each source is connected to each sink by a single arrow is covered by Proposition 4.8. Suppose next that Q is a subquiver of $K(n,n)$ having a relative invariant of weight θ (i.e. $K(n,n)$ has a perfect matching all of whose arrows belong to Q). The lattice polytope $\nabla(Q,\theta)$ can be identified with a subset of $\nabla(K(n,n),\theta)$: think of $m \in \mathbb{Z}^{Q_1}$ as $\tilde{m} \in \mathbb{Z}^{K(n,n)_1}$ where $\tilde{m}(a) = 0$ for $a \in K(n,n)_1 \setminus Q_1$ and $\tilde{m}(a) = m(a)$ for $a \in Q_1 \subseteq K(n,n)_1$. The surjection $\tilde{\varphi} : \mathbb{C}[t_m \mid m \in \nabla(K(n,n),\theta)] \to \mathcal{A}(K(n,n),\theta)$ restricts to $\varphi : \mathbb{C}[t_m \mid m \in \nabla(Q,\theta)] \to \mathcal{A}(Q,\theta)$. Denote by π the surjection of polynomial rings that sends to zero the variables t_m with $m \notin \nabla(Q,\theta)$. Then π maps the ideal $\ker(\tilde{\varphi})$ onto $\ker(\varphi)$, consequently generators of $\ker(\tilde{\varphi})$ are mapped onto generators of $\ker(\varphi)$. Since we know already that the first ideal is generated by elements of degree at most 3, the same holds for $\ker(\varphi)$. The case when Q is an arbitrary bipartite quiver with n sources and n sinks having possibly multiple arrows, and $\theta(v) = -1$ for each source v and $\theta(w) = 1$ for each sink w follows from the above case by a repeated application of Proposition 4.11 below.

Assume next that $\sum_{v \in Q_0}(|\theta(v)| - 1) \geq 1$, so there exists a vertex $w \in Q_0$ with $|\theta(w)| > 1$. By symmetry we may assume that w is a sink, so $\theta(w) > 1$. Construct a new quiver Q' as follows: add a new vertex w' to Q_0, for each arrow b with $b^+ = w$ add an extra arrow b' with

$(b')^+ = w'$ and $(b')^- = b^-$, and consider the weight θ' with $\theta'(w') = 1$, $\theta'(w) = \theta(w) - 1$, and $\theta'(v) = \theta(v)$ for all other vertices v. By Corollary 4.3 and (4.3) it is sufficient to show that if

$$m_1 + \cdots + m_k = n_1 + \cdots + n_k = s \in S := S(Q, \theta)$$

for some $k \geq 4$ and $m_1, \ldots, m_k, n_1, \ldots, n_k \in \nabla(Q, \theta) \cap \mathbb{Z}^{Q_1}$, then $m_i \sim_s n_j$ for some i, j. Set $S' := S(Q', \theta')$, and consider the semigroup homomorphism $\pi : S' \to S$ given by

$$\pi(m')(a) = \begin{cases} m'(a) + m'(a') & \text{if } a^+ = w; \\ m'(a) & \text{if } a^+ \neq w. \end{cases}$$

Take an arrow α with $\alpha^+ = w$ and $s(\alpha) > 0$. After a possible renumbering we may assume that $m_1(\alpha) > 0$ and $n_1(\alpha) > 0$. Define $m'_1 \in \mathbb{N}_0^{Q'_1}$ as $m'_1(\alpha) = m_1(\alpha) - 1$, $m'_1(\alpha') = 1$, and $m'_1(a) = m_1(a)$ for all other arrows $a \in Q'_1$. Similarly define $n'_1 \in \mathbb{N}_0^{Q'_1}$ as $n'_1(\alpha) = n_1(\alpha) - 1$, $n'_1(\alpha') = 1$, and $n'_1(a) = n_1(a)$ for all other arrows $a \in Q'_1$. Clearly $\pi(m'_1) = m_1$, $\pi(n'_1) = n_1$. Now we construct $s' \in S'$ with $\pi(s') = s$, $s' - m'_1 \in \mathbb{N}_0^{Q'_1}$ and $s' - n'_1 \in \mathbb{N}_0^{Q'_1}$ (thus m'_1 and n'_1 divide s' in S'). Note that $\sum_{a^+=w} s(a) = k\theta(w)$ and $\sum_{a^+=w} \max\{m_1(a), n_1(a)\} < \sum_{a^+=w}(m_1(a) + n_1(a)) = 2\theta(w)$ (since $m_1(\alpha) > 0$ and $n_1(\alpha) > 0$). The inequalities $\theta(w) \geq 2$ and $k \geq 4$ imply that $\sum_{a^+=w}(s(a) - \max\{m_1(a), n_1(a)\}) \geq k$. Consequently there exist non-negative integers $\{t(a) \mid a^+ = w\}$ such that $\sum_{a^+=w} t(a) = (\sum_{a^+=w} s(a)) - k$, $s(a) \geq t(a) \geq \max\{m_1(a), n_1(a)\}$ for all $a \neq \alpha$ with $a^+ = w$, and $s(\alpha) - 1 \geq t(\alpha) \geq \max\{m_1(\alpha), n_1(\alpha)\} - 1$. Consider $s' \in \mathbb{Z}^{Q'_1}$ given by $s'(a') = s(a) - t(a)$ and $s'(a) = t(a)$ for $a \in Q_1$ with $a^+ = w$ and $s'(b) = s(b)$ for all other $b \in Q'_1$. By construction s' has the desired properties, and so there exist $m'_i, n'_j \in \nabla(Q', \theta')$ with $s' = m'_1 + \cdots + m'_k = n'_1 + \cdots + n'_k$. Since $\sum_{v \in Q'_0}(|\theta'(v)| - 1)$ is one less than $\sum_{v \in Q_0}(|\theta(v)| - 1)$, by the induction hypothesis we have $m'_1 \sim_{s'} n'_1$. It is clear that $a \sim_t b$ implies $\pi(a) \sim_{\pi(t)} \pi(b)$, so we deduce $m_1 \sim_s n_1$. As we pointed out before, this shows by Corollary 4.3 that $\ker(\varphi)$ is generated by elements of degree at most 3. $\square$

The above proof refered to a general recipe to derive a minimal generating system of $\mathcal{I}(Q, \theta)$ from a minimal generating system for the quiver obtained by collapsing multiple arrows to a single arrow. Let us consider the following situation: let Q be a quiver with no oriented cycles, $\alpha_1, \alpha_2 \in Q_1$ with $\alpha_1^- = \alpha_2^-$ and $\alpha_1^+ = \alpha_2^+$. Denote by Q' the quiver obtained from Q by collapsing the α_i to a single arrow α. Take a weight $\theta \in \mathbb{Z}^{Q_0} = \mathbb{Z}^{Q'_0}$. The map $\pi : \nabla(Q, \theta) \to \nabla(Q', \theta)$ mapping $m \mapsto m'$ with $m'(\alpha) = m(\alpha_1) + m(\alpha_2)$ and $m'(\beta) = m(\beta)$ for all $\beta \in Q'_1 \setminus \{\alpha\} = Q_1 \setminus \{\alpha_1, \alpha_2\}$ induces a surjection from the monoid

$S := S(Q,\theta)$ onto the monoid $S' := S(Q',\theta')$. This extends to a surjection of semigroup algebras $\pi : \mathbb{C}[S] \to \mathbb{C}[S']$, which are identified with $\mathcal{A}(Q,\theta)$ and $\mathcal{A}(Q',\theta)$, respectively. Keep the notation π for the induced $\mathbb{C}$-algebra surjection $\mathcal{A}(Q,\theta) \to \mathcal{A}(Q',\theta)$. We have the commutative diagram of $\mathbb{C}$-algebra surjections

$$\begin{array}{ccc} F = \mathbb{C}[t_m \mid m \in \nabla(Q,\theta) \cap \mathbb{Z}^{Q_1}] & \xrightarrow{\varphi} & \mathcal{A}(Q,\theta) \\ \downarrow \pi & & \downarrow \pi \\ F' = \mathbb{C}[t_{m'} \mid m' \in \nabla(Q',\theta) \cap \mathbb{Z}^{Q_1'}] & \xrightarrow{\varphi'} & \mathcal{A}(Q',\theta) \end{array}$$

where the left vertical map (denoted also by π) sends the variable t_m to $t_{\pi(m)}$. For any monomial $u \in F'$ and any $s \in S$ with $\pi(x^s) = \varphi'(u) \in S'$ we choose a monomial $\psi_s(u) \in F$ such that $\pi(\psi_s(u)) = u$ and $\varphi(\psi_s(u)) = x^s$. This is clearly possible: let $u = t_{m_1} \dots t_{m_r}$, then we take for $\psi_s(u)$ an element $t_{n_1} \dots t_{n_r}$ where $\pi(n_j) = m_j$, such that $(n_1 + \dots + n_r)(\alpha_1) = s(\alpha_1)$. Denote by $\varepsilon_i \in \mathbb{N}_0^{Q_1}$ the characteristic function of $\alpha_i \in Q_1$ $(i = 1,2)$.

Proposition 4.11 *Let $u_\lambda - v_\lambda$ $(\lambda \in \Lambda)$ be a set of binomial relations generating the ideal* $\ker(\varphi')$. *Then* $\ker(\varphi)$ *is generated by* $\mathcal{G}_1 \bigcup \mathcal{G}_2$, *where*

$$\begin{aligned} \mathcal{G}_1 &:= \{\psi_s(u_\lambda) - \psi_s(v_\lambda) \mid \lambda \in \Lambda, \pi(x^s) = \varphi'(u_\lambda)\} \\ \mathcal{G}_2 &:= \{t_m t_n - t_{m+\varepsilon_2-\varepsilon_1} t_{n+\varepsilon_1-\varepsilon_2} \mid m,n \in \nabla(Q,\theta) \cap \mathbb{Z}^{Q_1}, m(\alpha_1) > 0, n(\alpha_2) > 0\}. \end{aligned}$$

Proof. Clearly $\mathcal{G}_1$ and $\mathcal{G}_2$ are contained in $\ker(\varphi)$. Denote by I the ideal generated by them in F, so $I \subseteq \ker(\varphi)$. In order to show the reverse inclusion, take any binomial relation $u - v \in \ker(\varphi)$, then $\varphi(u) = \varphi(v) = x^s$ for some $s \in S$. It follows that $\pi(u) - \pi(v) \in \ker(\varphi')$, whence there exist monomials w_i such that $\pi(u) - \pi(v) = \sum_{i=1}^k w_i(u_i - v_i)$, where $u_i - v_i \in \{u_\lambda - v_\lambda, \quad v_\lambda - u_\lambda \mid \lambda \in \Lambda\}$, $w_1 u_1 = \pi(u)$, $w_i v_i = w_{i+1} u_{i+1}$ for $i = 1, \dots, k-1$ and $w_k v_k = \pi(v)$. Moreover, for each i choose a divisor $r_i \mid s$ such that $\pi(x^{r_i}) = \varphi'(u_i)$ (this is clearly possible). Then I contains the element $\sum_{i=1}^k \psi_{s-r_i}(w_i)(\psi_{r_i}(u_i) - \psi_{r_i}(v_i))$, whose ith summand we shall denote by $y_i - z_i$ for notational simplicity. Then we have that $\pi(y_1) = \pi(u)$, $\pi(z_k) = \pi(v)$, $\pi(z_i) = \pi(y_{i+1})$ for $i = 1, \dots, k-1$, and $x^s = \varphi(y_i) = \varphi(z_i)$. It follows by Lemma 4.12 below $u - y_1$, $v - z_k$, and $y_{i+1} - z_i$ for $i = 1, \dots, k-1$ are all contained in the ideal J generated by $\mathcal{G}_2$. Whence $u - v$ is contained in I. □

Lemma 4.12 *Suppose that for monomials $u, v \in F$ we have $\varphi(u) = \varphi(v) \in \mathcal{A}(Q,\theta)$ and $\pi(u) = \pi(v) \in F'$. Then $u - v$ is contained in the ideal J generated by $\mathcal{G}_2$ (with the notation of Proposition 4.11).*

Proof. If u and v have a common variable t, then $u - v = t(u' - v')$, and u', v' satisfy the conditions of the lemma. By induction on the degree we may assume that $u' - v'$ belongs to the ideal J. Take $m_1 \in \nabla(Q,\theta) \cap \mathbb{Z}^{Q_1}$ such that t_{m_1} is a variable occurring in u. There exists an $m_2 \in \nabla(Q,\theta) \cap \mathbb{Z}^{Q_1}$ such that t_{m_2} occurs in v, and $\pi(m_1) = \pi(m_2)$. By symmetry we may assume that $m_1(\alpha_1) \geq m_2(\alpha_1)$, and apply induction on the non-negative difference $m_1(\alpha_1) - m_2(\alpha_1)$. If $m_1(\alpha_1) - m_2(\alpha_1) = 0$, then $m_1 = m_2$ and we are done by the above considerations. Suppose next that $m_1(\alpha_1) - m_2(\alpha_1) > 0$. By $\pi(m_1) = \pi(m_2)$ we have $m_2(\alpha_2) > 0$, and the condition $\varphi(u) = \varphi(v)$ implies that there exists an $m_3 \in \nabla(Q,\theta) \cap \mathbb{Z}^{Q_1}$ such that $t_{m_2} t_{m_3}$ divides v, and $m_3(\alpha_1) > 0$. Set $m_2' := m_2 + \varepsilon_1 - \varepsilon_2$, $m_3' := m_3 - \varepsilon_1 + \varepsilon_2$. Clearly $m_2', m_3' \in \nabla(Q,\theta) \cap \mathbb{Z}^{Q_1}$ and $t_{m_2} t_{m_3} - t_{m_2'} t_{m_3'} \in J$. So modulo J we may replace v by $t_{m_2'} t_{m_3'} v'$ where $v = t_{m_2} t_{m_3} v'$. Clearly $0 \leq m_1(\alpha_1) - m_2'(\alpha_1) < m_1(\alpha_1) - m_2(\alpha_1)$, and by induction we are finished. $\square$

We conclude this section by a graph theoretical reformulation of Proposition 4.8 and Corollary 4.3. For an undirected graph G, we define its *matching graph* $\mathrm{Match}(G)$ to be the graph whose vertex set is labelled by the perfect matchings of G and two vertices are connected by an edge in $\mathrm{Match}(G)$ if the corresponding matchings are edge disjoint in G.

Proposition 4.13 *If G is bipartite and k-regular for $k \geq 4$ then* $\mathrm{Match}(G)$ *is connected.*

Proof. Let Q be the quiver we obtain by orienting each edge of G from one side of its bipartition to the other and θ be the weight that is -1 on each source of Q and 1 on each sink. Now let s denote the lattice point in $\mathbb{Z}^{Q_1}$ that takes value 1 on each arrow. Since G is k-regular we have $s \in k\nabla(Q,\theta)$ and the statement follows immediately from Proposition 4.8 and Corollary 4.3. $\square$

A well-studied special case of matching graphs are the graphs $\mathrm{Match}(K(n,n))$, which are usually referred to as derangement graphs in the literature. They can be equivalently defined as having vertices labelled by the elements of the symmetric group S_n, and edges running between σ_1 and σ_2 whenever $\sigma_1 \sigma_2^{-1}$ has no fixed point. These graphs are known to be connected for $n \geq 4$ (see for example [39]).

4.3.1 The general case in [46]

In this section we give a short derivation of the main result of [46] from the special case Proposition 4.8. To reformulate the result in our context consider a bipartite quiver Q

with at least as many sinks as sources. By a *one-sided matching* of Q we mean an arrow set which has exactly one arrow incident to each source, and at most one arrow incident to each sink. By abuse of language the characteristic vector in $\mathbb{Z}^{Q_1}$ of a one-sided matching will also be called a one-sided matching. The convex hull of the one-sided matchings in $\mathbb{Z}^{Q_1}$ is a lattice polytope in $\mathbb{R}^{Q_1}$ which we will denote by $OSM(Q)$. Clearly the lattice points of $OSM(Q)$ are precisely the one-sided matchings. The normality of $OSM(Q)$ is explained in section 4.2 of [46] or it can be directly shown using the König-Hall Theorem for regular graphs and an argument similar to that in the proof below. Denote by $S(OSM(Q))$ the submonoid of $\mathbb{N}_0^{Q_1}$ generated by $OSM(Q) \cap \mathbb{Z}^{Q_1}$. This is graded, the generators have degree 1. Consider the ideal of relations among the generators $\{x^m \mid m \in OSM(Q) \cap \mathbb{Z}^{Q_1}\}$ of the semigroup algebra $\mathbb{C}[S(OSM(Q))]$. Theorem 2.1 from [46] can be stated as follows:

Theorem 4.14 *The ideal of relations of* $\mathbb{C}[S(OSM(Q))]$ *is generated by binomials of degree at most* 3.

Proof. Consider a quiver Q' that we obtain by adding enough new sources to Q so that it has the same number of sources and sinks, and adding an arrow from each new source to every sink. Let θ be the weight of Q' that is -1 on each source and 1 on each sink. Now the natural projection $\pi : \mathbb{R}^{Q'_1} \to \mathbb{R}^{Q_1}$ induces a surjective map from $\nabla(Q', \theta) \cap \mathbb{Z}^{Q'_1}$ onto $OSM(Q) \cap \mathbb{Z}^{Q_1}$ giving us a degree preserving surjection between the corresponding semigroup algebras. By Corollary 4.3 it is sufficient to prove that for any $k \geq 4$, any degree k element $s \in S(OSM(Q))$, and any $m, n \in OSM(Q) \cap \mathbb{Z}^{Q_1}$ with m, n dividing s we have $m \sim_s n$. In order to show this we shall construct an $s' \in \nabla(Q', k\theta) \cap \mathbb{Z}^{Q'_1}$ and $m', n' \in \nabla(Q', \theta) \cap \mathbb{Z}^{Q'_1}$ such that $m' \leq s'$, $n' \leq s'$, $\pi(m') = m$, $\pi(n') = n$ and $\pi(s') = s$. By Proposition 4.8 we have $m' \sim_{s'} n'$, hence the surjection π yields $m \sim_s n$. The desired s', m', n' can be obtained as follows: think of s as the multiset of arrows from Q, where the multiplicity of an arrow a is $s(a)$. Pairing off the new sources $Q'_0 \setminus Q_0$ with the sinks in Q not covered by m and adding the corresponding arrows to m we get a perfect matching m' of Q' with $\pi(m') = m$. Next do the same for n, with the extra condition that if none of n and m covers a sink in Q, then in n' it is connected with the same new source as in m'. Let $t \in \mathbb{N}_0^{Q'_1}$ be the multiset of arrows obtained from s by adding once each of the arrows $Q'_1 \setminus Q_1$ occuring in m' or n'. For a vertex $v \in Q'_1$ set $\deg_t(v) := \sum_{v \in \{c^-, c^+\}} t(c)$. Observe that $s - m$ and $s - n$ belong to $S(OSM(Q))_{k-1}$, hence $\deg_{s-m}(w) \leq k - 1$ and $\deg_{s-n}(w) \leq k - 1$ for any vertex w. If w is a sink not covered by m or n, then $\deg_s(w)$ agrees with $\deg_{s-m}(w)$ or $\deg_{s-n}(w)$, thus $\deg_s(w) \leq k - 1$, and hence $\deg_t(w) \leq k$. For the remaining sinks we have $\deg_t(w) = \deg_s(w) \leq k$ as well, moreover, $\deg_t(v) = k$ for the sources $v \in Q_0 \setminus Q'_0$, whereas $\deg_t(v) \leq 2$ for the new sources $v \in Q'_0 \setminus Q_0$. Consequently successively adding further new arrows from $Q'_1 \setminus Q_1$ to t we obtain $s' \geq t$ with $\deg_{s'}(v) = k$ for all $v \in Q'_0$. Moreover, $m' \leq t \leq s'$, $n' \leq t \leq s'$, and $\pi(s') = s$, so we are done. $\square$

4.4 Quiver cells

The aim of this section is to compile a full list of quiver polytopes whose toric ideals are not generated in degree 2 up to dimension 4. For an undirected graph G we will denote by G^* the quiver we obtain from G by putting a valency 2 sink on each edge. Recall from Theorem 3.22 that any d-dimensional prime quiver polytope can be realized as $\nabla(G^*, \theta)$ for a 3-regular graph G. The following proposition shows, that to obtain a full list of quiver polytopes in a given dimension such that their toric ideals are not generated in degree 2, one only has to consider prime quivers and products of lower dimensional examples. The statement is likely well-known but we provide a short proof for the sake of completeness.

Proposition 4.15 *Let $\nabla_1 \subset \mathbb{R}^{d_1}$ and $\nabla_2 \subset \mathbb{R}^{d_2}$ be normal lattice polytopes and $\nabla = \nabla_1 \times \nabla_2 \subset \mathbb{R}^{d_1+d_2}$. Denote the corresponding toric ideals by $\mathcal{I}(\nabla)$, $\mathcal{I}(\nabla_1)$ and $\mathcal{I}(\nabla_2)$. Then $\mathcal{I}(\nabla)$ is generated in degree 2 if and only if both $\mathcal{I}(\nabla_1)$ and $\mathcal{I}(\nabla_2)$ are generated in degree 2.*

Proof. Recall that $S(\nabla)$ denotes the graded semigroup corresponding to ∇. By replacing a polytope $\nabla \subset \mathbb{R}^d$ with the integral-affinely equivalent polytope $\nabla \times \{1\} \subset \mathbb{R}^{d+1}$ it can be always assumed that $k\nabla \cap l\nabla = \emptyset$ for positive integers $k \neq l$, hence it makes sense to identify the elements of $k\nabla$ with the degree k part of $S(\nabla)$.

First assume that $\mathcal{I}(\nabla)$ is generated in degree 2 and pick any $s_1 \in k\nabla_1 \cap \mathbb{Z}^{d_1}$ for $k \geq 3$. We need to show that $\sim_{s_1}$ has only one equivalence class. Indeed choose any $m_1, m_2 \in \nabla_1 \cap \mathbb{Z}^{d_1}$ such that $m_1, m_2 \mid_{S(\nabla_1)} s_1$, moreover choose a $s_2 \in k\nabla_2 \cap \mathbb{Z}^{d_2}$ and $n \in \nabla_2 \cap \mathbb{Z}^{d_2}$ such that $n \mid_{S(\nabla_2)} s_2$. By the assumption we have $(m_1, n) \sim_{(s_1,s_2)} (m_2, n)$, hence there is a sequence as in the definition of $\sim_s$ starting from (m_1, n) and ending in (m_2, n). Projecting this sequence to the coordinates that correspond to ∇_1 we obtain that $m_1 \sim_{s_1} m_2$.

For the other direction let $(s_1, s_2) \in k\nabla \cap \mathbb{Z}^{d_1+d_2}$ for $k \geq 3$ and $(m_1, n_1), (m_2, n_2) \in \nabla \cap \mathbb{Z}^{d_1+d_2}$ such that $(m_1, n_1), (m_2, n_2) \mid_{S(\nabla)} (s_1, s_2)$. By the assumption that $\mathcal{I}(\nabla_1)$ is generated in degree 2 there is a sequence $w_1, \ldots, w_j \in \nabla_1 \cap \mathbb{Z}^{d_1}$ such that $w_1 = m_1$, $w_j = m_2$ and $w_i +$

$w_{i+1} \mid_{S(\nabla_1)} s_1$ for all $i = 1, \dots, j-1$. Note that we can always assume $j \geq 3$. By normality of ∇_2 we have $s_2 = n_1 + u_1 + \dots u_{k-1}$ for some $u_1, \dots, u_{k-1} \in \nabla_2 \cap \mathbb{Z}^{d_2}$. Since $k, j \geq 3$ it is possible to choose a sequence $p_2, \dots, p_{j-1}$ such that $p_i \in \{n_1, u_1, \dots, u_{k-1}\}$, $p_i \neq p_{i+1}$ and $n_1 \notin \{p_2, p_{j-1}\}$. Now the sequence $(m_1, n_1), (w_2, p_2), \dots, (w_{j-1}, p_{j-1}), (m_2, n_1)$ satisfies the conditions in the definition of $\sim_{(s_1,s_2)}$ hence we obtained $(m_1, n_1) \sim_{(s_1,s_2)} (m_2, n_1)$. Applying the same argument for n_1 instead of m_1 we have $(m_1, n_1) \sim_{(s_1,s_2)} (m_2, n_2)$ completing the proof. □

The following proposition shows that for our purpose it is enough to deal with the cases when G is a simple graph (i.e. it contains no multiple edges).

Lemma 4.16 *Let G be a graph, containing two edges - e_1 and e_2 - running between the same vertices and denote by v_1, v_2 the valency 2 sinks of G^* that are placed on e_1 and e_2 respectively. Let H be the graph we obtain from G by collapsing e_1 and e_2 into a single edge e and denote by w the valency 2 sink of H^* that is placed on e. Let θ be a weight on G^*, such that $\nabla(Q, \theta)$ is non-empty, and θ' the weight on H^* we obtain by setting $\theta'(w) = \theta(v_1) + \theta(v_2)$ and $\theta' = \theta$ on the rest of the vertices. We have that $\mathcal{I}(G^*, \theta)$ is generated by its elements of degree 2 if and only $\mathcal{I}(H^*, \theta')$ is generated by its elements of degree 2.*

Proof. Let us denote the arrows of G^* and H^* incident to v_1, v_2 and w as in the picture below.

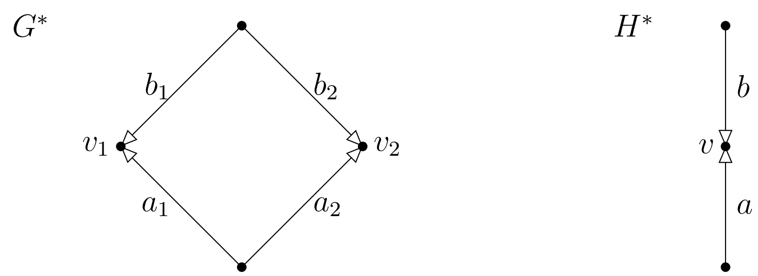

Let us identify the arrows in $G_1^* \setminus \{a_1, a_2, b_1, b_2\}$ with the arrows of $H_1^* \setminus \{a, b\}$ and define the linear map $\varphi : \mathbb{R}^{G_1^*} \to \mathbb{R}^{H_1^*}$, as $\varphi(x)(a) = x(a_1) + x(a_2)$, $\varphi(x)(b) = x(b_1) + x(b_2)$ and $\varphi(x)(c) = x(c)$ for $c \in G_1^* \setminus \{a_1, a_2, b_1, b_2\}$. The morphism φ maps $k\nabla(G^*, \theta)$ onto $k\nabla(H^*, \theta)$ for any positive integer k. Moreover for $m \in \nabla(G^*, \theta)$ and $s \in k\nabla(G^*, \theta)$ the inequality $m \leq s$ implies $\varphi(m) \leq \varphi(s)$ and it can be easily checked that φ maps the set $\{m \in \nabla(G^*, \theta) \mid m \leq s\}$ onto $\{m \in \nabla(H^*, \theta) \mid m \leq \varphi(s)\}$. It follows that sequences that appear in the definition of $\sim_s$ map onto sequences that appear in the definition of $\sim_{\varphi(s)}$, hence $\sim_{\varphi(s)}$ has at most as many equivalence classes as $\sim_s$. It follows by Corollary 4.3 that if $\mathcal{I}(G^*,\theta)$ is generated by elements of degree 2 then so is $\mathcal{I}(H^*,\theta)$.

For the other direction first record that for any $x \in \nabla(G^*,\theta)$ we have that $x(b_1) = \theta(v_1) - x(a_1)$, $x(b_2) = \theta(v_1) - x(a_2)$, moreover $x(a_1), x(b_1) \leq \theta(v_1)$ and $x(a_2), x(b_2) \leq \theta(v_2)$. Similarly for $x \in \nabla(H^*,\theta)$ we have that $x(b) = \theta'(v) - x(a)$ and $x(a), x(b) \leq \theta'(v)$. Hence for any lattice point $m \in \nabla(H^*,\theta) \cap \mathbb{Z}^{H_1^*}$ and $\alpha_1, \alpha_2 \in \mathbb{N}$ satisfying $\alpha_1 \leq \theta(v_1)$, $\alpha_2 \leq \theta(v_2)$ and $\alpha_1 + \alpha_2 = m(a)$, there is a unique lattice point $n \in \nabla(G^*,\theta) \cap \mathbb{Z}^{G_1^*}$ such that $n(a_1) = \alpha_1$, $n(a_2) = \alpha_2$ and $\varphi(n) = m$. We will denote this preimage of m by $m(\alpha_1, \alpha_2)$. Now assume that $\mathcal{I}(H^*,\theta')$ is generated in degree 2 and let s be a lattice point in $3\nabla(G^*,\theta)$. By Corollary 4.3 we need to show that $\sim_s$ has precisely one equivalence class.

First we claim that if for some $m \in \nabla(H^*,\theta)$ and $\alpha_1, \alpha_2, \beta_1, \beta_2 \in \mathbb{N}$ we have that $m(\alpha_1, \alpha_2), m(\beta_1, \beta_2) \leq s$ then $m(\alpha_1, \alpha_2) \sim_s m(\beta_1, \beta_2)$. We may assume that $\alpha_1 < \beta_1$ and $\alpha_2 > \beta_2$. By applying induction it is enough to deal with the case when $\beta_1 = \alpha_1 + 1$ and $\beta_2 = \alpha_2 - 1$. By normality of $\nabla(G^*,\theta)$ we can write

$$s = m(\alpha_1, \alpha_2) + m'(\alpha_1', \alpha_2') + m''(\alpha_1'', \alpha_2'')$$

for some $m', m'' \in \nabla(H^*,\theta) \cap \mathbb{Z}^{H_1^*}$ and $\alpha_1', \alpha_2', \alpha_1'', \alpha_2'' \in \mathbb{N}$. Since $m(\alpha_1 + 1, \alpha_2 - 1) \leq s$ one of α_1' and α_1'' need to be positive. Without loss of generality we may assume that $\alpha_1' > 0$. If $\alpha_2' < \theta(v_2)$ then we have

$$m(\alpha_1, \alpha_2) + m(\alpha_1', \alpha_2') = m(\alpha_1 + 1, \alpha_2 - 1) + m(\alpha_1' - 1, \alpha_2' + 1),$$

hence $m(\alpha_1, \alpha_2) \sim_s m(\alpha_1'', \alpha_2'') \sim_s m(\alpha_1 + 1, \alpha_2 - 1)$. We are done unless $\alpha_2' = \theta(v_2)$.

Now if $\alpha_1'' > 0$ we can argue similarly and see that if $\alpha_2'' < \theta(v_2)$ then $m(\alpha_1, \alpha_2) \sim_s m(\alpha_1 + 1, \alpha_2 - 1)$. However if $\alpha_2' = \alpha_2'' = \theta(v_2)$ then $s(a_2) = \alpha_2 + 2\theta(v_2)$ and consequently $s(b_2) = 3\theta(v_2) - s(a_2) = \theta(v_2) - \alpha_2$. On the other hand since $m(\alpha_1 + 1, \alpha_2 - 1) \leq s$ we have that $s(b_2) \geq \theta(v_2) - \alpha_2 + 1$ a contradiction.

Hence we are only left to deal with the case when $\alpha_2' = \theta(v_2)$ and $\alpha_1'' = 0$, and then we have

$$s(a_1) = \alpha_1 + \alpha_1' \leq \alpha_1 + \theta(v_1)$$

and

$$s(a_2) = \alpha_2 + \theta(v_2) + \alpha_2'' \geq \theta(v_2) + 1.$$

Summarizing we obtained that either $m(\alpha_1, \alpha_2) \sim_s m(\alpha_1 + 1, \alpha_2 - 1)$ or $s(a_1) \leq \alpha_1 + \theta(v_1)$

and $s(a_2) \geq \theta(v_2) + \alpha_2$. Now repeating the entire argument for $m(\alpha_1 + 1, \alpha_2 - 1)$ instead of $m(\alpha_1, \alpha_2)$ we obtain that either $m(\alpha_1, \alpha_2) \sim_s m(\alpha_1+1, \alpha_2-1)$ or $s(a_1) \geq \theta(v_1)+\alpha_1+1$ and $s(a_2) \leq \alpha_2 - 1 + \theta(v_2)$. Now it is clear that both pairs of inequalities can not hold at the same time hence we have $m(\alpha_1, \alpha_2) \sim_s m(\alpha_1 + 1, \alpha_2 - 1)$ proving the claim.

Now let $m(\alpha_1, \alpha_2)$ and $m'(\alpha'_1, \alpha'_2)$ be two arbitrary lattice points in $\nabla(G^*, \theta) \cap \{x \in \mathbb{Z}^{G^*} \mid x \leq s\}$. Since $\mathcal{I}(H^*, \theta')$ is generated by its elements of degree 2 we have $m \sim_{\varphi(s)} m'$, whence there is a series of lattice points $m_1, \ldots m_k \in \nabla(H^*, \theta')$ satisfying $m_1 = m$, $m_2 = m'$ and $m_i + m_{i+1} \leq \varphi(s)$. It is not difficult to see that integers $\gamma_1^i, \gamma_1'^i, \gamma_i^2, \gamma_i'^2$ for $i = 1, \ldots, k$ can be chosen, such that the expressions $m_i(\gamma_1^i, \gamma_2^i), m_i(\gamma_1'^i, \gamma_2'^i)$ make sense and satisfy $m_i(\gamma_1'^i, \gamma_2'^i) + m_{i+1}(\gamma_1^{i+1}, \gamma_2^{i+1}) \leq s$ and hence $m_i(\gamma_1'^i, \gamma_2'^i) \sim_s m_{i+1}(\gamma_1^{i+1}, \gamma_2^{i+1})$ for all $i \in \{1, \ldots, k-1\}$. Comparing this with the claim proven with the previous paragraph we obtain $m(\alpha_1, \alpha_2) \sim_s m'(\alpha'_1, \alpha'_2)$, completing the proof of the Lemma. □

Note that while we showed in Section 3.1 that it is possible to list every toric quiver variety in a given dimension, we estimate these lists in dimension 4 and higher to be extremely long. More importantly to each toric quiver variety there are an infinite number of quiver polyhedra associated, hence there did not seem to be any obvious way to achieve our goal via direct computation. Instead we follow an approach similar to that of [22], where it was shown that the amongst the 3×3 transportation polytopes, only the Birkhoff polytope B_3 (cf. the end of Section 2.4 for its definition) yields relations that are not generated in degree 2. This is a special case of our result, since 3×3 transportation polytopes are quiver polytopes of the bipartite quiver $K_{3,3}$. The key tool in their proof was to use a hyperplane subdivision to decompose the polytopes into "cells", which are subpolytopes of facet width 1, and then carry out a case by case analysis of the - finitely many - cells that occur.

For a lattice polytope $\nabla \subset \mathbb{R}^d$ and an integer vector $\underline{k} \in \mathbb{Z}^d$ we define the $\underline{k}$-*cell* of ∇ to be

$$\nabla_{\underline{k}} = \{x \in \nabla \mid \underline{k}(i) \leq x(i) \leq \underline{k}(i) + 1 \, \forall (i : 1 \leq i \leq d)\}.$$

By the $\underline{0}$-cell of ∇ we just mean the k-cell for $k = (0, \ldots, 0)$. For a quiver polytope $\nabla(Q, \theta)$ and a non-negative integer vector $k \in \mathbb{N}^{Q_1}$ we will denote by $\theta_{\underline{k}}$ the weight defined by

$$\theta_{\underline{k}}(v) = \sum_{a^+=v} \underline{k}(a) - \sum_{a^-=v} \underline{k}(a),$$

i.e. $\theta_{\underline{k}}$ is the unique weight θ' such that $\underline{k} \in \nabla(Q, \theta')$.

We will call the non-empty cells that can be obtained from quiver polytopes *quiver cells*. Since quiver polytopes always lie in the positive quadrant it suffices to consider cells defined by $\underline{k} \in \mathbb{N}^{Q_1}$. As we will show in the next proposition, one only has to consider a finite set of weights to obtain a complete list of quiver cells (up to translation) associated to a fixed quiver.

Proposition 4.17 *Let Q be a quiver and $\underline{k} \in \mathbb{N}^{Q_1}$.*

(i) For any integer weight θ we have that $\nabla(Q,\theta)_{\underline{k}} = \nabla(Q,\theta-\theta_{\underline{k}})_{\underline{0}} + \underline{k}$.

(ii) There are only finitely many different weights θ, such that $\nabla(Q,\theta)_{\underline{0}}$ is non-empty.

Proof. (i) follows immediately from the definition of quiver polytopes. For (ii) consider that the lattice points of $\nabla(Q,\theta)_{\underline{0}}$ take values $\{0,1\}$ on each edge, for the polytope to be non-empty, θ needs to satisfy

$$-|\{a \in Q_1 \mid a^- = v\}| \leq \theta(v) \leq |\{a \in Q_1 \mid a^+ = v\}|,$$

for each vertex $v \in Q_0$. □

It follows from Proposition 2.21 that quiver cells are also quiver polytopes, in particular they are normal and their toric ideals are generated in degree at most 3.

By an *alternating cycle* of the quiver Q we mean a lattice point $c \in M^Q = \mathcal{F}^{-1}(0) \cap \mathbb{Z}^{Q_1}$, such that $c(a) \in \{0,1,-1\}$ for all $a \in Q_1$ and $\mathrm{supp}(c)$ is a primitive (undirected) cycle. Recall from Proposition 2.12 that the alternating cycles generate M^Q. Moreover a simple inductive argument shows that any lattice point $m \in M^Q$ can be (greedily) decomposed as a sum of alternating cycles $c_1,\ldots,c_l$, such that the coordinates of the c_i are either zero or have the same sign as m. (Alternatively one can derive this statement from Theorem 21.2 from [42] along with the discussion that follows it and the fact that vertex-arrow incidence matrices of quivers are totally unimodular.) The following proposition shows us why cells play an important role in studying the generators of toric ideals. We note that an alternative proof could be derived from Theorem 6.2 in [9], which in turn is proven using several facts about the defining ideals of so-called monoidal complexes, but we preferred to show that there is a direct combinatorial argument.

Proposition 4.18 *For a quiver polytope $\nabla(Q,\theta)$, and a relation $b = t^{m_1}t^{m_2}t^{m_3} - t^{n_1}t^{n_2}t^{n_3} \in \mathcal{I}(Q,\theta)$ let $\underline{k} \in \mathbb{N}^{Q_1}$ be such that $(m_1+m_2+m_3)/3 \in \nabla(Q,\theta)_{\underline{k}}$. Then there exist*

$m_1', m_2', m_3', n_1', n_2', n_3' \in \nabla(Q,\theta)_{\underline{k}}$ such that $b' = t^{m_1'}t^{m_2'}t^{m_3'} - t^{n_1'}t^{n_2'}t^{n_3'} \in \mathcal{I}(Q,\theta)$ and b is contained in an ideal generated by b' and some degree 2 elements of $\mathcal{I}(Q,\theta)$.

Proof. We need to show that we can transform m_1, m_2, m_3 into m_1', m_2', m_3' as in the theorem by successively replacing pairs of lattice points with new lattice points that have the same sum. For an integer n and an arrow $a \in Q_1$ denote by $d_a(n)$ the distance of n from the set $\{\underline{k}(a), \underline{k}(a)+1\}$ and set

$$D(m_i) = \sum_{a \in Q_1} d_a(m_i(a)).$$

If $D(m_1)+D(m_2)+D(m_3) = 0$ we are done since $D(m_i) = 0$ if and only if $m_i \in \nabla(Q,\theta)_{\underline{k}}$. Otherwise we will show that we can replace two of m_1, m_2, m_3 with new lattice points from $\nabla(Q,\theta)$ having the same sum, such that the value of $D(m_1)+D(m_2)+D(m_3)$ strictly decreases, and - since the $D(m_i)$ are all integers - in finitely many steps the sum will be 0. If $D(m_1)+D(m_2)+D(m_3) \neq 0$ one of the m_j does not lie in the cell $\nabla(Q,\theta)_{\underline{k}}$, so we can assume that say $m_1(a) < \underline{k}(a)$ for some $a \in Q_1$. Since $\underline{k}(a) \leq (m_1(a)+m_2(a)+m_3(a))/3$ it follows that one of $m_2(a)$ and $m_3(1)$ has to be at least $k+1$. Assume that $m_2(a) \geq k+1$. Now since $m_2 - m_1 \in M^Q$ it decomposes as a sum of alternating cycles $c_1 + \dots + c_l$. Since $m_2(a) > m_1(a)$ we have that $c_j(a) = 1$ for some j. Since $c_1, \dots, c_l$ can be chosen such that their coordinates are either zero or have the same sign as $m_2 - m_1$, the lattice points $m_1 + c_j, m_2 - c_j$ have non-negative entries, hence $m_1 + c_j, m_2 - c_j \in \nabla(Q,\theta)$. We claim that for any $b \in Q_1$ we have that

$$d_b(m_1(b)) + d_b(m_2(b)) \geq d_b(m_1(b)+c_j(b)) + d_b(m_2(b)-c_j(b))$$

with strict inequality when $b = a$. If $c_j(b) = 0$ then the claim is trivial. If $c_j(b) \neq 0$ then $m_1(b) \neq m_2(b)$ and $c_j(b) = sgn(m_2(b) - m_1(b))$. It follows that

$$|d_b(m_1(b)+c_j(b)) - d_b(m_1(b))| \leq 1$$

and whenever

$$d_b(m_1(b)+c_j(b)) = d_b(m_1(b)) + 1$$

we also have

$$d_b(m_2(b)+c_j(b)) = d_b(m_2(b)) - 1$$

(and similarly in the other direction), proving the inequality in the claim. Moreover in the case $b = a$, we have that $c_j(a) = 1$ and

$$d_a(m_1(a)) > d_a(m_1(a) + 1),$$

since $m_1(a) < \underline{k}(a)$, and

$$d_a(m_2(a)) \geq d_a(m_2(a) - 1),$$

since $m_2(a) \geq \underline{k}(a)$. It follows that

$$D(m_1) + D(m_2) + D(m_3) > D(m_1 + c_j) + D(m_2 - c_j) + D(m_3)$$

and induction completes the proof. □

Note that our argument differs from Section 2 of [22] in that they used the fact that in the case of the 3×3 transportation polytopes the cells that do not yield cubic relations, also possess a quadratic Gröbner basis (in fact they were simplices).

Corollary 4.19 *If $s \in 3\nabla(Q, \theta)_{\underline{k}}$ is such that $\sim_s$ as a relation of the semigroup $S(\nabla(Q, \theta))$ has more than one equivalence class, then $\sim_s$ as a relation of the semigroup $S(\nabla(Q, \theta)_{\underline{k}})$ also has more than one equivalence class. In particular if $\mathcal{I}(\nabla(Q, \theta))$ is not generated in degree 2, then for some $\underline{k}$ the toric ideal $\mathcal{I}(\nabla(Q, \theta)_{\underline{k}})$ is not generated in degree 2 either.*

Note that the relation $\sim_s$ as a relation of the semigroup $S(\nabla(Q, \theta)_{\underline{k}})$ is not the same as the relation we get by considering as a relation of the semigroup $S(\nabla(Q, \theta))$ and restricting it to the cell $\nabla(Q, \theta)_{\underline{k}}$. Instead it is the relation we obtain by only considering sequences of vertices that are all in $\nabla(Q, \theta)_{\underline{k}}$ in the definition of $\sim_s$.

Proposition 4.17 provides us with a way to calculate a complete list of quiver cells that can be obtained from a fixed quiver. One needs to check for a finite set of weights whether the polytope $\nabla(Q, \theta)_{\underline{0}}$ is non-empty. In fact the number of weights to be considered can be further reduced if we are only interested in full-dimensional cells. We implemented this idea as an algorithm in SAGE to calculate cells of the quivers G^* where G is a 3-regular simple graph on $2d - 2$ vertices for $d = 3, 4$. The list of weights to be considered can be obtained from the bounds in Proposition 4.17. To calculate the lattice points of any particular cell one just needs to find the 0-1 solutions of a set of inequalities similarly to Example 3.39. After obtaining the weights we used the "Toric Varieties" module of SAGE, written by A. Novoseltsev and V. Braun, and the "Lattice and reflexive poltyopes" module, written by

A. Novoseltsev, to calculate the dimension of our polytopes and to decide whether they are smooth. Finally we used the "Toric Ideals" module, written by V. Braun, to find generators for the toric ideal of each polytope (note that by Theorem 4.24 one only has to check the singular cases). In the case $d = 3$, there is only one 3-regular simple graph on 4 vertices, the complete graph K_4. For $d = 4$ there are two 3-regular simple graphs on 6 vertices, the complete bipartite graph $K_{3,3}$ and the prism graph Y_3, shown in the picture below:

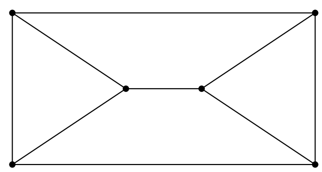

The following proposition collects the results we obtained from our SAGE calculations.

Proposition 4.20 *(i) The toric ideal of any cell* $\nabla(K_4^*, \theta)_{\underline{0}}$ *is generated by its elements of degree* 2.
(ii) The toric ideal of any cell $\nabla(Y_3^*, \theta)_{\underline{0}}$ *is generated by its elements of degree* 2*. The only cell* $\nabla(K_{3,3}^*, \theta)_{\underline{0}}$ *with toric ideal not generated in degree* 2 *is the Birkhoff polytope* B_3*.* $\nabla(K_{3,3}^*, \theta)_{\underline{0}}$ *is integral-affinely equivalent to* B_3 *precisely when* θ *is* -1 *on the vertices that belong to one of the classes in the bipartition of* $K_{3,3}$*,* -2 *on the vertices that belong to the other and* 1 *on the valency* 2 *sinks.*

We are ready to prove the main result of this section.

Theorem 4.21 *Let* Q *be a quiver without oriented cycles.*

(i) If $dim(\nabla(Q, \theta)) \leq 3$ *then* $\mathcal{I}(Q, \theta)$ *is generated in degree* 2*.*

(ii) If $dim(\nabla(Q, \theta)) = 4$ *then either* $\mathcal{I}(Q, \theta)$ *is generated in degree* 2*, or* $\nabla(Q, \theta)$ *is integral-affinely equivalent to the Birkhoff polytope* B_3*.*

Proof. The one dimensional case is trivial - in fact the only one dimensional projective toric variety is $\mathbb{P}_1$. In dimension 2 we can obtain the result without directly calculating the cells in the following way: It follows from Proposition 3.24 that every projective toric quiver variety is smooth, so in particular every cell is smooth. Hence by Theorem 4.24 the toric ideal of any cell is generated by its elements of degree 2 and then by Proposition 4.18 we see that the toric ideal of any 2-dimensional projective toric quiver variety is generated by its elements of degree 2.

For the higher dimensional cases recall from Theorem 3.22 that we only need to consider quiver polytopes $\nabla(G^*,\theta)$ for a 3-regular graph G on $2d-2$ vertices. In dimension 3 if G has multiple edges, then we can apply Lemma 4.16 and the fact that in the lower dimensional cases every toric ideal is generated in degree 2 to obtain the result. Note that the dimension of G^* will be strictly less than 3 after collapsing edges as in Lemma 4.16. If G has no multiple edges then it has to be K_4 and we are done by applying Propositions 4.20 and 4.18.

In dimension 4 the cases when G has multiple edges or is isomorphic to Y_3 can be dealt with similarly. The argument for the case when G is $K_{3,3}$ goes similarly to the one in Section 2 of [22]. If for some θ we have that $\mathcal{I}(K^*_{3,3},\theta)$ is not generated by its elements of degree 2 then, denoting by ω the weight defined in (ii) of Proposition 4.20 we have that $\nabla(K^*_{3,3},\theta)_{\underline{k}} = \underline{k} + \nabla(K^*_{3,3},\omega)_{\underline{0}}$ and that there is an $s \in 3\nabla(K^*_{3,3},\theta)_{\underline{k}}$, such that $\sim_s$ - considered as a relation on the vertices of $\nabla(K^*_{3,3},\theta)$ - has more than one equivalence class. Let us denote by $a_{i,j}$ for $i,j \in 1,2,3$ the arrows of $K^*_{3,3}$ that are incident to the vertices on which ω takes values -1. The six vertices of $\nabla(K^*_{3,3},\omega)_{\underline{0}} \simeq B_3$ can be indexed by the elements of the symmetric group S_3 with σ_{pqr} denoting the vertex defined by $\sigma_{pqr}(a_{1,p}) = \sigma_{pqr}(a_{2,q}) = \sigma_{pqr}(a_{3,r}) = 1$ and $\sigma_{pqr}(a_{i,j}) = 0$ for $(i,j) \notin \{(1,p),(2,q),(3,r)\}$. With this notation the unique element of $3\nabla(K^*_{3,3},\omega)_{\underline{0}}$ that corresponds to a relation that is not generated by elements of degree 2 is $s = \sigma_{123}+\sigma_{312}+\sigma_{231} = \sigma_{132}+\sigma_{312}+\sigma_{213}$. Set $s' = s+\underline{k}$. It remains to find the weights θ for which $\sim_{s'}$ has more than one equivalence class as a relation on the vertices of $\nabla(K^*_{3,3},\theta)$. It is easy to check that if $\underline{k}(a_{i,j}) = 0$ for all $i,j \in \{1,2,3\}$ then $\nabla(K^*_{3,3},\theta) = \nabla(K^*_{3,3},\theta)_{\underline{k}}$ and hence it is intergral-affinely equivalent to the Birkhoff polytope B_3, which indeed has the above cubic relation that is not generated in lower degree. Otherwise without loss of generality we may assume that $\underline{k}(a_{1,1}) > 0$. In this case however we have that $\sigma_{123}+\underline{k}+\sigma_{132}+\underline{k} \leq s'$, hence $\sigma_{123}+\underline{k} \sim_{s'} \sigma_{132}+\underline{k}$ and hence by Corollary 4.19 we see that $\sim'_s$ has only one equivalence class. □

4.5 Ideals of binary polytopes of toric GIT quotients

Our primary motivation for the results in this section was to study the toric ideals of quiver polytopes in the special case when the coordinates of the lattice points in the polytope take values 1 or 0. Up to integral-affine equivalence these are exactly the quiver cells which we studied in Section 4.4. However it turned out to be convenient to generalize our results for a wider class of polytopes, which all arise from toric GIT constructions.

Let $A \in \mathbb{Z}^{n\times d}$ denote an integer matrix and consider the action of $(\mathbb{C}^*)^n$ on $\mathbb{C}^d$ defined as

$$(t_1, \dots, t_n) \cdot (x_1, \dots, x_d) := (x_1 \prod_{i=1}^{n} t_i^{A_{i,1}}, \dots, x_d \prod_{i=1}^{n} t_i^{A_{i,d}}).$$

To a weight vector $\theta \in \mathbb{Z}^n$ we assign a character of $(\mathbb{C}^*)^n$ given by $(t_1, \dots, t_n) \to \prod_{i=1}^{n}(t_i)^{\theta_i}$. We define

$$\nabla(A, \theta) := \{x \in \mathbb{R}^d \mid x \geq 0, \quad Ax = \theta\}. \tag{4.8}$$

Our previous notation $\nabla(Q, \theta)$ can be considered a special case of the above if we identify Q with its signed vertex-arrow incidence matrix. Now, similarily to the case of toric quiver varieties, when $\nabla(A, \theta)$ is a normal lattice polyhedron one can verify without difficulty that the toric variety $X_{\nabla(A,\theta)}$ is isomorphic to the GIT quotient $\mathbb{C}^d /\!\!/_\theta (\mathbb{C}^*)^n$.

We note that in general $\nabla(A, \theta)$ need not be a lattice polyhedron, however the variety $\mathbb{C}^d /\!\!/_\theta (\mathbb{C}^*)^n$ is always toric and can be embedded by a suitably large integer multiple of $\nabla(A, \theta)$ (see Chapter 14 of [12] for details). A well-known sufficient condition for $\nabla(A, \theta)$ to be a normal lattice polyhedron is that A is totally unimodular, i.e. all of its square submatrices have determinants -1, 1 or 0 (see Theorems 19.2 and 19.4 in [42]).

We will call lattice polytopes that arise as $\nabla(A, \theta)$ as in (4.8) *standard* polytopes. Note that $\nabla(A, \theta)$ is a polytope if and only if it is non-empty and $\nabla(A, 0) = \{0\}$.

We will call a lattice polytope whose lattice points are all 0-1 vectors (i.e. vectors with all entries in the set $\{0, 1\}$) a *binary polytope*. The primary goal of this section is to show that toric ideals of standard normal binary polytopes are generated by elements of degree 2, under some assumptions on the arrangement of singular points. Note that if ∇ is a binary polytope then every lattice point in ∇ is a vertex. For a point $x \in \mathbb{R}^d$ we will denote by $\mathrm{supp}(x)$ the set $\{i \mid x(i) \neq 0\}$. We say that two vertices of a polytope are neighbours if they lie on the same edge with no intermediate lattice point, note that this happens if and only if there is a hyperplane that intersects $\nabla \cap \mathbb{Z}^d$ in precisely those two points. To prove our main results we will need two lemmas that hold for any standard binary polytope (normality is not required).

Lemma 4.22 *Let ∇ be a standard binary polytope. The vertices $v_1, v_2 \in \nabla \cap \mathbb{Z}^d$ are neighbouring if and only if there are no other vertices whose support is a subset of* $\mathrm{supp}(v_1) \cup \mathrm{supp}(v_2)$.

Proof. If v_1 and v_2 are the only vertices whose support is a subset of $\mathrm{supp}(v_1) \cup \mathrm{supp}(v_2)$ then for the hyperplane $H = \{x \mid \sum_{i \notin \mathrm{supp}(v_1) \cup \mathrm{supp}(v_2)} x(i) = 0\}$ we have $(\nabla \cap \mathbb{Z}^d) \cap H = \{v_1, v_2\}$.

For the other direction let w_1 be a vertex with $\mathrm{supp}(w_1) \subseteq \mathrm{supp}(v_1) \cup \mathrm{supp}(v_2)$. Then we have that $w_2 := v_1 + v_2 - w_1 \in \nabla \cap \mathbb{Z}^d$ since $w_1 \leq v_1 + v_2$ and $v_1 + v_2 - w_1$ clearly satisfies the linear equations defining ∇. Note that $w_2 \neq w_1$ since no vertex can lie on the line segment connecting v_1 and v_2. Now we have that $v_1 + v_2 = w_1 + w_2$, showing that the line segment connecting v_1 and v_2 is not an edge of ∇, thus v_1 and v_2 are not neighbouring. □

Lemma 4.23 *Let ∇ be a standard binary polytope. Let s be an element of $S(\nabla)$ and $v_1, v_2 \leq s$ two vertices of ∇. Then there is a series of vertices $w_1, \dots, w_k \in \nabla$, such that $w_1, \dots, w_k \leq s$, $v_1 = w_1$, $v_2 = w_k$ and w_i is a neighbour of w_{i+1} for all $i:\ 1 \leq i \leq k-1$.*

Proof. First note that $\nabla' := \{x \in \nabla \mid x \leq s\} = \{x \in \nabla \mid \mathrm{supp}(x) \subseteq \mathrm{supp}(s)\}$ is a polytope, since it is the intersection of a polytope and a linear subspace. We show that the set of vertices of ∇' is a subset of the set of vertices of ∇ (hence ∇' itself is a standard binary polytope). Assume that w is a vertex of ∇'. If w is not a vertex of ∇ then it can be written as a convex combination with non-zero coefficients of some vertices of ∇: $u_1 \dots u_l$, $l \geq 2$. However this would imply that $\mathrm{supp}(u_i) \subseteq \mathrm{supp}(w)$ for all $1 \leq i \leq l$, and consequently $u_i \in \nabla'$, so w itself can not be a vertex of ∇'. Now the lemma is proven by picking a series of neighbouring vertices in ∇' that connect v_1 to v_2. □

Recall from Section 2.1, that we call a vertex of a polytope smooth when the affine open set U_v is smooth (or equivalently an affine space), and we call it singular otherwise. Now we can state the main results of this section:

Theorem 4.24 *Let ∇ be a normal standard binary polytope. If ∇ contains no neighbouring singular vertices then the toric ideal of ∇ is generated by its elements of degree 2.*

Proof. By Corollary 4.3 we have to show that for any $k \geq 3$ and $s \in k\nabla \cap \mathbb{Z}^d$ the relation $\sim_s$ has precisely one equivalence class. Let us first treat the case when $v_1, w_1 \leq s$ are neighbouring vertices of ∇. By the assumption at least one of them - say v_1 - is a smooth vertex. By normality of ∇ there are - not necessarily distinct - vertices $v_2, \dots, v_k, w_2, \dots, w_k$ such that $s = \sum_{i=1}^k v_i = \sum_{i=1}^k w_i$. By the assumption that v_1 and w_1 are neighbours one of the ray generators of $\mathrm{Cone}(\nabla - v_1)$ is $w_1 - v_1$. Let $u_1 - w_1, \dots, u_{dim(\nabla)-1} - v_1$ denote the rest of the ray generators, that together with $w_1 - v_1$ form a $\mathbb{Z}$-basis of the lattice $\mathbb{Z}^d \cap \mathrm{Span}(\mathrm{Cone}(\nabla - v_1))$. Consider the "localized" equation $\sum_{i=2}^k (v_i - v_1) = \sum_{i=1}^k (w_i - v_1)$. Each $v_i - v_1$ and $w_i - w_1$ decomposes uniquely as a linear combination of the ray generators of $\text{Cone}(\nabla - v_1)$ with non-negative integer coefficients, so for the two sides to be equal the decomposition of one of the $v_i - v_1$ has to contain $w_1 - v_1$ with a positive coefficient. It follows that for some i we have an equation $v_i + jv_1 = w_1 + \sum_{l=1}^{j} u_{k_j}$, implying that $v_i + jv_1 \geq w_1$. On the other hand since all of the lattice points of ∇ are 0-1 vectors it follows that $v_i + v_1 \geq w_1$, and consequently that for any $j \notin \{1, i\}$ we have $v_j + w_1 \leq s$ and hence $v_j \sim_s w_1$. Since $v_j \sim_s v_1$ for all j, we also have $v_1 \sim_s w_1$. We have established that for any s of degree at least 3 if $v \leq s$ and $w \leq s$ are neighbouring vertices then we have $v \sim_s w$. Now we are done by Lemma 4.23. □

Example 4.25 It can also be deduced from the proof of Theorem 4.24 that for a standard binary polytope ∇ whenever $s \in k\nabla$ for $k \geq 3$ is such that there is a smooth vertex $v \in \nabla$ satisfying $v \leq s$, the relation $\sim_s$ has only one equivalence class. However as shown by the following example, for any $n \in \mathbb{N}$ there is a standard binary polytope with at least n smooth vertices, such that its toric ideal is not generated in degree 2.

Consider the complete bipartite quiver $K(3,3)$ and let us write $v_{1,2,3}$ for its sinks, $u_{1,2,3}$ for its sources and $a_{i,j}$ for the arrow pointing from v_i to u_j. Let Q be a quiver we obtain from $K(3,3)$ after adding a new vertex w, arrows $b_{1,2,3}$ from $v_{1,2,3}$ to w, and arrows $c_{1,2,3}$ from w to $u_{1,2,3}$. Set $\theta(v_{1,2,3}) = -1$, $\theta(u_{1,2,3}) = 1$ and $\theta(w) = 0$. The polytope $\nabla(Q, \theta)$ is binary since every arrow is incident to a source of weight -1 or a sink of weight 1. Denoting by σ_{ijk} the lattice points of $\nabla(Q, \theta)$ that correspond to perfect matchings of $K(3,3)$ as in the proof of Theorem 4.21, and setting $s = \sigma_{123} + \sigma_{312} + \sigma_{231} = \sigma_{132} + \sigma_{312} + \sigma_{213}$, we see that $\sim_s$ has two equivalence classes since $\nabla(Q, \theta) \cap \{x \leq s\}$ contains no lattice points other than the σ_{ijk}. On the other hand consider the vertex $m \in \nabla(Q, \theta)$ defined as $m(b_{1,2,3}) = m(c_{1,2,3}) = 1$ and $m(a_{i,j}) = 0$ for all $i, j \in \{1, 2, 3\}$. The quiver Q^m consists of a single vertex with loops, hence $U_m \cong \mathcal{M}(Q^m, 0)$ is smooth, so m is a smooth vertex. Now adding multiple copies of the arrows $b_{1,2,3}$ and $c_{1,2,3}$ we can obtain an arbitrarily large amount of smooth vertices in $\nabla(Q, \theta)$.

Theorem 4.26 *Let ∇ be a normal standard binary polytope. If ∇ has at most one singular vertex then the toric ideal of ∇ has a quadratic Gröbner basis.*

Proof. Let us denote by $v_1, \dots, v_k$ the vertices of ∇, let $R = \mathbb{C}[t^{v_1}, \dots, t^{v_k}]$ and φ denote the surjection $R \twoheadrightarrow \mathbb{C}[S(\nabla)]$ defined by $\varphi(t^{v_i}) = x^{v_i}$. Let $\leq$ denote the lexicographical order on the vertices of ∇, i.e. $v_i > v_k$ if and only if for the smallest j such that $v_i(j) \neq v_k(j)$ we have $v_i(j) > v_k(j)$. If ∇ contains a singular vertex v, after a possible renumbering of the coordinates we can can assume that $supp(v) = \{1, \dots, b\}$ for some integer b. Since in a standard binary polytope the supports of distinct vertices can not contain each other, this assumption implies that v is maximal with respect to $\leq$. Now we define the monomial ordering $\preceq$ on R, such that for two monomials $m_1 = \prod_i (t^{v_i})^{l_i}$ and $m_2 = \prod_i (t^{v_i})^{j_i}$ we have $m_2 \preceq m_1$ if and only if either $\deg(m_1) > \deg(m_2)$ or $\deg(m_1) = \deg(m_2)$ and for the smallest v_i (with respect to the ordering $\leq$) such that $l_i \neq j_i$ we have $j_i > l_i$.

We claim that the degree 2 elements of the ideal of relations $\ker(\varphi)$ are a Gröbner basis under the ordering $\preceq$. We call a monomial *initial* if it is the initial monomial with respect to $\preceq$ of some element of $\ker(\varphi)$. The remaining monomials are called *standard*. Record that a monomial m is standard if and only if it is minimal with respect to $\preceq$ amongst the monomials n with $\varphi(m) = \varphi(n)$. To prove the claim it is sufficient to show that if m is an initial monomial with $\deg(m) \geq 2$, then m is divisible by an initial monomial of degree 2. We will prove this by induction on the degree. The case $\deg(m) = 2$ is trivial.

Now assume that $m = \prod_{i=1}^{j} t^{w_i}$ is an initial monomial, where the w_i are not necessarily distinct vertices of ∇ satisfying $t^{w_1} \succeq t^{w_2} \succeq \cdots \succeq t^{w_j}$ and $j \geq 3$. Note that w_j is a smooth vertex, since if it was the unique singular vertex then - since we set it to be maximal with respect to the ordering $\leq$ - we would have $w_1 = \cdots = w_j$ and since $supp(w_j)$ does not contain the support of any other vertex $m = (t^{w_j})^j$ would be the unique monomial that maps to $\varphi(m)$ contradicting that it is initial. If t^{w_j} is minimal with respect to $\preceq$ in the set $\{t^{v_i} : x^{v_i} | \varphi(m)\}$ then $\prod_{i=1}^{j-1} t^{w_i}$ can not be standard, otherwise – by the definition of $\preceq$ and the characterization of standard monomials mentioned above – m would also be standard. Thus $\prod_{i=1}^{j-1} t^{w_i}$ is an initial monomial, hence is divisible by a degree 2 initial monomial by the induction hypothesis, implying in turn that m is divisible by a degree 2 initial monomial. It remains to deal with the case when there is a vertex $v' \in \nabla$ such that $t^{v'} \prec t^{w_j}$ and $x^{v'} | \varphi(m)$. Let c be the smallest integer such that $w_j(c) > v'(c)$. Denoting the neighbouring vertices of w_j by $u_1, \ldots, u_d$, by smoothness of w_j we have that $v' - w_j$ can be uniquely written as $\sum_{i=1}^{d} \alpha_i (u_i - w_j)$ where the coefficients α_i are non-negative integers. Since all the vertices are 0-1 vectors it follows that whenever $\alpha_i \neq 0$ we have that $\mathrm{supp}(u_i) \subseteq \mathrm{supp}(w_j) \cup \mathrm{supp}(v')$. Moreover since $w_j(c) = 1$ and $v'(c) = 0$ there is an l such that $\alpha_l \neq 0$ and $u_l(c) = 0$. Since for $b < c$ we have that $u_l(b) = w_j(b) = v'(b)$, it follows that $t^{u_l} \prec t^{w_j}$, and by $\mathrm{supp}(u_i) \subseteq \mathrm{supp}(w_j) \cup \mathrm{supp}(v')$ we have that $x^{u_l} | \varphi(m)$. Note that u_l and w_j are neighbouring vertices. Applying the argument from the proof of Proposition 4.24 for the pair (w_j, u_l) instead of (v_1, w_1) we conclude that for some $r < j$ we have $x^{u_l} | \varphi(t^{w_j} t^{w_r})$. Now $t^{u_l} \prec t^{w_j} \preceq t^{w_r}$, implying that the degree 2 monomial $t^{w_j} t^{w_r}$ is not standard. We are done as $t^{w_j} t^{w_r} | m$. □

A finitely generated graded k-algebra A is called a *Koszul algebra* if the ground field k has a linear graded free resolution over A. By a result of Priddy from [38], the existence of a quadratic Gröbner basis of the ideal of relations is a satisfactory condition for the Koszul property of the algebra, so we also have the following corollary:

Corollary 4.27 *Let ∇ be a normal standard binary polytope. If ∇ has at most one singular vertex then $\mathbb{C}[S(\nabla)]$ is Koszul.*

Example 4.28 Consider the following quiver Q:

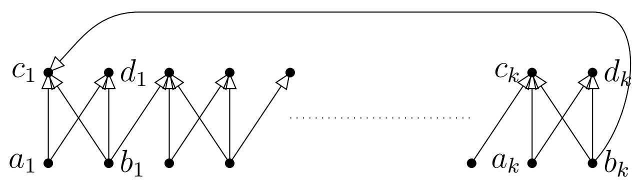

Setting the weight θ to be -1 on each source and 1 on each sink, the vertices of $\nabla(Q,\theta)$ correspond to the perfect matchings of the underlying undirected graph of Q. We can obtain 2^k perfect matchings by matching the vertices a_i and b_i with c_i and d_i, let us denote the corresponding vertices in $\nabla(Q,\theta)$ by $v_{i_1,\dots,i_k}$ where $i_j = 1$ if a_j is matched with c_j and $i_j = 2$ if a_j is matched with d_j. Aside of these there is only one more perfect matching where all of the a_i are matched with d_i and b_i is matched with c_{i+1} for $i < k$ and b_k is matched with c_1, let us denote the vertex corresponding to this matching by w. $\nabla(Q,\theta)$ is a $k+1$ dimensional binary polytope and applying Proposition 3.26 it is easy to verify that w is the only singular vertex. Therefore by Theorem 4.26 it admits a quadratic Gröbner basis and in particular its ideal of relations is generated by degree 2 elements. Indeed one can check that all relations are generated by the $\binom{k}{2}$ binomials of the type $t^{v_{1,1,i_3,\dots,i_k}}t^{v_{2,2,i_3,\dots,i_k}} - t^{v_{1,2,i_3,\dots,i_k}}t^{v_{2,1,i_3,\dots,i_k}}$.

In the light of Proposition 4.8 it is natural to ask if there is a general degree bound for the toric ideals of normal standard binary lattice polytopes. The answer to this question is negative, we conclude this section by giving an example for each positive integer n of a normal standard lattice polytope, such that the corresponding ideal of relations is not generated by its elements of degree less than n.

Fix $n \geq 2$ and let $K(n,n)$ denote the complete bipartite graph on $2n$ vertices. Write the bipartition of $K(n,n)$ as $K(n,n)_0 = V \coprod W$, where $|V| = |W| = n$ and every $v \in V$ and $w \in W$ is connected by an edge. Index the coordinates of $\mathbb{R}^{n^2}$ by $K(n,n)_1$. Let Π

denote the set of the $n!$ perfect matchings of $K(n,n)$. Consider the polytope defined as

$$\mathcal{P}_n := \{x \in \mathbb{R}^{n^2} \mid x \geq 0 \quad \forall P \in \Pi : \sum_{e \in P} x(e) = 1\}$$

For each $v \in K(n,n)_0$ define the lattice point m_v:

$$m_v(e) = \begin{cases} 1 & \text{if } e \textit{ is incident to } v \\ 0 & \textit{otherwise} \end{cases}$$

Lemma 4.29 *For any integer $k \geq 1$ and point $x \in k\mathcal{P}_n$ there is a $v \in K(n,n)_0$ such that* $\mathrm{supp}(m_v) \subseteq \mathrm{supp}(x)$.

Proof. First note that for any $x \in k\mathcal{P}_n$ and cycle of $K(n,n)$ with edges $c_1, \dots c_{2i}$, indexed consecutively along the cycle, we have that $x(c_1) + x(c_3) + \dots x(c_{2i-1}) = x(c_2) + x(c_4) + \dots x(c_{2i})$. Otherwise consider a perfect matching P_1 of $K(n,n)$ containing the edges $c_1, c_3, \dots c_{2i-1}$ and set P_2 be the perfect matching we obtain from P_1 by replacing c_j by c_{j+1} for each $j < 2i$. For P_1 and P_2 we have $\sum_{e \in P_1} x(e) \neq \sum_{e \in P_2} x(e)$ contradicting $x \in k\mathcal{P}_n$. Now suppose that the conclusion of the lemma does not hold for some x and set $F := K(n,n)_1 \setminus \mathrm{supp}(x)$. Note that every vertex is incident to at least one edge in F. Pick an edge $e \in \mathrm{supp}(x)$ with endpoints $v \in V$ and $w \in W$. Choose edges $f_v, f_w \in F$ that are incident to v and w respectively and denote by e' the edge that connects the other endpoints of f_v and f_w. Now the edges e, f_v, e', f_w form a cycle of length 4 such that the second and the fourth edge are not in $\mathrm{supp}(x)$ but the first edge is in $\mathrm{supp}(x)$, contradicting the observation made at the beginning of the proof. $\square$

Proposition 4.30 *(i) $\mathcal{P}_n$ is a normal standard binary polytope with vertex set $\{m_v \mid v \in K(n,n)_0\}$.*

(ii) The toric ideal of $\mathcal{P}_n$ is not generated by its elements of degree at most $n-1$.

Proof. Suppose x is a vertex of $\mathcal{P}_n$. By Lemma 4.29 we have a vertex v with $\mathrm{supp}(m_v) \subseteq \mathrm{supp}(x)$. Set $\lambda := \min\{x(e) \mid e \in \mathrm{supp}(m_v)\}$. If $x \neq m_v$ then $0 < \lambda < 1$, and we have that $(1-\lambda)^{-1}(x - \lambda m_v) \in \nabla$. Since x is then an interior point of the line segment between m_v and $(1-\lambda)^{-1}(x - \lambda m_v)$ it can not be a vertex. This shows us that $\mathcal{P}_n$ is indeed a lattice polytope with vertex set $\{m_v \mid v \in V(K(n,n))\}$ and it follows from the definition that $\mathcal{P}_n$ is standard and binary. For showing normality take a lattice point $m \in k\mathcal{P}_n$ for $k \geq 2$.

By Lemma 4.29 we have a vertex v with $\mathrm{supp}(m_v) \subseteq \mathrm{supp}(m)$. Since m is a lattice point we have that $m(e) \geq 1$ for all $e \in \mathrm{supp}(m_v)$, so $m - m_v \in (k-1)\mathcal{P}_n$ and we are done by induction, completing the proof of (i).

For (ii) set the notation $V = \{v_1, \dots, v_n\}$ and $W = \{w_1, \dots, w_n\}$ for the vertices in the two sides of the bipartition of $K(n, n)$. We have the equality $s = m_{v_1} + \cdots + m_{v_n} = m_{w_1} + \cdots + m_{w_n} \in n\mathcal{P}_n$. Clearly for any $i, j \in \{1, \dots, n\}$ we have $m_{v_i} + m_{w_j} \not\leq s$, hence $\sim_s$ has two equivalence classes: $m_{v_1}, \dots, m_{v_n}$ and $m_{w_1}, \dots, m_{w_n}$. Now we are done by Corollary 4.3. $\square$

Appendix A

List of 3-dimensional reflexive quiver polytopes

The following tables contain data of the prime 3-dimensional reflexive quiver polytopes, which we obtained from the computations explained in Section 3.4.

Q^I

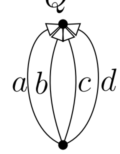

Val. 2 sinks	#Lattice Points	#Facets	#Vertices	#Smooth Vertices
none	35	4	4	4
a	31	5	6	6
a,b	27	6	7	6
a,b,c	23	7	7	4
a,b,c,d	19	8	6	0

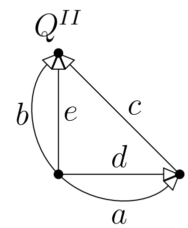

Val. 2 sinks	#Lattice Points	#Facets	#Vertices	#Smooth Vertices
none	30	5	6	6
a	26	6	8	8
b	28	6	8	8
c	25	6	8	8
a,b	24	7	9	8
a,c	23	7	10	10
a,d	22	7	9	8
b,e	26	7	10	10
a,b,c	21	8	11	10
a,b,d	20	8	9	6
a,b,e	22	8	10	8
a,c,d	21	8	12	12
a,b,c,d	19	9	12	10
a,b,d,e	18	9	9	4
a,b,c,d,e	17	10	12	8

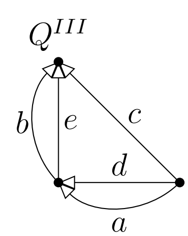

Val. 2 sinks	#Lattice Points	#Facets	#Vertices	#Smooth Vertices
none	30	5	5	4
a	26	6	7	6
c	29	6	8	8
a,b	23	7	8	6
a,c	25	7	10	10
a,b,c	22	8	11	10

Q^{IV}

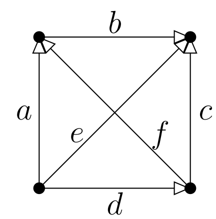

Val. 2 sinks	#Lattice Points	#Facets	#Vertices	#Smooth Vertices
none	26	6	7	6
a	24	7	9	8
b	22	7	8	6
d	25	7	9	8
e	22	7	9	8
a,b	21	8	10	8
a, c	22	8	10	8
a,d	20	8	10	8
a,e	23	8	11	10
b, d	19	8	8	4
e,f	21	8	10	8
a,b,c	19	9	11	8
a,b,e	20	9	12	10
a,c,e	21	9	12	10
a,c,f	18	9	10	6
a,d,e	19	9	12	10
a,d,f	18	9	11	8
a,b,c,d	17	10	11	6
a,b,c,e	18	10	13	10
a,b,d,e	17	10	12	8
a,b,c,d,e	16	11	13	8
a,b,c,d,e,f	15	12	14	8

Bibliography

[1] J. Adriaenssens and L. Le Bruyn, Local quivers and stable representations, Comm. Alg. 31 (2003), 1777-1797.

[2] K. Altmann and L. Hille, Strong exceptional sequences provided by quivers, Algebras and Represent. Theory 2 (1) (1999), 1-17.

[3] K. Altmann, B. Nill, S. Schwentner, and I. Wiercinska, Flow polytopes and the graph of reflexive polytopes, Discrete Math. 309 (2009), no. 16, 4992-4999.

[4] K. Altmann and D. van Straten, Smoothing of quiver varieties, Manuscripta Math. 129 (2009), no. 2, 211-230.

[5] V. V. Batyrev, Dual polyhedra and mirror symmetry for Calabi-Yau hypersurfaces in toric varieties, J. Algebr. Geom. 3, 493-535 (1994)

[6] Bocklandt, R., Quiver quotient varieties and complete intersections, Algebras and Representation Theory (2005) 8: 127-145.

[7] R. Bocklandt, Smooth quiver representation spaces, J. Algebra 253(2) (2002), 296-313.

[8] K. Bongartz, Some geometric aspects of representation theory, Algebras and Modules I, CMS Conf. Proc. 23 (1998), 1-27, (I. Reiten, S.O. Smalø, Ø. Solberg, editors).

[9] W. Bruns, The quest for counterexamples in toric geometry, Proc. CAAG 2010, Ramanujan Math. Soc. Lect. Notes Series No. 17 (2013), 1–17.

[10] W. Bruns and J. Gubeladze, Semigroup algebras and discrete geometry, Séminaires Congrès 6 (2002), 43-127.

[11] A. T. Carroll, C. Chindris and Z. Lin, Quiver representations of constant Jordan type and vector bundles, arXiv: 1402.2568

[12] D. Cox, J. Little, and H. Schenck, Toric Varieties, Amer. Math. Soc., Providence, Rhode Island, 2010.

[13] A. Craw and G. G. Smith, Projective toric varieties as fine moduli spaces of quiver representations, Amer. J. Math. 130 (2008), 1509-1534.

[14] P. Diaconis and N. Eriksson, Markov bases for noncommutative Fourier analysis of ranked data, J. of Symb. Comp. 41 (2006), 182-195.

[15] M. Domokos, On singularities of quiver moduli, Glasgow Math. J. 53 (2011), 131-139.

[16] M. Domokos and D. Joó, On the equations and classification of toric quiver varieties, Proceedings A of The Royal Society of Edinburgh, accepted

[17] M. Domokos and A. N. Zubkov, Semi-invariants of quivers as determinants, Transform. Groups 6, No. 1 (2001), 9-24.

[18] M. Domokos and A. N. Zubkov, Semisimple representations of quivers in characteristic p, Algebr. Represent. Theory 5 (2002), 305-317.

[19] N. Epstein and H. D. Nguyen, Algebra retracts and Stanley-Reisner rings, J. Pure Appl. Algebra 218 (2014), 1665-1682.

[20] J. Fei, Moduli of representations I. Projections from quivers, arXiv:1011.6106v3

[21] S. Greco and R. Strano: Complete Intersections, Lecture Notes in Math. 1092, Springer, Berlin, 1984.

[22] C. Haase and A. Paffenholz, Quadratic Gröbner bases for smooth 3×3 transportation polytopes, J. Algebr. Comb. 30 (2009), 477-489.

[23] M. Hatanaka, Uniqueness of the direct decomposition of toric manifolds, Osaka J. Math. 52 (2015), 439-453.

[24] T. Hibi, K. Matsuda and H. Ohsugi, Strongly Koszul Edge Rings, Acta Mathematica Vietnamica (2014), 1-8.

[25] L. Hille, Moduli spaces of thin sincere representations of quivers, preprint, Chemnitz, 1995.

[26] L. Hille, Toric quiver varieties, pp. 311-325, Canad. Math. Soc. Conf. Proc. 24, Amer. Math. Soc., Providence, RI, 1998.

[27] L. Hille, Quivers, cones and polytopes, Linear Alg. Appl. 365 (2003), 215-237.

[28] A. Ishii and K. Ueda, On moduli spaces of quiver representations associated with dimer models. Higher dimensional algebraic varieties and vector bundles, Res. Inst. Math. Sci. (2008), 127-141.

[29] D. Joó, Complete intersection quiver settings with one dimensional vertices, Algebr. Represent. Theory 16 (2013), 1109-1133.

[30] V. G. Kac, Infinite root systems, representations of graphs and invariant theory, Invent. Math. 56 (1980), 57-92.

[31] A. D. King, Moduli of representations of finite dimensional algebras, Quart. J. Math. Oxford (2), 45 (1994), 515-530.

[32] M. Lason and M. Michalek, On the toric ideal of a matroid, Advances in Mathematics 259 (2014), 1-12.

[33] L. Le Bruyn and C. Procesi, Semisimple representations of quivers, Trans. Amer. Math. Soc. 317 (1990), 585-598.

[34] M. Lenz, Toric ideals of flow polytopes, 22nd International Conference on Formal Power Series and Algebraic Combinatorics (FPSAC 2010), 889-896, Discrete Math. Theor. Comput. Sci. Proc., AN, Assoc. Discrete Math. Theor. Comput. Sci., Nancy, 2010.

[35] M. Lenz, Toric ideals of flow polytopes arXiv:0801.0495v3

[36] P. E. Newstead, Introduction to Moduli Problems and Orbit Spaces, Tata Institute Lecture Notes, Springer-Verlag, 1978.

[37] H. Ohsugi, Toric Ideals and an Infinite Family of Normal (0, 1)-Polytopes without Unimodular Regular Triangulations, Disc.and Comp. Geom. 27 (2002), 551–565

[38] S. B. Priddy, Koszul resolutions, Trans. Amer. Math. Soc. 152 (1970), 39-60.

[39] P. Renteln, On the Spectrum of the Derangement Graph, Electron. J. Combin. 14 (2007), Research Paper 82, 17 pp. electronic

[40] A. Schofield, General representations of quivers, Proc. London. Math. Soc. (3) 65 (1992), 46–64.

[41] A. Schrijver, Combinatorial Optimization – Polyhedra and Efficiency. Number 24A in Algorithms and Combinatorics, Springer, Berlin, 2003.

[42] A. Schrijver, Theory of Linear and Integer Programming, Wiley-Interscience Series in Discrete Mathematics, Chichester, 1986.

[43] A. Skowroński and J. Weyman, The algebras of semi-invariants of quivers, Transformation Groups, 5 (2000), No. 4., 361-402.

[44] B. Sturmfels, Gröbner Bases and Convex Polytopes, University Lecture Series 8, AMS, Providence, Rhode Island, 1996.

[45] M. Thaddeus, Geometric invariant theory and flips, J. Amer. Math. Soc. 9 (1996), 691-723.

[46] T. Yamaguchi, M. Ogawa and A. Takemura, Markov degree of the Birkhoff model, J. Alg. Comb. 40 (2014), 293-311.