Affine Mirković-Vilonen polytopes

Pierre Baumann, Joel Kamnitzer and Peter Tingley

To the memory of Andrei V. Zelevinsky

Abstract

Each integrable lowest weight representation of a symmetrizable Kac-Moody Lie algebra $\mathfrak{g}$ has a crystal in the sense of Kashiwara, which describes its combinatorial properties. For a given $\mathfrak{g}$, there is a limit crystal, usually denoted by $B(-\infty)$, which contains all the other crystals. When $\mathfrak{g}$ is finite dimensional, a convex polytope, called the Mirković-Vilonen polytope, can be associated to each element in $B(-\infty)$. This polytope sits in the dual space of a Cartan subalgebra of $\mathfrak{g}$, and its edges are parallel to the roots of $\mathfrak{g}$. In this paper, we generalize this construction to the case where $\mathfrak{g}$ is a symmetric affine Kac-Moody algebra. The datum of the polytope must however be complemented by partitions attached to the edges parallel to the imaginary root δ. We prove that these decorated polytopes are characterized by conditions on their normal fans and on their 2-faces. In addition, we discuss how our polytopes provide an analog of the notion of Lusztig datum for affine Kac-Moody algebras. Our main tool is an algebro-geometric model for $B(-\infty)$ constructed by Lusztig and by Kashiwara and Saito, based on representations of the completed preprojective algebra Λ of the same type as $\mathfrak{g}$. The underlying polytopes in our construction are described with the help of Buan, Iyama, Reiten and Scott's tilting theory for the category Λ-mod. The partitions we need come from studying the category of semistable Λ-modules of dimension-vector a multiple of δ.

1 Introduction

Let A be a symmetrizable generalized Cartan matrix, with rows and columns indexed by a set I. We denote by $\mathfrak{g}$ the Kac-Moody algebra defined by A. It comes with a triangular decomposition $\mathfrak{g} = \mathfrak{n}_- \oplus \mathfrak{h} \oplus \mathfrak{n}_+$, with a root system Φ, and with a Weyl group W. The simple roots α_i are indexed by I and the group W is a Coxeter system, generated by the simple reflections s_i. We denote the length function of W by $\ell : W \to \mathbb{N}$ and the set of positive (respectively, negative) roots by Φ_+ (respectively, Φ_-). The root lattice is denoted by $\mathbb{Z}I = \bigoplus_{i \in I} \mathbb{Z}\alpha_i$ and we set $\mathbb{R}I = \mathbb{Z}I \otimes_{\mathbb{Z}} \mathbb{R}$. The canonical pairing between $\mathbb{R}I$ and its dual $(\mathbb{R}I)^*$ will be denoted by angle brackets. Lastly, we denote by $\mathbb{R}_{\geq 0}I$ the set of linear combinations of the simple roots with nonnegative coefficients and we set $\mathbb{N}I = \mathbb{Z}I \cap \mathbb{R}_{\geq 0}I$.

1.1 Crystals

The combinatorics of the representation theory of $\mathfrak{g}$ is captured by Kashiwara's theory of crystals. Let us summarize quickly this theory; we refer the reader to the nice survey [37] for detailed explanations.

A $\mathfrak{g}$-crystal is a set B endowed with maps wt, ε_i, φ_i, $\tilde{e}_i$ and $\tilde{f}_i$, for each $i \in I$, that satisfy certain axioms. This definition is of combinatorial nature and the axioms stipulate the local behavior of the structure maps around an element $b \in B$. This definition is however quite permissive, so one wants to restrict to crystals that actually come from representations.

In this respect, an important object is the crystal $B(-\infty)$, which contains the crystals of all the irreducible lowest weight integrable representations of $\mathfrak{g}$ (see Theorem 8.1 in [37]). This crystal contains a lowest weight element $u_{-\infty} \in B(-\infty)$ annihilated by all the lowering operators $\tilde{f}_i$, and any element of $B(-\infty)$ can be obtained by applying a sequence of raising operators $\tilde{e}_i$ to $u_{-\infty}$.

The crystal $B(-\infty)$ itself is defined as a basis of the quantum group $U_q(\mathfrak{n}_+)$ in the limit $q \to 0$. Working with this algebraic construction is cumbersome, and there exist other, more handy, algebro-geometric or combinatorial models for $B(-\infty)$.

One of these combinatorial models is Mirković-Vilonen (MV) polytopes. In this model, proposed by Anderson [2], one associates a convex polytope $\mathrm{Pol}(b) \subseteq \mathbb{R}I$ to each element $b \in B(-\infty)$. The construction of $\mathrm{Pol}(b)$ is based on the geometric Satake correspondence. More precisely, the affine Grassmannian of the Langlands dual of $\mathfrak{g}$ contains remarkable subvarieties, called MV cycles after Mirković and Vilonen [47]. There is a natural bijection $b \mapsto Z_b$ from $B(-\infty)$ onto the set of all MV cycles [11, 12, 23], and $\mathrm{Pol}(b)$ is simply the image of Z_b by the moment map.

Using Berenstein and Zelevinsky's work [9], the second author showed in [36] that these MV polytopes can be described in a completely combinatorial fashion: these are the convex lattice polytopes whose normal fan is a coarsening of the Weyl fan in the dual of $\mathbb{R}I$, and whose 2-faces have a shape constrained by the tropical Plücker relations. In addition, the length of the edges of $\mathrm{Pol}(b)$ is given by the Lusztig data of b, which indicate how b, viewed as a basis element of $U_q(\mathfrak{g})$ at the limit $q \to 0$, compares with the PBW bases.

1.2 Generalization to the affine case

This paper aims at generalizing this model of MV polytopes to the case where $\mathfrak{g}$ is an affine Kac-Moody algebra.

Obstacles pop up when one tries to generalize the above constructions of $\mathrm{Pol}(b)$ to the affine case. Despite difficulties in defining the double-affine Grassmannian, the algebro-geometric model of $B(-\infty)$ using MV cycles still exists in the affine case, thanks to Braverman, Finkelberg and Gaitsgory's work [11]; however, there is no obvious way to go from MV cycles to MV polytopes.

On the algebraic side, several PBW bases for $U_q(\mathfrak{n}_+)$ have been defined in the affine case by Beck [6], Beck and Nakajima [8], and Ito [30], but the relationship between the different Lusztig data they provide has not been studied[1].

As recalled above, in finite type, the normal fan of an MV polytope is a coarsening of the Weyl fan, so the facets of an MV polytope are orthogonal to the rays in the Weyl fan. Therefore an MV polytope is determined just by the position of theses facets, which form a set of numerical values dubbed "Berenstein-Zelevinsky (BZ) data". In the case $\mathfrak{g} = \widehat{\mathfrak{sl}}_n$, a combinatorial model for an analog of these BZ data was introduced by Naito, Sagaki, and Saito [51, 52]. Later, Muthiah [49] related this combinatorial model to the geometry of the MV cycles. However, the complete relationship between this combinatorial model and our affine MV polytopes is not yet clear.

1.3 The preprojective model

Due to the difficulties in the MV cycle and PBW bases models, we are led to use a third construction of $\mathrm{Pol}(b)$, recently obtained by the first two authors for the case of a finite dimensional $\mathfrak{g}$ [5]. This construction uses a geometric model for $B(-\infty)$ based on quiver varieties, which we now recall.

This model exists for any Kac-Moody algebra $\mathfrak{g}$ (not necessarily of finite or affine type) but only when the generalized Cartan matrix A is symmetric. Then $2\,\mathrm{id} - A$ is the incidence matrix of the Dynkin graph (I, E); here our index set I serves as the set of vertices and E is the set of edges. Choosing an orientation of this graph yields a quiver Q, and one can then define the completed preprojective algebra Λ of Q.

[1] Recently, Muthiah and Tingley [50] have considered this problem in the case $\mathfrak{g} = \widehat{\mathfrak{sl}}_2$. They have shown that the resulting combinatorics matches that produced in the present paper, in the sense that the MV polytopes coming from the Lusztig data provided by the PBW bases match those defined here. It should be easy to extend this result to the case of an arbitrary symmetric affine Kac-Moody algebra.

A Λ-module is an I-graded vector space equipped with linear maps. If the dimension-vector is given, we can work with a fixed vector space; the datum of a Λ-module then amounts to the family of linear maps, which can be regarded as a point of an algebraic variety. This variety is called Lusztig's nilpotent variety; we denote it by $\Lambda(\nu)$, where $\nu \in \mathbb{N}I$ is the dimension-vector. Abusing slightly the language, we often view a point $T \in \Lambda(\nu)$ as a Λ-module.

For $\nu \in \mathbb{N}I$, let $\mathfrak{B}(\nu)$ be the set of irreducible components of $\Lambda(\nu)$. We set $\mathfrak{B} = \bigsqcup_{\nu\in\mathbb{N}I} \mathfrak{B}(\nu)$. In [42], Lusztig endows $\mathfrak{B}$ with a crystal structure, and in [38], Kashiwara and Saito show the existence of an isomorphism of crystals $b \mapsto \Lambda_b$ from $B(-\infty)$ onto $\mathfrak{B}$. This isomorphism is unique since $B(-\infty)$ has no non-trivial automorphisms.

Given a finite-dimensional Λ-module T, we can consider the dimension-vectors of the Λ-submodules of T; they are finitely many, since they belong to a bounded subset of the lattice $\mathbb{Z}I$. The convex hull in $\mathbb{R}I$ of these dimension-vectors will be called the Harder-Narasimhan (HN) polytope of T and will be denoted by $\operatorname{Pol}(T)$.

The main result of [5] is equivalent to the following statement: if $\mathfrak{g}$ is finite dimensional, then for each $b \in B(-\infty)$, the set $\{T \in \Lambda_b \mid \operatorname{Pol}(T) = \operatorname{Pol}(b)\}$ contains a dense open subset of Λ_b. In other words, $\operatorname{Pol}(b)$ is the general value of the map $T \mapsto \operatorname{Pol}(T)$ on Λ_b.

This result obviously suggests a general definition for MV polytopes. We will however see that for $\mathfrak{g}$ of affine type, another piece of information is needed to have an complete model for $B(-\infty)$; namely, we need to equip each polytope with a family of partitions. Our task now is to explain what our polytopes look like, and where these partitions come from.

1.4 Faces of HN polytopes

Choose a linear form $\theta : \mathbb{R}I \to \mathbb{R}$ and let $\psi_{\operatorname{Pol}(T)}(\theta)$ denote the maximum value of θ on $\operatorname{Pol}(T)$. Then $P_\theta = \{x \in \operatorname{Pol}(T) \mid \langle\theta, x\rangle = \psi_{\operatorname{Pol}(T)}(\theta)\}$ is a face of $\operatorname{Pol}(T)$. Moreover, the set of submodules $X \subseteq T$ whose dimension-vectors belong to P_θ has a smallest element $T_\theta^{\min}$ and a largest element $T_\theta^{\max}$.

The existence of $T_\theta^{\min}$ and $T_\theta^{\max}$ follows from general considerations: if we define the slope of a finite dimensional Λ-module X as $\langle\theta, \underline{\dim}\, X\rangle / \dim X$, then $T_\theta^{\max}/T_\theta^{\min}$ is the semistable subquotient of slope zero in the Harder-Narasimhan filtration of T. Introducing the abelian subcategory $\mathscr{R}_\theta$ of semistable Λ-modules of slope zero, it follows that, for each submodule $X \subseteq T$,

$$\underline{\dim}\,(X) \in P_\theta \iff \Big(T_\theta^{\min} \subseteq X \subseteq T_\theta^{\max} \quad \text{and} \quad X/T_\theta^{\min} \in \mathscr{R}_\theta\Big).$$

In other words, the face P_θ coincides with the HN polytope of $T_\theta^{\max}/T_\theta^{\min}$, computed relative to the category $\mathscr{R}_\theta$, and shifted by $\underline{\dim}\, T_\theta^{\min}$.

Our aim now is to describe the normal fan to $\mathrm{Pol}(T)$, that is, to understand how $T_\theta^{\min}$, $T_\theta^{\max}$ and $\mathscr{R}_\theta$ depend on θ. For that, we need tools that are specific to preprojective algebras.

1.5 Tits cone and tilting theory

One of these tools is Buan, Iyama, Reiten and Scott's tilting ideals for Λ [14]. Let S_i be the simple Λ-module of dimension-vector α_i and let I_i be its annihilator, a one-codimensional two-sided ideal of Λ. The products of these ideals I_i are known to satisfy the braid relations, so to each w in the Weyl group of $\mathfrak{g}$, we can attach a two-sided ideal I_w of Λ by the rule $I_w = I_{i_1} \cdots I_{i_\ell}$, where $s_{i_1} \cdots s_{i_\ell}$ is any reduced decomposition of w. Given a finite-dimensional Λ-module T, we denote the image of the evaluation map $I_w \otimes_\Lambda \mathrm{Hom}_\Lambda(I_w, T) \to T$ by T^w.

Recall that the dominant Weyl chamber C_0 and the Tits cone C_T are the convex cones in the dual of $\mathbb{R}I$ defined as

$$C_0 = \{\theta \in (\mathbb{R}I)^* \mid \forall i \in I,\ \langle \theta, \alpha_i \rangle > 0\} \quad \text{and} \quad C_T = \bigcup_{w \in W} w\,\overline{C_0}.$$

We will show the equality $T_\theta^{\min} = T_\theta^{\max} = T^w$ for any finite dimensional Λ-module T, any $w \in W$ and any linear form $\theta \in wC_0$. This implies that $\underline{\dim}\,T^w$ is a vertex of $\mathrm{Pol}(T)$ and that the normal cone to $\mathrm{Pol}(T)$ at this vertex contains wC_0. This also implies that $\mathrm{Pol}(T)$ is contained in

$$\{x \in \mathbb{R}I \mid \forall \theta \in wC_0,\ \langle \theta, x \rangle \leq \langle \theta, \underline{\dim}\,T^w \rangle\} = \underline{\dim}\,T^w - w\big(\mathbb{R}_{\geq 0} I\big).$$

When θ runs over the Tits cone, it generically belongs to a chamber, and we have just seen that in this case, the face P_θ is a vertex. When θ lies on a facet, P_θ is an edge (possibly degenerate). More precisely, if θ lies on the facet that separates the chambers wC_0 and ws_iC_0, with say $\ell(ws_i) > \ell(w)$, then $(T_\theta^{\min}, T_\theta^{\max}) = (T^{ws_i}, T^w)$. Results in [1] and [27] moreover assert that T^w/T^{ws_i} is the direct sum of a finite number of copies of the Λ-module $I_w \otimes_\Lambda S_i$.

There is a similar description when θ is in $-C_T$; here the submodules T_w of T that come into play are the kernels of the coevaluation maps $T \to \mathrm{Hom}_\Lambda(I_w, I_w \otimes_\Lambda T)$, where again $w \in W$.

1.6 Imaginary edges and partitions (in affine type)

From now on in this introduction, we focus on the case where $\mathfrak{g}$ is of symmetric affine type, which in particular implies $\mathfrak{g}$ is of untwisted affine type.

The root system for $\mathfrak{g}$ decomposes into real and imaginary roots $\Phi = \Phi^{\mathrm{re}} \sqcup (\mathbb{Z}_{\neq 0}\,\delta)$; the real roots are the conjugate of the simple roots under the Weyl group action, whereas the imaginary roots are fixed under this action. The Tits cone is $C_T = \{\theta : \mathbb{R}I \to \mathbb{R} \mid \langle \theta, \delta \rangle > 0\} \cup \{0\}$.

We set $\mathfrak{t}^* = \mathbb{R}I/\mathbb{R}\delta$. The projection $\pi : \mathbb{R}I \to \mathfrak{t}^*$ maps Φ^{re} onto the "spherical" root system Φ^s, whose Dynkin diagram is obtained from that of $\mathfrak{g}$ by removing an extending vertex. The rank of Φ^s is $r = \dim \mathfrak{t}^*$, which is also the multiplicity of the imaginary roots.

The vector space $\mathfrak{t}$ identifies with the hyperplane $\{\theta : \mathbb{R}I \to \mathbb{R} \mid \langle \theta, \delta \rangle = 0\}$ of the dual of $\mathbb{R}I$. The root system $\Phi^s \subseteq \mathfrak{t}^*$ defines an hyperplane arrangement in $\mathfrak{t}$, called the spherical Weyl fan. The open cones in this fan will be called the spherical Weyl chambers. Together, this fan and the hyperplane arrangement that the real roots define in $C_T \cup (-C_T)$ make up a (non locally finite) fan in the dual of $\mathbb{R}I$, which we call the affine Weyl fan and which we denote by $\mathscr{W}$.

Each set of simple roots in Φ^s is a basis of $\mathfrak{t}^*$; we can then look at the dual basis in $\mathfrak{t}$, whose elements are the corresponding fundamental coweights. We denote by Γ the set of all fundamental coweights, for all possible choices of simple roots. Elements in Γ are called spherical chamber coweights; the rays they span are the rays of the spherical Weyl fan.

Now take a Λ-module T. As we saw in the previous section, the normal cone to $\mathrm{Pol}(T)$ at the vertex $\underline{\dim}\, T^w$ (respectively, $\underline{\dim}\, T_w$) contains wC_0 (respectively, $-w^{-1}C_0$). Altogether, these cones form a dense subset of the dual of $\mathbb{R}I$: this leaves no room for other vertices. This analysis also shows that the normal fan to $\mathrm{Pol}(T)$ is a coarsening of $\mathscr{W}$.

Thus, the edges of $\mathrm{Pol}(T)$ point in directions orthogonal to one-codimensional faces of $\mathscr{W}$, that is, parallel to roots. In the previous section, we have described the edges that point in real root directions. We now need to understand the edges that are parallel to δ. We call these the imaginary edges.

More generally, we are interested in describing the faces parallel to δ. Let us pick $\theta \in \mathfrak{t}$ and let us look at the face $P_\theta = \{x \in \mathrm{Pol}(T) \mid \langle \theta, x \rangle = \psi_{\mathrm{Pol}(T)}(\theta)\}$. As we saw in section 1.4, this face is the HN polytope of $T_\theta^{\max}/T_\theta^{\min}$, computed relative to the category $\mathscr{R}_\theta$. It turns out that $T_\theta^{\max}/T_\theta^{\min}$ and $\mathscr{R}_\theta$ only depend on the face F of the spherical Weyl fan to which θ belongs. We record this fact in the notation by writing $\mathscr{R}_F$ for $\mathscr{R}_\theta$.

We need one more definition: for $\gamma \in \Gamma$, we say that a Λ-module is a γ-core if it belongs to $\mathscr{R}_\theta$ for all $\theta \in \mathfrak{t}$ sufficiently close to γ. In other words, the category of γ-cores is the intersection of the categories $\mathscr{R}_C$, taken over all spherical Weyl chambers C such that $\gamma \in \overline{C}$.

For each $\nu \in \mathbb{N}I$, the set of indecomposable modules is a constructible subset of $\Lambda(\nu)$. It thus makes sense to ask if the general point of an irreducible subset of $\Lambda(\nu)$ is indecomposable. Similarly, the set of modules that belong to $\mathscr{R}_C$ is an open subset of $\Lambda(n\delta)$, so we may ask if the general point of an irreducible subset of $\Lambda(n\delta)$ is in $\mathscr{R}_C$. In section 7.4, we will show the following theorems.

Theorem 1.1 *For each integer $n \geq 1$ and each $\gamma \in \Gamma$, there is a unique irreducible component of $\Lambda(n\delta)$ whose general point is an indecomposable γ-core.*

We denote by $I(\gamma, n)$ this component.

Theorem 1.2 *Let n be a positive integer and let C be a spherical Weyl chamber. There are exactly r irreducible components of $\Lambda(n\delta)$ whose general point is an indecomposable module in $\mathscr{R}_C$. These components are the $I(\gamma, n)$, for $\gamma \in \Gamma \cap \overline{C}$.*

In Theorem 1.2, the multiplicity r of the root $n\delta$ materializes as a number of irreducible components.

Now let $b \in B(-\infty)$ and pick θ in a spherical Weyl chamber C. Let T be a general point of Λ_b and let $X = T_\theta^{\max}/T_\theta^{\min}$, an object in $\mathscr{R}_C$. Write the Krull-Schmidt decomposition of X as $X_1 \oplus \cdots \oplus X_\ell$, with $X_1, \ldots, X_\ell$ indecomposable; then each X_k is in $\mathscr{R}_C$, so $\underline{\dim}\, X_k = n_k\delta$ for a certain integer $n_k \geq 1$. Moreover, it follows from Crawley-Boevey and Schröer's theory of canonical decomposition [18] that each X_k is the general point of an irreducible component $Z_k \subseteq \Lambda(n_k\delta)$. Using Theorem 1.2, we then see that each Z_k is a component $I(\gamma_k, n_k)$ for a certain $\gamma_k \in \Gamma \cap \overline{C}$. Gathering the integers n_k according to the coweights γ_k, we get a tuple of partitions $(\lambda_\gamma)_{\gamma \in \Gamma \cap \overline{C}}$. In this context, we will show that the partition λ_γ depends only on b and γ, and not on the Weyl chamber C.

We are now ready to give the definition of the MV polytope of b: it is the datum $\widetilde{\mathrm{Pol}}(b)$ of the HN polytope $\mathrm{Pol}(T)$, for T general in Λ_b, together with the family of partitions $(\lambda_\gamma)_{\gamma \in \Gamma}$ defined above.

1.7 2-faces of MV polytopes

Let us now consider the 2-faces of our polytopes $\mathrm{Pol}(T)$. Such a face is certainly of the form

$$P_\theta = \{x \in \mathrm{Pol}(T) \mid \langle \theta, x \rangle = \psi_{\mathrm{Pol}(T)}(\theta)\},$$

where θ belongs to a 2-codimensional face of $\mathscr{W}$. There are three possibilities, whether θ belongs to C_T, $-C_T$ or $\mathfrak{t}$.

Suppose first that $\theta \in C_T$. Then the root system $\Phi_\theta = \Phi \cap (\ker \theta)$ is finite of rank 2, of type $A_1 \times A_1$ or type A_2. More precisely, let $w \in W$ be of minimal length such that $\theta \in w\,\overline{C_0}$ and let

$J = \{i \in I \mid \langle w^{-1}\theta, \alpha_i\rangle = 0\}$; then the element w^{-1} maps Φ_θ onto the root system $\Phi_J = \Phi \cap \mathbb{R}J$. The full subgraph of (I, E) defined by J gives rise to a preprojective algebra Λ_J. The obvious surjective morphism $\Lambda \to \Lambda_J$ induces an inclusion Λ_J-mod $\hookrightarrow$ Λ-mod, whose image is the category $\mathscr{R}_{w^{-1}\theta}$. Further, the tilting ideals I_w provide an equivalence of categories

$$\mathscr{R}_{w^{-1}\theta} \underset{\mathrm{Hom}_\Lambda(I_w,?)}{\overset{I_w \otimes_\Lambda ?}{\rightleftarrows}} \mathscr{R}_\theta \ ,$$

whose action on the dimension-vectors is given by w. Putting all this together, we see that P_θ is the image under w of the HN polytope of the Λ_J-module $X = \mathrm{Hom}_\Lambda(I_w, T_\theta^{\max}/T_\theta^{\min})$. In addition, genericity is preserved in this construction: if T is a general point in an irreducible component of a nilpotent variety for Λ, then X is a general point in an irreducible component of a nilpotent variety for Λ_J. When Φ_J is of type A_2, this implies that the 2-face P_θ obeys the tropical Plücker relation from [36].

A similar analysis can be done in the case where θ is in $-C_T$. It then remains to handle the case where $\theta \in \mathfrak{t}$, that is, where θ belongs to a face F of codimension one in the spherical Weyl fan. Here $\Phi_\theta = \Phi \cap (\ker\theta)$ is an affine root system of type $\widetilde{A}_1$. The face F separates two spherical Weyl chambers of $\mathfrak{t}$, say C' and C'', and there are spherical chamber coweights γ' and γ'' such that $\Gamma \cap \overline{C'} = \big(\Gamma \cap \overline{F}\big) \sqcup \{\gamma'\}$ and $\Gamma \cap \overline{C''} = \big(\Gamma \cap \overline{F}\big) \sqcup \{\gamma''\}$.

Choose $\theta' \in C'$ and $\theta'' \in C''$. Assume that T is the general point of an irreducible component. As we saw in section 1.6, the modules $T_{\theta'}^{\max}/T_{\theta'}^{\min}$ and $T_{\theta''}^{\max}/T_{\theta''}^{\min}$ are then described by tuples of partitions $(\lambda_\gamma)_{\gamma\in\Gamma\cap\overline{C'}}$ and $(\lambda_\gamma)_{\gamma\in\Gamma\cap\overline{C''}}$, respectively. Both these modules are subquotients of $T_\theta^{\max}/T_\theta^{\min}$, so this latter contains the information about the partitions λ_γ for all $\gamma \in (\Gamma \cap \overline{F}) \sqcup \{\gamma', \gamma''\}$.

Theorem 1.3 *Let $\overline{P_\theta}$ be the polytope obtained by shortening each imaginary edge of the 2-face P_θ by $\Big(\sum_{\gamma\in\Gamma\cap\overline{F}} |\lambda_\gamma|\Big)\delta$. Then $\overline{P_\theta}$, equipped with the two partitions $\lambda_{\gamma'}$ and $\lambda_{\gamma''}$, is an MV polytope of type $\widetilde{A}_1$.*

A partition can be thought of as an MV polytope of type $\widetilde{A}_0$, since the generating function for the number of partitions equals the graded dimension of the upper half of the Heisenberg algebra. Thus, the family of partitions $(\lambda_\gamma)_{\gamma\in\Gamma\cap\overline{F}}$ can be thought of as an MV polytope of type $(\widetilde{A}_0)^{r-1}$. We can therefore regard the datum of the face P_θ and of the partitions $(\lambda_\gamma)_{\gamma\in(\Gamma\cap\overline{F})\sqcup\{\gamma',\gamma''\}}$ as an MV polytope of type $\widetilde{A}_1 \times \widetilde{A}_0^{r-1}$.

Theorem 1.3 will be proved in section 7.5. Our method is to construct an embedding of Π-mod into $\mathscr{R}_\theta$, where Π is the completed preprojective algebra of type $\widetilde{A}_1$; this embedding depends on F and its essential image is large enough to capture a dense open subset in the relevant irreducible component of Lusztig's nilpotent variety. In this construction, we were inspired by the work of I. Frenkel et al. [20] who produced analogous embeddings in the quiver setting.

So the final picture is the following. Let $\mathcal{MV}$ be the set of all lattice convex polytopes P in $\mathbb{R}I$, equipped with a family of partitions $(\lambda_\gamma)_{\gamma\in\Gamma}$, such that:

- The normal fan to P is a coarsening of the Weyl fan $\mathscr{W}$.
- To each spherical Weyl chamber C corresponds a imaginary edge of P; the difference between the two endpoints of this edge is equal to $\left(\sum_{\gamma\in\Gamma\cap\overline{C}}|\lambda_\gamma|\right)\delta$.
- A 2-face of P is an MV polytope of type $A_1\times A_1$, A_2 or $\widetilde{A}_1\times\widetilde{A}_0^{r-1}$; in the type A_2 case, this means that its shape obeys the tropical Plücker relation.

At the end of section 1.6, we associated an element $\widetilde{\mathrm{Pol}}(b)$ of $\mathcal{MV}$ to each $b\in B(-\infty)$.

Theorem 1.4 *The map* $\widetilde{\mathrm{Pol}}:B(-\infty)\to\mathcal{MV}$ *is bijective.*

Theorem 1.4 will be proved in section 7.7. In a companion paper [4], we will provide a combinatorial description of MV polytopes of type $\widetilde{A}_1$. With that result in hand, the above conditions provide an explicit characterization of the polytopes $\widetilde{\mathrm{Pol}}(b)$.

1.8 Lusztig data

As explained at the end of section 1.1, for a finite dimensional $\mathfrak{g}$, the MV polytope $\mathrm{Pol}(b)$ of an element $b\in B(-\infty)$ geometrically encodes all the Lusztig data of b.

In more detail, let N be the number of positive roots. Each reduced decomposition of the longest element w_0 of W provides a PBW basis of the quantum group $U_q(\mathfrak{n}_+)$, which goes to the basis $B(-\infty)$ at the limit $q\to 0$. To an element $b\in B(-\infty)$, one can therefore associate many PBW monomials, one for each PBW basis. In other words, one can associate to b many elements of $\mathbb{N}^N$, one for each reduced decomposition of w_0. These elements in $\mathbb{N}^N$ are called the Lusztig data of b. A reduced decomposition of w_0 specifies a path in the 1-skeleton of $\mathrm{Pol}(b)$ that connects the top vertex to the bottom one, and the corresponding Lusztig datum materializes as the lengths of the edges of this path.

With this in mind, we now explain that when $\mathfrak{g}$ is of affine type, our MV polytopes $\widetilde{\mathrm{Pol}}(b)$ provide a fair notion of Lusztig data.

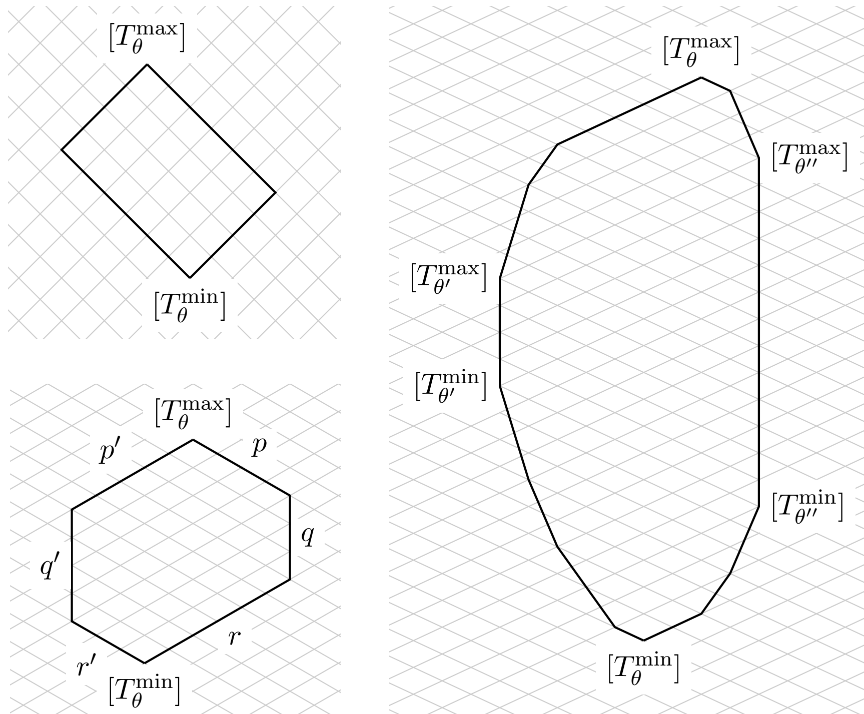

Figure 1: Examples of 2-faces of affine MV polytopes. These faces are of type $A_1 \times A_1$ (top left), A_2 (bottom left), and $\widetilde{A}_1$ (right). In type A_2, the tropical Plücker relation is $q' = \min(p, r)$. Note that the edges are parallel to root directions. In type $\widetilde{A}_1$, the edges parallel to the imaginary root δ (displayed here vertically) must be decorated with partitions $\lambda_{\gamma'}$ and $\lambda_{\gamma''}$ which encode the structure of the modules $T_{\theta'}^{\max}/T_{\theta'}^{\min}$ and $T_{\theta''}^{\max}/T_{\theta''}^{\min}$ (notation as in the text above). The lengths of all the edges and the partitions are subject to relations which are the analogs of the tropical Plücker relations; the displayed polygon endowed with the partitions $\lambda_{\gamma'} = (3, 1)$ and $\lambda_{\gamma''} = (4, 2, 1, 1, 1, 1, 1, 1, 1)$ satisfies these relations.

To this aim, we first note that a reasonable analog of the reduced decompositions of w_0 is certainly the notion of "total reflection order" (Dyer) or "convex order" (Ito), see [15, 28]. By definition, this is a total order $\preccurlyeq$ on Φ_+ such that

$$\big(\alpha + \beta \in \Phi_+ \ \text{ and } \ \alpha \preccurlyeq \beta\big) \implies \alpha \preccurlyeq \alpha + \beta \preccurlyeq \beta.$$

(Unfortunately, the convexity relation implies that $m\delta \preccurlyeq n\delta$ for any positive integers m and n. We therefore have to accept that $\preccurlyeq$ is only a preorder; this blemish is however limited to the imaginary roots.)

A convex order $\preccurlyeq$ splits the positive real roots in two parts: those that are greater than δ and those that are smaller. One easily shows that the projection $\pi : \mathbb{R}I \to \mathfrak{t}^*$ maps $\{\beta \in \Phi_+ \mid \beta \succ \delta\}$ onto a positive system of Φ^s. Thus, there exists $\theta \in \mathfrak{t}$ such that

$$\forall \beta \in \Phi_+^{\mathrm{re}}, \quad \beta \succ \delta \iff \langle \theta, \beta \rangle > 0. \tag{1.1}$$

Given such a convex order $\preccurlyeq$, we will construct a functorial filtration $(T_{\succcurlyeq \alpha})_{\alpha \in \Phi_+}$ on each finite dimensional Λ-module T, such that each $\underline{\dim}\, T_{\succcurlyeq \alpha}$ is a vertex of $\mathrm{Pol}(T)$. The family of dimension-vectors $(\underline{\dim}\, T_{\succcurlyeq \alpha})_{\alpha \in \Phi_+}$ are the vertices along a path in the 1-skeleton of $\mathrm{Pol}(T)$ connecting the top vertex and bottom vertices. The lengths of the edges in this path form a family of natural numbers $(n_\alpha)_{\alpha \in \Phi_+}$, defined by the relation $\underline{\dim}\, T_{\succcurlyeq \alpha}/T_{\succ \alpha} = n_\alpha \alpha$. Further, if we choose $\theta \in \mathfrak{t}$ satisfying (1.1), then $T_{\succcurlyeq \delta} = T_\theta^{\max}$ and $T_{\succ \delta} = T_\theta^{\min}$.

Fix $b \in B(-\infty)$ and take a general point T in Λ_b. Besides the family $(n_\alpha)_{\alpha \in \Phi_+}$ of natural numbers mentioned just above, we can construct a tuple of partitions $(\lambda_\gamma)_{\gamma \in \Gamma \cap \overline{C}}$ by applying the analysis carried after Theorem 1.2 to the module $T_\theta^{\max}/T_\theta^{\min}$, where C is the spherical Weyl chamber containing θ. To b, we can thus associate the pair $\Omega_\preccurlyeq(b)$ consisting in the two families $(n_\alpha)_{\alpha \in \Phi_+^{\mathrm{re}}}$ and $(\lambda_\gamma)_{\gamma \in \Gamma \cap \overline{C}}$. All this information can be read from $\widetilde{\mathrm{Pol}}(b)$. We call $\Omega_\preccurlyeq(b)$ the Lusztig datum of b in direction $\preccurlyeq$.

Let us denote by $\mathcal{P}$ the set of all partitions and by $\mathbb{N}^{(\Phi_+^{\mathrm{re}})}$ the set of finitely supported families $(n_\alpha)_{\alpha \in \Phi_+^{\mathrm{re}}}$ of non-negative integers.

Theorem 1.5 *The map $\Omega_\preccurlyeq : B(-\infty) \to \mathbb{N}^{(\Phi_+^{\mathrm{re}})} \times \mathcal{P}^{\Gamma \cap \overline{C}}$ is bijective.*

Theorem 1.5 will be proved in section 7.6. Let us conclude by a few remarks.

(i) The MV polytope $\widetilde{\mathrm{Pol}}(b)$ contains the information of all Lusztig data of b, for all convex orders. This is in complete analogy with the situation in the case where $\mathfrak{g}$ is finite dimensional. The conditions on the 2-faces given in the definition of $\mathcal{MV}$ say how the Lusztig datum of b varies when the convex order changes; they can be regarded as the analog in the affine type case of Lusztig's piecewise linear bijections.

(ii) The knowledge of a single Lusztig datum of b, for just one convex order, allows one to reconstruct the irreducible component Λ_b. This fact is indeed an ingredient of the proof of injectivity in Theorem 1.4.

(iii) Through the bijective map $\widetilde{\mathrm{Pol}}$, the set $\mathcal{MV}$ acquires the structure of a crystal, isomorphic to $B(-\infty)$. This structure can be read from the Lusztig data. Specifically, if α_i is the smallest element of the order $\preccurlyeq$, then $\varphi_i(b)$ is the α_i-coordinate of $\Omega_\preccurlyeq(b)$, and the operators $\tilde{e}_i$ and $\tilde{f}_i$ act by incrementing or decrementing this coordinate.

(iv) As mentioned at the beginning of section 1.3, Beck in [6], Beck and Nakajima in [8], and Ito in [30] construct PBW bases of $U_q(\mathfrak{n}_+)$ for $\mathfrak{g}$ of affine type. An element in one of these bases is a monomial in root vectors, the product being computed according to a convex order $\preccurlyeq$. To describe a monomial, one needs an integer for each real root α and a r-tuple of integers for each imaginary root $n\delta$, so in total, monomials in a PBW basis are indexed by $\mathbb{N}^{(\Phi_+^{\mathrm{re}})} \times \mathcal{P}^r$. Moreover, such a PBW basis goes to $B(-\infty)$ at the limit $q \to 0$. (This fact has been established in [7] for Beck's bases, and the result can probably be extended to Ito's more general bases by using [46] or [57].) In the end, we get a bijection between $B(-\infty)$ and $\mathbb{N}^{(\Phi_+^{\mathrm{re}})} \times \mathcal{P}^r$. We expect that this bijection is our map $\Omega_{\preccurlyeq}$.

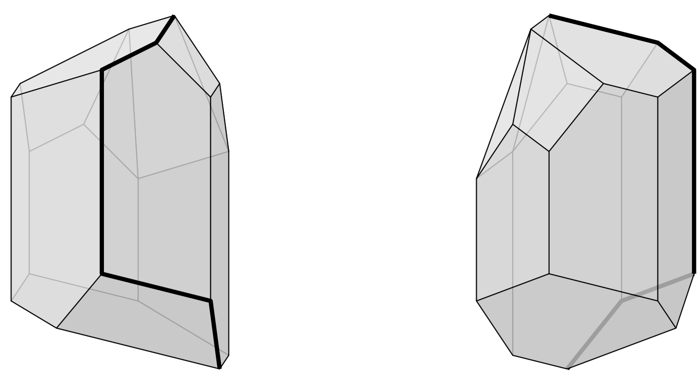

Figure 2: Two views of the same affine MV polytope $\widetilde{P}(b)$ of type $\widetilde{A}_2$. The thick line goes successively through the points $\nu_0 = 0$, $\nu_1 = \nu_0 + (\alpha_0 + 2\alpha_1 + 2\alpha_2)$, $\nu_2 = \nu_1 + 2\alpha_1$, $\nu_3 = \nu_2 + 5\delta$, $\nu_4 = \nu_3 + (\alpha_0 + \alpha_2)$, $\nu_5 = \nu_4 + 2\alpha_0$. The length of the edges of this line, together with the two partitions $\lambda_{\varpi_1} = (1,1)$ and $\lambda_{\varpi_2} = (2,1)$, form the Lusztig datum $\Omega_{\preccurlyeq}(b)$ relative to any convex order $\preccurlyeq$ such that $\alpha_0 \prec \alpha_0 + \alpha_2 \prec \delta \prec \alpha_1 \prec \alpha_0 + 2\alpha_1 + 2\alpha_2$. The other vertices were calculated using the conditions on the 2-faces. The MV polytope $\widetilde{P}(b)$ includes the data of λ_γ for all spherical chamber coweights ϖ, and in this example, the rest of this decoration is given by $\lambda_{s_1\varpi_1} = (1,1)$, $\lambda_{s_2 s_1 \varpi_1} = (0)$, $\lambda_{s_2\varpi_2} = (2,1)$ and $\lambda_{s_1 s_2 \varpi_2} = (2,1)$.

1.9 Plan of the paper

Section 2 recalls combinatorial notions and facts related to root systems. We emphasize the notion of biconvex subsets, which is crucial to the study of convex orders and to the definition of the functorial filtration $(T_{\succcurlyeq\alpha})_{\alpha\in\Phi_+}$ mentioned in section 1.8.

Section 3 is devoted to generalities about HN polytopes in abelian categories.

In section 4, we recall known facts about preprojective algebras and Lusztig's nilpotent varieties. We also prove that cutting a Λ-module according to a torsion pair is an operation that preserves genericity.

In section 5, we exploit the tilting theory on Λ-mod to define and study the submodules T^w and T_w mentioned in section 1.5. An important difference with the works of Iyama, Reiten et al. and of Geiß, Leclerc and Schröer is the fact that we are interested not only in the small slices that form the categories $\mathrm{Sub}(\Lambda/I_w)$ (notation of Iyama, Reiten et al.) or $\mathcal{C}_w$ (notation of Geiß, Leclerc and Schröer), but also at controlling the remainder. Moreover, we track the tilting theory at the level of the irreducible components of Lusztig's nilpotent varieties and interpret the result in term of crystal operations.

In section 6, we construct embeddings of Π-mod into Λ-mod, where Π is the completed preprojective algebra of type $\widetilde{A}_1$. The data needed to define such an embedding is a pair (S, R) of rigid orthogonal bricks in Λ-mod satisfying $\dim \mathrm{Ext}^1_\Lambda(S, R) = \dim \mathrm{Ext}^1_\Lambda(R, S) = 2$. The key ingredient in the construction is the 2-Calabi-Yau property of Λ-mod.

The final section 7 deals with the specifics of the affine type case. All the results concerning the imaginary edges, the cores, or the partitions are stated and proved there.

1.10 Thanks

We thank Claire Amiot for suggesting to us that the reflection functors in [5] are related to those in [1], which allowed us to take the current literature [14, 27, 59] into account in section 5. We also thank Thomas Dunlap for sharing his ideas about affine MV polytopes and for providing us with his PhD thesis [19]. We thank Alexander Braverman, Bernhard Keller, Bernard Leclerc, and Dinakar Muthiah for very helpful discussions. Finally, we thank two anonymous referees for thorough and insightful reports which led to significant improvements in the presentation.

P. B. acknowledges support from the ANR, project ANR-09-JCJC-0102-01. J. K. acknowledges support from NSERC. P. T. acknowledges support from the NSF, grants DMS-0902649, DMS-1162385 and DMS-1265555.

1.11 Summary of the main notations

$\mathbb{N} = \{0, 1, 2, \dots\}$.
$\mathcal{P}$ the set of partitions.
$\mathbf{K}(\mathscr{A})$ the Grothendieck group of $\mathscr{A}$, an essentially small abelian category.
$\mathrm{Irr}\,\mathscr{A}$ the set of isomorphism classes of simple objects in $\mathscr{A}$.

$\mathrm{Pol}(T) \subseteq \mathbf{K}(\mathscr{A})_{\mathbb{R}}$ the HN polytope of an object $T \in \mathscr{A}$.
$P_\theta = \{x \in P \mid \langle \theta, x \rangle = m\}$ the face of a HN polytope P, where $\theta \in (\mathbf{K}(\mathscr{A})_{\mathbb{R}})^*$ and m is the maximal value of θ on P.
K the base field for representation of quivers and preprojective algebras.
(I, E) a finite graph, wihout loops (encoding a symmetric generalized Cartan matrix).
$\mathfrak{g}$ the corresponding symmetric Kac-Moody algebra.
$\mathfrak{n}_+$ the upper nilpotent subalgebra of $\mathfrak{g}$.
Ω an orientation of (I, E) (thus $Q = (I, \Omega)$ is a quiver).
$H = \Omega \sqcup \overline{\Omega}$ the set of edges of the double quiver $\overline{Q}$.
$s, t : H \to I$ the source and target maps.
Λ the completed preprojective algebra of Q.
Λ-mod the category of finite dimensional left Λ-modules.
Φ the root system of $\mathfrak{g}$.
$\{\alpha_i \mid i \in I\}$ the standard basis of Φ.
$\Phi = \Phi_+ \sqcup \Phi_-$ the positive and negative roots with respect to this basis.
$W = \langle s_i \mid i \in I \rangle$ the Weyl group.
$\ell : W \to \mathbb{N}$ the length function.
$\mathbb{Z}I = \bigoplus_{i \in I} \mathbb{Z}\alpha_i$ the root lattice.
$\mathbb{R}I = \mathbb{Z}I \otimes_{\mathbb{Z}} \mathbb{R}$, the $\mathbb{R}$-vector space with basis $(\alpha_i)_{i \in I}$.
$(\mathbb{R}I)^*$ the dual vector space.
$\omega_i \in (\mathbb{R}I)^*$ the i-th coordinate on $\mathbb{R}I$; thus $(\omega_i)_{i \in I}$ is the dual basis of $(\alpha_i)_{i \in I}$.
$(\ ,\ ) : \mathbb{Z}I \times \mathbb{Z}I \to \mathbb{Z}$ the W-invariant symmetric bilinear form (real roots have square length 2).
$C_0 = \{\theta \in (\mathbb{R}I)^* \mid \langle \theta, \alpha_i \rangle > 0\}$ the dominant Weyl chamber.
$C_T = \bigcup_{w \in W} \overline{C_0}$ the Tits cone.
$F_J = \{\theta \in (\mathbb{R}I)^* \mid \forall j \in J,\ \langle \theta, \alpha_j \rangle = 0 \text{ and } \forall i \in I \setminus J,\ \langle \theta, \alpha_i \rangle > 0\}$, for $J \subseteq I$.
Φ_J and W_J, the root subsystem and the parabolic subgroup defined by $J \subseteq I$.
w_J the longest element in W_J, when the latter is finite.
$\mathrm{ht} : \mathbb{Z}I \to \mathbb{Z}$ the linear form such that $\mathrm{ht}(\alpha_i) = 1$ for each $i \in I$.
$N_w = \Phi_+ \cap w\Phi_-$, for $w \in W$; thus $N_w = \{s_{i_1} \cdots s_{i_{k-1}} \alpha_{i_k} \mid 1 \leq k \leq \ell\}$ for any reduced decomposition $w = s_{i_1} \cdots s_{i_\ell}$.
Π the preprojective algebra of type $\widetilde{A}_1$.

In the case of an affine root system:
δ the positive primitive imaginary root.
$\mathfrak{t}^* = \mathbb{R}I / \mathbb{R}\delta$.
$\pi : \mathbb{R}I \to \mathfrak{t}^*$ the projection modulo $\mathbb{R}\delta$.
$\Phi^s = \pi(\Phi^{\mathrm{re}})$ the spherical (finite) root system.
$\iota : \Phi^s \to \Phi_+^{\mathrm{re}}$ the "minimal" lift, a right inverse of π.
$r = \dim \mathfrak{t}^*$ the rank of Φ^s.
$\mathfrak{t} = \{\theta \in (\mathbb{R}I)^* \mid \langle \theta, \delta \rangle = 0\}$ the dual of $\mathfrak{t}^*$.

$\Gamma \subseteq \mathfrak{t}$ the set of all spherical chamber coweights.
$\mathscr{W}$ the Weyl fan on $(\mathbb{R}I)^*$, completed on $\mathfrak{t}$ by the spherical Weyl fan.
$Q^\vee \subseteq \mathfrak{t}$ the coroot lattice, spanned over $\mathbb{Z}$ by the elements $(\alpha_i, ?)$.
$t_\lambda \in W$ the translation, for $\lambda \in Q^\vee$; thus $t_\lambda(\nu) = \nu - \langle \lambda, \nu \rangle \delta$ for each $\nu \in \mathbb{R}I$.
$W_0 = W/Q^\vee$ the image of W in $\mathrm{GL}(\mathfrak{t}^*)$.
$\mathscr{V} = \{A \subseteq \Phi_+ \mid A \text{ biconvex}\}$.

And after having chosen an extending vertex 0 in the extended Dynkin diagram:
$I_0 = I \setminus \{0\}$ the vertices of the (finite type) Dynkin diagram.
$\{\pi(\alpha_i) \mid i \in I_0\}$ a preferred system of simple roots for Φ^s.
$(\varpi_i)_{i \in I_0}$ the spherical fundamental coweights, a basis of $\mathfrak{t}$.
$C_0^s = \sum_{i \in I_0} \mathbb{R}_{>0} \varpi_i$ the dominant spherical Weyl chamber.

Geometry:
The set of irreducible components of a topological space X is denoted by $\operatorname{Irr} X$. If Z is an irreducible topological space, then we say that a propriety $P(x)$ depending on a point $x \in Z$ holds for x general in Z if the set of points of Z at which P holds true contains a dense open subset of Z. We sometimes extend this vocabulary by simply saying "let x be a general point in Z"; in this case, it is understood that we plan to impose finitely many such conditions P.

2 Combinatorics of root systems and of MV polytopes

In this section, we introduce our notations and recall general results about root systems and biconvex subsets. Starting from section 2.3 onwards, we focus on the case of an affine root system.

2.1 General setup

Let (I, E) be a finite graph, without loops: here I is the set of vertices and E is the set of edges. We denote by $\mathbb{Z}I$ the free abelian group on I and we denote its canonical basis by $\{\alpha_i \mid i \in I\}$. We endow it with the symmetric bilinear form $(\,,\,) : \mathbb{Z}I \times \mathbb{Z}I \to \mathbb{Z}$, given by $(\alpha_i, \alpha_i) = 2$ for any i, and for $i \neq j$, (α_i, α_j) is the negative of the number of edges between the vertices i and j in the graph (I, E). The Weyl group W is the subgroup of $\mathrm{GL}(\mathbb{Z}I)$ generated by the simple reflections $s_i : \alpha_j \mapsto \alpha_j - (\alpha_j, \alpha_i)\alpha_i$; this is in fact a Coxeter system, whose length function is denoted by ℓ. Lastly, we denote by $\mathbb{N}I$ the set of all linear combinations of the α_i with coefficients in $\mathbb{N}$ and we denote by $\mathrm{ht} : \mathbb{Z}I \to \mathbb{Z}$ the linear form that maps each α_i to 1.

The matrix with entries (α_i, α_j) is a symmetric generalized Cartan matrix, hence it gives rise to a Kac-Moody algebra $\mathfrak{g}$ and a root system Φ. The latter is a W-stable subset of $\mathbb{Z}I$, which can be split into positive and negative roots $\Phi = \Phi_+ \sqcup \Phi_-$ and into real and imaginary roots $\Phi = \Phi^{\mathrm{re}} \sqcup \Phi^{\mathrm{im}}$.

Given a subset $J \subseteq I$, we can look at the root system $\Phi_J = \Phi \cap \operatorname{span}_{\mathbb{Z}}\{\alpha_j \mid j \in J\}$. Its Weyl group is the parabolic subgroup $W_J = \langle s_j \mid j \in J \rangle$ of W. If W_J is finite, then it has a longest element, which we denote by w_J. An element $u \in W$ is called J-reduced on the right if $\ell(us_j) > \ell(u)$ for each $j \in J$. If u is J-reduced on the right, then $\ell(uv) = \ell(u) + \ell(v)$ for all $v \in W_J$. Each right coset of W_J in W contains a unique element that is J-reduced on the right.

The Weyl group acts on $\mathbb{R}I$ and on its dual $(\mathbb{R}I)^*$. The dominant chamber C_0 and the Tits cone C_T are the convex cones in $(\mathbb{R}I)^*$ defined as

$$C_0 = \{\theta \in (\mathbb{R}I)^* \mid \forall i \in I,\ \langle \theta, \alpha_i \rangle > 0\} \quad \text{and} \quad C_T = \bigcup_{w \in W} w\,\overline{C_0}.$$

The closure $\overline{C}_0$ is the disjoint union of faces

$$F_J = \{\theta \in (\mathbb{R}I)^* \mid \forall j \in J,\ \langle \theta, \alpha_j \rangle = 0 \text{ and } \forall i \in I \setminus J,\ \langle \theta, \alpha_i \rangle > 0\},$$

for $J \subseteq I$. The stabilizer of any point in F_J is precisely the parabolic subgroup W_J. Thus

$$\overline{C}_0 = \bigsqcup_{J \subseteq I} F_J \quad \text{and} \quad C_T = \bigsqcup_{J \subseteq I} \bigsqcup_{w \in W/W_J} wF_J.$$

The disjoint union on the right endows C_T with the structure of a (non locally finite) fan, which we call the Tits fan.

To an element $w \in W$, we associate the subset $N_w = \Phi_+ \cap w\Phi_-$. If $w = s_{i_1} \cdots s_{i_\ell}$ is a reduced decomposition, then

$$N_w = \{s_{i_1} \cdots s_{i_{k-1}} \alpha_{i_k} \mid 1 \le k \le \ell\}.$$

The following result is well-known (see for instance Remark ♣ in [15]).

Lemma 2.1 *For $(u, v) \in W^2$, the following three properties are equivalent:*

$$\ell(u) + \ell(v) = \ell(uv), \qquad N_u \subseteq N_{uv}, \qquad N_{u^{-1}} \cap N_v = \varnothing.$$

Corollary 2.2 *Let $J \subseteq I$ and let $w \in W$. If w is J-reduced on the right, then $N_{w^{-1}} \cap \Phi_J = \varnothing$.*

Proof. Let $J \subseteq I$ and let $w \in W$ be such that $N_{w^{-1}} \cap \Phi_J \neq \varnothing$. Then there exists $\beta \in \Phi_J \cap \Phi_+$ such that $w\beta \in \Phi_-$. Since β is a nonnegative linear combination of the roots α_j for $j \in J$, it follows that there exists $j \in J$ such that $w\alpha_j \in \Phi_-$. For this j, we have $N_{s_j} \subseteq N_{w^{-1}}$, whence $\ell(s_j) + \ell(s_j w^{-1}) = \ell(w^{-1})$ by Lemma 2.1. Therefore w is not J-reduced on the right. □

2.2 Biconvex sets (general type)

A subset $A \subseteq \Phi$ is said to be clos if the conditions $\alpha \in A$, $\beta \in A$, $\alpha + \beta \in \Phi$ imply $\alpha + \beta \in A$ (see [10], chapitre 6, §1, n° 7, Définition 4). A subset $A \subseteq \Phi_+$ is said to be biconvex if both A and $\Phi_+ \setminus A$ are clos. We denote by $\mathscr{V}$ the set of all biconvex subsets of Φ_+ and endow it with the inclusion order.

Examples 2.3. (i) An increasing union or a decreasing intersection of biconvex subsets is itself biconvex.

(ii) Each finite biconvex subset of Φ_+ consists of real roots and is a N_w, with $w \in W$ (see [15], Proposition 3.2). For convenience, we will say that a biconvex set A is cofinite if its complement $\Phi_+ \setminus A$ is finite. Given $w \in W$, we set $A_w = N_{w^{-1}}$ and $A^w = \Phi_+ \setminus N_w$. Thus the map $w \mapsto A_w$ (respectively, $w \mapsto A^w$) is a bijection from W onto the set of finite (respectively, cofinite) biconvex subsets of Φ_+.

(iii) Each $\theta \in (\mathbb{R}I)^*$ gives birth to two biconvex subsets

$$A_\theta^{\min} = \{\alpha \in \Phi_+ \mid \langle \theta, \alpha \rangle > 0\} \quad \text{and} \quad A_\theta^{\max} = \{\alpha \in \Phi_+ \mid \langle \theta, \alpha \rangle \geq 0\}.$$

Remark 2.4. Define a positive root system as a subset $X \subseteq \Phi$ such that $\Phi = X \sqcup (-X)$ and that the convex cone spanned in $\mathbb{R}I$ by X is acute. Denote by $\widetilde{\mathscr{V}}$ the set of all positive root systems in Φ. If X is a positive root system, then $A = X \cap \Phi_+$ is biconvex and X can be recovered from A by the formula $X = A \sqcup (-(\Phi_+ \setminus A))$. Thus the map $X \mapsto X \cap \Phi_+$ from $\widetilde{\mathscr{V}}$ to $\mathscr{V}$ is well-defined and injective. The image of this map is the set $\mathscr{V}'$ of all subsets $A \subseteq \Phi_+$ such that the convex cones spanned by A and by $\Phi_+ \setminus A$ intersect only at the origin. Lemma 2.11 below shows that $\mathscr{V}' = \mathscr{V}$ whenever Φ is of finite or affine type.

Proposition 2.5 *Let $J \subseteq I$, let $\theta \in F_J$, and let $w \in W$. Assume w is J-reduced on the right. Then $A^w = A_{w\theta}^{\max}$ and $A_w = A_{-w^{-1}\theta}^{\min}$. In addition, if W_J is finite, then $A^{ww_J} = A_{w\theta}^{\min}$ and $A_{w_J w} = A_{-w^{-1}\theta}^{\max}$.*

Proof. Let J, θ, w as in the statement of the proposition.

We have $A^w = \{\alpha \in \Phi_+ \mid w^{-1}\alpha \in \Phi_+\}$ and $A_{w\theta}^{\max} = \{\alpha \in \Phi_+ \mid \langle \theta, w^{-1}\alpha \rangle \geq 0\}$. The inclusion $A^w \subseteq A_{w\theta}^{\max}$ is straightforward. To show the reverse inclusion, we take $\alpha \in \Phi_+ \setminus A^w$, that is, $\alpha \in N_w$. Then $\beta = -w^{-1}\alpha$ is in $N_{w^{-1}}$, in particular $\beta \in \Phi_+$, but $\beta \notin \Phi_J$ by Corollary 2.2, and so $\langle \theta, \beta \rangle > 0$, which means that $\alpha \notin A_{w\theta}^{\max}$. We conclude that $A^w = A_{w\theta}^{\max}$.

Suppose now that W_J is finite. Then $A^{ww_J} = \{\alpha \in \Phi_+ \mid w_J w^{-1}\alpha \in \Phi_+\}$ and $A_{w\theta}^{\min} = \{\alpha \in \Phi_+ \mid \langle \theta, w_J w^{-1}\alpha \rangle > 0\}$. The inclusion $A_{w\theta}^{\min} \subseteq A^{ww_J}$ is straightforward. To show the reverse inclusion, we take $\alpha \in A^{ww_J} \setminus A^{\min}_{w\theta}$, if possible. Then $w_J w^{-1}\alpha$ necessarily belongs to $\Phi_+ \cap \Phi_J$, and so does $\beta = -w^{-1}\alpha$. Then $\beta \in N_{w^{-1}} \cap \Phi_J$, which contradicts Corollary 2.2. We conclude that $A^{ww_J} = A^{\min}_{w\theta}$.

The last two equalities $A_w = A^{\min}_{-w^{-1}\theta}$ and $A_{w_J w} = A^{\max}_{-w^{-1}\theta}$ are proved in a similar fashion. □

2.3 Setup in the affine type

In the rest of section 2, we will focus on the case where the root system Φ is of affine type. Then there exists $\delta \in \mathbb{Z}I$ such that $\Phi^{\mathrm{im}}_+ = \mathbb{Z}_{>0}\,\delta$. We set $\mathfrak{t}^* = \mathbb{R}I/\mathbb{R}\delta$ and we denote the natural projection by $\pi : \mathbb{R}I \to \mathfrak{t}^*$. Then $\Phi^s = \pi(\Phi^{\mathrm{re}})$ is a finite root system in $\mathfrak{t}^*$, called the spherical root system.

The dual vector space $\mathfrak{t}$ of $\mathfrak{t}^*$ is identified with $\{\theta \in (\mathbb{R}I)^* \mid \langle \theta, \delta\rangle = 0\}$. The linear forms $x \mapsto (\alpha_i, x)$ on $\mathbb{R}I$ span a lattice in $\mathfrak{t}$, called the coroot lattice; we denote it by $Q^\vee$. The Weyl group leaves δ invariant, hence acts on $\mathfrak{t}^*$. The kernel of this action consists of translations t_λ, for $\lambda \in Q^\vee$. The translation t_λ acts on $\mathbb{R}I$ by $t_\lambda(\nu) = \nu - \langle \lambda, \nu\rangle\,\delta$ and acts on $(\mathbb{R}I)^*$ by $t_\lambda(\theta) = \theta + \langle \theta, \delta\rangle \lambda$. We denote by W_0 the quotient of W by this subgroup of translations; it can be viewed as a subgroup of $\mathrm{GL}(\mathfrak{t}^*)$, indeed as the Weyl group of Φ^s.

A basis of the root system Φ^s is in particular a basis of $\mathfrak{t}^*$, and the elements in the dual basis are called the spherical fundamental coweights. We define a spherical chamber coweight as an element of $\mathfrak{t}$ that is conjugate under W_0 to a spherical fundamental coweight, and we denote by Γ the set of all spherical chamber coweights. The root system Φ^s defines a hyperplane arrangement in $\mathfrak{t}$, called the spherical Weyl fan. The open cones in this fan will be called the spherical Weyl chambers. The rays of this fan are spanned by the spherical chamber coweights.

The Tits cone is $C_T = \{\theta \in (\mathbb{R}I)^* \mid \langle \theta, \delta\rangle > 0\} \cup \{0\}$. Thus $(\mathbb{R}I)^*$ is covered by C_T, $-C_T$ and $\mathfrak{t}$. Gathering the faces of the Tits fan, their opposite, and the faces of the spherical Weyl fan, we get a (non-locally finite) fan on $(\mathbb{R}I)^*$. We call it the affine Weyl fan and denote it by $\mathscr{W}$. The cones of this fan are the equivalence classes of the relation $\sim$ on $(\mathbb{R}I)^*$ defined in the following way: one says that $x \sim y$ if for each root $\alpha \in \Phi$, the two real numbers $\langle x, \alpha\rangle$ and $\langle y, \alpha\rangle$ have the same sign.

For each $\alpha \in \Phi^s$, we denote by $\iota(\alpha) \in \Phi^{\mathrm{re}}_+$ the unique positive real root such that $\pi(\iota(\alpha)) = \alpha$ and $\iota(\alpha) - \delta \notin \Phi^{\mathrm{re}}_+$. Thus $\iota : \Phi^s \to \Phi^{\mathrm{re}}_+$ is the "minimal" right inverse to π.

It is often convenient to embed the spherical root system Φ^s in the affine root system Φ. To do that, we choose an extending vertex 0 in I and we set $I_0 = I \setminus \{0\}$. Then the spherical Weyl group W_0 can be identified with the parabolic subgroup $\langle s_i \mid i \in I_0\rangle$ of W. Further, $\{\pi(\alpha_i) \mid i \in I_0\}$ is a basis of the spherical root system Φ^s, whence a dominant spherical Weyl

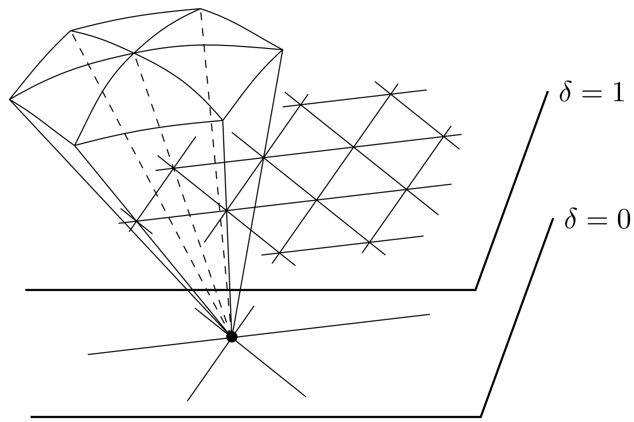

Figure 3: The upper half of the affine Weyl fan of type $\widetilde{A}_2$. The intersection with the affine hyperplane $\{\theta \in (\mathbb{R}I)^* \mid \langle \theta, \delta \rangle = 1\}$ draws the familiar pattern of alcoves. On the hyperplane $\{\theta \in (\mathbb{R}I)^* \mid \langle \theta, \delta \rangle = 0\}$, one can see the spherical Weyl fan.

chamber $C_0^s = \{\theta \in \mathfrak{t} \mid \forall i \in I_0,\ \langle \theta, \alpha_i \rangle > 0\}$. The highest root $\widetilde{\alpha}$ of Φ^s relative to this set of simple roots satisfies $\delta = \alpha_0 + \iota(\widetilde{\alpha})$. We denote by $\{\varpi_i \mid i \in I_0\}$ the basis of $\mathfrak{t}$ dual to the basis $\{\pi(\alpha_i) \mid i \in I_0\}$ of $\mathfrak{t}^*$.

2.4 Biconvex sets (affine type)

One nice feature of the affine type is the following key result, which is a direct application of Theorem 3.12 in [15].

Proposition 2.6 *Let $A \subseteq \Phi_+$ be a biconvex subset. If $\delta \notin A$, then A is the union of an increasing sequence of finite biconvex subsets. If $\delta \in A$, then A is the intersection of a decreasing sequence of cofinite biconvex subsets.*

Example 2.7. Let $\lambda \in Q^\vee$. With the notation of Remark 2.3 (iii), we then have

$$A_\lambda^{\min} = \bigcup_{n \in \mathbb{N}} A_{t_{n\lambda}} \quad \text{and} \quad A_\lambda^{\max} = \bigcap_{n \in \mathbb{N}} A^{t_{n\lambda}};$$

these equalities readily follow from the formula $t_{n\lambda}(\alpha) = \alpha - n\langle \lambda, \alpha \rangle \delta$.

Lemma 2.8 *Let $A \subseteq \Phi_+$ be a biconvex subset such that $\delta \notin A$ and let $X = \pi(A)$. Then X is contained in a positive root system of Φ^s and $\iota(X) \subseteq A$.*

Proof. Since A is clos, so is X. Since $\delta \notin A$, we furthermore have $X \cap (-X) = \varnothing$. By [10], chapitre 6, §1, n° 7, Proposition 22, X is contained in a positive root system of Φ^s. Lastly, let $\alpha \in X$, and choose $\beta \in A \cap \pi^{-1}(\alpha)$ of minimal height. Then $\beta - \delta$ is not in A. It is not in $\Phi_+ \setminus A$ either, for this latter is clos and contains δ but not β. Therefore $\beta - \delta \notin \Phi_+$, which means that $\beta = \iota(\alpha)$. We have shown that $\iota(X) \subseteq A$. □

Lemma 2.9 *Let $\alpha \in \Phi_+^{\mathrm{re}}$ and let A and B be two biconvex subsets such that $B = A \sqcup \{\alpha\}$. We assume that $\delta \notin A$. Then, for each finite subset $X \subseteq A$, there are finite biconvex subsets $A' \subseteq A$ and $B' \subseteq B$ such that $X \subseteq A'$ and $B' = A' \sqcup \{\alpha\}$.*

Proof. Since A is the increasing union of finite biconvex subsets, one can find a biconvex $A_0 \subseteq A$ that contains X. Similarly, one can find a finite biconvex $B_0 \subseteq B$ that contains $A_0 \cup \{\alpha\}$. By Example 2.3 (ii), we can write $A_0 = N_u$ and $B_0 = N_{uv}$, with $(u, v) \in W^2$. Lemma 2.1 says then that $\ell(uv) = \ell(u) + \ell(v)$. Let us write a reduced decomposition $s_{i_1} \cdots s_{i_\ell}$ for v. There exists k such that $\alpha = us_{i_1} \cdots s_{i_{k-1}} \alpha_{i_k}$. We then take $A' = N_{us_{i_1} \cdots s_{i_{k-1}}}$ and $B' = N_{us_{i_1} \cdots s_{i_k}}$. □

One can of course obtain a similar statement in the case where $\delta \in B$ by taking complements in Φ_+.

Lemma 2.10 *(i) Let A and B be two biconvex subsets such that $B = A \sqcup (\mathbb{Z}_{>0}\delta)$. Then there is a positive system $X \subseteq \Phi^s$ such that $A = \{\alpha \in \Phi_+ \mid \pi(\alpha) \in X\}$.*

(ii) Let $A \subseteq B$ be two biconvex subsets. Suppose that A is finite and that B is cofinite. Then there is a positive system $X \subseteq \Phi^s$ such that $A \subseteq \{\alpha \in \Phi_+ \mid \pi(\alpha) \in X\} \subseteq B$.

Proof. Let us first show (i). We take A and B as in the statement to be proved. Let $X = \pi(A)$ and $Y = \pi(\Phi_+ \setminus B)$. Certainly, $\Phi^s = \pi(\Phi_+^{\mathrm{re}}) = X \cup Y$. In addition, X and Y are disjoint, because otherwise the inclusions $\iota(X) \subseteq A$ and $\iota(Y) \subseteq \Phi_+ \setminus B$ given by Lemma 2.8 would force A and $\Phi_+ \setminus B$ to share a common element. Lastly, X and Y are clos. By [10], chapitre 6, §1, n° 7, Corollaire 1 to Proposition 20, X is a positive system in Φ^s. Assertion (i) then follows from the observation that $A \subseteq \pi^{-1}(X)$ and $\pi(\Phi_+ \setminus A) \subseteq Y \cup \{0\}$.

Now we consider (ii). By Example 2.3 (ii), there are $(u, v) \in W^2$ such that $A = A_u = N_{u^{-1}}$ and $B = A^v = \Phi_+ \setminus N_v$. The condition $A \subseteq B$ means that $N_{u^{-1}} \cap N_v = \varnothing$, so $\ell(uv) = \ell(u) + \ell(v)$ by Lemma 2.1. By Lemma 2.8, $\pi(N_{uv})$ is contained in a positive root system of Φ^s, say Y.

Since $\pi : \mathbb{R}I \to \mathfrak{t}^*$ is W-equivariant, $\pi(u^{-1}N_{uv})$ is contained in $u^{-1}Y$. Let $X = -u^{-1}Y$. From the equality $u^{-1}N_{uv} = u^{-1}(N_u \sqcup uN_v) = (-N_{u^{-1}}) \sqcup N_v$, we deduce that $\pi(N_{u^{-1}}) \subseteq X$ and that $\pi(N_v) \cap X = \varnothing$. Therefore $N_{u^{-1}} \subseteq \{\alpha \in \Phi_+ \mid \pi(\alpha) \in X\} \subseteq \Phi_+ \setminus N_v$. □

Lemma 2.11 *For any biconvex set A, the convex cones spanned by A and by $\Phi_+ \setminus A$ intersect only at the origin.*

Proof. Taking complements in Φ_+, we can assume that $\delta \notin A$. Suppose that the convex cones spanned by A and $\Phi_+ \setminus A$ share a common nonzero element x. This x can be expressed as a non-negative linear combination of a finite family of elements of A, so using Proposition 2.6, we can find a finite biconvex subset $B \subseteq A$ such that x belongs to the convex cone spanned by B. Further there exists $\theta \in (\mathbb{R}I)^*$ such that $B = A_\theta^{\min}$, by Example 2.3 (ii) and Proposition 2.5. We then have $\langle \theta, x \rangle > 0$. On the other hand, x belongs to the convex cone spanned by $\Phi_+ \setminus B = A_{-\theta}^{\max}$, and so $\langle -\theta, x \rangle \geq 0$. This contradiction shows that the convex cones spanned by A and $\Phi_+ \setminus A$ do not share any nonzero element. □

2.5 Convex orders

One motivation for studying biconvex subsets comes from the notion of "convex order" on Φ_+. Specifically, a preorder $\preccurlyeq$ on Φ_+ is called a convex order if for all $(\alpha, \beta) \in \Phi_+^2$, the three following conditions hold:

$$\alpha \preccurlyeq \beta \ \text{ or } \ \beta \preccurlyeq \alpha,$$

$$\big(\alpha + \beta \in \Phi_+ \ \text{ and } \ \alpha \preccurlyeq \beta\big) \implies \alpha \preccurlyeq \alpha + \beta \preccurlyeq \beta,$$

$$\big(\alpha \preccurlyeq \beta \ \text{ and } \ \beta \preccurlyeq \alpha\big) \iff \alpha \text{ and } \beta \text{ are proportional.}$$

In this section, we restrict to affine type. In this case, in the last condition above, α and β are proportional if and only if they are equal or they are both imaginary.

A terminal section for a convex order $\preccurlyeq$ is a subset $A \subseteq \Phi_+$ such that

$$\big(\alpha \in A \ \text{ and } \ \alpha \preccurlyeq \beta\big) \implies \beta \in A.$$

We denote the set of terminal sections of $\preccurlyeq$ by $\mathscr{U}(\preccurlyeq)$. The following result, implicit in [15, 28], provides the link between biconvex subsets and convex orders. We leave its (routine) proof as an exercise for the reader.

Lemma 2.12 *For each convex order $\preccurlyeq$, the set $\mathscr{U}(\preccurlyeq)$ is a maximal totally ordered subset of $\mathscr{V}$. The datum of $\mathscr{U}(\preccurlyeq)$ completely determines $\preccurlyeq$.*

Remarks 2.13. (i) It is known that any biconvex subset is the terminal section of a convex order (Corollary 3.13 in [15]).

(ii) Let us say that a pair (A, B) of biconvex subsets is adjacent if $A \subsetneq B$ and if there does not exist a biconvex subset C such that $A \subsetneq C \subsetneq B$. Each $(w, i) \in W \times I$ with $\ell(ws_i) > \ell(w)$ gives such a pair, namely (N_w, N_{ws_i}); indeed, one here observes that $N_{ws_i} = N_w \sqcup \{w\alpha_i\}$, so there is no room between N_w and N_{ws_i}. Using Lemma 2.1, one easily shows that any adjacent pair of finite biconvex subsets is of this form, so the notion of adjacent biconvex subset generalizes the covering relation for the weak Bruhat order.

(iii) Given a real positive root $\alpha \in \Phi_+^{\mathrm{re}}$, let us say that a pair (A, B) of biconvex subsets is α-adjacent if $B = A \sqcup \{\alpha\}$. Let us say that (A, B) is δ-adjacent if $B = A \sqcup \big(\mathbb{Z}_{>0}\delta\big)$. We conjecture that a pair (A, B) of biconvex subsets is adjacent (in the sense of (ii) above) if and only if there is a root β (real or imaginary) such that (A, B) is β-adjacent. This conjecture seems reasonable in view of our current understanding of biconvex subsets, but we were not able to extract it from the papers [15, 29]. If it is correct, then Lemma 2.12 admits a converse, and "maximal totally ordered subset of $\mathscr{V}$" would be a notion equivalent to that of "convex order". In any case, Zorn's lemma shows that any totally ordered subset of $\mathscr{V}$ can be completed to a maximal one.

(iv) Let $\preccurlyeq$ be a convex order. The terminal sections

$$\{\beta \in \Phi_+ \mid \beta \succ \delta\} \quad \text{and} \quad \{\beta \in \Phi_+ \mid \beta \succcurlyeq \delta\}$$

satisfy the assumptions of Lemma 2.10 (i), so there is a positive system $X \subseteq \Phi^s$ such that $\{\beta \in \Phi_+ \mid \beta \succ \delta\} = \{\beta \in \Phi_+ \mid \pi(\beta) \in X\}$. This fact was announced in section 1.8, see equation (1.1).

Examples 2.14. (i) Let us consider the type $\widetilde{A}_1$. As is customary, we use $I = \{0, 1\}$. Then

$$\Phi_+ = \{\alpha_0 + n\delta, \alpha_1 + n\delta, (n+1)\delta \mid n \in \mathbb{N}\}.$$

There are exactly two convex orders on Φ_+. One of them is

$$\alpha_1 \prec \alpha_1 + \delta \prec \alpha_1 + 2\delta \prec \cdots \prec \delta \prec \cdots \prec \alpha_0 + 2\delta \prec \alpha_0 + \delta \prec \alpha_0,$$

the other is the opposite order.

(ii) A linear form $\theta \in (\mathbb{R}I)^*$ defines a preorder on Φ_+, as follows: we say that $\alpha \preccurlyeq \beta$ if $\langle \theta, \alpha \rangle / \operatorname{ht}(\alpha) \leq \langle \theta, \beta \rangle / \operatorname{ht}(\beta)$. For θ general enough (outside countably many hyperplanes), this preorder is a convex order.

2.6 GGMS polytopes in affine type

To a non-empty compact convex subset $K \subseteq \mathbb{R}I$, one associates its support function $\psi_K : (\mathbb{R}I)^* \to \mathbb{R}$, defined by $\psi_K(\theta) = \max(\theta(K))$. One can reconstruct K from the datum of ψ_K. If P is a convex polytope, then ψ_P is piecewise linear; the maximal regions of linearity are closed cones that cover $(\mathbb{R}I)^*$. The relative interiors of these cones and of their faces form the normal fan $\mathscr{N}_P$ of P. In addition, each face of P is of the form $P_\theta = \{x \in P \mid \langle \theta, x\rangle = \psi_P(\theta)\}$ for some $\theta \in (\mathbb{R}I)^*$.

We say that a convex lattice polytope $P \subseteq \mathbb{R}I$ is GGMS if each open cone of the affine Weyl fan $\mathscr{W}$ is contained in an open cone of the normal fan $\mathscr{N}_P$. In other words, we ask that for each open cone $C \in \mathscr{W}$, there is a vertex x of P such that $P_\theta = \{x\}$ for all $\theta \in C$. It follows that an edge of a GGMS polytope always points in a root direction. As we will see later in this section, a convex lattice polytope P is GGMS if and only if $\mathscr{N}_P$ is a coarsening of $\mathscr{W}$, in the sense that each cone of $\mathscr{N}_P$ is the union of cones of $\mathscr{W}$.

Let us fix a GGMS polytope P. If A is a finite or cofinite biconvex subset, then there is a unique open cone $C \in \mathscr{W}$ such that $A = A_\theta^{\min} = A_\theta^{\max}$ for each $\theta \in C$, by Example 2.3 and Proposition 2.5, and we denote by $\mu_P(A)$ the vertex x of P such that $P_\theta = \{x\}$ for each $\theta \in C$.

We now want to extend this definition to all biconvex subsets A. As mentioned above, the edges of a GGMS polytope P point in root directions. Let us denote by $E_P \subseteq (\Phi_+^{\mathrm{re}} \sqcup \{\delta\})$ the (finite) set of all these directions. Furthermore, let us denote the symmetric difference between two sets A and B by $A\Delta B = (A \setminus B) \cup (B \setminus A)$.

Lemma 2.15 *Let P be a GGMS polytope and let A and B be finite or cofinite biconvex subsets. Then there is a family of nonnegative integers $(n_\alpha)_{\alpha\in A\Delta B}$ such that $n_\alpha = 0$ if $\alpha \notin E_P$ and*

$$\mu_P(B) - \mu_P(A) = \sum_{\alpha\in B\setminus A} n_\alpha \alpha - \sum_{\alpha \in A\setminus B} n_\alpha \alpha.$$

Proof. We choose θ_0 and θ_1 in open cones of $\mathscr{W}$ such that $A = \{\alpha \in \Phi_+ \mid \langle\theta_0, \alpha\rangle > 0\}$ and $B = \{\alpha \in \Phi_+ \mid \langle\theta_1, \alpha\rangle > 0\}$. By moving θ_0 and θ_1 if necessary, we may assume that the segment $[\theta_0, \theta_1]$ does not meet any cone of codimension 2 in the normal fan $\mathscr{N}_P$. We consider $\theta(t) = (1-t)\theta_0 + t\theta_1$. As t varies from 0 to 1, the face $P_{\theta(t)}$ is generally a vertex of P, occasionally an edge, but never a face of higher dimension. The vertices and edges found in this way form a path in the 1-skeleton of P from $\mu_P(A)$ to $\mu_P(B)$. Each edge traversed by this path points in the direction of a root α such that $\langle\theta_0, \alpha\rangle < 0 < \langle\theta_1, \alpha\rangle$, so that either $\alpha \in \Phi_+ \cap (B \setminus A)$ or $-\alpha \in \Phi_+ \cap (A \setminus B)$, and moreover the length of this edge is an integral multiple of α, because P is a lattice polytope. $\square$

With the notation of the lemma, we have $\mu_P(A) = \mu_P(B)$ as soon as $A\Delta B$ does not meet E_P. We can thus extend μ_P to a map from all of $\mathscr{V}$ to the set of vertices of P as follows. If A is a biconvex subset such that $\delta \notin A$, then we set $\mu_P(A) = \mu_P(B)$, where B is any finite biconvex subset such that $A \cap E_P \subseteq B \subseteq A$; the result does not depend on the choice of B, because the set of all possible B is filtered. Similarly, if A is a biconvex subset such that $\delta \in A$, then we set $\mu_P(A) = \mu_P(B)$, where B is any cofinite biconvex subset such that $A \subseteq B \subseteq (A \cup (\Phi_+ \setminus E_P))$. With these conventions, Lemma 2.15 trivially extends to any biconvex subsets A and B.

Recall the biconvex subsets $A_\theta^{\min}$ and $A_\theta^{\max}$ from Example 2.3 (iii).

Proposition 2.16 *The support function Ψ_P of a GGMS polytope P is given by*

$$\psi_P(\theta) = \langle \theta, \mu_P(A_\theta^{\min})\rangle = \langle \theta, \mu_P(A_\theta^{\max})\rangle,$$

for all $\theta \in (\mathbb{R}I)^*$.

Proof. Let P be a GGMS polytope, let $\theta \in (\mathbb{R}I)^*$, and let B be any biconvex subset. By definition, $\langle \theta, \alpha\rangle > 0$ for each $\alpha \in A_\theta^{\min} \setminus B$, and $\langle \theta, \alpha\rangle \leq 0$ for each $\alpha \in B \setminus A_\theta^{\min}$. Lemma 2.15 then implies that $\langle \theta, \mu_P(A_\theta^{\min}) - \mu_P(B)\rangle \geq 0$. Since each vertex of P can be written as a $\mu_P(B)$, it follows that $\langle \theta, \mu_P(A_\theta^{\min})\rangle$ is the supremum of θ on P, whence the first equality. The second equality also directly follows Lemma 2.15. $\square$

The biconvex subsets $A_\theta^{\min}$ and $A_\theta^{\max}$ only depend on the cone of $\mathscr{W}$ to which θ belongs. From Proposition 2.16, we then deduce that the support function of a GGMS polytope P is linear on each cone of $\mathscr{W}$. Therefore the normal fan $\mathscr{N}_P$ is a coarsening of $\mathscr{W}$, as announced earlier in this section.

The fact that $\mathscr{N}_P$ is a coarsening of $\mathscr{W}$ restricts the shape of the 2-faces of P. Specifically, given θ in a codimension 2 face of $\mathscr{W}$, two cases can happen: either $\pm\theta \in C_T$, and then $\Phi \cap (\ker\theta)$ is a finite root system of type $A_1 \times A_1$ or A_2; or θ belongs to a facet of the spherical Weyl fan, and then $\Phi \cap (\ker\theta)$ is an affine root system of type $\widetilde{A}_1$. In both cases, P_θ is a GGMS polytope of the same type as $\Phi \cap (\ker\theta)$. In the former case, we thus say that P_θ is a 2-face of finite type; in the latter, we say that P_θ is a 2-face of affine type.

Now fix a GGMS polytope P and a convex order $\preccurlyeq$ on Φ_+. For $\alpha \in \Phi_+^{\mathrm{re}} \sqcup \{\delta\}$, look at

$$A = \{\beta \in \Phi_+ \mid \beta \succ \alpha\} \quad \text{and} \quad B = \{\beta \in \Phi_+ \mid \beta \succcurlyeq \alpha\}.$$

These are biconvex subsets such that $B = A \sqcup \{\alpha\}$, if α is real, or $B = A \sqcup \big(\mathbb{Z}_{>0}\delta\big)$, if $\alpha = \delta$. In either case, $[\mu_P(A), \mu_P(B)]$ is an edge of P (possibly degenerate) which points in the direction of α, so we may write $\mu_P(B) - \mu_P(A) = n_\alpha \alpha$, where n_α is a non-negative integer (nonzero only if $\alpha \in E_P$). For any $A \in \mathscr{U}(\preccurlyeq)$, we then have

$$\mu_P(A) = \sum_{\alpha \in A \cap (\Phi_+^{\mathrm{re}} \sqcup \{\delta\})} n_\alpha \alpha.$$

The collection of numbers (n_α) will be called the Lusztig datum of P in direction $\preccurlyeq$. We will however later decorate our GGMS polytopes in order to refine the information carried by n_δ, taking into account all the imaginary roots and their multiplicities.

3 Torsion pairs and Harder-Narasimhan polytopes

In this section we will study general facts about torsion pairs and Harder-Narasimhan polytopes. We consider an essentially small abelian category $\mathscr{A}$ such that all objects have finite length. This assumption ensures that the Grothendieck group $\mathbf{K}(\mathscr{A})$ is a free abelian group, with basis the set of isomorphism classes of simple objects. As usual, we denote by $[T]$ the class in $\mathbf{K}(\mathscr{A})$ of an object $T \in \mathscr{A}$. Our subcategories will always be full subcategories.

3.1 Torsion pairs

Following [3], a torsion pair in $\mathscr{A}$ is a pair $(\mathscr{T}, \mathscr{F})$ of two subcategories, called the torsion class and the torsion-free class, that satisfy the following two axioms:

(T1) $\mathrm{Hom}_{\mathscr{A}}(X, Y) = 0$ for each $(X, Y) \in \mathscr{T} \times \mathscr{F}$.

(T2) Each object $T \in \mathscr{A}$ has a subobject X such that $(X, T/X) \in \mathscr{T} \times \mathscr{F}$.

Axiom (T1) forces the subobject X in (T2) to be the largest subobject of T that belongs to $\mathscr{T}$, and a fortiori to be unique; this X is called the torsion subobject of T with respect to the torsion pair $(\mathscr{T}, \mathscr{F})$.

An equivalent set of axioms are the two requirements:

(T'1) $\mathscr{T} = \{X \in \mathscr{A} \mid \forall Y \in \mathscr{F},\ \mathrm{Hom}(X, Y) = 0\}$.

(T'2) $\mathscr{F} = \{Y \in \mathscr{A} \mid \forall X \in \mathscr{T},\ \mathrm{Hom}(X, Y) = 0\}$.

With this second formulation, it is clear that $\mathscr{T}$ is closed under taking quotients and extensions and that $\mathscr{F}$ is closed under taking subobjects and extensions.

Given two torsion pairs $(\mathscr{T}', \mathscr{F}')$ and $(\mathscr{T}'', \mathscr{F}'')$, we write $(\mathscr{T}', \mathscr{F}') \preccurlyeq (\mathscr{T}'', \mathscr{F}'')$ if the following three equivalent conditions hold:

$$\mathscr{T}' \subseteq \mathscr{T}'', \qquad \mathscr{F}' \supseteq \mathscr{F}'', \qquad \mathscr{T}' \cap \mathscr{F}'' = \{0\}.$$

In this case, each object $T \in \mathscr{A}$ is endowed with a three-step filtration $0 \subseteq X' \subseteq X'' \subseteq T$, where X' and X'' are the torsion subobjects of T with respect to $(\mathscr{T}', \mathscr{F}')$ and $(\mathscr{T}'', \mathscr{F}'')$, respectively. Since $\mathscr{F}'$ is stable under taking subobjects and $\mathscr{T}''$ is stable under taking quotients, we have $(X', X''/X', T/X'') \in (\mathscr{T}', \mathscr{F}' \cap \mathscr{T}'', \mathscr{F}'')$.

A typical example of torsion pair is obtained by the following construction, directly translated from the well-known theories of Harder-Narasimhan filtrations and stability conditions [58, 53, 56]. Fix a group homomorphism $\theta : \mathbf{K}(\mathscr{A}) \to \mathbb{R}$ and define five subcategories $\mathscr{I}_\theta$, $\overline{\mathscr{I}}_\theta$, $\mathscr{P}_\theta$, $\overline{\mathscr{P}}_\theta$ and $\mathscr{R}_\theta$ of $\mathscr{A}$:

- An object T is in $\mathscr{I}_\theta$ (respectively, $\overline{\mathscr{I}}_\theta$) if any nonzero quotient X of T satisfies $\theta([X]) > 0$ (respectively, $\theta([X]) \geq 0$).
- An object T is in $\mathscr{P}_\theta$ (respectively, $\overline{\mathscr{P}}_\theta$) if any nonzero subobject X of T satisfies $\theta([X]) < 0$ (respectively, $\theta([X]) \leq 0$).
- An object T is in $\mathscr{R}_\theta$ if $\theta([T]) = 0$ and any nonzero subobject X of T satisfies $\theta([X]) \leq 0$.

The objects in the category $\mathscr{R}_\theta$ are called θ-semistable [40]. Note that $\mathscr{R}_\theta = \overline{\mathscr{I}}_\theta \cap \overline{\mathscr{P}}_\theta$.

Proposition 3.1 *Both $(\mathscr{I}_\theta, \overline{\mathscr{P}}_\theta)$ and $(\overline{\mathscr{I}}_\theta, \mathscr{P}_\theta)$ are torsion pairs in $\mathscr{A}$. The category $\mathscr{R}_\theta$ is an abelian subcategory of $\mathscr{A}$.*

Proof. Let us first prove that $(\mathscr{I}_\theta, \overline{\mathscr{P}}_\theta)$ is a torsion pair. The axiom (T1) is obvious, so we have to prove the axiom (T2).

We first show that $\mathscr{I}_\theta$ is closed under extensions. Let $0 \to T' \to T \xrightarrow{f} T'' \to 0$ be a short exact sequence with T' and T'' in $\mathscr{I}_\theta$ and let $g : T \to X$ be an epimorphism. The pushout of (f, g) then exhibits X as the extension of a quotient X'' of T'' by a quotient X' of T'. By assumption, $\theta([X'])$ and $\theta([X''])$ are both nonnegative, so $\theta([X]) \geq 0$. Moreover, equality holds only if both X' and X'' are zero, thus only if $X = 0$.

Now let $T \in \mathscr{A}$. Our assumption of finite length allows us to pick a maximal element X among the subobjects of T that belong to $\mathscr{I}_\theta$. Suppose that T/X is not in $\overline{\mathscr{P}}_\theta$. Then it contains a subobject Y such that $\theta([Y]) > 0$, and we may assume that Y has been chosen minimal with this property. Certainly, Y does not belong to $\mathscr{I}_\theta$; otherwise, the extension of Y by X inside T would belong to $\mathscr{I}_\theta$, contradicting the maximality of X. So Y has a nonzero quotient Y/Z such that $\theta([Y/Z]) \leq 0$. Since Z is a subobject of T/X properly contained in Y, the minimality of Y requires $\theta([Z]) \leq 0$. We thus reach a contradiction, namely $0 \geq \theta([Z]) + \theta([Y/Z]) = \theta([Y]) > 0$. Therefore $T/X \in \overline{\mathscr{P}}_\theta$, which establishes (T2).

We have thus shown that $(\mathscr{I}_\theta, \overline{\mathscr{P}}_\theta)$ is a torsion pair. The proof for $(\overline{\mathscr{I}}_\theta, \mathscr{P}_\theta)$ is similar. The fact that $\mathscr{R}_\theta$ is an abelian subcategory is well-known. □

Since $(\mathscr{I}_\theta, \overline{\mathscr{P}}_\theta) \preccurlyeq (\overline{\mathscr{I}}_\theta, \mathscr{P}_\theta)$, these two torsion pairs endow each object $T \in \mathscr{A}$ with a three-step filtration $0 \subseteq T_\theta^{\min} \subseteq T_\theta^{\max} \subseteq T$. The quotient $T_\theta^{\max}/T_\theta^{\min}$ belongs to $\mathscr{R}_\theta = \overline{\mathscr{I}}_\theta \cap \overline{\mathscr{P}}_\theta$.

Proposition 3.2 *Let $\theta : \mathbf{K}(\mathscr{A}) \to \mathbb{R}$ be a group homomorphism and let $T \in \mathscr{A}$. Then*

$$\theta([T_\theta^{\min}]) = \theta([T_\theta^{\max}]) \geq \theta([X])$$

for any subobject $X \subseteq T$. Equality holds if only if $T_\theta^{\min} \subseteq X \subseteq T_\theta^{\max}$ and $X/T_\theta^{\min}$ is θ-semistable.

Proof. We adopt the notation of the statement. Let X be a subobject of T. Since $T_\theta^{\min} \in \mathscr{I}_\theta$, we have $\theta([T_\theta^{\min}/(X \cap T_\theta^{\min})]) \geq 0$, with equality only if $T_\theta^{\min} \subseteq X$. Since $T/T_\theta^{\max} \in \mathscr{P}_\theta$, we have $\theta([(X + T_\theta^{\max})/T_\theta^{\max}]) \leq 0$, with equality only if $X \subseteq T_\theta^{\max}$. Lastly, we note that $(X \cap T_\theta^{\max})/(X \cap T_\theta^{\min})$ is a subobject of $T_\theta^{\max}/T_\theta^{\min}$; since the latter is in $\overline{\mathscr{P}}_\theta$, we have $\theta([(X \cap T_\theta^{\max})/(X \cap T_\theta^{\min})]) \leq 0$, with equality if and only if $(X \cap T_\theta^{\max})/(X \cap T_\theta^{\min})$ is θ-semistable. The result now follows from the relation $[X + T_\theta^{\max}] + [X \cap T_\theta^{\max}] = [X] + [T_\theta^{\max}]$. □

3.2 Harder-Narasimhan polytopes

We set $\mathbf{K}(\mathscr{A})_\mathbb{R} = \mathbf{K}(\mathscr{A}) \otimes_\mathbb{Z} \mathbb{R}$. We view this $\mathbb{R}$-vector space as the inductive limit of its finite dimensional subspaces; it is thus a locally convex topological vector space. Linear forms on this vector space are automatically continuous. We denote the canonical pairing between $\mathbf{K}(\mathscr{A})_\mathbb{R}$ and its dual $(\mathbf{K}(\mathscr{A})_\mathbb{R})^*$ by angle brackets. We may regard a linear form $\theta \in (\mathbf{K}(\mathscr{A})_\mathbb{R})^*$ as a group homomorphism $\theta : \mathbf{K}(\mathscr{A}) \to \mathbb{R}$, whence the torsion pairs $(\mathscr{I}_\theta, \overline{\mathscr{P}}_\theta)$ and $(\overline{\mathscr{I}}_\theta, \mathscr{P}_\theta)$ and the abelian subcategory $\mathscr{R}_\theta$.

Given an object $T \in \mathscr{A}$, there are finitely many classes $[X]$ of subobjects $X \subseteq T$. The convex hull in $\mathbf{K}(\mathscr{A})_\mathbb{R}$ of all these points is a convex lattice polytope. We call it the Harder-Narasimhan polytope of T and we denote it by $\mathrm{Pol}(T)$. The support function $\psi_{\mathrm{Pol}(T)}$ of this polytope is defined as the function on $(\mathbf{K}(\mathscr{A})_{\mathbb{R}})^*$ that maps a linear form θ to its maximum on $\mathrm{Pol}(T)$. As in section 2.6, $P_\theta = \{x \in \mathrm{Pol}(T) \mid \langle \theta, x\rangle = \psi_{\mathrm{Pol}(T)}(\theta)\}$ is a face of $\mathrm{Pol}(T)$.

The inclusion $i : \mathscr{R}_\theta \subseteq \mathscr{A}$ is an exact functor, so it induces a group homomorphism $\mathbf{K}(i) : \mathbf{K}(\mathscr{R}_\theta) \to \mathbf{K}(\mathscr{A})$ and a corresponding linear map $\mathbf{K}(i)_{\mathbb{R}}$. Proposition 3.2 then receives the following interpretation.

Corollary 3.3 *Let $\theta \in (\mathbf{K}(\mathscr{A})_{\mathbb{R}})^*$. Let us denote by Q the HN polytope of $T_\theta^{\max}/T_\theta^{\min}$, regarded as an object of $\mathscr{R}_\theta$ (so that $Q \subseteq \mathbf{K}(\mathscr{R}_\theta)_{\mathbb{R}}$). Then*

$$P_\theta = [T_\theta^{\min}] + \mathbf{K}(i)_{\mathbb{R}}(Q).$$

Thus the face P_θ of $\mathrm{Pol}(T)$ is the HN polytope of $T_\theta^{\max}/T_\theta^{\min}$, computed relative to the category $\mathscr{R}_\theta$, and shifted by $[T_\theta^{\min}]$.

A consequence of this observation is that $[T_\theta^{\min}]$ and $[T_\theta^{\max}]$ are vertices of $\mathrm{Pol}(T)$. Another noteworthy consequence is the following rigidity property: if x is a vertex of $\mathrm{Pol}(T)$, then T has a unique subobject X such that $[X] = x$.

We may also interpret the categories $\mathscr{I}_\theta$, $\overline{\mathscr{I}}_\theta$, etc., in terms of HN polytopes. For instance, an object T belongs to $\overline{\mathscr{I}}_\theta$ if and only if $T = T_\theta^{\max}$, hence if and only if the top vertex $[T]$ of $\mathrm{Pol}(T)$ lies on the face defined by θ.

Given T and θ, the subfaces of P_θ are obtained by perturbing slightly θ. The following result states this formally.

Proposition 3.4 *Let $(\theta, \eta) \in \mathrm{Hom}_{\mathbb{Z}}(\mathbf{K}(\mathscr{A}), \mathbb{R})^2$, let $T \in \mathscr{A}$, let $X = T_\theta^{\max}/T_\theta^{\min}$, and let $i : \mathscr{R}_\theta \subseteq \mathscr{A}$ be the inclusion functor. Let $0 \subseteq X_{\eta\circ\mathbf{K}(i)}^{\min} \subseteq X_{\eta\circ\mathbf{K}(i)}^{\max} \subseteq X$ be the filtration of X, regarded as an object in $\mathscr{R}_\theta$, relative to the group homomorphism $\eta \circ \mathbf{K}(i) \in \mathrm{Hom}_{\mathbb{Z}}(\mathbf{K}(\mathscr{R}_\theta), \mathbb{R})$. Call $T_\theta^{\min} \subseteq T' \subseteq T'' \subseteq T_\theta^{\max}$ the pull-back of this filtration by the canonical epimorphism $T_\theta^{\max} \to X$. Then for m large enough, $T' = T_{m\theta+\eta}^{\min}$ and $T'' = T_{m\theta+\eta}^{\max}$.*

Proof. The classes in $\mathbf{K}(\mathscr{A})$ of subquotients of T are finitely many. Pick m large enough so that, for all subquotients Z of T,

$$\theta([Z]) > 0 \implies (m\theta + \eta)([Z]) > 0 \quad \text{and} \quad \theta([Z]) < 0 \implies (m\theta + \eta)([Z]) < 0.$$

Each nonzero quotient Y of $T_\theta^{\min}$ satisfies $\theta([Y]) > 0$, hence satisfies $(m\theta + \eta)([Y]) > 0$. Therefore $T_\theta^{\min} \in \mathscr{I}_{m\theta+\eta}$, and so $T_\theta^{\min} \subseteq T_{m\theta+\eta}^{\min}$. The quotient $U = T_{m\theta+\eta}^{\min}/T_\theta^{\min}$ belongs to $\mathscr{I}_{m\theta+\eta}$ and to $\overline{\mathscr{P}}_\theta$, so we have $(m\theta + \eta)([U]) \geq 0$ and $\theta([U]) \leq 0$, which forces $\theta([U]) = 0$.

In a similar fashion, we see that $T^{\max}_{m\theta+\eta} \subseteq T^{\max}_\theta$ and $\theta([T^{\max}_\theta / T^{\max}_{m\theta+\eta}]) = 0$. We conclude that

$$T^{\min}_\theta \subseteq T^{\min}_{m\theta+\eta} \subseteq T^{\max}_{m\theta+\eta} \subseteq T^{\max}_\theta$$

and that the subquotients of this filtration belong to $\mathscr{R}_\theta$.

Reducing modulo $T^{\min}_\theta$, we get a three-step filtration $0 \subseteq X' \subseteq X'' \subseteq X$ of $X = T^{\max}_\theta / T^{\min}_\theta$, viewed as an object of $\mathscr{R}_\theta$. Any nonzero quotient Y of X' in $\mathscr{R}_\theta$ is a nonzero quotient of $T^{\min}_{m\theta+\eta}$ in $\mathscr{A}$ such that $\theta([Y]) = 0$, and so $\eta([Y]) = (m\theta + \eta)([Y]) > 0$. Therefore X' belongs to the subcategory $\mathscr{I}_{\eta\circ\mathbf{K}(i)}$ of $\mathscr{R}_\theta$. One checks in a similar fashion that $X''/X' \in \mathscr{R}_{\eta\circ\mathbf{K}(i)}$ and $X/X'' \in \mathscr{P}_{\eta\circ\mathbf{K}(i)}$. We conclude that $X' = X^{\min}_{\eta\circ\mathbf{K}(i)}$ and $X'' = X^{\max}_{\eta\circ\mathbf{K}(i)}$. □

Let us now compare this construction to the more usual notion of Harder-Narasimhan filtration. To define the latter, we need to fix a pair $(\eta, \theta) \in \operatorname{Hom}_{\mathbb{Z}}(\mathbf{K}(\mathscr{A}), \mathbb{R})^2$ such that $\eta([T]) > 0$ for each nonzero object T. The slope of a nonzero object $T \in \mathscr{A}$ is defined as $\mu(T) = \theta([T])/\eta([T])$ and an object T is called semistable if it is zero or if $\mu(X) \leq \mu(T)$ for any nonzero subobject $X \subseteq T$. It can then be shown that any object $T \in \mathscr{A}$ has a finite filtration

$$0 = T_0 \subset T_1 \subset \cdots \subset T_{\ell-1} \subset T_\ell = T \tag{3.1}$$

whose subquotients are nonzero and semistable, with moreover $\mu(T_k/T_{k-1})$ decreasing with k (see for instance [53], close to the present context). This filtration is unique and is called the Harder-Narasimhan filtration of T.

Given $a \in \mathbb{R}$, the torsion subobject of T with respect to the torsion pair $(\mathscr{I}_{\theta-a\eta}, \overline{\mathscr{P}}_{\theta-a\eta})$ is T_k, where k is the largest index such that $\mu(T_k/T_{k-1}) > a$. In our former notation, this means that $T_k = T^{\min}_{\theta-a\eta}$; in particular, $[T_k]$ is a vertex of $\operatorname{Pol}(T)$. One can even be more precise: the linear map $\varphi : \mathbf{K}(\mathscr{A})_{\mathbb{R}} \to \mathbb{R}^2$ given in coordinates as (η, θ) projects $\operatorname{Pol}(T)$ to a convex polygon of the plane, and the upper ridge of this polygon is the polygonal line going successively through the points $\varphi([T_k])$, for $0 \leq k \leq \ell$. We leave the proof of this fact to the reader.

The polygonal line just obtained is what Shatz calls the HN polygon [58]. Thus our HN polytopes are a multidimensional analog of those HN polygons; they simply take into account the existence of a whole space of stability conditions. There may well exist sensible adaptations of this notion of HN polytope to other contexts where spaces of stability conditions have been defined (see for instance [13]).

Remarks 3.5. (i) In Corollary 3.3, the map $\mathbf{K}(i)_{\mathbb{R}}$ induces an actual loss of information. For example, in our study of imaginary edges (sections 1.6 and 7.4), the category $\mathscr{R}_\theta$ has infinitely many simple objects; since they all have the same dimension-vector, their classes have the same image by $\mathbf{K}(i)_{\mathbb{R}}$.

(ii) The HN polytope of the direct sum of two objects is the Minkowski sum of the HN polytopes of the two objects.

3.3 Nested families of torsion pairs

The subobjects of T that appear in the HN filtration (3.1) are the torsion subobjects with respect to the torsion pairs $(\mathscr{I}_{\theta+a\eta}, \overline{\mathscr{P}}_{\theta+a\eta})$, as a varies over $\mathbb{R}$. Observe that, in the notation of section 3.1,

$$\forall(a,b) \in \mathbb{R}^2, \qquad a \leq b \implies (\mathscr{I}_{\theta+a\eta}, \overline{\mathscr{P}}_{\theta+a\eta}) \preccurlyeq (\mathscr{I}_{\theta+b\eta}, \overline{\mathscr{P}}_{\theta+b\eta}).$$

This prompts the following definition: a nested family of torsion pairs is the datum of a family $(\mathscr{T}_a, \mathscr{F}_a)_{a\in A}$ of torsion pairs, indexed by a totally ordered set A, such that

$$\forall(a,b) \in A^2, \qquad a \leq b \implies (\mathscr{T}_a, \mathscr{F}_a) \preccurlyeq (\mathscr{T}_b, \mathscr{F}_b).$$

This definition is certainly less general than Rudakov's study [56] but is sufficient for our purposes.

A nested family of torsion pairs $(\mathscr{T}_a, \mathscr{F}_a)_{a\in A}$ in $\mathscr{A}$ induces a non-decreasing filtration $(T_a)_{a\in A}$ on any object $T \in \mathscr{A}$: simply define T_a as the torsion subobject of T with respect to $(\mathscr{T}_a, \mathscr{F}_a)$. As already shown in section 3.1, the object T_b/T_a is in $\mathscr{F}_a \cap \mathscr{T}_b$ whenever $a \leq b$.

4 Background on preprojective algebras

4.1 Basic definitions

We fix a base field K, which we assume for convenience to be algebraically closed of characteristic 0. As in section 2.1, we fix a graph (I, E), where I is the set of vertices and E the set of edges. We denote by H the set of oriented edges of this graph. Thus each edge in E gives birth to two oriented edges in H, and H comes with a source map $s : H \to I$, a target map $t : H \to I$ and a fixed-point free involution $a \mapsto \overline{a}$ such that $s(a) = t(\overline{a})$ for each $a \in H$.

An orientation is a subset $\Omega \subset H$ such that $H = \Omega \sqcup \overline{\Omega}$. Such an orientation yields a quiver $Q = (I, \Omega, s, t)$, and then $\overline{Q} = (I, H, s, t)$ is the double quiver of Q. We set $\varepsilon(a) = 1$ if $a \in \Omega$ and $\varepsilon(a) = -1$ if $a \notin \Omega$.

Let $K\overline{Q}$ be the path algebra of $\overline{Q}$. The linear span $\mathbf{S} = \operatorname{span}_K(e_i)_{i\in I}$ of the lazy paths is a commutative semisimple subalgebra of $K\overline{Q}$. The linear span $\mathbf{A} = \operatorname{span}_K(a)_{a\in H}$ of the paths of length one is an $\mathbf{S}$-$\mathbf{S}$-bimodule. Then $K\overline{Q}$ is the tensor algebra $T_{\mathbf{S}}\mathbf{A}$. For $i \in I$, set

$$\rho_i = \sum_{\substack{a \in H \\ s(a)=i}} \varepsilon(a)\overline{a}a,$$

the so-called preprojective relation at vertex i. The linear span $\mathbf{R} = \operatorname{span}_K(\rho_i)_{i\in I}$ is a $\mathbf{S}$-$\mathbf{S}$-subbimodule of $K\overline{Q}$.

By definition, the preprojective algebra of Q is the quotient of $K\overline{Q}$ by the ideal generated by $\mathbf{R}$. This is an augmented algebra over $\mathbf{S}$. Its completion with respect to the augmentation ideal is called the completed preprojective algebra and is denoted by Λ_Q. For brevity, we will generally drop the Q in the notation Λ_Q. Completing has the effect that the augmentation ideal becomes the Jacobson radical; thus Λ quotiented by its Jacobson radical is isomorphic to $\mathbf{S}$, and the simple Λ-modules are just the simple $\mathbf{S}$-modules, namely the one dimensional modules S_i. We denote by Λ-mod the category of finite dimensional left Λ-modules.

The involution $a \mapsto \overline{a}$ on the set H of oriented edges induces an anti-automorphism of Λ. If M is a finite dimensional left Λ-module, then we denote by M^* the dual module $\operatorname{Hom}_K(M, K)$, viewed as a left module by means of this anti-automorphism.

Occasionally, we will have to write Λ-modules in a concrete fashion. Our notation is as follows. An $\mathbf{S}$-module M is an I-graded vector space $M = \bigoplus_{i\in I} M_i$, where $M_i = e_iM$; we define the dimension-vector of M to be the element $\underline{\dim}\, M = \sum_{i\in I}(\dim M_i)\alpha_i$ in $\mathbb{N}I$. A $K\overline{Q}$-module M is the datum of an $\mathbf{S}$-module and of linear maps $M_a : M_{s(a)} \to M_{t(a)}$ for each $a \in H$.

Since the simple Λ-modules are the modules S_i, concentrated at a vertex of the quiver, it is natural to present a special notation designed to analyze a Λ-module M locally around a vertex i. Specifically, we break the datum of M in two parts: the first part consists of the vector spaces M_j for $j \neq i$ and of the linear maps between them; the second part consists of the vector spaces and of the linear maps that appear in the diagram

$$\bigoplus_{\substack{a \in H \\ s(a)=i}} M_{t(a)} \xrightarrow{(M_{\overline{a}})} M_i \xrightarrow{(\varepsilon(a)M_a)} \bigoplus_{\substack{a \in H \\ s(a)=i}} M_{t(a)}.$$

Here $(M_{\overline{a}})$ denotes a column-matrix, whose lines are indexed by $\{a \in H \mid s(a) = i\}$; likewise, $(\varepsilon(a)M_a)$ denotes a row-matrix.

For brevity, we will write this diagram as

$$\widetilde{M_i} \xrightarrow{M_{\mathrm{in}(i)}} M_i \xrightarrow{M_{\mathrm{out}(i)}} \widetilde{M_i}. \tag{4.1}$$

With this notation, the preprojective relation at i is $M_{\mathrm{in}(i)}M_{\mathrm{out}(i)} = 0$.

We define the i-socle of M as the largest submodule of M that is isomorphic to a direct sum of copies of S_i; we denote it by $\operatorname{soc}_i M$. It is concentrated at the vertex i and it can be identified with the vector space $\ker M_{\mathrm{out}(i)}$. Likewise, the largest quotient of M that is isomorphic to a direct sum of copies of S_i is called the i-head of M and is denoted by $\operatorname{hd}_i M$. It is concentrated at the vertex i and can be identified with the vector space $\operatorname{coker} M_{\mathrm{in}(i)}$.

The assignment $[M] \mapsto \underline{\dim}\, M$ provides an isomorphism between the Grothendieck group $\mathbf{K}(\Lambda\text{-mod})$ and the root lattice $\mathbb{Z}I$. Crawley-Boevey's formula (Lemma 1 in [17]) gives a module-theoretic meaning to the bilinear form on the root lattice:

$$\dim \operatorname{Hom}_\Lambda(M,N) + \dim \operatorname{Hom}_\Lambda(N,M) - \dim \operatorname{Ext}^1_\Lambda(M,N) = \big(\underline{\dim}\, M,\, \underline{\dim}\, N\big) \tag{4.2}$$

for any finite dimensional Λ-modules M and N.

Remark 4.1. The HN polytope $\operatorname{Pol}(T)$ of an object $T \in \Lambda$-mod lives in $\mathbf{K}(\Lambda\text{-mod})_\mathbb{R} \cong \mathbb{R}I$. Since the duality exchanges submodules and quotients and leaves the dimension-vector unchanged, $\operatorname{Pol}(T^*)$ is the image of $\operatorname{Pol}(T)$ under the involution $x \mapsto \underline{\dim}\, T - x$ of $\mathbb{R}I$. We leave it to the reader to check the equalities $\mathscr{I}_{-\theta} = (\mathscr{P}_\theta)^*$, $\overline{\mathscr{I}}_{-\theta} = (\overline{\mathscr{P}}_\theta)^*$ and $\mathscr{R}_{-\theta} = (\mathscr{R}_\theta)^*$, for any $\theta \in \operatorname{Hom}_\mathbb{Z}(\mathbf{K}(\Lambda\text{-mod}), \mathbb{R})$.

4.2 Projective resolutions

We now recall Geiß, Leclerc and Schröer's description of the extension groups in the category Λ-mod (see [26], section 8).

Consider the complex of Λ-bimodules

$$\Lambda \otimes_\mathbf{S} \mathbf{R} \otimes_\mathbf{S} \Lambda \xrightarrow{d_1} \Lambda \otimes_\mathbf{S} \mathbf{A} \otimes_\mathbf{S} \Lambda \xrightarrow{d_0} \Lambda \otimes_\mathbf{S} \mathbf{S} \otimes_\mathbf{S} \Lambda \to \Lambda \to 0, \tag{4.3}$$

where the map on the right is multiplication in Λ, where for each $a \in H$

$$d_0(1 \otimes a \otimes 1) = a \otimes e_{s(a)} \otimes 1 - 1 \otimes e_{t(a)} \otimes a,$$

and where for each $i \in I$

$$d_1(1 \otimes \rho_i \otimes 1) = \sum_{\substack{a \in H \\ s(a) = i}} \varepsilon(a)(\overline{a} \otimes a \otimes 1 + 1 \otimes \overline{a} \otimes a).$$

Then (4.3) is the beginning of a projective resolution of Λ, by [26], Lemma 8.1.1.

Given $M, N \in \Lambda$-mod, one can apply $\operatorname{Hom}_\Lambda(? \otimes_\Lambda M, N)$ to (4.3). One then obtains the complex

$$0 \to \bigoplus_{i \in I} \operatorname{Hom}_K(M_i, N_i) \xrightarrow{d^0_{M,N}} \bigoplus_{a \in H} \operatorname{Hom}_K(M_{s(a)}, N_{t(a)}) \xrightarrow{d^1_{M,N}} \bigoplus_{i \in I} \operatorname{Hom}_K(M_i, N_i), \qquad (4.4)$$

where

$$d^0_{M,N} : (f_i)_{i \in I} \mapsto \big(N_a f_{s(a)} - f_{t(a)} M_a\big)_{a \in H}$$

and

$$d^1_{M,N} : (g_a)_{a \in H} \mapsto \left(\sum_{\substack{a \in H \\ s(a) = i}} \varepsilon(a) \big(N_{\overline{a}} g_a + g_{\overline{a}} M_a\big) \right)_{i \in I}.$$

Thus for $k \in \{0, 1\}$, the extension group $\operatorname{Ext}^k_\Lambda(M, N)$ can be identified to the cohomology groups in degree k of the complex (4.4).

In [26], section 8.2, Geiß, Leclerc and Schröer explain that in this identification, the bilinear map

$$\tau_1 : \left(\bigoplus_{a \in H} \operatorname{Hom}_K(M_{s(a)}, N_{t(a)}) \right) \times \left(\bigoplus_{a \in H} \operatorname{Hom}_K(N_{s(a)}, M_{t(a)}) \right) \to K$$

defined by

$$\tau_1\big((g_a), (h_a)\big) = \sum_{i \in I} \operatorname{Tr} \left(\sum_{\substack{a \in H \\ s(a) = i}} \varepsilon(a) g_{\overline{a}} h_a \right)$$

induces a non-degenerate pairing between $\operatorname{Ext}^1_\Lambda(M, N)$ and $\operatorname{Ext}^1_\Lambda(N, M)$. Note that because of the cyclicity of the trace and of the presence of the signs $\varepsilon(a)$, this pairing τ_1 is antisymmetric.

One sees likewise that the bilinear map

$$\tau_2 : \left(\bigoplus_{i \in I} \operatorname{Hom}_K(M_i, N_i) \right) \times \left(\bigoplus_{i \in I} \operatorname{Hom}_K(N_i, M_i) \right) \to K$$

defined by

$$\tau_2\big((f_i), (h_i)\big) = \sum_{i \in I} \operatorname{Tr}(f_i h_i)$$

induces a non-degenerate pairing between $\operatorname{coker} d^1_{M,N}$ and $\operatorname{Hom}_\Lambda(N, M)$.

4.3 Lusztig's nilpotent varieties

Given a dimension-vector $\nu = \sum_{i \in I} \nu_i \alpha_i$ in $\mathbb{N}I$, we set

$$\operatorname{Rep}_K(\overline{Q}, \nu) = \bigoplus_{a \in H} \operatorname{Hom}_K(K^{\nu_{s(a)}}, K^{\nu_{t(a)}}).$$

A structure of Λ-module on the I-graded vector space $\bigoplus_{i \in I} K^{\nu_i}$ is specified by linear maps $T_a : K^{\nu_{s(a)}} \to K^{\nu_{t(a)}}$, for each $a \in H$, that is, by a point $T = (T_a)_{a \in H}$ in $\operatorname{Rep}_K(\overline{Q}, \nu)$. To have an action of the completed preprojective algebra, we must moreover impose the preprojective relations and the nilpotency condition. These equations define a subvariety

$$\Lambda(\nu) \subseteq \operatorname{Rep}_K\bigl(\overline{Q}, \nu\bigr),$$

called the affine variety of representations of Λ or Lusztig's nilpotent variety.

The group $G(\nu) = \prod_{i \in I} \mathrm{GL}_{\nu_i}(K)$ acts by conjugation on $\operatorname{Rep}_K\bigl(\overline{Q}, \nu\bigr)$. This action preserves the nilpotent variety $\Lambda(\nu)$. Then any isomorphism class of a Λ-module of dimension-vector ν can be regarded as a $G(\nu)$-orbit in $\Lambda(\nu)$.

In [43], Lusztig shows that $\Lambda(\nu)$ is a Lagrangian subvariety of $\operatorname{Rep}_K\bigl(\overline{Q}, \nu\bigr)$; in particular, all the irreducible components of $\Lambda(\nu)$ have dimension $\dim\bigl(\operatorname{Rep}_K\bigl(\overline{Q}, \nu\bigr)\bigr)/2$. A straightforward calculation shows that this common dimension is

$$\dim \Lambda(\nu) = \dim G(\nu) - (\nu, \nu)/2. \tag{4.5}$$

A Λ-module M is called rigid if $\operatorname{Ext}^1_\Lambda(M, M) = 0$. An immediate consequence of Crawley-Boevey's formula (see Corollary 3.15 in [25]) is that the closure of the $G(\nu)$-orbit through a point $T \in \Lambda(\nu)$ is an irreducible component if and only if T is a rigid module.

As in the introduction, we denote by $\mathfrak{B}(\nu) = \operatorname{Irr} \Lambda(\nu)$ the set of irreducible components of the nilpotent variety and we define $\mathfrak{B} = \bigsqcup_{\nu \in \mathbb{N}I} \mathfrak{B}(\nu)$. This set $\mathfrak{B}$ is endowed with a crystal structure, defined by Lusztig ([42], section 8), which we quickly recall.

The weight of an element $Z \in \mathfrak{B}(\nu)$ is $\operatorname{wt} Z = \nu$. The number $\varphi_i(Z)$ is the dimension of the i-head of a general point $T \in Z$. The number $\varepsilon_i(Z)$ can then be found be the general formula $\varphi_i(Z) - \varepsilon_i(Z) = \langle \alpha_i^\vee, \operatorname{wt} Z \rangle$. The operators $\tilde{e}_i$ and $\tilde{f}_i$ add or remove a copy of S_i in general position at the top of a module $T \in Z$. In other words, the relationship $Z' = \tilde{e}_i Z$ corresponds to extensions $0 \to T \to T' \to S_i \to 0$ as general as possible: if T runs over a dense open subset of Z, then T' will also run over a dense open subset of Z', and vice versa. The duality $*$ corresponds to the involution $Z \mapsto Z^*$ on $\mathfrak{B}$, which preserves the weight. We refer the reader to the literature for the formal definitions.

Kashiwara and Saito show in [38] that the crystal $\mathfrak{B}$ is isomorphic to the crystal $B(-\infty)$ of $U_q(\mathfrak{n}_+)$. This isomorphism is canonical, because the only endomorphism of the crystal $B(-\infty)$ is the identity. For $b \in B(-\infty)$, we write Λ_b for the image of b under this isomorphism. Moreover, the crystal $B(-\infty)$ is endowed with an involution, usually also denoted by $*$ (see [37], §8.3), and the proof in [38] shows that the isomorphism $B(-\infty) \to \mathfrak{B}$ commutes with both involutions — thus, the use of the same notation $*$ does not lead to any confusion.

4.4 The canonical decomposition of a component

In this section, we quickly recall Crawley-Boevey and Schröer's results [18] on the canonical decomposition of irreducible components of module varieties, specializing their results to the case of nilpotent varieties.

Let ν' and ν'' be two dimension-vectors. The function $(T', T'') \mapsto \dim \operatorname{Hom}_\Lambda(T', T'')$ on $\Lambda(\nu') \times \Lambda(\nu'')$ is upper semicontinuous. Given $Z' \in \operatorname{Irr} \Lambda(\nu')$ and $Z'' \in \operatorname{Irr} \Lambda(\nu'')$, we denote its minimum on $Z' \times Z''$ by $\hom_\Lambda(Z', Z'')$. Then $\hom_\Lambda(Z', Z'') = \dim \operatorname{Hom}_\Lambda(T', T'')$ for (T', T'') general in $Z' \times Z''$.

Likewise, the function $(T', T'') \mapsto \dim \operatorname{Ext}^1_\Lambda(T', T'')$ on $\Lambda(\nu') \times \Lambda(\nu'')$ is upper semicontinuous. Given $Z' \in \operatorname{Irr} \Lambda(\nu')$ and $Z'' \in \operatorname{Irr} \Lambda(\nu'')$, we denote its minimum on $Z' \times Z''$ by $\operatorname{ext}^1_\Lambda(Z', Z'')$. Then $\operatorname{ext}^1_\Lambda(Z', Z'') = \dim \operatorname{Ext}^1_\Lambda(T', T'')$ for (T', T'') general in $Z' \times Z''$.

Let $n \geq 1$, let $\nu_1, \ldots, \nu_n$ be dimension-vectors, and for $1 \leq j \leq n$, let $Z_j \in \operatorname{Irr} \Lambda(\nu_j)$. Set $\nu = \nu_1 + \cdots + \nu_n$ and denote by $Z_1 \oplus \cdots \oplus Z_n$ the set of all modules in $\Lambda(\nu)$ that are isomorphic to a direct sum $T_1 \oplus \cdots \oplus T_n$, with $T_j \in Z_j$ for all $1 \leq j \leq n$. This is an irreducible subset of $\Lambda(\nu)$. Its closure $\overline{Z_1 \oplus \cdots \oplus Z_n}$ is an irreducible component of $\Lambda(\nu)$ if and only if $\operatorname{ext}^1_\Lambda(Z_j, Z_k) = 0$ for all $j \neq k$.

Conversely, for any $Z \in \operatorname{Irr} \Lambda(\nu)$, there exists n, ν_j and Z_j as above such that the general point in Z_j is an indecomposable Λ-module and

$$Z = \overline{Z_1 \oplus \cdots \oplus Z_n}.$$

Furthermore, the Z_j are unique up to permutation. This is called the canonical decomposition of Z.

4.5 Torsion pairs in Λ-mod

Here we consider the constructions of section 3 in the category Λ-mod. Since we are primarily interested in the crystal $B(-\infty)$, we need to make sure that our constructions go down to the level of irreducible components of the nilpotent varieties.

Proposition 4.2 *Let $\nu \in \mathbb{N}I$ be a dimension-vector and let $\theta \in (\mathbb{R}I)^*$.*

(i) For each $\xi \in \mathbb{N}I$, the set of all $T \in \Lambda(\nu)$ that contain a submodule of dimension-vector ξ is closed.

(ii) There are finitely many polytopes $\operatorname{Pol}(T)$*, for $T \in \Lambda(\nu)$. For each polytope $P \subseteq \mathbb{R}I$, the set $\{T \in \Lambda(\nu) \mid \operatorname{Pol}(T) = P\}$ is constructible.*

(iii) For each category $\mathscr{C}$ among $\mathscr{I}_\theta$, $\overline{\mathscr{I}}_\theta$, $\mathscr{P}_\theta$, $\overline{\mathscr{P}}_\theta$ and $\mathscr{R}_\theta$, the subset $\{T \in \Lambda(\nu) \mid T \in \mathscr{C}\}$ is open in $\Lambda(\nu)$.

Proof. When we view a point $T \in \Lambda(\nu)$ as a Λ-module, we tacitly agree that the underlying I-graded vector space of this module is $\bigoplus_{i\in I} K^{\nu_i}$. Let X be the set of all its I-graded vector subspaces $\bigoplus_{i\in I} V_i$ of dimension-vector ξ; as a product of Grassmannians, X is naturally endowed with the structure of a smooth projective variety. The incidence variety Y consisting of all pairs $(T, V) \in \Lambda(\nu) \times X$ such that $T_a(V_{s(a)}) \subseteq V_{t(a)}$ for all $a \in H$ is closed. The first projection $Y \to \Lambda(\nu)$ is therefore a projective morphism, so is proper. Its image is therefore closed, which shows (i).

Let $R = (\mathbb{N}I) \cap (\nu - \mathbb{N}I)$; this is a finite set. If $T \in \Lambda(\nu)$, then the dimension-vectors of the submodules of T form a subset $S(T)$ of R. Assertion (i) says that $\{T \in \Lambda(\nu) \mid \xi \in S(T)\}$ is closed for each $\xi \in R$. This implies that $\{T \in \Lambda(\nu) \mid S(T) = S\}$ is locally closed for each subset $S \subseteq R$. Gathering these locally closed subsets according to the convex hull of S, we obtain assertion (ii).

According to a remark following Corollary 3.3, a point $T \in \Lambda(\nu)$ belongs to $\overline{\mathscr{I}}_\theta$ if and only if ν lies on the face of $\operatorname{Pol}(T)$ defined by θ. This condition means that $S(T)$ does not meet $\{\xi \in R \mid \langle\theta, \xi\rangle > \langle\theta, \nu\rangle\}$. Thus assertion (i) exhibits $\{T \in \Lambda(\nu) \mid T \in \overline{\mathscr{I}}_\theta\}$ as a finite intersection of open subsets of $\Lambda(\nu)$. This shows the case $\mathscr{C} = \overline{\mathscr{I}}_\theta$ in assertion (iii). The other cases are dealt with in a similar fashion. □

Now let us fix a torsion pair $(\mathscr{T}, \mathscr{F})$ in Λ-mod. All torsion submodules mentioned hereafter in this section are taken with respect to it. We make the following assumption:

(O) For each $\nu \in \mathbb{N}I$, both sets $\{T \in \Lambda(\nu) \mid T \in \mathscr{T}\}$ and $\{T \in \Lambda(\nu) \mid T \in \mathscr{F}\}$ are open.

Under this assumption, it is legitimate to consider the set $\mathfrak{T}(\nu)$ (respectively, $\mathfrak{F}(\nu)$) of all irreducible components of $\Lambda(\nu)$ whose general point belongs to $\mathscr{T}$ (respectively, $\mathscr{F}$). We then get two subsets

$$\mathfrak{T} = \bigsqcup_{\nu\in\mathbb{N}I} \mathfrak{T}(\nu) \quad \text{and} \quad \mathfrak{F} = \bigsqcup_{\nu\in\mathbb{N}I} \mathfrak{F}(\nu)$$

of $\mathfrak{B}$. Our aim now is to construct a bijection $\Xi : \mathfrak{T} \times \mathfrak{F} \to \mathfrak{B}$ that reflects at the component level the decomposition of Λ-modules provided by $(\mathscr{T}, \mathscr{F})$.

Let $(\nu_t, \nu_f) \in (\mathbb{N}I)^2$. We set $\nu = \nu_t + \nu_f$ and define $\Lambda^{\mathscr{T}}(\nu_t) = \{T_t \in \Lambda(\nu_t) \mid T_t \in \mathscr{T}\}$ and $\Lambda^{\mathscr{F}}(\nu_f) = \{T_f \in \Lambda(\nu_f) \mid T_f \in \mathscr{F}\}$. We define $\Theta(\nu_t, \nu_f)$ as the set of all tuples (T, T_t, T_f, f, g) such that $T \in \Lambda(\nu)$, $(T_t, T_f) \in \Lambda^{\mathscr{T}}(\nu_t) \times \Lambda^{\mathscr{F}}(\nu_f)$, and $0 \to T_t \xrightarrow{f} T \xrightarrow{g} T_f \to 0$ is an exact sequence in Λ-mod. This is a quasi-affine algebraic variety. We can then form the diagram

$$\Lambda^{\mathscr{T}}(\nu_t) \times \Lambda^{\mathscr{F}}(\nu_f) \xleftarrow{p} \Theta(\nu_t, \nu_f) \xrightarrow{q} \Lambda(\nu) \tag{4.6}$$

in which p and q are the obvious projections.

Lemma 4.3 *The map p is a locally trivial fibration with a smooth and connected fiber of dimension $\dim G(\nu) - (\nu_t, \nu_f)$. The image of q is the set of all points $T \in \Lambda(\nu)$ whose torsion submodule has dimension-vector ν_t. The non-empty fibers of q are isomorphic to $G(\nu_t) \times G(\nu_f)$.*

Proof. The statements concerning q are obvious, so we only have to deal with p.

The points (T_a), $(T_{t,a})$ and $(T_{f,a})$, chosen in the nilpotent varieties $\Lambda(\nu)$, $\Lambda^{\mathscr{T}}(\nu_t)$ and $\Lambda^{\mathscr{F}}(\nu_f)$, define Λ-module structures on the vector spaces $\bigoplus_{i \in I} K^{\nu_i}$, $\bigoplus_{i \in I} K^{\nu_{t,i}}$ and $\bigoplus_{i \in I} K^{\nu_{f,i}}$.

Let first consider the complex (4.4) from section 4.2, with the Λ-modules T_f and T_t in place of M and N. The maps $d^0_{T_f,T_t}$ and $d^1_{T_f,T_t}$ of the complex depend on the datum of the arrows $(T_{t,a}) \in \Lambda^{\mathscr{T}}(\nu_t)$ and $(T_{f,a}) \in \Lambda^{\mathscr{F}}(\nu_f)$, but the spaces of the complex depend only on ν_f and ν_t. The map $d^0_{T_f,T_t}$ has rank $\dim \operatorname{Hom}_{\mathbf{S}}(T_f, T_t) - \dim \operatorname{Hom}_\Lambda(T_f, T_t)$. In addition, $\operatorname{Ext}^1_\Lambda(T_f, T_t) \cong \ker d^1_{T_f,T_t} / \operatorname{im} d^0_{T_f,T_t}$. Using Crawley-Boevey's formula (4.2) and using the axiom (T1) of torsion pairs, we easily compute

$$\dim \ker d^1_{M,N} = \dim \operatorname{Ext}^1_\Lambda(T_f, T_t) + \operatorname{rk} d^0_{T_f,T_t} = \dim \operatorname{Hom}_{\mathbf{S}}(T_f, T_t) - (\nu_f, \nu_t).$$

Remarkably, this dimension depends only on ν_f and ν_t, and not on the datum of the arrows $(T_{t,a})$ and $(T_{f,a})$.

Let E be the set of all exact sequences $0 \to T_t \xrightarrow{f} T \xrightarrow{g} T_f \to 0$ of I-graded vector spaces. This is a homogeneous space for the group $G(\nu)$ and the stabilizer of a point (f, g) is $\{\mathrm{id} + fhg \mid h \in \operatorname{Hom}_{\mathbf{S}}(T_f, T_t)\}$. It is thus a smooth connected variety of dimension $\dim G(\nu) - \dim \operatorname{Hom}_{\mathbf{S}}(T_f, T_t)$.

The fiber of p over a point $((T_{t,a}), (T_{f,a})) \in \Lambda^{\mathscr{T}}(\nu_t) \times \Lambda^{\mathscr{F}}(\nu_f)$ consists of the datum of $(f, g) \in E$ and of $(T_a) \in \Lambda(\nu)$, with a compatibility condition between the two. The datum of (f, g) corresponds to a trivial fiber bundle over $\Lambda^{\mathscr{T}}(\nu_t) \times \Lambda^{\mathscr{F}}(\nu_f)$ with fiber E. Let us now examine how (T_a) can be chosen when $((T_{t,a}), (T_{f,a}))$ and (f, g) are given.

Once chosen an I-graded complementary subspace of $\ker g$ in the vector space T, the set of possible choices for (T_a) is isomorphic to $\ker d^1_{T_f,T_t}$; moreover, the isomorphism depends smoothly on (f,g). The linear map $d^1_{T_f,T_t}$ depends smoothly on $((T_{t,a}),(T_{f,a}))$ and has constant rank, as we have seen above, so its kernel depends smoothly on $((T_{t,a}),(T_{f,a}))$. In this fashion, we eventually see that the set of possible choices for (T_a) depends smoothly on $((T_{t,a}),(T_{f,a}),f,g)$. Choosing trivializations where needed, we conclude that p is a locally trivial fibration.

Finally, the dimension of the fibers of p is the sum of two contributions, namely $\dim E$ and $\dim\ker d^1_{T_f,T_t}$. We find $\dim G(\nu)-(\nu_f,\nu_t)$, as announced. $\square$

Let $(Z_t,Z_f)\in\operatorname{Irr}\Lambda^{\mathscr{T}}(\nu_t)\times\operatorname{Irr}\Lambda^{\mathscr{F}}(\nu_f)$. In view of Lemma 4.3, $p^{-1}(Z_t\times Z_f)$ is an irreducible component of $\Theta(\nu_t,\nu_f)$. Then $Z=q(p^{-1}(Z_t\times Z_f))$ is an irreducible subset of $\Lambda(\nu)$ which, by e.g. I, §8, Theorem 3 in [48], has dimension

$$\dim(Z_t\times Z_f)+(\dim G(\nu)-(\nu_t,\nu_f))-\dim(G(\nu_t)\times G(\nu_f)).$$

Equation (4.5) shows that this dimension is equal to that of $\Lambda(\nu)$, so $\overline{Z}\in\operatorname{Irr}\Lambda(\nu)$.

This construction defines a map $(\overline{Z_t},\overline{Z_f})\mapsto\overline{Z}$ from $\mathfrak{T}(\nu_t)\times\mathfrak{F}(\nu_f)$ to $\mathfrak{B}(\nu)$. Gluing these maps for all possible (ν_t,ν_f), we eventually get a map $\Xi:\mathfrak{T}\times\mathfrak{F}\to\mathfrak{B}$.

Theorem 4.4 *The map $\Xi:\mathfrak{T}\times\mathfrak{F}\to\mathfrak{B}$ is bijective.*

Proof. Given ν_t and ν_f, Ξ is a bijection from $\mathfrak{T}(\nu_t)\times\mathfrak{F}(\nu_f)$ onto the set of irreducible components of $\overline{q(\Theta(\nu_t,\nu_f))}$.

Now we take $\nu\in\mathbb{N}I$. We consider the diagrams (4.6) for all ν_t and ν_f such that $\nu_t+\nu_f=\nu$. In this fashion, we split $\Lambda(\nu)$ according to the dimension-vector of the torsion submodule:

$$\Lambda(\nu)=\bigsqcup_{\nu_t+\nu_f=\nu}q(\Theta(\nu_t,\nu_f)).$$

Each piece of this partition is constructible, therefore an irreducible component of $\Lambda(\nu)$ is contained in one and only one closure $\overline{q(\Theta(\nu_t,\nu_f))}$.

Taking the union, we see that Ξ defines a bijection from $\bigsqcup_{\nu_t+\nu_f=\nu}(\mathfrak{T}(\nu_t)\times\mathfrak{F}(\nu_f))$ onto $\mathfrak{B}(\nu)$. $\square$

The construction of Ξ implies that, if T is a general point of $\overline{Z}$, then the torsion submodule X of T has dimension-vector ν_t and the point $(X,T/X)$ is general in $\overline{Z_t}\times\overline{Z_f}$. To see this, take a $G(\nu_t)$-invariant dense open subset $U_t \subseteq Z_t$ and a $G(\nu_f)$-invariant dense open subset of $U_f \subseteq Z_f$. Then $p^{-1}(U_t \times U_f)$ is dense in $p^{-1}(Z_t \times Z_f)$. The subset $q(p^{-1}(U_t \times U_f))$ is thus dense in the irreducible set $\overline{Z}$, and is constructible by Chevalley's theorem, so it contains a dense open subset U of $\overline{Z}$. By construction, if T belongs to U, then $\underline{\dim}\, X = \nu_t$ and $(X, T/X)$ belongs to the prescribed open subset $U_t \times U_f$, as desired.

Suppose now that we are given two torsion pairs $(\mathscr{T}', \mathscr{F}')$ and $(\mathscr{T}'', \mathscr{F}'')$ in Λ-mod that both satisfy the openness condition (O). They give rise to subsets $\mathfrak{T}'$, $\mathfrak{F}'$, $\mathfrak{T}''$ and $\mathfrak{F}''$ of $\mathfrak{B}$ and to bijections $\Xi' : \mathfrak{T}' \times \mathfrak{F}' \to \mathfrak{B}$ and $\Xi'' : \mathfrak{T}'' \times \mathfrak{F}'' \to \mathfrak{B}$.

Proposition 4.5 *Assume that $(\mathscr{T}', \mathscr{F}') \preccurlyeq (\mathscr{T}'', \mathscr{F}'')$. Then the map Ξ' restricts to a bijection $\mathfrak{T}' \times (\mathfrak{F}' \cap \mathfrak{T}'') \to \mathfrak{T}''$, the map Ξ'' restricts to a bijection $(\mathfrak{F}' \cap \mathfrak{T}'') \times \mathfrak{F}'' \to \mathfrak{F}'$, and we have a commutative diagram*

$$\begin{tikzcd} \mathfrak{T}' \times (\mathfrak{F}' \cap \mathfrak{T}'') \times \mathfrak{F}'' \arrow[r, "\Xi' \times \mathrm{id}"] \arrow[d, "\mathrm{id} \times \Xi''"'] & \mathfrak{T}'' \times \mathfrak{F}'' \arrow[d, "\Xi''"] \\ \mathfrak{T}' \times \mathfrak{F}' \arrow[r, "\Xi'"'] & \mathfrak{B}. \end{tikzcd}$$

Proof. Let $(Z_1, Z_2) \in \mathfrak{T}' \times \mathfrak{F}'$ and set $Z = \Xi'(Z_1, Z_2)$. Let T be a general point of Z and let X be the torsion submodule of T with respect to $(\mathscr{T}', \mathscr{F}')$. Then $X \in \mathscr{T}''$ and the point $(X, T/X)$ is general in $Z_1 \times Z_2$. Since a torsion class is stable under taking quotients and extensions, T belongs to $\mathscr{T}''$ if and only is T/X does. This means that Z belongs to $\mathfrak{T}''$ if and only if Z_2 does. Thus Ξ' restricts to a bijection $\mathfrak{T}' \times (\mathfrak{F}' \cap \mathfrak{T}'') \to \mathfrak{T}''$, as announced.

One shows that Ξ'' restricts to a bijection $(\mathfrak{F}' \cap \mathfrak{T}'') \times \mathfrak{F}'' \to \mathfrak{F}'$ in a similar fashion.

Now let $(Z_1, Z_2, Z_3) \in \mathfrak{T}' \times (\mathfrak{F}' \cap \mathfrak{T}'') \times \mathfrak{F}''$. Set $Z_4 = \Xi'(Z_1, Z_2)$ and $Z = \Xi''(Z_4, Z_3)$. Let T be a general point of Z and let X' and X'' be the torsion submodules of T with respect to $(\mathscr{T}', \mathscr{F}')$ and $(\mathscr{T}'', \mathscr{F}'')$, respectively. Then the point $(X'', T/X'')$ is general in $Z_4 \times Z_3$. Since X' is the torsion submodule of X'' with respect to $(\mathscr{T}', \mathscr{F}')$, the point $(X', X''/X', T/X'')$ is general in $Z_1 \times Z_2 \times Z_3$.

A similar reasoning shows that $(X', X''/X', T/X'')$ is also general in $\widetilde{Z}_1 \times \widetilde{Z}_2 \times \widetilde{Z}_3$, where $(\widetilde{Z}_1, \widetilde{Z}_2, \widetilde{Z}_3) = (\Xi' \circ (\mathrm{id} \times \Xi''))^{-1}(Z)$. Therefore a point can be general in $Z_1 \times Z_2 \times Z_3$ and in $\widetilde{Z}_1 \times \widetilde{Z}_2 \times \widetilde{Z}_3$ at the same time. Then necessarily $(Z_1, Z_2, Z_3) = (\widetilde{Z}_1, \widetilde{Z}_2, \widetilde{Z}_3)$, which establishes the commutativity. □

The torsion pairs $(\mathscr{I}_\theta, \overline{\mathscr{P}}_\theta)$ and $(\overline{\mathscr{I}}_\theta, \mathscr{P}_\theta)$ satisfy the assumption (O), thanks to Proposition 4.2 (iii). Applying Proposition 4.5 to them, we get a bijection

$$\Xi_\theta : \mathfrak{I}_\theta \times \mathfrak{R}_\theta \times \mathfrak{P}_\theta \to \mathfrak{B},$$

where $\mathfrak{I}_\theta$, $\mathfrak{R}_\theta$ and $\mathfrak{P}_\theta$ are the subsets of $\mathfrak{B}$, consisting of components whose general points belong to $\mathscr{I}_\theta$, $\mathscr{R}_\theta$ and $\mathscr{P}_\theta$, respectively.

We also note that Proposition 4.5 can be generalized in an obvious fashion to any finite nested sequence

$$(\mathscr{T}_0, \mathscr{F}_0) \preccurlyeq \cdots \preccurlyeq (\mathscr{T}_\ell, \mathscr{F}_\ell)$$

of torsion pairs that satisfy (O).

5 Tilting theory on preprojective algebras

5.1 Reflection functors

Let $I_i = \Lambda(1 - e_i)\Lambda$; in other words, let I_i be the annihilator of the simple Λ-module S_i. Amplifying their previous work [31], Iyama and Reiten, in a joint paper [14] with Buan and Scott, show that I_i is a tilting Λ-module of projective dimension at most one and that $\operatorname{End}_\Lambda(I_i) \cong \Lambda$ (at least, when no connected component of (I, E) is of Dynkin type).

This certainly invites us to look at the endofunctors $\Sigma_i = \operatorname{Hom}_\Lambda(I_i, ?)$ and $\Sigma_i^* = I_i \otimes_\Lambda ?$ of the category Λ-mod. It turns out that these functors can be described in a very explicit fashion.

Recall that we locally depict a Λ-module M around the vertex i by the diagram (4.1).

Proposition 5.1 *(i) The module $\Sigma_i M$ is obtained by replacing* (4.1) *with*

$$\widetilde{M_i} \xrightarrow{\overline{M}_{\mathrm{out}(i)} M_{\mathrm{in}(i)}} \ker M_{\mathrm{in}(i)} \hookrightarrow \widetilde{M_i},$$

where the map $\overline{M}_{\mathrm{out}(i)} : M_i \to \ker M_{\mathrm{in}(i)}$ is induced by $M_{\mathrm{out}(i)}$.

(ii) The module $\Sigma_i^* M$ is obtained by replacing (4.1) *with*

$$\widetilde{M_i} \twoheadrightarrow \operatorname{coker} M_{\mathrm{out}(i)} \xrightarrow{M_{\mathrm{out}(i)} \overline{M}_{\mathrm{in}(i)}} \widetilde{M_i},$$

where the map $\overline{M}_{\mathrm{in}(i)} : \operatorname{coker} M_{\mathrm{out}(i)} \to M_i$ is induced by $M_{\mathrm{in}(i)}$.

Proof. Applying the functor $S_i \otimes_\Lambda ?$ to the resolution (4.3) and changing the right arrow by a sign, one finds the following exact sequence of right Λ-modules:

$$K\rho_i \otimes_{\mathbf{S}} \Lambda \xrightarrow{\partial_1} \bigoplus_{\substack{a \in H \\ s(a)=i}} K\overline{a} \otimes_{\mathbf{S}} \Lambda \xrightarrow{\partial_0} e_i I_i \to 0, \tag{5.1}$$

where

$$\partial_1(\rho_i \otimes 1) = \sum_{\substack{a\in H \\ s(a)=i}} \varepsilon(a)\overline{a} \otimes a \qquad \text{and} \qquad \partial_0(\overline{a} \otimes 1) = \overline{a}.$$

The sequence obtained by applying $? \otimes_\Lambda M$ to (5.1) can be identified with

$$M_i \xrightarrow{M_{\mathrm{out}(i)}} \widetilde{M_i} \to e_i I_i \otimes_\Lambda M \to 0.$$

Using the decomposition $I_i = (1-e_i)\Lambda \oplus e_i I_i$, one can subsequently identify $\Sigma_i^* M = I_i \otimes_\Lambda M$ with the vector space described in Statement (ii).

Let us check the equality $(\Sigma_i^* M)_{\mathrm{out}(i)} = M_{\mathrm{out}(i)} \overline{M}_{\mathrm{in}(i)}$. Let $x \in \operatorname{coker} M_{\mathrm{out}(i)}$. It can be represented by an element $(x_b) \in \widetilde{M_i}$, which, in the identification

$$\widetilde{M_i} \cong \left(\bigoplus_{\substack{b\in H \\ s(b)=i}} K\overline{b} \otimes_{\mathbf{S}} \Lambda \right) \otimes_\Lambda M, \quad \text{corresponds to} \quad \sum_{\substack{b\in H \\ s(b)=i}} \overline{b} \otimes x_b.$$

In the Λ-module $\Sigma_i^* M$, an arrow a that starts at i maps x to

$$\sum_{\substack{b\in H \\ s(b)=i}} (a\overline{b}) \otimes x_b = \sum_{\substack{b\in H \\ s(b)=i}} M_a M_{\overline{b}} x_b,$$

where the left-hand side lives in $(1-e_i)\Lambda \otimes_\Lambda M$. Therefore

$$(\Sigma_i^* M)_{\mathrm{out}(i)}(x) = \Big(\varepsilon(a) \big(\Sigma_i^* M\big)_a(x) \Big)_{\substack{a\in H \\ s(a)=i}} = \left(\sum_{\substack{b\in H \\ s(b)=i}} \varepsilon(a) M_a M_{\overline{b}} x_b \right)_{\substack{a\in H \\ s(a)=i}} = M_{\mathrm{out}(i)} \overline{M}_{\mathrm{in}(i)}(x).$$

One checks similarly that $(\Sigma_i^* M)_{\mathrm{in}(i)}$ is the canonical map $\widetilde{M_i} \to \operatorname{coker} M_{\mathrm{out}(i)}$, which concludes the proof of (ii). The proof of (i) is similar, with two differences: one starts with the exact sequence obtained by applying $?\otimes_\Lambda S_i$ to the resolution (4.3), and one has to change the position of the signs $\varepsilon(a)$ (Remark 2.4 in [5] explains that this change is without consequences). □

These mutually adjoint functors Σ_i and Σ_i^* are called reflection functors. The concrete description afforded by the proposition yields several important properties that they enjoy:

- Adjunction morphisms can be inserted in functorial short exact sequences

$$0 \to \mathrm{soc}_i \to \mathrm{id} \to \Sigma_i \Sigma_i^* \to 0 \quad \text{and} \quad 0 \to \Sigma_i^* \Sigma_i \to \mathrm{id} \to \mathrm{hd}_i \to 0 \tag{5.2}$$

(see Proposition 2.5 in [5]).

- They are exchanged by the $*$-duality; in other words, $\Sigma_i^* T^* \cong (\Sigma_i T)^*$ for all finite dimensional Λ-module T.

- They induce Kashiwara and Saito's crystal reflections S_i and S_i^* on $\mathfrak{B}$ (see section 5.5).

- The operation of restricting a representation of Λ to a representation of the quiver Q intertwines the functors Σ_i and Σ_i^* with the traditional Bernstein-Gelfand-Ponomarev reflection functors (see Proposition 7.1 in [5]).

Remark 5.2. These functors Σ_i and Σ_i^* were introduced by Iyama and Reiten in [32] by means of the ideals I_i, and also, independently, by the first two authors in [5] by the explicit description of Proposition 5.1. The link between the two constructions was suggested to us by Amiot.

5.2 The tilting ideals I_w

Reflection functors satisfy the braid relations, so it is natural to study products of reflection functors computed according to reduced decompositions of elements in W. We now look for an analog of the exact sequences (5.2) for such a product.

To simplify the presentation, we consider in this section the case where no connected component of the diagram (I, E) is of Dynkin type. This allows us to rely on the following result, due to Buan, Iyama, Reiten and Scott (section II.1 in [14]) We will however argue in section 5.6 that almost all the results presented here hold true in general.

Theorem 5.3 *(i) If $s_{i_1} \cdots s_{i_\ell}$ is a reduced decomposition, then the multiplication in Λ gives rise to an isomorphism of bimodules $I_{i_1} \otimes_\Lambda \cdots \otimes_\Lambda I_{i_\ell} \to I_{i_1} \cdots I_{i_\ell}$.*

(ii) Under the same assumption, the product $I_{i_1} \cdots I_{i_\ell}$ depends only on $w = s_{i_1} \cdots s_{i_\ell}$; we can thus denote it by I_w. It has finite codimension in Λ.

(iii) Each I_w is a tilting Λ-bimodule of projective dimension at most 1 and $\mathrm{End}_\Lambda(I_w) \cong \Lambda$.

(iv) If $\ell(ws_i) > \ell(w)$, then $\mathrm{Tor}_1^\Lambda(I_w, S_i) = 0$. If $\ell(s_i w) > \ell(w)$, then $\mathrm{Ext}^1_\Lambda(I_w, S_i) = 0$.

In view of Theorem 5.3 (iii), it is natural to apply Brenner and Butler's theorem. For that, we define categories

$$\mathscr{T}_w = \{T \mid I_w \otimes_\Lambda T = 0\}, \qquad \mathscr{T}^w = \{T \mid \operatorname{Ext}^1_\Lambda(I_w, T) = 0\},$$
$$\mathscr{F}_w = \{T \mid \operatorname{Tor}^1_\Lambda(I_w, T) = 0\}, \qquad \mathscr{F}^w = \{T \mid \operatorname{Hom}_\Lambda(I_w, T) = 0\}.$$

Theorem 5.4 *(i) The pair $(\mathscr{T}^w, \mathscr{F}^w)$ is a torsion pair in Λ-mod. For each Λ-module T, the evaluation map $I_w \otimes_\Lambda \operatorname{Hom}_\Lambda(I_w, T) \to T$ is injective and its image is the torsion submodule of T with respect to $(\mathscr{T}^w, \mathscr{F}^w)$.*

(ii) The pair $(\mathscr{T}_w, \mathscr{F}_w)$ is a torsion pair in Λ-mod. For each Λ-module T, the coevaluation map $T \to \operatorname{Hom}_\Lambda(I_w, I_w \otimes_\Lambda T)$ is surjective and its kernel is the torsion submodule of T with respect to $(\mathscr{T}_w, \mathscr{F}_w)$.

(iii) There are mutually inverse equivalences

$$\begin{tikzcd} \mathscr{F}_w \arrow[r, shift left, "I_w \otimes_\Lambda ?"] & \mathscr{T}^w \arrow[l, shift left, "\operatorname{Hom}_\Lambda(I_w, ?)"] \end{tikzcd}.$$

Proof. See [3], in particular the lemma in section 1.6, the corollary in section 1.9, and the theorem and its corollary in section 2.1. □

Given $w \in W$ and $T \in \Lambda$-mod, we will denote by T^w and T_w the torsion submodules of T with respect to $(\mathscr{T}^w, \mathscr{F}^w)$ and $(\mathscr{T}_w, \mathscr{F}_w)$, respectively.

If $u, v \in W$ are such that $\ell(u) + \ell(v) = \ell(uv)$, then $I_{uv} \cong I_u \otimes_\Lambda I_v$. It immediately follows that

$$\mathscr{F}^{uv} \supseteq \mathscr{F}^u \quad \text{and} \quad \mathscr{T}_v \subseteq \mathscr{T}_{uv},$$

which can be written

$$(\mathscr{T}^{uv}, \mathscr{F}^{uv}) \preccurlyeq (\mathscr{T}^u, \mathscr{F}^u) \quad \text{and} \quad (\mathscr{T}_v, \mathscr{F}_v) \preccurlyeq (\mathscr{T}_{uv}, \mathscr{F}_{uv}). \tag{5.3}$$

Remarks 5.5. (i) The reader may here object that we apply a theorem valid for finite dimensional algebras to an infinite dimensional framework. In fact, there is no difficulty. The non-obvious point is to show that the functors $\operatorname{Hom}_\Lambda(I_w, ?)$ and $I_w \otimes_\Lambda ?$ preserve the category of finite dimensional Λ-modules. In the case where w is a simple reflection, this follows from Proposition 5.1. The general case follows by composition.

(ii) By Theorem 5.4 (iii), any module $T \in \mathscr{T}^w$ is isomorphic to a module of the form $I_w \otimes_\Lambda X$. Conversely, one easily checks that if T is of the form $I_w \otimes_\Lambda X$, then the evaluation map $I_w \otimes_\Lambda \operatorname{Hom}_\Lambda(I_w, T) \to T$ is surjective, so T belongs to $\mathscr{T}^w$ by Theorem 5.4 (i). Therefore $\mathscr{T}^w$ is the essential image of the functor $I_w \otimes_\Lambda ?$. Likewise, one shows that $\mathscr{F}_w$ is the essential image of $\operatorname{Hom}_\Lambda(I_w, ?)$.

Examples 5.6. (i) The case where w is a simple reflection has been dealt with in the previous section: the functorial short exact sequences (5.2) imply that the torsion submodule of T with respect to $(\mathscr{T}_{s_i}, \mathscr{F}_{s_i})$ is the i-socle of T and that the torsion-free quotient of T with respect to $(\mathscr{T}^{s_i}, \mathscr{F}^{s_i})$ is the i-head of T. Therefore

$$\mathscr{T}_{s_i} = \operatorname{add} S_i, \qquad \mathscr{T}^{s_i} = \{T \mid \operatorname{hd}_i T = 0\},$$
$$\mathscr{F}_{s_i} = \{T \mid \operatorname{soc}_i T = 0\}, \qquad \mathscr{F}^{s_i} = \operatorname{add} S_i,$$

where $\operatorname{add} S_i$ is the additive closure of S_i in Λ-mod. These equalities were also obtained by Sekiya and Yamaura (Lemma 2.23 in [59]).

(ii) Let us generalize the first example. Let $J \subseteq I$ and let Q_J be the full subquiver Q with set of vertices J. The preprojective algebra Λ_J is a quotient of Λ. The kernel of the natural morphism $\Lambda \to \Lambda_J$ is the ideal $I_J = \Lambda\big(1 - \sum_{j \in J} e_j\big)\Lambda$. The pull-back functor allows to identify Λ_J-mod with the full subcategory

$$\{M \in \Lambda\text{-mod} \mid e_j M \neq 0 \;\Rightarrow j \in J\}$$

of Λ-mod. Moreover, a Λ-module T belongs to Λ_J-mod if and only if $\operatorname{Hom}_\Lambda(I_J, T) = 0$. In fact, if this equality holds, then $I_J T = 0$, hence T is a Λ/I_J-module. Conversely, if $T \in \Lambda_J$-mod, then

$$\operatorname{Ext}^1_\Lambda(\Lambda_J, T) \cong \operatorname{Ext}^1_{\Lambda_J}(\Lambda_J, T) = 0,$$

for Λ_J-mod is closed under extensions; since the first arrow in the short exact sequence

$$\operatorname{Hom}_\Lambda(\Lambda_J, T) \hookrightarrow \operatorname{Hom}_\Lambda(\Lambda, T) \to \operatorname{Hom}_\Lambda(I_J, T) \to \operatorname{Ext}^1_\Lambda(\Lambda_J, T) = 0$$

is an isomorphism, we get $\operatorname{Hom}_\Lambda(I_J, T) = 0$.

Assume now that Q_J is of Dynkin type. Then Theorem II.3.5 in [14] says that I_J is the ideal I_{w_J}, where w_J is the longest element in the parabolic subgroup $W_J = \langle s_j \mid j \in J\rangle$ of W. We conclude that $\mathscr{F}^{w_J} = \Lambda_J$-mod. Further, by duality, we also have $\mathscr{T}_{w_J} = \Lambda_J$-mod.

Thus the torsion-free part of a module T with respect to the torsion pair $(\mathscr{T}^{w_J}, \mathscr{F}^{w_J})$ is the largest quotient of T that belongs to Λ_J-mod, and the torsion part of T with respect to the torsion pair $(\mathscr{T}_{w_J}, \mathscr{F}_{w_J})$ is the largest submodule of T that belongs to Λ_J-mod.

(iii) We have already mentioned that the reflection functors Σ_i and Σ_i^* are exchanged by the $*$-duality. By composition, we obtain $(I_w \otimes_\Lambda T)^* = \mathrm{Hom}_\Lambda(I_{w^{-1}}, T^*)$. This readily implies

$$\mathscr{F}^w = (\mathscr{T}_{w^{-1}})^* \quad \text{and} \quad \mathscr{T}^w = (\mathscr{F}_{w^{-1}})^*.$$

(iv) We will explain in Examples 5.13 and 5.14 that $\mathscr{F}^w$ is Buan, Iyama, Reiten and Scott's category $\mathrm{Sub}(\Lambda/I_w)$ [14] and that $\mathscr{T}_w$ is Geiß, Leclerc and Schröer's category $\mathcal{C}_w$ [27].

We conclude this section by showing that the equivalence of categories described in Theorem 5.4 (iii) can be broken into pieces according to any reduced decomposition of w.

Proposition 5.7 *Let $(u, v, w) \in W^3$ be such that $\ell(uvw) = \ell(u) + \ell(v) + \ell(w)$. Then one has mutually inverse equivalences*

$$\mathscr{F}_{uv} \cap \mathscr{T}^w \underset{\mathrm{Hom}_\Lambda(I_v,?)}{\overset{I_v \otimes_\Lambda ?}{\rightleftarrows}} \mathscr{F}_u \cap \mathscr{T}^{vw} .$$

Proof. Certainly, $I_v \otimes_\Lambda ?$ maps $\mathscr{T}^w$, the essential image of $I_w \otimes_\Lambda ?$, to $\mathscr{T}^{vw}$, the essential image of $I_{vw} \otimes_\Lambda ?$. On the other side, any module $T \in \mathscr{F}_{uv}$ belongs to the essential image of

$$\mathrm{Hom}_\Lambda(I_{uv}, ?) = \mathrm{Hom}_\Lambda(I_v, \mathrm{Hom}_\Lambda(I_u, ?)),$$

hence is isomorphic to $\mathrm{Hom}_\Lambda(I_v, X)$, with $X \in \mathscr{F}_u$. By Theorem 5.4 (i), $I_v \otimes_\Lambda T$ is a submodule of X, hence belongs to $\mathscr{F}_u$. To sum up, $I_v \otimes_\Lambda ?$ maps $\mathscr{T}^w$ to $\mathscr{T}^{vw}$ and maps $\mathscr{F}_{uv}$ to $\mathscr{F}_u$, hence maps $\mathscr{F}_{uv} \cap \mathscr{T}^w$ to $\mathscr{F}_u \cap \mathscr{T}^{vw}$.

One shows in a dual fashion that $\mathrm{Hom}_\Lambda(I_v, ?)$ maps $\mathscr{F}_u \cap \mathscr{T}^{vw}$ to $\mathscr{F}_{uv} \cap \mathscr{T}^w$. □

Under the assumptions of the proposition, one thus has a chain of equivalences of categories

$$\mathscr{F}_{uvw} \underset{\mathrm{Hom}_\Lambda(I_w,?)}{\overset{I_w \otimes_\Lambda ?}{\rightleftarrows}} \mathscr{F}_{uv} \cap \mathscr{T}^w \underset{\mathrm{Hom}_\Lambda(I_v,?)}{\overset{I_v \otimes_\Lambda ?}{\rightleftarrows}} \mathscr{F}_u \cap \mathscr{T}^{vw} \underset{\mathrm{Hom}_\Lambda(I_u,?)}{\overset{I_u \otimes_\Lambda ?}{\rightleftarrows}} \mathscr{T}^{uvw} .$$

Corollary 5.8 *Let $w \in W$ and let $T \in \mathscr{F}_w$. Then*

$$\underline{\dim}\, I_w \otimes_\Lambda T = w(\underline{\dim}\, T).$$

Proof. The particular case where w is a simple reflection is Lemma 2.5 in [1]; it can also be proved in a more elementary way with the help of Proposition 5.1. The general case is obtained by writing w as a product of $\ell(w)$ simple reflections. □

5.3 Layers and stratification

In section 10 of [27], Geiß, Leclerc and Schröer construct a filtration of the objects in $\mathscr{T}_w$. In section 2 of [1], Amiot, Iyama, Reiten and Todorov construct a filtration of the finite dimensional module Λ/I_w by layers. Our aim in this section is to show that these two constructions are $*$-dual to each other and are defined by the torsion pairs associated to the tilting ideals I_w. These results will help us later to identify the categories $\mathscr{T}^w$, $\mathscr{F}^w$, $\mathscr{T}_w$ and $\mathscr{F}_w$ with the categories $\mathscr{I}_\theta$ and $\mathscr{P}_\theta$ from section 3.1.

We begin by introducing the layers. The next lemma is identical to Corollary 9.3 in [27], to Theorem 2.6 in [1], and to Lemma 5.11 in [59]. We cannot resist offering a fourth proof of this basic result.

Lemma 5.9 *Let $(w,i) \in W \times I$ be such that $\ell(ws_i) > \ell(w)$. Then the simple module S_i belongs to $\mathscr{F}_w$. Moreover, the module $I_w \otimes_\Lambda S_i$ belongs to $\mathscr{F}^{ws_i}$ and has dimension-vector $w\alpha_i$.*

Proof. We proceed by induction on the length of w. The case $w = 1$ is obvious. Assume $\ell(w) > 0$ and write $w = s_j v$ with $\ell(v) < \ell(w)$. Applying $\operatorname{Hom}_\Lambda(I_v, ?)$ to the short exact sequence

$$0 \to \operatorname{soc}_j(I_v \otimes_\Lambda S_i) \to I_v \otimes_\Lambda S_i \to \Sigma_j\Sigma_j^*(I_v \otimes_\Lambda S_i) \to 0$$

leads to

$$0 \to \operatorname{Hom}_\Lambda(I_v, S_j)^{\oplus n} \to \operatorname{Hom}_\Lambda(I_v, I_v \otimes_\Lambda S_i) \to \operatorname{Hom}_\Lambda(I_w, I_w \otimes_\Lambda S_i) \to 0$$

with $n = \dim \operatorname{soc}_j(I_v \otimes_\Lambda S_i)$, by Theorem 5.3 (iv). Applied to (v,i), the inductive hypothesis says that the middle term is isomorphic to S_i, hence is simple. Applied to (v^{-1}, j), the inductive hypothesis, together with Example 5.6 (iii) and Corollary 5.8, implies that the left term has dimension-vector $n\,v^{-1}(\alpha_j)$. Moreover, the assumption $\ell(ws_i) > \ell(w)$ means that $v^{-1}(\alpha_j) \neq \alpha_i$. We conclude that necessarily $n = 0$ and that the right term is isomorphic to S_i. Therefore S_i belongs to the essential image of $\operatorname{Hom}_\Lambda(I_w, ?)$, that is, to $\mathscr{F}_w$. We also see that

$$\operatorname{Hom}_\Lambda(I_{ws_i}, I_w \otimes_\Lambda S_i) \cong \Sigma_i S_i = 0,$$

which means that $I_w \otimes_\Lambda S_i$ belongs to $\mathscr{F}^{ws_i}$. The last claim follows from Corollary 5.8. $\square$

The following result is in essence due to Amiot, Iyama, Reiten and Todorov [1] and to Geiß, Leclerc and Schröer [27].

Theorem 5.10 *(i) Let $(w,i) \in W \times I$ be such that $\ell(ws_i) > \ell(w)$ and let $T \in \Lambda\text{-mod}$. Then T^w/T^{ws_i} is the largest quotient of T^w isomorphic to a direct sum of copies of the module $I_w \otimes_\Lambda S_i$.*

(ii) *Let $w = s_{i_1} \cdots s_{i_\ell}$ be a reduced decomposition. Then a module $T \in \Lambda$-mod belongs to $\mathscr{F}^w$ if and only if it has a filtration whose subquotients are modules of the form $I_{s_{i_1} \cdots s_{i_{k-1}}} \otimes_\Lambda S_{i_k}$ with $1 \le k \le \ell$.*

(iii) *Let $w = s_{i_1} \cdots s_{i_\ell}$ be a reduced decomposition. Then a module $T \in \Lambda$-mod belongs to $\mathscr{T}^w$ if and only if there is no epimorphism $T \twoheadrightarrow I_{s_{i_1} \cdots s_{i_{k-1}}} \otimes_\Lambda S_{i_k}$ with $1 \le k \le \ell$.*

Proof. Let w, i and T as in (i) and set $X = \mathrm{Hom}_\Lambda(I_w, T)$. Applying $I_w \otimes_\Lambda ?$ to the short exact sequence

$$0 \to \Sigma_i^* \Sigma_i X \to X \to \mathrm{hd}_i X \to 0$$

and using $\mathrm{Tor}_1^\Lambda(I_w, S_i) = 0$ (Theorem 5.3 (iv)), we get

$$0 \to I_{ws_i} \otimes_\Lambda \mathrm{Hom}_\Lambda(I_{ws_i}, T) \to I_w \otimes_\Lambda \mathrm{Hom}_\Lambda(I_w, T) \to I_w \otimes_\Lambda S_i^{\oplus m} \to 0,$$

where $m = \dim \mathrm{hd}_i X$. Therefore T^w / T^{ws_i} has the desired form.

To finish the proof of (i), it remains to show the maximality. Before that, we look at Statements (ii) and (iii). For $1 \le k \le \ell$, we set $L_k = I_{s_{i_1} \cdots s_{i_{k-1}}} \otimes_\Lambda S_{i_k}$.

Let $T \in \Lambda$-mod. What has already been showed from Statement (i) implies that T/T^w has a filtration whose subquotients are all isomorphic to some L_k. If T belongs to $\mathscr{F}^w$, then $T = T/T^w$ has such a filtration: this shows the necessity of the condition proposed in Statement (ii). If T has no quotient isomorphic to a module L_k, then $T = T^w$, and therefore T belongs to $\mathscr{T}^w$: this shows the sufficiency of the condition proposed in Statement (iii).

By Lemma 5.9, L_k belongs to $\mathscr{F}^{s_{i_1} \cdots s_{i_k}}$, hence to $\mathscr{F}^w$. Any iterated extension of modules in the set $\{L_k \mid 1 \le k \le \ell\}$ therefore belongs to $\mathscr{F}^w$, for $\mathscr{F}^w$ is closed under extensions: this shows the sufficiency of the condition in Statement (ii). On the other hand, a module in $\mathscr{T}^w$ cannot have a nonzero map to a module in $\mathscr{F}^w$, hence has no quotient isomorphic to a module L_k: this shows the necessity of the condition in Statement (iii).

Now let us go back to Statement (i), resuming the proof where we left it. We choose a reduced decomposition $w = s_{i_1} \cdots s_{i_{\ell-1}}$ and set $i_\ell = i$. For $1 \le k \le \ell$, we set $L_k = I_{s_{i_1} \cdots s_{i_{k-1}}} \otimes_\Lambda S_{i_k}$.

Assume the existence of a short exact sequence of the form $0 \to Z \to T^w \to L_\ell^{\oplus n} \to 0$, where n is a positive integer and Z does not belong to $\mathscr{T}^w$. Then Z has a quotient Z/Y isomorphic to some L_k, with $k < \ell$, and we have an extension $0 \to L_k \to T^w/Y \to L_\ell^{\oplus n} \to 0$. By Lemma 5.9, $\underline{\dim}\, L_k$ and $\underline{\dim}\, L_\ell$ are distinct real roots, hence are not proportional; therefore $\underline{\dim}\, T^w/Y$ is not a multiple of $\underline{\dim}\, L_\ell$. Let us set $U = T^w/Y$. This module belongs to $\mathscr{T}^w$, because T^w does, so $U = U^w$; it also belongs to $\mathscr{F}^{ws_i}$, because L_k and L_ℓ do and because $\mathscr{F}^{ws_i}$ is closed under extensions, so $U^{ws_i} = 0$. Therefore $U = U^w / U^{ws_i}$, which implies (by the first part of the proof applied to U) that U is a direct sum of copies of L_ℓ. We thus reach a contradiction, and conclude that the assumption at the beginning of this paragraph is wrong.

Consider a short exact sequence $0 \to Z \to T^w \to L_\ell^{\oplus n} \to 0$. The preceding paragraph says that Z belongs to $\mathscr{T}^w$. Further, we note that $\mathrm{Hom}_\Lambda(I_w, T^w) = \mathrm{Hom}_\Lambda(I_w, T) = X$, by Theorem 5.4 and Remark 5.5 (ii). Applying $\mathrm{Hom}_\Lambda(I_w, ?)$ to the short exact sequence, we then get $0 \to \mathrm{Hom}_\Lambda(I_w, Z) \to X \to S_i^{\oplus n} \to 0$. Therefore, by definition of the i-head, $\mathrm{Hom}_\Lambda(I_w, Z)$ contains $\Sigma_i^* \Sigma_i X$, and moreover the quotient $\mathrm{Hom}_\Lambda(I_w, Z)/\Sigma_i^*\Sigma_i X$ is a direct sum of copies of S_i. Applying $I_w \otimes_\Lambda ?$ to the inclusion $\Sigma_i^*\Sigma_i X \hookrightarrow \mathrm{Hom}_\Lambda(I_w, Z)$ and using Theorem 5.3 (iv), we conclude that $Z = I_w \otimes_\Lambda \mathrm{Hom}_\Lambda(I_w, Z)$ contains $T^{ws_i} = I_w \otimes_\Lambda \Sigma_i^* \Sigma_i X$. $\square$

This theorem admits a dual version, which we state for reference.

Theorem 5.11 *(i) Let $(w, i) \in W \times I$ be such that $\ell(s_i w) > \ell(w)$ and let $T \in \Lambda$-*mod*. Then $T_{s_i w}/T_w$ is the largest submodule of T/T_w isomorphic to a direct sum of copies of the module* $\mathrm{Hom}_\Lambda(I_w, S_i)$.

*(ii) Let $w = s_{i_\ell} \cdots s_{i_1}$ be a reduced decomposition. Then a module $T \in \Lambda$-*mod *belongs to $\mathscr{T}_w$ if and only if it has a filtration whose subquotients are modules of the form* $\mathrm{Hom}_\Lambda(I_{s_{i_{k-1}} \cdots s_{i_1}}, S_{i_k})$ *with $1 \leq k \leq \ell$.*

*(iii) Let $w = s_{i_\ell} \cdots s_{i_1}$ be a reduced decomposition. Then a module $T \in \Lambda$-*mod *belongs to $\mathscr{F}_w$ if and only if there is no monomorphism* $\mathrm{Hom}_\Lambda(I_{s_{i_{k-1}} \cdots s_{i_1}}, S_{i_k}) \hookrightarrow T$ *with $1 \leq k \leq \ell$.*

As already mentioned, most of the results presented above in this section can be found in papers by Iyama, Reiten et al. and by Geiß, Leclerc and Schröer. The aim of the next three examples is to explain some connections in more detail.

Example 5.12. Let us denote by $\{\omega_i \mid i \in I\}$ the set of fundamental weights of the root system Φ; these weights are elements of a representation of W which contains $\mathbb{R}I$, and for all $(i, j) \in I^2$, we have $s_j \omega_i = \omega_i - \delta_{ij} \alpha_i$, where δ_{ij} is Kronecker's symbol. Let us now fix $(i, w) \in I \times W$. By [5], Theorem 3.1 (ii), there exists a unique Λ-module $N(-w\omega_i)$ whose dimension-vector is $\omega_i - w\omega_i$ and whose socle is 0 or S_i. (The paper [5] mainly deals with the case where $\mathfrak{g}$ is finite dimensional, but the constructions in sections 2 and 3 are valid in general, with the exception of Proposition 3.6.) The aim of this example is to show that $N(-w\omega_i) \cong ((\Lambda/I_w) \otimes_\Lambda \Lambda e_i)^*$, where e_i is the lazy path at vertex i (see section 4.1). For that, we consider a reduced decomposition $w = s_{j_1} \cdots s_{j_\ell}$. The filtration $\Lambda \supset I_{s_{j_1}} \supset I_{s_{j_1} s_{j_2}} \supset \cdots \supset I_{s_{j_1} \cdots s_{j_{\ell-1}}} \supset I_w$ induces a filtration of Λ/I_w, whose k-th subquotient is the layer

$$I_{s_{j_1} \cdots s_{j_{k-1}}} / I_{s_{j_1} \cdots s_{j_k}} \cong I_{s_{j_1} \cdots s_{j_{k-1}}} \otimes_\Lambda (\Lambda/I_{j_k}) \cong I_{s_{j_1} \cdots s_{j_{k-1}}} \otimes_\Lambda S_{j_k}$$

studied by Amiot, Iyama, Reiten and Todorov [1]. Tensoring this filtration on the right with the projective Λ-module Λe_i kills all the subquotients with $j_k \neq i$, so by Lemma 5.9,

$$\underline{\dim}\,(\Lambda/I_w) \otimes_\Lambda \Lambda e_i = \sum_{\substack{1\leq k\leq \ell \\ j_k=i}} s_{j_1}\cdots s_{j_{k-1}}\alpha_{j_k} = \omega_i - w\omega_i.$$

In addition, the head of Λe_i is equal to S_i, so the head of $(\Lambda/I_w) \otimes_\Lambda \Lambda e_i$ is 0 or S_i. The dual of $(\Lambda/I_w) \otimes_\Lambda \Lambda e_i$ therefore satisfies the two conditions that characterize $N(-w\omega_i)$.

Example 5.13. Theorem 5.10 (ii) and Corollary 2.9 in [32] imply that our category $\mathscr{F}^w$ is equal to $\mathrm{Sub}(\Lambda/I_w)$, the full subcategory of Λ-mod whose objects are the modules that can be embedded in a finite direct sum of copies of Λ/I_w. (To make sure that the assumptions of the statements in [32] are fulfilled, observe that a module $T \in \mathscr{F}^w$ satisfies $\mathrm{Hom}_\Lambda(I_w, T) = 0$ by definition, whence $I_wT = 0$, so T can be seen as a Λ/I_w-module.)

Example 5.14. We now compare Theorem 5.10 to Geiß, Leclerc and Schröer's stratification. We fix a reduced decomposition $w = s_{i_\ell}\cdots s_{i_1}$. Then Geiß, Leclerc and Schröer define modules V_k and M_k for $1 \leq k \leq \ell$ (sections 2.4 and 9 of [27]). By construction, V_k is a submodule of the injective hull of S_{i_k}, hence has socle 0 or S_{i_k}; moreover, the dimension-vector of V_k is $\omega_{i_k} - s_{i_1}\cdots s_{i_k}\omega_{i_k}$, by Corollary 9.2 in [27]. Comparing with Example 5.12, we see that

$$V_k \cong N(-s_{i_1}\cdots s_{i_k}\omega_{i_k}) \cong ((\Lambda/I_{s_{i_1}\cdots s_{i_k}}) \otimes_\Lambda \Lambda e_{i_k})^*.$$

If $\{s \in \{1,\ldots,k-1\} \mid i_s = i_k\}$ is not empty, then k^- is defined to be the largest element in this set and M_k is defined to be the quotient V_k/V_{k^-}; note here that

$$V_{k^-} \cong N(-s_{i_1}\cdots s_{i_{k^-}}\omega_{i_{k^-}}) = N(-s_{i_1}\cdots s_{i_{k-1}}\omega_{i_k}) \cong ((\Lambda/I_{s_{i_1}\cdots s_{i_{k-1}}}) \otimes_\Lambda \Lambda e_{i_k})^*.$$

Otherwise, M_k is defined to be V_k. In both case, M_k is the dual of

$$(I_{s_{i_1}\cdots s_{i_{k-1}}}/I_{s_{i_1}\cdots s_{i_k}}) \otimes_\Lambda \Lambda e_{i_k} \cong I_{s_{i_1}\cdots s_{i_{k-1}}} \otimes_\Lambda (\Lambda/I_{i_k}) \otimes_\Lambda \Lambda e_{i_k} \cong I_{s_{i_1}\cdots s_{i_{k-1}}} \otimes_\Lambda S_{i_k}.$$

By Example 5.6 (iii), this gives $M_k \cong \mathrm{Hom}_\Lambda(I_{s_{i_{k-1}}\cdots s_{i_1}}, S_{i_k})$. Further, Geiß, Leclerc and Schröer consider the module $V = V_1\oplus\cdots V_\ell$ and the category $\mathcal{C}_w = \mathrm{Fac}(V)$, the full subcategory of Λ-mod whose objects are the homomorphic images of a direct sum of copies of V. Comparing Lemma 10.2 in [27] with Theorem 5.11 (ii), we see that $\mathcal{C}_w$ is our category $\mathscr{T}_w$ and that the stratification constructed in section 10 of [27] on a Λ-module $T \in \mathcal{C}_w$ coincides with the filtration by the submodules $T_{s_{i_k}\cdots s_{i_1}}$. Combining this with the conclusion of Example 5.13, we additionally obtain that $\mathcal{C}_w = \mathrm{Fac}((\Lambda/I_{w^{-1}})^*)$, in agreement with Theorem 2.8 (iv) in [27].

Theorem 5.10 has the following noteworthy consequence.

Proposition 5.15 *Let $(u, v) \in W^2$ such that $\ell(u)+\ell(v) = \ell(uv)$. Then $(\mathscr{T}_u, \mathscr{F}_u) \preccurlyeq (\mathscr{T}^v, \mathscr{F}^v)$.*

Proof. We want to show that $\mathscr{T}_u \cap \mathscr{F}^v = \{0\}$. Assume the contrary and choose $T \neq 0$ of minimal dimension in the intersection. Write $u = s_{i_\ell} \cdots s_{i_1}$ and $v = s_{j_1} \cdots s_{j_m}$. By Theorem 5.10 (ii), T has a filtration with subquotients of the form $I_{s_{j_1} \cdots s_{j_{q-1}}} \otimes_\Lambda S_{j_q}$. Any subquotient of this filtration belongs to $\mathscr{F}^v$, and the top one also belongs to $\mathscr{T}_u$ since a torsion class is closed under taking quotients. The minimality of $\dim T$ imposes then that the filtration has just one step. Dually, T has a filtration with subquotients of the form $\operatorname{Hom}_\Lambda(I_{s_{i_{p-1}} \cdots s_{i_1}}, S_{i_p})$, and minimality impose again that this filtration has just one step. We end up with an isomorphism

$$\operatorname{Hom}_\Lambda(I_{s_{i_{p-1}} \cdots s_{i_1}}, S_{i_p}) \cong I_{s_{j_1} \cdots s_{j_{q-1}}} \otimes_\Lambda S_{j_q}.$$

Taking dimension-vectors, we get $s_{i_1} \cdots s_{i_{p-1}} \alpha_{i_p} = s_{j_1} \cdots s_{j_{q-1}} \alpha_{j_q}$, by Lemma 5.9. This contradicts the assumption $\ell(u) + \ell(v) = \ell(uv)$. $\square$

We conclude this section with a proposition that slightly refines Proposition 3.2 in [32].

Proposition 5.16 *Assume that $\ell(u) + \ell(v) = \ell(uv)$. Then one has equivalences of categories*

$$\mathscr{F}^v \underset{\operatorname{Hom}_\Lambda(I_u,?)}{\overset{I_u \otimes_\Lambda ?}{\rightleftarrows}} \mathscr{F}^{uv} \cap \mathscr{T}^u \qquad \textit{and} \qquad \mathscr{T}_{uv} \cap \mathscr{F}_v \underset{\operatorname{Hom}_\Lambda(I_v,?)}{\overset{I_v \otimes_\Lambda ?}{\rightleftarrows}} \mathscr{T}_u .$$

Proof. If $T \in \mathscr{F}^v$, then $T \in \mathscr{F}_u$, by Proposition 5.15. Thus $T = T/T_u \cong \operatorname{Hom}_\Lambda(I_u, I_u \otimes_\Lambda T)$, whence

$$\operatorname{Hom}_\Lambda(I_{uv}, I_u \otimes_\Lambda T) \cong \operatorname{Hom}_\Lambda(I_v, \operatorname{Hom}_\Lambda(I_u, I_u \otimes_\Lambda T)) \cong \operatorname{Hom}_\Lambda(I_v, T) = 0,$$

which shows that $I_u \otimes_\Lambda T \in \mathscr{F}^{uv}$. Conversely, if $T \in \mathscr{F}^{uv}$, then

$$\operatorname{Hom}_\Lambda(I_v, \operatorname{Hom}_\Lambda(I_u, T)) = \operatorname{Hom}_\Lambda(I_{uv}, T) = 0,$$

hence $\operatorname{Hom}_\Lambda(I_u, T) \in \mathscr{F}^v$. The first pair of equivalences then follows from Theorem 5.4 (iii).

The second equivalence is proved in a similar fasion (or follows by duality). $\square$

5.4 Tilting structure and HN polytopes in Λ-mod

We now relate all this material about the reflection functors and the torsion pairs $(\mathscr{T}^w, \mathscr{F}^w)$ and $(\mathscr{T}_w, \mathscr{F}_w)$ to the categories $\mathscr{I}_\theta$, $\mathscr{R}_\theta$, etc., and to the HN polytopes defined in section 3.

The following result is almost identical to Theorem 3.5 in [59]; we however prove the key point in a different fashion.

Theorem 5.17 *Let $\theta : \mathbb{Z}I \to \mathbb{R}$ be a group homomorphism and let $i \in I$. If $\langle \theta, \alpha_i \rangle > 0$, then Σ_i and Σ_i^* induce mutually inverse equivalences*

$$\mathscr{R}_\theta \underset{\Sigma_i}{\overset{\Sigma_i^*}{\rightleftarrows}} \mathscr{R}_{s_i\theta} .$$

Proof. The assumption $\langle \theta, \alpha_i \rangle > 0$ forbids a module $T \in \mathscr{R}_\theta$ to have a submodule isomorphic to S_i, so $\mathscr{R}_\theta \subseteq \mathscr{F}_{s_i}$ by Example 5.6 (i). Likewise, $\mathscr{R}_{s_i\theta} \subseteq \mathscr{T}^{s_i}$.

Let $T \in \mathscr{R}_\theta$. Then $T \in \mathscr{F}_{s_i}$ and $\langle \theta, \underline{\dim}\, T \rangle = 0$. Using Corollary 5.8 with $w = s_i$, we obtain $\langle s_i\theta, \underline{\dim}\, \Sigma_i^* T \rangle = 0$.

To show that $\Sigma_i^* T$ is $s_i\theta$-semistable, it remains to show that $\langle s_i\theta, \underline{\dim}\, X \rangle \geq 0$ for any quotient X of $\Sigma_i^* T$. In this aim, consider a surjective morphism $f : \Sigma_i^* T \to X$. The functor Σ_i modifies only the vector space at vertex i, so the cokernel of $\Sigma_i f$ is concentrated at this vertex. There exists thus a natural integer n such that $\underline{\dim}\ \mathrm{coker}(\Sigma_i f) = n\alpha_i$. Note now that not only $\Sigma_i^* T$, but also X belong to $\mathscr{T}^{s_i}$, for a torsion class is closed under taking quotients. By Corollary 5.8 again, we have $\underline{\dim}\, \Sigma_i X = s_i\, \underline{\dim}\, X$. We therefore have $s_i\, \underline{\dim}\, X = \underline{\dim}\ \mathrm{im}(\Sigma_i f) + n\alpha_i$. Since T is θ-semistable, $\langle \theta, \underline{\dim}\ \mathrm{im}(\Sigma_i f) \rangle \geq 0$. We eventually find that $\langle s_i\theta, \underline{\dim}\, X \rangle \geq n \langle \theta, \alpha_i \rangle \geq 0$, as desired.

We thus see that Σ_i^* maps $\mathscr{R}_\theta$ to $\mathscr{R}_{s_i\theta}$. A dual reasoning shows that Σ_i maps $\mathscr{R}_{s_i\theta}$ to $\mathscr{R}_\theta$. The theorem now follows from Theorem 5.4 (iii) and from the inclusions $\mathscr{R}_\theta \subseteq \mathscr{F}_{s_i}$ and $\mathscr{R}_{s_i\theta} \subseteq \mathscr{T}^{s_i}$. $\square$

Recall from section 2.1 the definition of the dominant cone $\overline{C}_0$ and of the Tits cone $C_T = \bigcup_{w \in W} w\overline{C}_0$. A subset $J \subseteq I$ gives rise to a face $F_J \subseteq \overline{C}_0$, to a parabolic subgroup $W_J = \langle s_j \mid j \in J \rangle$ of W, and to a root system $\Phi_J \subseteq \Phi$. In the case W_J is finite, we denote its longest element by w_J. Recall also that for $w \in W$, we set $N_w = \Phi_+ \cap w\Phi_-$ and that for any reduced decomposition $w = s_{i_1} \cdots s_{i_\ell}$, we have

$$N_w = \{ s_{i_1} \cdots s_{i_{k-1}} \alpha_{i_k} \mid 1 \leq k \leq \ell \}.$$

Theorem 5.18 *Let $J \subseteq I$, let $\theta \in F_J$, and let $w \in W$. We assume that w is J-reduced on the right, that is, $\ell(ws_j) > \ell(w)$ for each $j \in J$.*

(i) The category $\mathscr{R}_\theta$ coincides with the subcategory Λ_J-mod. Moreover, there are mutually inverse equivalences

$$\mathscr{R}_\theta \underset{\operatorname{Hom}_\Lambda(I_w,?)}{\overset{I_w \otimes_\Lambda ?}{\rightleftarrows}} \mathscr{R}_{w\theta} .$$

(ii) We have $(\mathscr{T}^w, \mathscr{F}^w) = (\overline{\mathscr{I}}_{w\theta}, \mathscr{P}_{w\theta})$.

(iii) If W_J is finite, then have $(\mathscr{T}^{ww_J}, \mathscr{F}^{ww_J}) = (\mathscr{I}_{w\theta}, \overline{\mathscr{P}}_{w\theta})$.

Proof. Let J, θ, w as in the statement of the theorem.

Given $T \in \Lambda$-mod, the condition $\langle \theta, \underline{\dim}\, T\rangle = 0$ is necessary for T to be in $\mathscr{R}_\theta$. It is also sufficient, because any quotient module X of T satisfies $\langle \theta, \underline{\dim}\, X\rangle \geq 0$ by the dominance of θ. The equality $\mathscr{R}_\theta = \Lambda_J$-mod then follows from Example 5.6 (ii).

Since θ is in F_J, it takes a positive value at each root in $\Phi_+ \setminus \Phi_J$. Corollary 2.2 then ensures that θ take positive values on $N_{w^{-1}}$. Choosing a reduced decomposition $w = s_{i_1} \cdots s_{i_\ell}$, we obtain $\langle s_{i_{k+1}} \cdots s_{i_\ell}\theta, \alpha_{i_k}\rangle > 0$ for each $1 \leq k \leq \ell$. Using Theorem 5.17, we get a chain of equivalences of categories

$$\mathscr{R}_\theta \underset{\Sigma_{i_\ell}}{\overset{\Sigma^*_{i_\ell}}{\rightleftarrows}} \mathscr{R}_{s_{i_\ell}\theta} \underset{\Sigma_{i_{\ell-1}}}{\overset{\Sigma^*_{i_{\ell-1}}}{\rightleftarrows}} \cdots \underset{\Sigma_{i_2}}{\overset{\Sigma^*_{i_2}}{\rightleftarrows}} \mathscr{R}_{s_{i_\ell}\cdots s_{i_2}\theta} \underset{\Sigma_{i_1}}{\overset{\Sigma^*_{i_1}}{\rightleftarrows}} \mathscr{R}_{w\theta} .$$

By composition, we get assertion (i).

Let $T \in \mathscr{T}^w$ and let X be a quotient of T; then $X \in \mathscr{T}^w$. By Theorem 5.4 (iii) and Corollary 5.8, $\underline{\dim}\, X$ is of the form $w\nu$ with $\nu \in \mathbb{N}I$, and so $\langle w\theta, \underline{\dim}\, X\rangle = \langle \theta, \nu\rangle \geq 0$. This proves that $T \in \overline{\mathscr{I}}_{w\theta}$.

Let $T \in \mathscr{F}^w$ and let $X \subseteq T$ be a nonzero submodule; then X is in $\mathscr{F}^w$. Theorem 5.10 (ii) and Lemma 5.9 then imply that $\underline{\dim}\, X$ is a nontrivial $\mathbb{N}$-linear combination of elements in N_w. In addition, $w\theta$ takes negative values on N_w, for θ takes positive values on $N_{w^{-1}} = -w^{-1}N_w$, as we have seen during the course of the proof of assertion (i). Therefore $\langle w\theta, \underline{\dim}\, X\rangle < 0$. This reasoning shows that $T \in \mathscr{P}_{w\theta}$.

We have established that $\mathscr{T}^w \subseteq \overline{\mathscr{I}}_{w\theta}$ and that $\mathscr{F}^w \subseteq \mathscr{P}_{w\theta}$. This implies assertion (ii).

We now prove assertion (iii), assuming that W_J is finite.

Consider first $T \in \mathscr{T}^{ww_J}$ and take a nonzero quotient X of T. Then X also belongs to $\mathscr{T}^{ww_J}$, and by Proposition 5.7, we can write $X = I_w \otimes_\Lambda Y$, with $Y \in \mathscr{T}^{w_J}$. Since $X \neq 0$, we have $Y \neq 0$, hence $Y \notin \mathscr{F}^{w_J}$. By Example 5.6 (ii), this means that $\underline{\dim}\, Y$ is not in $\mathbb{N}J$, the set of $\mathbb{N}$-linear combinations of elements in $\{\alpha_j \mid j \in J\}$. Therefore $\langle w\theta, \underline{\dim}\, X\rangle = \langle \theta, \underline{\dim}\, Y\rangle > 0$. This proves that $T \in \mathscr{I}_{w\theta}$.

Now let $T \in \mathscr{F}^{ww_J}$ and take a submodule $X \subseteq T$. Then $X \in \mathscr{F}^{ww_J}$ as well. Theorem 5.10 (ii) and Lemma 5.9 then imply that $\underline{\dim}\, X$ is a $\mathbb{N}$-linear combination of elements in N_{ww_J}. Certainly, $ww_J\theta$ takes nonpositive values on N_{ww_J}, so $\langle ww_J\theta, \underline{\dim}\, X\rangle \leq 0$. Observing that $w_J\theta = \theta$, we conclude that $T \in \overline{\mathscr{P}}_{w\theta}$.

We have established that $\mathscr{T}^{ww_J} \subseteq \mathscr{I}_{w\theta}$ and that $\mathscr{F}^{ww_J} \subseteq \overline{\mathscr{P}}_{w\theta}$, whence assertion (iii). □

Remarks 5.19. (i) In the context of assertions (ii) and (iii) of Theorem 5.18, we have $T^w = T^{\max}_{w\theta}$ and $T^{ww_J} = T^{\min}_{w\theta}$ for any Λ-module T. This shows in particular that each $\underline{\dim}\, T^w$ is a vertex of the HN polytope $\mathrm{Pol}(T)$.

(ii) With the help of Remark 4.1 and Example 5.6 (iii), we can complement the statement of Theorem 5.18 as follows:

$$(\mathscr{T}_{w^{-1}}, \mathscr{F}_{w^{-1}}) = (\mathscr{I}_{-w\theta}, \overline{\mathscr{P}}_{-w\theta}) \quad \text{and} \quad (\mathscr{T}_{w_Jw^{-1}}, \mathscr{F}_{w_Jw^{-1}}) = (\overline{\mathscr{I}}_{-w\theta}, \mathscr{P}_{-w\theta}).$$

(iii) Combining (ii) with Proposition 4.2 (iii) and Remark 5.14, we get another proof of the openness statement in Remark 14.2 of [27].

(iv) In the case where W_J is finite, assertion (i) of Theorem 5.18 is a particular case of Proposition 5.16, since $\mathscr{R}_{w\theta} = \overline{\mathscr{I}}_{w\theta} \cap \overline{\mathscr{P}}_{w\theta} = \mathscr{T}^w \cap \mathscr{F}^{ww_J}$ and $\mathscr{R}_\theta = \Lambda_J\text{-mod} = \mathscr{F}^{w_J}$.

Our first corollary to Theorem 5.18 shows that the position of the facets of our polytopes can be computed as the dimension of homomorphism spaces.

Corollary 5.20 *Let J, θ and w be as in Theorem 5.18. Suppose that θ is integral, that is, each $\langle\theta, \alpha_i\rangle$ is an integer. Set $N(w\theta) = \bigoplus_{i\in I}(I_w \otimes_\Lambda \Lambda e_i)^{\oplus\langle\theta,\alpha_i\rangle}$ and denote by $N(-w\theta)$ the dual of $\bigoplus_{i\in I}((\Lambda/I_w) \otimes_\Lambda \Lambda e_i)^{\oplus\langle\theta,\alpha_i\rangle}$. Then for any Λ-module T,*

$$\dim \mathrm{Hom}_\Lambda(N(\pm w\theta), T) = \psi_{\mathrm{Pol}(T)}(\pm w\theta).$$

Proof. The case $+w\theta$ comes from the computation

$$\begin{aligned}\dim \operatorname{Hom}_\Lambda(N(w\theta), T) &= \sum_{i\in I} \langle \theta, \alpha_i\rangle \dim \operatorname{Hom}_\Lambda(\Lambda e_i, \operatorname{Hom}_\Lambda(I_w, T)) \\ &= \sum_{i\in I} \langle \theta, \alpha_i\rangle \dim e_i \operatorname{Hom}_\Lambda(I_w, T) \\ &= \langle \theta, \underline{\dim} \operatorname{Hom}_\Lambda(I_w, T)\rangle \\ &= \langle w\theta, \underline{\dim}\, T^w\rangle \\ &= \langle w\theta, \underline{\dim}\, T^{\max}_{w\theta}\rangle \\ &= \psi_{\operatorname{Pol}(T)}(w\theta).\end{aligned}$$

Here the first equality is adjunction, the fourth is Corollary 5.8, and the fifth is Remark 5.19 (i).

Now applying the functor $\operatorname{Hom}_\Lambda(?, T^*)$ to the short exact sequence $0 \to I_w \otimes_\Lambda \Lambda e_i \to \Lambda e_i \to (\Lambda/I_w) \otimes_\Lambda \Lambda e_i \to 0$, we get a long exact sequence

$$\begin{aligned}0 \to \operatorname{Hom}_\Lambda((\Lambda/I_w) \otimes_\Lambda \Lambda e_i, T^*) &\to \operatorname{Hom}_\Lambda(\Lambda e_i, T^*) \\ &\to \operatorname{Hom}_\Lambda(I_w \otimes_\Lambda \Lambda e_i, T^*) \to \operatorname{Ext}^1_\Lambda((\Lambda/I_w) \otimes_\Lambda \Lambda e_i, T^*) \to 0.\end{aligned}$$

Taking dimensions, and using Crawley-Boevey's formula (4.2) and Example 5.12, we obtain

$$\begin{aligned}&\dim \operatorname{Hom}_\Lambda(T^*, (\Lambda/I_w) \otimes_\Lambda \Lambda e_i) \\ &\quad = \dim \operatorname{Hom}_\Lambda(I_w \otimes_\Lambda \Lambda e_i, T^*) - \dim \operatorname{Hom}_\Lambda(\Lambda e_i, T^*) + (\underline{\dim}\, (\Lambda/I_w) \otimes_\Lambda \Lambda e_i, \underline{\dim}\, T^*) \\ &\quad = \dim \operatorname{Hom}_\Lambda(I_w \otimes_\Lambda \Lambda e_i, T^*) - \dim e_i T^* + (\omega_i - w\omega_i, \underline{\dim}\, T^*).\end{aligned}$$

Multiplying by $\langle \theta, \alpha_i\rangle$ and summing over $i \in I$ gives then

$$\begin{aligned}\dim \operatorname{Hom}_\Lambda(N(-w\theta), T) &= \sum_{i\in I} \langle \theta, \alpha_i\rangle \dim \operatorname{Hom}_\Lambda(T^*, (\Lambda/I_w) \otimes_\Lambda \Lambda e_i) \\ &= \dim \operatorname{Hom}_\Lambda(N(w\theta), T^*) - \langle \theta, \underline{\dim}\, T^*\rangle + \langle \theta - w\theta, \underline{\dim}\, T^*\rangle \\ &= \psi_{\operatorname{Pol}(T^*)}(w\theta) - \langle w\theta, \underline{\dim}\, T\rangle.\end{aligned}$$

Remark 4.1 shows that the right-hand side is $\psi_{\operatorname{Pol}(T)}(-w\theta)$. $\square$

Our second corollary compares the normal fan of an HN polytope to the Tits fan.

Corollary 5.21 *The support function of the HN polytope of a Λ-module is linear on each face wF_J of the Tits cone.*

Proof. Let T be a Λ-module. The support function of $\operatorname{Pol}(T)$ is given by $\theta \mapsto \langle \theta, \underline{\dim}\, T_\theta^{\max} \rangle$. However, $T_\theta^{\max} = T^w$ if $\theta \in wF_J$, with w chosen J-reduced on the right. □

Using Remark 4.1, we see that the support function of an HN polytope is also linear on each face of the opposite $-C_T$ of the Tits cone.

The third corollary describes the face of an HN polytope defined by a linear form in the Tits cone.

Corollary 5.22 *Let J, θ and w as in the statement of Theorem 5.18 and let $T \in \Lambda\text{-mod}$. Set $X = \operatorname{Hom}_\Lambda(I_w, T_{w\theta}^{\max}/T_{w\theta}^{\min})$. Then X is in the subcategory $\Lambda_J\text{-mod}$ and*

$$\{x \in \operatorname{Pol}(T) \mid \langle w\theta, x\rangle = \psi_{\operatorname{Pol}(T)}(w\theta)\} = \underline{\dim}\, T_{w\theta}^{\min} + w\operatorname{Pol}(X).$$

Proof. This follows by combining Corollary 3.3, Theorem 5.18 (i), Corollary 5.8 and Example 5.6 (ii). □

In the case $J = \varnothing$, that is, if θ belongs to the open cone C_0, then $T^w = T_{w\theta}^{\min} = T_{w\theta}^{\max}$, $\mathscr{I}_{w\theta} = \overline{\mathscr{I}}_{w\theta} = \mathscr{T}^w$, $\mathscr{P}_{w\theta} = \overline{\mathscr{P}}_{w\theta} = \mathscr{F}^w$, and $\mathscr{R}_{w\theta} \cong \mathscr{R}_\theta = \Lambda_J\text{-mod} = \{0\}$. This case corresponds of course to a vertex of $\operatorname{Pol}(T)$.

In the case where J has just one element, say i, then $w_J = s_i$, $(T_{w\theta}^{\min}, T_{w\theta}^{\max}) = (T^{ws_i}, T^w)$, $\mathscr{R}_\theta = \Lambda_J\text{-mod} = \operatorname{add} S_i$, and $\mathscr{R}_{w\theta} = \operatorname{add}(I_w \otimes_\Lambda S_i)$. This case corresponds to an edge of $\operatorname{Pol}(T)$ that points in the direction $\underline{\dim}\,(I_w \otimes_\Lambda S_i) = w\alpha_i$.

The case where J contains just two vertices i and j linked by a single edge is more interesting. Here $w_J = s_is_js_i = s_js_is_j$, $(T_{w\theta}^{\min}, T_{w\theta}^{\max}) = (T^{ww_J}, T^w)$, and $\mathscr{R}_\theta = \Lambda_J\text{-mod}$ has four indecomposables. This case corresponds to a 2-face of type A_2. We will come back to this case soon: if T is a general point in an irreducible component Z, then the shape of this 2-face will be constrained by the tropical Plücker relations.

5.5 Tilting structure and crystal operations

Thanks to Theorem 4.4, we know that under a suitable openness condition (O), each torsion pair $(\mathscr{T}, \mathscr{F})$ gives rise to a bijection $\Xi : \mathfrak{T} \times \mathfrak{F} \to \mathfrak{B}$. The aim of this section is to show that in the case of the torsion pair $(\mathscr{T}^w, \mathscr{F}^w)$, this bijection can be described by elementary operations on the crystal $\mathfrak{B}$. Given $\nu \in \mathbb{N}I$, both sets $\{T \in \Lambda(\nu) \mid T \in \mathscr{T}^w\}$ and $\{T \in \Lambda(\nu) \mid T \in \mathscr{F}^w\}$ are open (Proposition 4.2 (iii) and Theorem 5.18 (ii)), so we can define the subsets $\mathfrak{T}^w$ and $\mathfrak{F}^w$ of $\mathfrak{B}$ formed by irreducible components whose general point belongs to $\mathscr{T}^w$ and $\mathscr{F}^w$. We define $\mathfrak{T}_w$ and $\mathfrak{F}_w$ in a similar fashion.

The first ingredient is the particular case where w is a simple reflection s_i. By Example 5.6 (i), $\mathfrak{F}^{s_i} = \bigsqcup_{n\in\mathbb{N}} \mathfrak{B}(n\alpha_i)$ is in bijection with $\mathbb{N}$; moreover, by the definition of the crystal structure on $\mathfrak{B}$ (see section 4.3), we have $\mathfrak{T}^{s_i} = \{Z \in \mathfrak{B} \mid \varphi_i(Z) = 0\}$. Under the identification $\mathfrak{F}^{s_i} \cong \mathbb{N}$, the bijection $\mathfrak{T}^{s_i} \times \mathfrak{F}^{s_i} \to \mathfrak{B}$ becomes the map $(Z, n) \mapsto \tilde{e}_i^n Z$. The inverse of this map is $Z \mapsto (\tilde{f}_i^{\max} Z, \varphi_i(Z))$, where $\tilde{f}_i^{\max} b = (\tilde{f}_i)^{\varphi_i(b)}\, b$.

To go further, we need to understand how the equivalence of categories provided by Theorem 5.4 (iii) relates to crystal operations. Twisting the usual crystal operators by the involution $*$, one defines the starred operators on $B(-\infty)$ as follows:

$$\tilde{e}_i^* : b \mapsto (\tilde{e}_i\, b^*)^*, \qquad \tilde{f}_i^* : b \mapsto (\tilde{f}_i\, b^*)^*, \qquad (\tilde{f}_i^*)^{\max} : b \mapsto (\tilde{f}_i^{\max}\, b^*)^*.$$

In [57], Corollary 3.4.8 (see also [38], section 8.2), Saito defines mutually inverse bijections

$$\{b \in B(-\infty) \mid \varphi_i(b) = 0\} \underset{S_i^*}{\overset{S_i}{\rightleftarrows}} \{b \in B(-\infty) \mid \varphi_i(b^*) = 0\}$$

by the rules $S_i(b) = (\tilde{e}_i)^{\varepsilon_i(b^*)}(\tilde{f}_i^*)^{\max} b$ and $S_i^*(b) = (\tilde{e}_i^*)^{\varepsilon_i(b)}\tilde{f}_i^{\max} b$. This definition ensures that if $\varphi_i(b) = 0$, then $\operatorname{wt} S_i(b) = s_i(\operatorname{wt} b)$. For convenience, we extend S_i and S_i^* on $B(-\infty)$ by setting $\sigma_i b = S_i(\tilde{f}_i^{\max} b)$ and $\sigma_i^* b = S_i^*((\tilde{f}_i^*)^{\max} b)$. By transport through the bijection $B(-\infty) \cong \mathfrak{B}$, we can view the maps σ_i and σ_i^* as maps from $\mathfrak{B}$ to itself. In view of Example 5.6 (i), σ_i and σ_i^* restrict to mutually inverse bijections

$$\mathfrak{T}^{s_i} \underset{\sigma_i^*}{\overset{\sigma_i}{\rightleftarrows}} \mathfrak{F}_{s_i} .$$

Proposition 5.23 *(i) Let $\nu \in \mathbb{N}I$, let $i \in I$, let $Z \in \mathfrak{F}_{s_i}(\nu)$ and let $Z' = \sigma_i^*(Z)$, an element in $\mathfrak{T}^{s_i}(s_i\nu)$. Let $U = \{T \in Z \mid \operatorname{soc}_i T = 0\}$ and $U' = \{T' \in Z' \mid \operatorname{hd}_i T' = 0\}$. Let Θ be the set of triples (T, T', h), such that $(T, T') \in U \times U'$ and $h : T' \to \Sigma_i^* T$ is an isomorphism. Then the first projection $\Theta \to U$ and the second one $\Theta \to U'$ are locally trivial fibrations with a smooth and connected fiber.*

(ii) Let $(u, v, i) \in W^2 \times I$ be such that $\ell(us_iv) = \ell(u) + \ell(v) + 1$. Then $\mathfrak{F}^v \subseteq \mathfrak{F}_{us_i}$ and $\mathfrak{F}^{s_iv} \subseteq \mathfrak{F}_u$. In addition, σ_i and σ_i^* restrict to mutually inverse bijections

$$\mathfrak{F}_u \cap \mathfrak{T}^{s_iv} \underset{\sigma_i^*}{\overset{\sigma_i}{\rightleftarrows}} \mathfrak{F}_{us_i} \cap \mathfrak{T}^v \quad \text{and} \quad \mathfrak{F}^{s_iv} \cap \mathfrak{T}^{s_i} \underset{\sigma_i^*}{\overset{\sigma_i}{\rightleftarrows}} \mathfrak{F}^v .$$

Proof. Assertion (i) is [5], Theorem 5.3. It is the precise way of stating that if T is a general point in Z, then $\Sigma_i^* T$ "belongs" to Z' and is general in Z'. (The quotes around "belongs" reflects the fact that a point of Z' can only be isomorphic to $\Sigma_i^* T$, and not equal to it.) Assertion (ii) follows then from Propositions 5.7, 5.15 and 5.16. □

Let $(w,i) \in W \times I$ be such that $\ell(ws_i) > \ell(w)$. By Proposition 5.16 and Example 5.6 (i), we have $\mathscr{F}^{ws_i} \cap \mathscr{T}^w = \operatorname{add}(I_w \otimes_\Lambda S_i)$, so $\mathfrak{F}^{ws_i} \cap \mathfrak{T}^w$ is in bijection with $\mathbb{N}$: to an integer n corresponds the closure in $\Lambda(w\alpha_i)$ of the orbit representing the module $I_w \otimes_\Lambda S_i^{\oplus n}$.

Now choose a finite sequence $\mathbf{i} = (i_1, \dots, i_\ell)$ such that $s_{i_1} \cdots s_{i_\ell}$ is a reduced decomposition. Set $(\mathscr{T}_0, \mathscr{F}_0) = (\Lambda\text{-mod}, \{0\})$, and for $1 \leq k \leq \ell$, set $(\mathscr{T}_k, \mathscr{F}_k) = (\mathscr{T}^{s_{i_1} \cdots s_{i_k}}, \mathscr{F}^{s_{i_1} \cdots s_{i_k}})$. Then

$$(\mathscr{T}_\ell, \mathscr{F}_\ell) \preccurlyeq \cdots \preccurlyeq (\mathscr{T}_0, \mathscr{F}_0).$$

The generalization of Proposition 4.5 to finite sequences provides a bijection

$$\Omega_{\mathbf{i}} : \mathfrak{B} \to \prod_{k=1}^{\ell} (\mathfrak{F}_k \cap \mathfrak{T}_{k-1}) \times \mathfrak{T}_\ell.$$

(Here we have used the inverse map to that defined in Proposition 4.5 and have reversed the order of the factors.) Under the identification $\mathfrak{F}_k \cap \mathfrak{T}_{k-1} \cong \mathbb{N}$, this bijection $\Omega_{\mathbf{i}}$ can be expressed in terms of the crystal operations in the following way.

Proposition 5.24 *Let $Z \in \mathfrak{B}$. Set $Z' = (\sigma^*_{i_1} \cdots \sigma^*_{i_\ell} \sigma_{i_\ell} \cdots \sigma_{i_1})(Z)$ and $n_k = \varphi_{i_k}(\sigma_{i_{k-1}} \cdots \sigma_{i_1} Z)$ for $1 \leq k \leq \ell$. Then $\Omega_{\mathbf{i}}(Z) = (n_1, \dots, n_\ell, Z')$.*

Proof. Set $i = i_1$ and $\mathbf{i}' = (i_2, \dots, i_\ell)$. For $2 \leq k \leq \ell$, set $(\mathscr{T}'_k, \mathscr{F}'_k) = (\mathscr{T}^{s_{i_2} \cdots s_{i_k}}, \mathscr{F}^{s_{i_2} \cdots s_{i_k}})$. By Propositions 5.7 and 5.16, Σ_i and Σ_i^* restrict to equivalences of categories

$$\mathscr{T}^{s_i} \cap \mathscr{T}_k \underset{\Sigma_i^*}{\overset{\Sigma_i}{\rightleftarrows}} \mathscr{F}_{s_i} \cap \mathscr{T}'_k \quad \text{and} \quad \mathscr{T}^{s_i} \cap \mathscr{F}_k \underset{\Sigma_i^*}{\overset{\Sigma_i}{\rightleftarrows}} \mathscr{F}_{s_i} \cap \mathscr{F}'_k\ .$$

Let $T \in Z$ be a general point and let X be its torsion submodule with respect to $(\mathscr{T}_1, \mathscr{F}_1)$. Then X is the top step in the filtration of T defined by our nested sequence of torsion pairs. The other modules define a filtration of X, whose image by Σ_i is the filtration on $\Sigma_i X$ defined by the nested sequence

$$(\mathscr{T}'_\ell, \mathscr{F}'_\ell) \preccurlyeq \cdots \preccurlyeq (\mathscr{T}'_2, \mathscr{F}'_2).$$

Now X is a general point of $\tilde{f}_i^{\max} Z$, so $\Sigma_i X$ is a general point of $\sigma_i(Z)$. By induction, we then have $\Omega_{\mathbf{i}'}(\sigma_i(Z)) = (n_2, \dots, n_\ell, Z'')$, where $Z'' = \sigma^*_{i_2} \cdots \sigma^*_{i_\ell} \sigma_{i_\ell} \cdots \sigma_{i_2}(\sigma_i Z)$. On the other hand, we have $T/X \cong S_i^{\oplus n_1}$, because X is the kernel of the canonical map $T \to \operatorname{hd}_i T$ and n_1 is the dimension of the i-head of T. The result now follows from Proposition 5.23 (ii). $\square$

Remarks 5.25. (i) Set $w = s_{i_1} \cdots s_{i_\ell}$. Obviously, $\Omega_{\mathbf{i}}$ induces a bijection between $\mathfrak{F}^w$ and $\mathbb{N}^l$. Up to duality, this bijection corresponds to the parameterization defined by Geiß, Leclerc and Schröer ([27], Proposition 14.5), as one sees from our discussion in Example 5.14. In this context, Proposition 5.24 has been independently proved by Jiang [33].

(ii) Consider the case where $\mathfrak{g}$ is finite dimensional and $s_{i_1} \cdots s_{i_\ell} = w_0$. Then $\Omega_{\mathbf{i}}$ induces a bijection between $\mathfrak{B}$ and $\mathbb{N}^l$. For $b \in B(-\infty)$, the procedure to compute $\Omega_{\mathbf{i}}(\Lambda_b)$ given in Proposition 5.24 coincides with Saito's method to determine the Lusztig datum of b in direction $\mathbf{i}$ (compare with the proof of [57], Lemma 4.1.3). So in this case, $\Omega_{\mathbf{i}}$ gives the usual Lusztig data. This gives us an incentive to use $\Omega_{\mathbf{i}}$ in order to define Lusztig data in the affine type case as well; we will pursue this road in section 7.6.

We now examine how $\Omega_{\mathbf{i}}$ depends on $\mathbf{i}$, when $s_{i_1} \cdots s_{i_\ell}$ is fixed in W. Matsumoto's lemma instructs us to look what happens under a braid move, so let us locate a subword of the form (i, j, i) in $\mathbf{i}$, where i and j are linked with a single edge in the graph (I, E). Let us denote by $m + 1$ the index at which this subword begins and let us denote by $\mathbf{j}$ the result of the substitution of (i, j, i) by (j, i, j) in $\mathbf{i}$.

Proposition 5.26 *Let $Z \in \mathfrak{B}$ and write $\Omega^{\mathbf{i}}(Z) = (n_1, \ldots, n_m, p, q, r, n_{m+4}, \ldots, n_\ell, Z')$. Define p', q' and r' by the formulas*

$$q' = \min(p, r), \quad p' + q' = q + r, \quad q' + r' = p + q.$$

Then $\Omega^{\mathbf{j}}(Z) = (n_1, \ldots, n_m, p', q', r', n_{m+4}, \ldots, n_\ell, Z')$.

Proof. We adopt the notation of the statement of the proposition. Let us set $J = \{i, j\}$.

The nested sequences of torsion pairs defined by $\mathbf{i}$ and $\mathbf{j}$ differ only in the places $m + 1$, $m + 2$ and $m+3$, so $\Omega_{\mathbf{i}}(Z)$ and $\Omega_{\mathbf{j}}(Z)$ differ only there. Moreover, thanks to Proposition 5.24, we can reduce to the case where $m = 0$ by replacing Z by $\sigma_{i_m} \cdots \sigma_{i_1}(Z)$. Finally, we can also assume without loss of generality that $\ell = 3$ and that $Z \in \mathfrak{F}^{w_J}$.

By Example 5.6 (ii), the category $\mathscr{F}^{w_J}$ is isomorphic to the category of representations of the preprojective algebra Λ_J, of type A_2. It is well-known that this category has four indecomposable objects, namely two simple objects S_i and S_j and two objects T_i and T_j of Loewy length 2, obtained as the middle terms of non-split extensions $0 \to S_j \to T_i \to S_i \to 0$ and $0 \to S_i \to T_j \to S_j \to 0$; the structure of these modules T_i and T_j can be represented pictorially as follows:

$$T_i = \begin{pmatrix} i & & \\ & \searrow & \\ & & j \end{pmatrix} \quad \text{and} \quad T_j = \begin{pmatrix} & & j \\ & \swarrow & \\ i & & \end{pmatrix}.$$

Thus, there exists $(a, b, c, d) \in \mathbb{N}^4$ such that a general point in Z is isomorphic to

$$S_i^{\oplus a} \oplus S_j^{\oplus b} \oplus T_i^{\oplus c} \oplus T_j^{\oplus d},$$

with moreover $\min(a, b) = 0$, because a module which has $S_i \oplus S_j$ as summand cannot be general. One then easily computes

$$\Omega^{(i,j,i)}(Z) = (a + c, d, b + c, \{0\}) \qquad \text{and} \qquad \Omega^{(j,i,j)}(Z) = (b + d, c, a + d, \{0\}).$$

Setting $p = a + c$, $q = d$ and $r = b + c$, one checks that $p' = b + d$, $q' = c$, $r' = a + d$, which shows the desired result. $\square$

Likewise, let us locate a subword of the form (i, j) in $\mathbf{i}$, where i and j are not linked in the graph (I, E). Let us denote by $m+1$ the index at which this subword begins and let us denote by $\mathbf{j}$ the result of the substitution of (i, j) by (j, i) in $\mathbf{i}$.

Proposition 5.27 *Let $Z \in \mathfrak{B}$ and write $\Omega^{\mathbf{i}}(Z) = (n_1, \ldots, n_m, p, q, n_{m+3}, \ldots, n_\ell, Z')$. Then $\Omega^{\mathbf{j}}(Z) = (n_1, \ldots, n_m, q, p, n_{m+3}, \ldots, n_\ell, Z')$.*

The proof is similar to that of Proposition 5.26, save that we now deal with a preprojective algebra Λ_J of type $A_1 \times A_1$.

For a fixed $w \in W$, let $\mathscr{X}(w)$ be the set of all tuples $\mathbf{i}$ that represent a reduced decomposition of w. Lusztig's piecewise linear bijections $R_{\mathbf{i}}^{\mathbf{j}} : \mathbb{N}^\ell \to \mathbb{N}^\ell$ ([41], sections 2.1 and 2.6) can be defined here just as in the case where w is the longest element in a finite W. Since any two elements in $\mathscr{X}(w)$ can be related by a sequence of braid and of commutation relations, Propositions 5.26 and 5.27 say that the numerical parts of $\Omega_{\mathbf{i}}$ and $\Omega_{\mathbf{j}}$ are related by $R_{\mathbf{i}}^{\mathbf{j}}$.

To conclude, let us consider again Proposition 5.24 and choose a general point $T \in Z$. The integers $n_1, \ldots, n_\ell$ are equal to the lengths of the edges of $\mathrm{Pol}(T)$ along the path

$$\underline{\dim}\, T, \quad \underline{\dim}\, T^{s_{i_1}}, \quad \underline{\dim}\, T^{s_{i_1} s_{i_2}}, \quad \ldots, \quad \underline{\dim}\, T^{s_{i_1} \cdots s_{i_\ell}}.$$

Now look at Proposition 5.26, set $J = \{i, j\}$, choose $\theta \in F_J$, and set $u = s_{i_1} \cdots s_{i_m}$; then u is J-reduced on the right. Using Theorem 5.18 (i) and Corollary 5.21, one can show that the vertices on the face of $\mathrm{Pol}(T)$ defined by $u\theta$ are the six vertices $\underline{\dim}\, T^{uv}$, with $v \in W_J$. The relation given in Proposition 5.26 constrains the lengths of the edges of this face; it is equivalent to the tropical Plücker relations of [36].

5.6 The finite type case

Our main focus of interest in this paper concerns the case where Φ is of affine type, to which section 5.2 directly applies. In the finite type case, the ideals I_w are not tilting of projective dimension at most 1 anymore. They nevertheless exist, so we can define the full subcategories

$$\begin{aligned}\mathscr{T}^w &= \text{essential image of } I_w \otimes_\Lambda ?,\\ \mathscr{F}^w &= \text{kernel of } \operatorname{Hom}_\Lambda(I_w, ?),\\ \mathscr{T}_w &= \text{kernel of } I_w \otimes_\Lambda ?,\\ \mathscr{F}_w &= \text{essential image of } \operatorname{Hom}_\Lambda(I_w, ?)\end{aligned}$$

of Λ-mod. With this definition, all the results in sections 5.1 hold unchanged, except of course Theorem 5.3 (iii).

To show this, one can adopt the method of Iyama, Reiten and their collaborators, namely, one chooses an embedding of the Dynkin diagram into a non-Dynkin one. One then get a natural surjective morphism from the preprojective algebra $\widehat{\Lambda}$ of non-Dynkin type onto the preprojective algebra Λ of Dynkin type, and thus a natural embedding of Λ-mod as a full subcategory of $\widehat{\Lambda}$-mod. This subcategory is abelian and closed under extensions. Each $i \in I$ yields then an ideal I_i of Λ and an ideal $\widehat{I}_i$ of $\widehat{\Lambda}$. By Proposition 5.1, the functors $I_i \otimes_\Lambda ?$ and $\widehat{I}_i \otimes_{\widehat{\Lambda}} ?$ (respectively, $\operatorname{Hom}_\Lambda(I_i, ?)$ and $\operatorname{Hom}_{\widehat{\Lambda}}(\widehat{I}_i, ?)$) coincide on Λ-mod.

Moreover, the Weyl group W of the Dynkin diagram embeds as a parabolic subgroup of the Weyl group $\widehat{W}$ of the larger diagram. Given $w \in W$, one can then define I_w as the bimodule $I_{i_1} \otimes_\Lambda \cdots \otimes_\Lambda I_{i_\ell}$ by choosing a reduced decomposition $w = s_{i_1} \cdots s_{i\ell}$. Since the functors $I_w \otimes_\Lambda ?$ and $\widehat{I}_w \otimes_{\widehat{\Lambda}} ?$ coincide on Λ-mod, we deduce the Dynkin case of Theorem 5.3 (iv) from the non-Dynkin case. The proof of [14], Proposition II.1.5 then immediately implies Theorem 5.3 (i). Using that moreover the functors $\operatorname{Hom}_\Lambda(I_w, ?)$ and $\operatorname{Hom}_{\widehat{\Lambda}}(\widehat{I}_w, ?)$ coincide on Λ-mod, we deduce the Dynkin case from the non-Dynkin case in Theorem 5.4, Examples 5.6 (i) and (ii), Propositions 5.7 and Corollary 5.8.

Finally, to see that Examples 5.6 (iii) and (iv) hold true in the Dynkin case, one can apply Example 5.13. One may here moreover note that the surjection $\widehat{\Lambda} \to \Lambda$ induces an isomorphism $\widehat{\Lambda}/\widehat{I}_w \cong \Lambda/I_w$ (this follows, for instance, by an obvious dimension argument based on Amiot, Iyama, Reiten and Todorov's filtration, see section 2 of [1]).

Another pecularity of the finite type case is the fact that the Tits cone C_T fills the whole dual space of $\mathbb{R}I$. By Theorem 5.18, any torsion pair $(\overline{\mathscr{I}}_\theta, \mathscr{P}_\theta)$ or $(\mathscr{I}_\theta, \overline{\mathscr{P}}_\theta)$ is of the form $(\mathscr{T}_w, \mathscr{F}_w)$, so these latter are enough to completely describe the polytopes $\operatorname{Pol}(T)$. This fact, combined with Corollary 5.20, shows that the definition of $\operatorname{Pol}(T)$ given in [5] is identical to the definition used in the present paper.

In addition, the intersection between C_T and its opposite is not $\{0\}$ anymore, so it is possible to write $-u^{-1}\eta = v\theta$ with $(u, v) \in W^2$ and both η and θ in C_0. In this case, we have

$$(\mathscr{T}^v, \mathscr{F}^v) = (\mathscr{I}_{v\theta}, \mathscr{I}_{v\theta}) = (\mathscr{I}_{-u^{-1}\eta}, \mathscr{P}_{-u^{-1}\eta}) = (\mathscr{T}_u, \mathscr{F}_u), \tag{5.4}$$

by Theorem 5.18 (ii) and Remark 5.19 (ii).

A last peculiarity of the finite type case was mentioned above in Remark 5.25 (ii).

6 The Hall functors

In this section, we assume we are given two orthogonal rigid bricks S and R in Λ-mod such that $\dim \operatorname{Ext}^1(S, R) = 2$. In other words, we assume that S and R are two finite dimensional Λ-modules such that

$$\begin{aligned} \operatorname{End}_\Lambda(S) = \operatorname{End}_\Lambda(R) = K, \qquad & \operatorname{Hom}_\Lambda(S, R) = \operatorname{Hom}_\Lambda(R, S) = 0, \\ \operatorname{Ext}^1_\Lambda(S, S) = \operatorname{Ext}^1_\Lambda(R, R) = 0, \qquad & \dim \operatorname{Ext}^1_\Lambda(S, R) = 2. \end{aligned} \tag{6.1}$$

We fix ξ and η in $\bigoplus_{a\in H} \operatorname{Hom}_K(S_{s(a)}, R_{t(a)})$ and $\hat{\xi}$ and $\hat{\eta}$ in $\bigoplus_{a\in H} \operatorname{Hom}_K(R_{s(a)}, S_{t(a)})$ such that, in the notation of section 4.2,

$$d^1_{S,R}(\xi) = d^1_{S,R}(\eta) = 0, \qquad d^1_{R,S}(\hat{\xi}) = d^1_{R,S}(\hat{\eta}) = 0, \tag{6.2}$$

$$\tau_1(\xi, \hat{\xi}) = \tau_1(\eta, \hat{\eta}) = 1, \qquad \tau_1(\xi, \hat{\eta}) = \tau_1(\eta, \hat{\xi}) = 0. \tag{6.3}$$

Thus (ξ, η) can be regarded as a basis of $\operatorname{Ext}^1_\Lambda(S, R)$ and $(\hat{\xi}, \hat{\eta})$ can be regarded as the dual basis of $\operatorname{Ext}^1_\Lambda(R, S)$.

6.1 A combinatorial lemma

We denote by Π the completed preprojective algebra of type $\widetilde{A}_1$. This is the completed preprojective algebra of the Kronecker quiver

$$\begin{tikzcd} 0 \arrow[r, bend left, "\alpha"] \arrow[r, bend right, "\beta"'] & 1. \end{tikzcd}$$

It contains orthogonal idempotents e_0 and e_1 and arrows $\alpha, \beta \in e_1 \Pi e_0$ and $\overline{\alpha}, \overline{\beta} \in e_0 \Pi e_1$. We denote its augmentation ideal by J.

For the following, it is useful to think of the algebra Π as the quotient

$$\mathbf{S}\langle\langle\, \alpha, \beta, \overline{\alpha}, \overline{\beta}\,\rangle\rangle/(\alpha\overline{\alpha}+\beta\overline{\beta}, \overline{\alpha}\alpha+\overline{\beta}\beta)$$

of the ring of non-commutative formal power series in four variables α, β, $\overline{\alpha}$, $\overline{\beta}$ with coefficients in the commutative semisimple algebra $\mathbf{S} = Ke_0 \oplus Ke_1$.

Lemma 6.1 *The image of the linear map*

$$C : \big(e_1\Pi e_0\big)^2 \times \big(e_0\Pi e_1\big)^2 \to e_0\Pi e_0 \times e_1\Pi e_1$$

given by

$$C(x, y, \hat{x}, \hat{y}) = (-\hat{x}\alpha - \hat{y}\beta - \overline{\alpha}x - \overline{\beta}y,\ x\overline{\alpha} + y\overline{\beta} + \alpha\hat{x} + \beta\hat{y})$$

contains:

- *any element of the form $(uv - vu, 0)$, where $(u, v) \in \big(e_0\Pi e_0\big)^2$;*
- *any element of the form $(0, uv - vu)$, where $(u, v) \in \big(e_1\Pi e_1\big)^2$;*
- *any element of the form $(uv, -vu)$, where $(u, v) \in e_0\Pi e_1 \times e_1\Pi e_0$.*

Proof. Let u and v be two words of even length in the alphabet $\{\alpha, \beta, \overline{\alpha}, \overline{\beta}\}$, in which barred and non-barred letters alternate, and which start with a barred letter. Thus u and v define elements in $e_0\Pi e_0$. We write $u = \overline{c_1}c_2 \cdots \overline{c_{2\ell-1}}c_{2\ell}$, with $c_k \in \{\alpha, \beta\}$. For $1 \le k \le \ell$, we set

$$m_k = c_{2k}\overline{c_{2k+1}} \cdots c_{2\ell}\ v\ \overline{c_1}c_2 \cdots c_{2k-2} \qquad \text{and} \qquad n_k = \overline{c_{2k+1}}c_{2k+2} \cdots c_{2\ell}\ v\ \overline{c_1}c_2 \cdots \overline{c_{2k-1}}.$$

Then $(uv - vu, 0)$ is the image of the element

$$\left(-\sum_{\substack{1\le k\le\ell \\ \overline{c_{2k-1}}=\overline{\alpha}}} m_k,\quad -\sum_{\substack{1\le k\le\ell \\ \overline{c_{2k-1}}=\overline{\beta}}} m_k,\quad \sum_{\substack{1\le k\le\ell \\ c_{2k}=\alpha}} n_k,\quad \sum_{\substack{1\le k\le\ell \\ c_{2k}=\beta}} n_k \right)$$

by our linear map. This shows that the elements of the first kind belong to the image of our map. The two other cases are similar. $\square$

6.2 A universal lifting

Given a K-vector space V, two elements X and Y in $V \otimes_K \Pi$, and an integer $k \geq 0$, generalizing a standard notation, we will write $X \equiv Y \bmod J^k$ to express that $X - Y$ belongs to $V \otimes_K J^k$.

For $a \in H$, we set $\dot{S}_a = S_a \otimes e_0$, where S_a is the linear map attached to the arrow a in the Λ-module S; this $\dot{S}_a$ is an element in $\operatorname{Hom}_K(S_{s(a)}, S_{t(a)}) \otimes_K e_0 \Pi e_0$. Likewise, we set $\dot{R}_a = R_a \otimes e_1$; this is an element in $\operatorname{Hom}_K(R_{s(a)}, R_{t(a)}) \otimes_K e_1 \Pi e_1$.

We set $P = \xi \otimes \alpha + \eta \otimes \beta$ and $Q = \hat{\xi} \otimes \overline{\alpha} + \hat{\eta} \otimes \overline{\beta}$; these are elements in

$$\bigoplus_{a \in H} \operatorname{Hom}_K(S_{s(a)}, R_{t(a)}) \otimes_K e_1 \Pi e_0 \qquad \text{and} \qquad \bigoplus_{a \in H} \operatorname{Hom}_K(R_{s(a)}, S_{t(a)}) \otimes_K e_0 \Pi e_1,$$

respectively. For $a \in H$, we write P_a for the component of P in $\operatorname{Hom}_K(S_{s(a)}, R_{t(a)}) \otimes_K e_1 \Pi e_0$ and Q_a for the component of Q in $\operatorname{Hom}_K(R_{s(a)}, S_{t(a)}) \otimes_K e_0 \Pi e_1$.

Lemma 6.2 *There are elements*

$$S^{(\infty)} \in \bigoplus_{a \in H} \operatorname{Hom}_K(S_{s(a)}, S_{t(a)}) \otimes_K e_0 \Pi e_0, \qquad R^{(\infty)} \in \bigoplus_{a \in H} \operatorname{Hom}_K(R_{s(a)}, R_{t(a)}) \otimes_K e_1 \Pi e_1,$$

$$P^{(\infty)} \in \bigoplus_{a \in H} \operatorname{Hom}_K(S_{s(a)}, R_{t(a)}) \otimes_K e_1 \Pi e_0, \qquad Q^{(\infty)} \in \bigoplus_{a \in H} \operatorname{Hom}_K(R_{s(a)}, S_{t(a)}) \otimes_K e_0 \Pi e_1$$

such that for each $a \in H$,

$$\begin{pmatrix} S_a^{(\infty)} & Q_a^{(\infty)} \\ P_a^{(\infty)} & R_a^{(\infty)} \end{pmatrix} \equiv \begin{pmatrix} \dot{S}_a & Q_a \\ P_a & \dot{R}_a \end{pmatrix} \mod \begin{pmatrix} J^2 & J^3 \\ J^3 & J^2 \end{pmatrix},$$

and for each $i \in I$,

$$\sum_{\substack{a \in H \\ s(a)=i}} \varepsilon(a) \begin{pmatrix} S_{\overline{a}}^{(\infty)} & Q_{\overline{a}}^{(\infty)} \\ P_{\overline{a}}^{(\infty)} & R_{\overline{a}}^{(\infty)} \end{pmatrix} \begin{pmatrix} S_a^{(\infty)} & Q_a^{(\infty)} \\ P_a^{(\infty)} & R_a^{(\infty)} \end{pmatrix} = 0.$$

Proof. The desired elements will be constructed as the limit in the J-adic topology of elements $S^{(k)}$, $R^{(k)}$, $P^{(k)}$ and $Q^{(k)}$ such that

$$\begin{pmatrix} S^{(k+1)} & Q^{(k+1)} \\ P^{(k+1)} & R^{(k+1)} \end{pmatrix} \equiv \begin{pmatrix} S^{(k)} & Q^{(k)} \\ P^{(k)} & R^{(k)} \end{pmatrix} \mod \begin{pmatrix} J^{2k+2} & J^{2k+3} \\ J^{2k+3} & J^{2k+2} \end{pmatrix}.$$

For $k=0$, we define

$$S^{(0)} \in \bigoplus_{a \in H} \operatorname{Hom}_K(S_{s(a)}, S_{t(a)}) \otimes_K e_0 \Pi e_0$$

by gathering the elements $\dot{S}_a$. We similarly define $R^{(0)}$ and we set $P^{(0)} = P$ and $Q^{(0)} = Q$. The conditions we impose at step k are

$$\sum_{\substack{a \in H \\ s(a)=i}} \varepsilon(a) \begin{pmatrix} S_{\overline{a}}^{(k)} & Q_{\overline{a}}^{(k)} \\ P_{\overline{a}}^{(k)} & R_{\overline{a}}^{(k)} \end{pmatrix} \begin{pmatrix} S_a^{(k)} & Q_a^{(k)} \\ P_a^{(k)} & R_a^{(k)} \end{pmatrix} \equiv 0 \mod \begin{pmatrix} J^{2k+2} & J^{2k+3} \\ J^{2k+3} & J^{2k+2} \end{pmatrix} \tag{6.4}$$

for each $i \in I$,

$$\dot{\tau}_1\Big(S^{(k)}, S^{(k)}\Big) + \dot{\tau}_1\Big(Q^{(k)}, P^{(k)}\Big) \equiv 0 \mod J^{2k+4} \tag{6.5}$$

and

$$\dot{\tau}_1\Big(R^{(k)}, R^{(k)}\Big) + \dot{\tau}_1\Big(P^{(k)}, Q^{(k)}\Big) \equiv 0 \mod J^{2k+4}. \tag{6.6}$$

In (6.5) and (6.6), the symbol $\dot{\tau}_1$ represents $\tau_1 \otimes \mathrm{id}_\Pi$, the map obtained from the bilinear form τ_1 defined in section 4.2 by extending the scalars from K to Π. Similarly, we will denote the map $\tau_2 \otimes \mathrm{id}_\Pi$ by the symbol $\dot{\tau}_2$.

Thanks to the preprojective relations for the Λ-modules S and R and to equations (6.2) and (6.3), the conditions (6.4)–(6.6) are fulfilled at step $k=0$.

Let us assume that $S^{(k)}$, $R^{(k)}$, $P^{(k)}$ and $Q^{(k)}$ have been constructed. We set

$$f_i = \sum_{\substack{a \in H \\ s(a)=i}} \varepsilon(a) \Big(S_{\overline{a}}^{(k)} S_a^{(k)} + Q_{\overline{a}}^{(k)} P_a^{(k)} \Big)$$

and we regard (f_i) as a formal series in α, β, $\overline{\alpha}$, $\overline{\beta}$ with coefficients in $\bigoplus_{i \in I} \operatorname{Hom}_K(S_i, S_i)$ and valuation at least $2k+2$. Then, thanks to (6.5), we have

$$\dot{\tau}_2\big((f_i), (\mathrm{id}_{S_i} \otimes e_0)\big) \in J^{2k+4};$$

in other words, the coefficient in (f_i) of any monomial of degree less than $2k+4$ is τ_2-orthogonal to $\operatorname{Hom}_\Lambda(S, S)$. We conclude that modulo J^{2k+4}, (f_i) belongs $\operatorname{im} d^1_{S,S}$. Therefore there exists

$$\widetilde{S} \in \bigoplus_{a \in H} \operatorname{Hom}_K(S_{s(a)}, S_{t(a)}) \otimes_K e_0 \Pi e_0$$

of valuation at least $2k+2$ such that $d^1_{S,S}\big(\widetilde{S}\big) \equiv (f_i) \mod J^{2k+4}$. We then set $S^{(k+1)} = S^{(k)} - \widetilde{S}$, and the upper left corner of (6.4) is satisfied at step $k+1$. One similarly finds first $R^{(k+1)}$, and next $P^{(k+1)}$ and $Q^{(k+1)}$, that satisfy (6.4) at step $k+1$. However, the elements $P^{(k+1)}$ and $Q^{(k+1)}$ obtained in this way are not the final ones, since they do not yet satisfy (6.5) and (6.6).

The left-hand sides of (6.5) and (6.6) are the components of

$$D = \left(\dot{\tau}_1\left(S^{(k+1)}, S^{(k+1)}\right) + \dot{\tau}_1\left(Q^{(k+1)}, P^{(k+1)}\right),\ \dot{\tau}_1\left(R^{(k+1)}, R^{(k+1)}\right) + \dot{\tau}_1\left(P^{(k+1)}, Q^{(k+1)}\right)\right).$$

Thus $D \in e_0\Pi e_0 \times e_1\Pi e_1$ and its two components have valuations at least $2k+4$. Because of the cyclicity of the trace and of the presence of the sign $\varepsilon(a)$ in

$$\dot{\tau}_1\left(S^{(k+1)}, S^{(k+1)}\right) = \sum_{i\in I} \operatorname{Tr}\left(\sum_{\substack{a\in H\\ s(a)=i}} \varepsilon(a)\left(S_{\overline{a}}^{(k+1)} S_a^{(k+1)}\right)\right),$$

the first term in the first component of D is a linear combination of elements of the kind $uv - vu$, with $(u,v) \in \left(e_0\Pi e_0\right)^2$. Likewise, we see that the contributions to D of $R^{(k+1)}$, and of $P^{(k+1)}$ and $Q^{(k+1)}$, are linear combination of elements of the second and third kind in the statement of Lemma 6.1. Therefore D belongs to the image of the map C, so we can find $(x,y) \in (e_1\Pi e_0)^2$ and $(\hat{x},\hat{y}) \in (e_0\Pi e_1)^2$ of valuation at least $2k+3$ such that $D = C(x,y,\hat{x},\hat{y})$. If we now correct $P^{(k+1)}$ and $Q^{(k+1)}$ by subtracting $\xi\otimes x+\eta\otimes y$ from the former and $\hat{\xi}\otimes\hat{x}+\hat{\eta}\otimes\hat{y}$ from the latter, then D will vanish modulo J^{2k+6}, thanks to (6.3), while the condition (6.4) remains satisfied, thanks to (6.2). □

6.3 Construction of the Hall functors

We are now in a position to define a fully faithful exact functor $\mathscr{H} : \Pi\text{-mod} \to \Lambda\text{-mod}$.

Let V be a Π-module. We define an I-graded vector space $M = \bigoplus_{i\in I} M_i$ by

$$M_i = \left(S_i \otimes_K V_0\right) \oplus \left(R_i \otimes_K V_1\right).$$

Now let $a \in H$. Then $S_a^{(\infty)}$ is an element of $\operatorname{Hom}_K(S_{s(a)}, S_{t(a)}) \otimes_K e_0\Pi e_0$, hence can be seen as a matrix with entries in $e_0\Pi e_0$. We can evaluate these entries in the representation V; the result belongs to

$$\operatorname{Hom}_K(S_{s(a)}, S_{t(a)}) \otimes_K \operatorname{Hom}_K(V_0, V_0) = \operatorname{Hom}_K(S_{s(a)} \otimes_K V_0, S_{t(a)} \otimes_K V_0).$$

Evaluating $R_a^{(\infty)}$, $P_a^{(\infty)}$ and $Q_a^{(\infty)}$ in a similar way, we map

$$\begin{pmatrix} S_a^{(\infty)} & Q_a^{(\infty)} \\ P_a^{(\infty)} & R_a^{(\infty)} \end{pmatrix}$$

to an element in $\mathrm{Hom}_K(M_{s(a)}, M_{t(a)})$. The conditions imposed in Lemma 6.2 assert that we then get a Λ-module M.

Since we have worked with a universal formula (the same for all Π-modules V), the assignment $V \mapsto M$ defines a functor $\mathscr{H}$, which moreover is exact. Let $\langle S, R\rangle$ be the smallest abelian, closed under extensions, subcategory of Λ-mod that contains the isomorphism classes of S and R.

Theorem 6.3 *The functor $\mathscr{H}$ induces an equivalence of categories between* Π-mod *and* $\langle S, R\rangle$.

Proof. The category $\langle S, R\rangle$ has only (up to isomorphism) two simple objects, namely S and R, for these latter are orthogonal bricks. In view of [22], Lemma 11.7, it thus suffices to show that for any simple Π-modules L and L', the induced homomorphism $\mathrm{Ext}^k_\Pi(L, L') \to \mathrm{Ext}^k_{\langle S,R\rangle}(\mathscr{H}(L), \mathscr{H}(L'))$ is bijective for $k \in \{0, 1\}$ and injective for $k = 2$. We can here replace the extension spaces in $\langle S, R\rangle$ by the extension spaces in Λ-mod: this does not change the Ext^0 nor the Ext^1, for $\langle S, R\rangle$ is full and closed under extensions; and if the injectivity condition holds for Ext^2_Λ, it will a fortiori holds for $\mathrm{Ext}^2_{\langle S,R\rangle}$.

Let us call W_0 and W_1 the two simple Π-modules, concentrated at vertices 0 and 1 respectively; then $\mathscr{H}(W_0) = S$ and $\mathscr{H}(W_1) = R$. Obviously,

$$\mathrm{End}(W_0) = \mathrm{End}(W_1) = K \qquad \text{and} \qquad \mathrm{Hom}(W_0, W_1) = \mathrm{Hom}(W_1, W_0) = 0,$$

so the condition is fulfilled for $k = 0$.

The Π-modules T_α and T_β with dimension-vector $(1, 1)$ obtained by letting the arrows of Π act by

$$(\alpha, \beta, \overline{\alpha}, \overline{\beta}) \mapsto (1, 0, 0, 0) \qquad \text{and} \qquad (\alpha, \beta, \overline{\alpha}, \overline{\beta}) \mapsto (0, 1, 0, 0)$$

are extensions of W_0 by W_1. We denote their extension classes in $\mathrm{Ext}^1_\Pi(W_0, W_1)$ by α and β, respectively. The extension classes of $\mathscr{H}(T_\alpha)$ and $\mathscr{H}(T_\beta)$ are ξ and η. Thus, the induced homomorphism $\mathrm{Ext}^1_\Pi(W_0, W_1) \to \mathrm{Ext}^1_\Lambda(S, R)$ maps the basis (α, β) of the first space to the basis (ξ, η) of the second space; it is therefore bijective. We check in a similar way the other cases for $k = 1$.

The equality $\tau_1(\xi, \hat{\xi}) = 1$ implies that the Yoneda product $\xi\hat{\xi} \in \mathrm{Ext}^2_\Lambda(R, R)$ does not vanish. The induced homomorphism $\mathrm{Ext}^2_\Pi(W_1, W_1) \to \mathrm{Ext}^2_\Lambda(R, R)$ maps $\alpha\overline{\alpha}$ to $\xi\hat{\xi}$, so it cannot be zero. It is thus injective, for $\mathrm{Ext}^2_\Pi(W_1, W_1)$ is one dimensional. The other cases for $k = 2$ are treated in like manner. □

We here note that a proof for Lemma 11.7 in [22] can be found in [39], Proposition 3.4.3.

6.4 Irreducible components

We now study the consequences of the existence of a Hall functor at the level of irreducible components of the nilpotent varieties.

Let $\mu = (\mu_0, \mu_1)$ be a dimension-vector for Π and set $\nu = \mu_0 \underline{\dim} S + \mu_1 \underline{\dim} R$. We denote by $\Pi(\mu)$ the nilpotent variety for Π and by $\Lambda_{\langle S,R\rangle}(\nu)$ the set of all points in $\Lambda(\nu)$ that belong to $\langle S, R\rangle$. In addition, we define $\Omega(\mu)$ to be the set of all triples (V, M, f) such that $V \in \Pi(\mu)$, $M \in \Lambda(\nu)$ and $f : \mathscr{H}(V) \to M$ is an isomorphism of Λ-modules. We can then form the diagram

$$\Pi(\mu) \xleftarrow{p} \Omega(\mu) \xrightarrow{q} \Lambda(\nu) \tag{6.7}$$

in which p and q are the first and second projection. Obviously, p is a principal $G(\nu)$-bundle, the image of q is $\Lambda_{\langle S,R\rangle}(\nu)$, and each non-empty fiber of q is isomorphic to $G(\mu)$.

Proposition 6.4 *The subset $\Lambda_{\langle S,R\rangle}(\nu)$ is constructible and all its irreducible components have full dimension in $\Lambda(\nu)$. The diagram (6.7) induces a bijection between the irreducible components of $\Pi(\mu)$ and the irreducible components of $\Lambda(\nu)$ whose general point belongs to $\Lambda_{\langle S,R\rangle}(\nu)$.*

Proof. Combining the conditions (6.1) with equation (4.2), we get

$$(\mu, \mu) = (\nu, \nu),$$

where $(\, , \,)$ in the left-hand side is the bilinear form on $\mathbf{K}(\Pi\text{-mod})$ and $(\, , \,)$ in the right-hand side is the bilinear form on $\mathbf{K}(\Lambda\text{-mod}) \cong \mathbb{Z}I$. In view of (4.5), this translates to

$$\dim G(\mu) - \dim \Pi(\mu) = \dim G(\nu) - \dim \Lambda(\nu).$$

The proposition now results from general results of algebraic geometry, similar to those used in section 4.5. □

7 Preprojective algebras of affine type

From now on, $\mathfrak{g}$ is of symmetric affine type. We will apply the theory we have been developing to finally obtain our affine MV polytopes. We will use the notation concerning affine roots systems discussed in section 2.3.

7.1 Torsion pairs associated to biconvex subsets

From Corollary 5.21, it follows that the HN polytope $P = \mathrm{Pol}(T)$ of a finite dimensional Λ-module T is GGMS. The construction in section 2.6 then assigns a vertex $\mu_P(A)$ of P to each biconvex subset $A \subseteq \Phi_+$. As noticed after Corollary 3.3, $\mu_P(A)$ is the dimension vector of a unique submodule T_A of T. Our aim in this section is to construct a torsion pair $(\mathscr{T}(A), \mathscr{F}(A))$ with respect to which T_A is the torsion submodule of T, for all Λ-module T.

As in section 2.2, each $w \in W$ gives rise to two biconvex subsets: a finite one, namely $A_w = N_{w^{-1}}$, and a cofinite one, $A^w = \Phi_+ \setminus N_w$ (see Example 2.3 (ii)). We set

$$(\mathscr{T}(A_w), \mathscr{F}(A_w)) = (\mathscr{T}_w, \mathscr{F}_w) \quad \text{and} \quad (\mathscr{T}(A^w), \mathscr{F}(A^w)) = (\mathscr{T}^w, \mathscr{F}^w).$$

In this fashion, we associate a torsion pair in Λ-mod to each finite or cofinite biconvex set. (By (5.4), this construction is unambiguous in the finite type case, where biconvex sets are both finite and cofinite.)

Using Lemma 2.1, (5.3) and Proposition 5.15, we deduce the following monotonicity property: if $A \subseteq B$, then $(\mathscr{T}(A), \mathscr{F}(A)) \preccurlyeq (\mathscr{T}(B), \mathscr{F}(B))$. We also note the following interpretation of Example 5.6 (iii): if $B = \Phi_+ \setminus A$, then $(\mathscr{T}(B), \mathscr{F}(B)) = (\mathscr{F}(A)^*, \mathscr{T}(A)^*)$.

By Proposition 2.6, every biconvex subset A is either the increasing union of finite biconvex subsets or the decreasing intersection of cofinite biconvex subsets. In the former case, we set

$$\mathscr{T}(A) = \bigcup_{\substack{B \text{ finite biconvex} \\ B \subseteq A}} \mathscr{T}(B) \qquad \text{and} \qquad \mathscr{F}(A) = \bigcap_{\substack{B \text{ finite biconvex} \\ B \subseteq A}} \mathscr{F}(B).$$

Then $(\mathscr{T}(A), \mathscr{F}(A))$ is a torsion pair. The axiom (T1) is indeed easily verified. To check (T2), we take $T \in \Lambda$-mod. Each finite biconvex subset $B \subseteq A$ provides a submodule $T_B \subseteq T$ such that $(T_B, T/T_B) \in \mathscr{T}(B) \times \mathscr{F}(B)$. The monotonicity property implies that the map $B \mapsto T_B$ is non-decreasing, so for dimension reasons, the family (T_B) has an maximal element, say T_{B_0}. Then for any $B_1 \supseteq B_0$, we have $T_{B_1} = T_{B_0}$, which shows that $(T_{B_0}, T/T_{B_0}) \in \mathscr{T}(B_1) \times \mathscr{F}(B_1)$. We conclude that $(T_{B_0}, T/T_{B_0}) \in \mathscr{T}(A) \times \mathscr{F}(A)$, as desired.

When A is the intersection of cofinite biconvex subsets, we set

$$\mathscr{T}(A) = \bigcap_{\substack{B \text{ finite biconvex} \\ B \supseteq A}} \mathscr{T}(B) \qquad \text{and} \qquad \mathscr{F}(A) = \bigcup_{\substack{B \text{ finite biconvex} \\ B \supseteq A}} \mathscr{F}(B).$$

We then have a torsion pair $(\mathscr{T}(A), \mathscr{F}(A))$ associated in a monotonous way to each biconvex subset A, and for $B = \Phi_+ \setminus A$, we have $(\mathscr{T}(B), \mathscr{F}(B)) = (\mathscr{F}(A)^*, \mathscr{T}(A)^*)$.

Proposition 7.1 *Let T be a finite-dimensional Λ-module T, let P be its HN polytope, let A be a biconvex subset, and let T_A be the torsion submodule of T with respect to the torsion pair $(\mathscr{T}(A), \mathscr{F}(A))$. Then $\mu_P(A) = \underline{\dim}\, T_A$.*

Proof. Suppose first that A is cofinite. Let us choose $\theta \in C_0$ and let w be the unique element in W such that $A = A^w$ (Proposition 2.5). We have $T^w = T_{w\theta}^{\min} = T_{w\theta}^{\max}$ by Theorem 5.18. Therefore the face $P_{w\theta}$ of P is reduced to the vertex $\underline{\dim}\, T^w$. The definition in section 2.6 then leads to $\mu_P(A) = \underline{\dim}\, T^w$, which is the desired equality in our case. The case where A is finite is proved in a similar fashion, using Remark 5.19 (ii). The case of a general A then follows by monotonous approximation. □

Recall from Example 2.3 (iii) that any $\theta \in (\mathbb{R}I)^*$ defines two biconvex subsets

$$A_\theta^{\min} = \{\alpha \in \Phi_+ \mid \langle \theta, \alpha \rangle > 0\} \quad \text{and} \quad A_\theta^{\max} = \{\alpha \in \Phi_+ \mid \langle \theta, \alpha \rangle \geq 0\}.$$

Proposition 7.2 *For each $\theta \in (\mathbb{R}I)^*$, we have*

$$(\mathscr{T}(A_\theta^{\min}), \mathscr{F}(A_\theta^{\min})) = (\mathscr{I}_\theta, \overline{\mathscr{P}}_\theta) \quad \textit{and} \quad (\mathscr{T}(A_\theta^{\max}), \mathscr{F}(A_\theta^{\max})) = (\overline{\mathscr{I}}_\theta, \mathscr{P}_\theta).$$

Proof. Let $\theta \in (\mathbb{R}I)^*$.

Let T be a finite dimensional Λ-module and let P be its HN polytope. By definition, $T \in \overline{\mathscr{P}}_\theta$ if and only if $T_\theta^{\min} = 0$. By Proposition 3.2, this is equivalent to the equation $\psi_P(\theta) = 0$, which holds if and only if $\langle \theta, \mu_P(A_\theta^{\min}) \rangle = 0$ by Proposition 2.16. Since $\mu_P(A_\theta^{\min})$ is a non-negative linear combination of roots α such that $\langle \theta, \alpha \rangle > 0$ (Lemma 2.15), this equation amounts to $\mu_P(A_\theta^{\min}) = 0$. By Proposition 7.1, this holds if and only if $T \in \mathscr{F}(A_\theta^{\min})$.

Therefore $\overline{\mathscr{P}}_\theta = \mathscr{F}(A_\theta^{\min})$, whence the first equality. The second follows by $*$-duality. □

The biconvex subsets $A_\theta^{\min}$ and $A_\theta^{\max}$ only depend on the face of $\mathscr{W}$ to which θ belongs. By Proposition 7.2, the same is true for the torsion pairs $(\mathscr{I}_\theta, \overline{\mathscr{P}}_\theta)$ and $(\overline{\mathscr{I}}_\theta, \mathscr{P}_\theta)$, and also for the category $\mathscr{R}_\theta = \overline{\mathscr{I}}_\theta \cap \overline{\mathscr{P}}_\theta$. For each face F of $\mathscr{W}$, we may thus define categories $\mathscr{I}_F$, $\overline{\mathscr{I}}_F$, etc. so that $\mathscr{I}_F = \mathscr{I}_\theta$, $\overline{\mathscr{I}}_F = \overline{\mathscr{I}}_\theta$, etc. for any $\theta \in F$. Further, we denote by $\mathfrak{I}_F$, $\mathfrak{R}_F$, and $\mathfrak{P}_F$ the subsets of $\mathfrak{B}$ consisting of irreducible components whose general point belong to the categories $\mathscr{I}_F$, $\mathscr{R}_F$, and $\mathscr{P}_F$, respectively.

Lemma 7.3 *Let T be a finite dimensional Λ-module and let $A \subseteq B$ be two biconvex subsets. Denote by T_A and T_B the torsion submodules of T with respect to the torsion pairs $(\mathscr{T}(A), \mathscr{F}(A))$ and $(\mathscr{T}(B), \mathscr{F}(B))$. If $\dim T < \operatorname{ht} \alpha$ for each $\alpha \in B \setminus A$, then $T_A = T_B$.*

Proof. Let P be the HN polytope of T. By Proposition 7.1, $\underline{\dim}\, T_B/T_A = \mu_P(B) - \mu_P(A)$, whence $\operatorname{ht}(\mu_P(B) - \mu_P(A)) = \dim T_B/T_A \leq \dim T$. On the other hand, Lemma 2.15 expresses the weight $\mu_P(B) - \mu_P(A)$ as a non-negative linear combination of roots in $B \setminus A$. The last assumption in the statement forces this linear combination to be trivial, which implies that $T_B/T_A = 0$. □

Proposition 7.4 *Let A and B be two biconvex subsets and let $\alpha \in \Phi_+^{\mathrm{re}}$. Assume that $B = A \sqcup \{\alpha\}$. Then there is a rigid indecomposable Λ-module $L(A, B)$ of dimension-vector α such that $\mathscr{F}(A) \cap \mathscr{T}(B) = \operatorname{add} L(A, B)$.*

Proof. The case where A and B are both finite follows from Theorem 5.11 (i): if $A = A_w$ and $B = A_{s_i w}$, then $L(A, B) = \operatorname{Hom}_\Lambda(I_w, S_i)$. This module is rigid, because S_i is rigid and the equivalence in Theorem 5.4 (iii) preserves rigidity.

Assume now that A and B are infinite and that $\delta \notin A$. By Lemma 2.9, we can then find $A' \subseteq A$ and $B' \subseteq B$ finite biconvex subsets such that $B' = A' \sqcup \{\alpha\}$. Certainly, A' and B' are not unique subject to these requirements. We claim however that $L(A', B')$ does not depend on the choice of A' and B'.

To see this, consider A'' and B'' finite biconvex subsets with $B'' = A'' \sqcup \{\alpha\}$ and $A'' \subseteq A$. We want to show that $L(A', B') \cong L(A'', B'')$. Without loss of generality, we may assume that $A'' \supseteq A'$. We write $A' = A_u$, $B' = A_{s_i u}$, $A'' = A_{vu}$, $B'' = A_{s_j vu}$ with $\ell(vu) = \ell(v) + \ell(u)$ and $\alpha = u^{-1}\alpha_i = (vu)^{-1}\alpha_j$. Then $L(A', B') = \operatorname{Hom}_\Lambda(I_u, S_i)$ and $L(A'', B'') = \operatorname{Hom}_\Lambda(I_{vu}, S_j) = \operatorname{Hom}_\Lambda(I_u, \operatorname{Hom}_\Lambda(I_v, S_j))$. Observing that $\operatorname{Hom}_\Lambda(I_v, S_j)$ has dimension-vector $v^{-1}\alpha_j = \alpha_i$, we obtain $L(A', B') \cong L(A'', B'')$, as announced.

Certainly, $L(A', B') \in \mathscr{T}(B)$. Let A_0 be a finite biconvex subset contained in A. By Lemma 2.9, there is a finite biconvex subset $A'' \subseteq A$ that contains A_0 and such that $B'' = A'' \sqcup \{\alpha\}$ is biconvex. Then $L(A', B') \cong L(A'', B'')$ belongs to $\mathscr{F}(A'')$, hence to $\mathscr{F}(A_0)$. Since A_0 was arbitrary, we get $L(A', B') \in \mathscr{F}(A)$. It follows that $\operatorname{add} L(A', B') \subseteq \mathscr{F}(A) \cap \mathscr{T}(B)$.

Conversely, let $T \in \mathscr{F}(A) \cap \mathscr{T}(B)$. By Lemma 2.9, there is a biconvex subset $A'' \subseteq A$ that contains $\{\beta \in A \mid \operatorname{ht}\beta \leq \dim T\}$ and such that $B'' = A'' \sqcup \{\alpha\}$ is biconvex. By Lemma 7.3, we have $T \in \mathscr{T}(B'')$. Therefore T belongs to $\mathscr{F}(A'') \cap \mathscr{T}(B'') = \operatorname{add} L(A'', B'')$.

Setting $L(A, B) = L(A', B')$, we thus have $\mathscr{F}(A) \cap \mathscr{T}(B) = \operatorname{add} L(A, B)$, as desired.

It remains to deal with the case where $\delta \in A$. When A and B are both cofinite, the result follows from Theorem 5.10 (i): if $A = A^{ws_i}$ and $B = A^w$, then $L(A, B) = I_w \otimes_\Lambda S_i$. The general case follows by approximation, in a similar fashion as above. □

We now claim that for any biconvex subset A, the torsion pair $(\mathscr{T}(A), \mathscr{F}(A))$ satisfies the openness condition (O) of section 4.5. When A is finite or cofinite, or more generally when A is of the form $A_\theta^{\min}$ or $A_\theta^{\max}$, this follows from Proposition 4.2 (iii). The general case is deduced from this particular case with the help of Lemma 7.3. For instance in the case $\delta \notin A$, for each dimension-vector $\nu \in \mathbb{N}I$, we can find a finite biconvex subset $A_0 \subseteq A$ that contains $\{\alpha \in A \mid \operatorname{ht} \alpha \leq \operatorname{ht} \nu\}$; then

$$\{T \in \Lambda(\nu) \mid T \in \mathscr{T}(A)\} = \{T \in \Lambda(\nu) \mid T \in \mathscr{T}(A_0)\}$$

and

$$\{T \in \Lambda(\nu) \mid T \in \mathscr{F}(A)\} = \{T \in \Lambda(\nu) \mid T \in \mathscr{F}(A_0)\};$$

and thus condition (O) for $(\mathscr{T}(A), \mathscr{F}(A))$ in dimension-vector ν follows from the condition (O) for $(\mathscr{T}(A_0), \mathscr{F}(A_0))$.

We may thus apply the results of section 4.5: each biconvex subset A defines subsets $\mathfrak{T}(A)$ and $\mathfrak{F}(A)$ of $\mathfrak{B}$, and we have a bijection

$$\Xi(A) : \mathfrak{T}(A) \times \mathfrak{F}(A) \to \mathfrak{B}.$$

More generally if $\mathbf{A} = (A_0, \ldots, A_\ell)$ is a nondecreasing list of biconvex subsets, then we have nested torsion pairs

$$(\mathscr{T}(A_0), \mathscr{F}(A_0)) \preccurlyeq \cdots \preccurlyeq (\mathscr{T}(A_\ell), \mathscr{F}(A_\ell)),$$

whence a bijection

$$\Xi(\mathbf{A}) : \mathfrak{T}(A_0) \times \prod_{k=1}^{\ell} \Big(\mathfrak{F}(A_{k-1}) \cap \mathfrak{T}(A_k) \Big) \times \mathfrak{F}(A_\ell) \to \mathfrak{B}.$$

We define the character of a subset $\mathfrak{X} \subseteq \mathfrak{B}$ as the formal series

$$P_{\mathfrak{X}}(t) = \sum_{\nu \in \mathbb{N}I} \operatorname{Card} \mathfrak{X}(\nu)\, t^\nu$$

in $\mathbb{Z}[[\,(t^{\alpha_i})_{i \in I}\,]]$, where $\mathfrak{X}(\nu) = \mathfrak{X} \cap \mathfrak{B}(\nu)$ is the set of elements of weight ν in $\mathfrak{X}$. We denote the multiplicity of a root α by m_α; thus $m_\alpha = 1$ if α is real and $m_\alpha = r$ if α is imaginary.

Proposition 7.5 *Let $A \subseteq B$ be two biconvex subsets. Then*

$$P_{\mathfrak{F}(A) \cap \mathfrak{T}(B)} = \prod_{\alpha \in B \setminus A} \frac{1}{(1 - t^\alpha)^{m_\alpha}}. \tag{7.1}$$

Proof. Let $A \subseteq B$ be two biconvex subsets. The bijection $\Xi((A, B))$ leads to the equation

$$P_{\mathfrak{T}(A)} \times P_{\mathfrak{F}(A) \cap \mathfrak{T}(B)} \times P_{\mathfrak{F}(B)} = P_{\mathfrak{B}}.$$

Since $\mathfrak{B}$ indexes a basis for $U(\mathfrak{n}_+)$, the series $P_{\mathfrak{B}}$ is given by the Kostant partition function

$$P_{\mathfrak{B}}(t) = \prod_{\alpha \in \Phi_+} \frac{1}{(1 - t^\alpha)^{m_\alpha}}. \tag{7.2}$$

For any subset $S \subseteq \Phi_+$ consider the formal power series

$$Q_S = \prod_{\alpha \in S} \frac{1}{(1 - t^\alpha)^{m_\alpha}},$$

and notice that a monomial t^ν can only occur in Q_S if ν belongs to the convex cone spanned by S. Our aim is to show that $P_{\mathfrak{F}(A) \cap \mathfrak{T}(B)} = Q_{B \setminus A}$.

To simplify the notation, let us set $Q' = Q_A$, $Q'' = Q_{B \setminus A}$, $Q''' = Q_{\Phi_+ \setminus B}$, $R' = P_{\mathfrak{T}(A)}$, $R'' = P_{\mathfrak{F}(A) \cap \mathfrak{T}(B)}$ and $R''' = P_{\mathfrak{F}(B)}$. We then have

$$Q'Q''Q''' = R'R''R'''$$

because each side is equal to $P_{\mathfrak{B}}$.

Given a formal power series

$$P = \sum_{\nu \in \mathbb{N}I} a_\nu t^\nu$$

and a nonnegative integer n, we define

$$P_n = \sum_{\substack{\nu \in \mathbb{N}I \\ \operatorname{ht} \nu = n}} a_\nu t^\nu \qquad \text{and} \qquad P_{\leq n} = P_0 + P_1 + \cdots + P_n.$$

We now show the equations

$$Q'_{\leq n} = R'_{\leq n}, \qquad Q''_{\leq n} = R''_{\leq n}, \qquad Q'''_{\leq n} = R'''_{\leq n} \tag{E_n}$$

by induction on n.

Certainly, (E_0) holds, because all the series involved have constant term equal to 1. Assume that (E_{n-1}) holds. Then

$$Q'Q''Q''' - Q'_{\leq n-1}Q''_{\leq n-1}Q'''_{\leq n-1} = R'R''R''' - R'_{\leq n-1}R''_{\leq n-1}R'''_{\leq n-1}.$$

After appropriate truncation, this leads to

$$Q'_n + Q''_n + Q'''_n = R'_n + R''_n + R'''_n. \tag{7.3}$$

Now consider a monomial t^ν that occurs in the right-hand side of (7.3). If it appears in R'_n, then there is a Λ-module in $\mathscr{T}(A)$ of dimension-vector ν, and it follows from Lemma 2.15 and Proposition 7.1 that ν belong to the convex cone spanned by A. By Lemma 2.11, ν does not belong to the convex cones spanned by $B \setminus A$ or by $\Phi_+ \setminus B$, hence t^ν cannot occur in Q'' nor in Q'''. Therefore t^ν can only appear in Q'_n. Likewise, we see that if t^ν appears in R''_n (respectively, R'''_n), then it can only appear in Q''_n (respectively, Q'''_n). We conclude that equation (7.3) splits into the three equations $Q'_n = R'_n$, $Q''_n = R''_n$ and $Q'''_n = R'''_n$, and thus that (E_n) holds.

Therefore (E_n) holds for each $n \in \mathbb{N}$, whence $Q'' = R''$. $\square$

In view of its later use, the particular case $(A, B) = (A_\theta^{\min}, A_\theta^{\max})$ with $\theta \in \mathfrak{t}$ deserves a special mention.

Corollary 7.6 *Let $\theta \in \mathfrak{t}$. Then*

$$P_{\mathfrak{R}_\theta} = \left(\prod_{\substack{\alpha \in \Phi_+^{\mathrm{re}} \\ \langle \theta, \alpha \rangle = 0}} \frac{1}{1 - t^\alpha} \right) \left(\prod_{n \geq 1} \frac{1}{1 - t^{n\delta}} \right)^r .$$

7.2 Simple regular modules

We come back to the description of the abelian categories $\mathscr{R}_F$. Our aim in this section is to get information on their simple objects.

We begin with a general remark: let F and G be two faces such that $F \subseteq \overline{G}$. If we pick $\theta \in F$ and $\eta \in G$, then

$$A_\theta^{\min} \subseteq A_\eta^{\min} \subseteq A_\eta^{\max} \subseteq A_\theta^{\max},$$

hence

$$(\mathscr{I}_\theta, \overline{\mathscr{P}}_\theta) \preccurlyeq (\mathscr{I}_\eta, \overline{\mathscr{P}}_\eta) \preccurlyeq (\overline{\mathscr{I}}_\eta, \mathscr{P}_\eta) \preccurlyeq (\overline{\mathscr{I}}_\theta, \mathscr{P}_\theta);$$

in other words,

$$(\mathscr{I}_F, \overline{\mathscr{P}}_F) \preccurlyeq (\mathscr{I}_G, \overline{\mathscr{P}}_G) \preccurlyeq (\overline{\mathscr{I}}_G, \mathscr{P}_G) \preccurlyeq (\overline{\mathscr{I}}_F, \mathscr{P}_F).$$

Therefore

$$\mathscr{I}_F \subseteq \mathscr{I}_G, \quad \mathscr{P}_F \subseteq \mathscr{P}_G \quad \text{and} \quad \mathscr{R}_F \supseteq \mathscr{R}_G, \tag{7.4}$$

and for any Λ-module T, we have a filtration $0 \subseteq T_\theta^{\min} \subseteq T_\eta^{\min} \subseteq T_\eta^{\max} \subseteq T_\theta^{\max} \subseteq T$. The three subquotients

$$T_\eta^{\min}/T_\theta^{\min} \in \overline{\mathscr{P}}_\theta \cap \mathscr{I}_\eta, \quad T_\eta^{\max}/T_\eta^{\min} \in \overline{\mathscr{P}}_\eta \cap \overline{\mathscr{I}}_\eta \quad \text{and} \quad T_\theta^{\max}/T_\eta^{\max} \in \mathscr{P}_\eta \cap \overline{\mathscr{I}}_\theta$$

all belong to $\mathscr{R}_\theta$; in particular, a simple object of $\mathscr{R}_F$ belongs either to $\mathscr{I}_G$, $\mathscr{R}_G$ or $\mathscr{P}_G$.

We denote by $\operatorname{Irr}\mathscr{R}_F$ the set of simple objects in $\mathscr{R}_F$. Recall that two objects T and U in $\operatorname{Irr}\mathscr{R}_F$ are said to be linked if there is a finite sequence $T = X_0, X_1, \ldots, X_n = U$ of objects in $\operatorname{Irr}\mathscr{R}_F$ such that $\operatorname{Ext}^1_\Lambda(X_{k-1}, X_k) \neq 0$ for each $k \in \{1, \ldots, n\}$. (Note here that the groups Ext^1 are the same computed in Λ-mod and in $\mathscr{R}_F$, for the latter is closed under extensions, and that $\operatorname{Ext}^1_\Lambda(X,Y)$ and $\operatorname{Ext}^1_\Lambda(Y,X)$ are K-dual to each other.) The linkage relation is an equivalence relation. Finally, recall the map $\iota : \Phi^s \to \Phi_+^{\mathrm{re}}$ constructed in section 2.3 as a right inverse to the projection π.

Theorem 7.7 *Let F be a face of the spherical Weyl fan.*

(i) If the dimension-vector of a simple object $T \in \mathscr{R}_F$ is a multiple of δ, then $\{T\}$ is a linkage class in $\operatorname{Irr}\mathscr{R}_F$, and T belongs to $\mathscr{R}_C$ for each spherical Weyl chamber C such that $F \subseteq \overline{C}$.

(ii) The other objects in $\operatorname{Irr}\mathscr{R}_F$ are rigid. Their dimension-vectors belong to $\iota(\Phi^s)$. Given $\alpha \in \iota(\Phi^s)$, there is at most one simple object in $\mathscr{R}_F$ of dimension-vector α, up to isomorphism.

Proof. Let T be as in (i). For any $X \in \operatorname{Irr}\mathscr{R}_F$ different from T, we have $\operatorname{Hom}_\Lambda(T,X) = \operatorname{Hom}_\Lambda(X,T) = 0$, by Schur's lemma applied in the category $\mathscr{R}_F$, so $\operatorname{Ext}^1_\Lambda(T,X) = 0$ by Crawley-Boevey's formula (4.2). Therefore $\{T\}$ is a linkage class. Let C be a spherical Weyl chamber such that $F \subseteq \overline{C}$. The assumption on $\underline{\dim}\,T$ rules out the possibility that $T \in \mathscr{I}_C$ or $\mathscr{P}_C$. We conclude that $T \in \mathscr{R}_C$. Thus assertion (i) is true.

We now turn to (ii). Let $T \in \operatorname{Irr}\mathscr{R}_F$, whose dimension-vector α is not a multiple of δ. Then (α,α) is a positive even integer. By Schur's lemma, the endomorphism algebra of T has dimension 1. Using Crawley-Boevey's formula (4.2), we then see that $(\alpha,\alpha) = 2$ and $\dim \operatorname{Ext}^1_\Lambda(T,T) = 0$. Thus T is a rigid Λ-module, and, by Proposition 5.10 in [34], α is a real root.

In this context, assume that $\alpha - \delta$ is a positive root. Then, by Corollary 7.6, $\Lambda(\alpha - \delta)$ has an irreducible component whose general point belongs to $\mathscr{R}_F$. In particular, there exists $X \in \mathscr{R}_F$ of dimension-vector $\alpha - \delta$. But then $(\underline{\dim}\,X, \underline{\dim}\,T) = 2$, and (4.2) gives that $\operatorname{Hom}_\Lambda(T,X)$ or $\operatorname{Hom}_\Lambda(X,T)$ is nonzero. Since $\dim T > \dim X$, this forbids T to be simple in $\mathscr{R}_F$, which contradicts our choice of T. Therefore $\alpha - \delta \notin \Phi_+$, which means that $\alpha \in \iota(\Phi^s)$.

Lastly, let T' and T'' be two simple objects in $\mathscr{R}_F$ with the same dimension-vector $\alpha \in \iota(\Phi^s)$. Since $(\underline{\dim}\, T', \underline{\dim}\, T'') = 2$, (4.2) gives that $\operatorname{Hom}_\Lambda(T', T'')$ or $\operatorname{Hom}_\Lambda(T'', T')$ is nonzero. By Schur's lemma, T' and T'' are isomorphic. $\square$

One can show that the simple objects of $\mathscr{R}_F$ described in Theorem 7.7 (i) always have dimension-vector δ (see Corollary 7.21).

7.3 The type $\widetilde{A}_1$

Let Π be the completed preprojective algebra of the Kronecker quiver, as in section 6. We identify the Grothendieck group of Π-mod with $\mathbb{Z}^2$ by writing the dimension-vector of a Π-module V as the pair $(\dim V_0, \dim V_1)$.

$$\begin{tikzcd} 0 \arrow[r, bend left=40, dotted, "\alpha"] \arrow[r, bend left=20, "\beta"'] & 1 \arrow[l, bend left=20, dotted, "\overline{\alpha}"'] \arrow[l, bend left=40, "\overline{\beta}"] \end{tikzcd}$$

As in section 6.4, we denote by $\Pi(\mu)$ the nilpotent variety of type $\widetilde{A}_1$ for a given dimension-vector $\mu = (\mu_0, \mu_1)$. A point in $\Pi(\mu)$ is a 4-tuple of matrices $T = (T_\alpha, T_\beta, T_{\overline{\alpha}}, T_{\overline{\beta}})$ which satisfy the equations $T_\alpha T_{\overline{\alpha}} + T_\beta T_{\overline{\beta}} = 0$ and $T_{\overline{\alpha}} T_\alpha + T_{\overline{\beta}} T_\beta = 0$ and the nilpotency condition.

We denote the root system of type $\widetilde{A}_1$ by Δ, so $\Delta_+ = \Delta_+^{\mathrm{re}} \sqcup (\mathbb{Z}_{>0}\delta)$, where $\delta = (1,1)$ is the primitive imaginary root and

$$\Delta_+^{\mathrm{re}} = \{(1,0) + n\delta, (0,1) + n\delta, \mid n \in \mathbb{N}\}$$

(Note that we use the same letter δ to denote the primitive imaginary root in both Φ_+ and Δ_+; this will not lead to confusion.)

There are two opposite spherical chamber coweights, namely

$$\gamma'(\mu_0, \mu_1) = \mu_0 - \mu_1 \quad \text{and} \quad \gamma''(\mu_0, \mu_1) = \mu_1 - \mu_0.$$

The spherical Weyl fan has three faces, namely $\{0\}$ and the two chambers $\mathbb{R}_{>0}\gamma'$ and $\mathbb{R}_{>0}\gamma''$.

Given $n \in \mathbb{N}$, we denote by $\Pi(n\delta)^\times$ the open subset of all points $T \in \Pi(n\delta)$ such that the $n \times n$ matrix T_α is invertible. Obviously, the dimension-vector μ of a submodule of T satisfies $\mu_0 \le \mu_1$, so T is γ'-semistable, and is even γ'-stable if $n = 1$.

We now describe $\Pi(n\delta)^\times$ with the help of an auxiliary variety. Let

$$Z_n = \{(X, Y) \in M_n(K)^2 \mid X \text{ is nilpotent, } XY = YX\},$$

where $M_n(K)$ denotes the algebra of $n \times n$ matrices over K. The group $\mathrm{GL}_n(K)$ acts by conjugation on Z_n. Given a partition λ of size n, let us denote by $\mathscr{O}_\lambda \subseteq M_n(K)$ the adjoint orbit of nilpotent matrices of Jordan type λ. The following lemma is due to I. Frenkel and Savage (it is a particular case of [21], Proposition 2.9).

Lemma 7.8 *(i) The map $f : \Pi(n\delta)^\times \to Z_n$ defined by $f(T) = (T_{\overline{\beta}}T_\alpha, T_\alpha^{-1}T_\beta)$ is a principal $\mathrm{GL}_n(K)$-bundle.*

(ii) The first projection $Z_n \to M_n(K)$ identifies Z_n with the disjoint union of the conormal bundles $T^*_{\mathscr{O}_\lambda} M_n(K)$, where λ runs over the set of all partitions of n.

Proof. We begin with (i). Let $T \in \Pi(n\delta)^\times$. By definition, $T_{\overline{\beta}}T_\alpha$ is nilpotent. In addition,

$$(T_{\overline{\beta}}T_\alpha)(T_\alpha^{-1}T_\beta) = -T_{\overline{\alpha}}T_\alpha = -T_\alpha^{-1}(T_\alpha T_{\overline{\alpha}})T_\alpha = (T_\alpha^{-1}T_\beta)(T_{\overline{\beta}}T_\alpha),$$

thanks to the preprojective equations. So f is well defined. Now $\mathrm{GL}_n(K)$ acts on $\Pi(n\delta)^\times$ by

$$U \cdot (T_\alpha, T_\beta, T_{\overline{\alpha}}, T_{\overline{\beta}}) = (UT_\alpha, UT_\beta, T_{\overline{\alpha}}U^{-1}, T_{\overline{\beta}}U^{-1}).$$

This action is free, for T_α is invertible, and the orbits of this action are the fibers of f. Thus (i) holds true.

For a matrix $X \in M_n(K)$, the tangent space to its adjoint orbit identifies with the subspace $\{[W, X] \mid W \in M_n(K)\}$. Under the standard trace duality, the orthogonal of this subspace identifies with the Lie algebra centralizer of X because for any $Y \in M_n(K)$,

$$\begin{aligned}[X, Y] = 0 &\iff \Big(\forall W \in M_n(K),\ \mathrm{Tr}\big(W[X, Y]\big) = 0\Big)\\ &\iff \Big(\forall W \in M_n(K),\ \mathrm{Tr}\big([W, X]Y\big) = 0\Big).\end{aligned}$$

Assertion (ii) follows. □

Thanks to Lemma 7.8, we see that the irreducible components of $\Pi(n\delta)^\times$ are of the form $f^{-1}\big(\overline{T^*_{\mathscr{O}_\lambda} M_n(K)}\big)$. We denote by $I(\lambda)$ the closure in $\Pi(n\delta)$ of this set; since $\Pi(n\delta)^\times$ is open, this is an irreducible component of $\Pi(n\delta)$, whose general point is γ'-semistable. When λ has just one nonzero part, here n, we write simply $I(n)$ instead of $I((n))$.

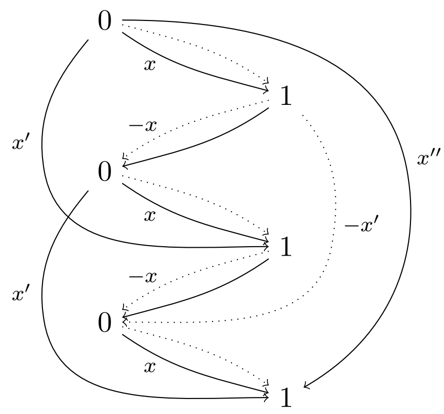

Figure 4: The Π-module afforded by a general point T of $I(3)$. On the picture, a digit 0 or 1 represents a basis vector of the corresponding degree, each dotted arrow indicates a nonzero entry (equal to 1 unless otherwise indicated) in the matrix that represents T_α or $T_{\overline{\alpha}}$, and each plain arrow indicates a nonzero entry in the matrix that represents T_β or $T_{\overline{\beta}}$. Things have been arranged so that $T_{\overline{\beta}}T_\alpha = \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{pmatrix}$ is a Jordan block of size 3 and $T_\alpha^{-1}T_\beta = \begin{pmatrix} x & x' & x'' \\ 0 & x & x' \\ 0 & 0 & x \end{pmatrix}$ is a point in the commutant of $T_{\overline{\beta}}T_\alpha$. The parameters x, x' and x'' can take any value in K.

We denote the set of partitions by $\mathcal{P}$. Recalling the well-known formula

$$\sum_{\lambda \in \mathcal{P}} t^{|\lambda|} = \prod_{n \geq 1} \frac{1}{1 - t^n}$$

and applying Corollary 7.6 to Π and γ', we see that $\{I(\lambda) \mid \lambda \in \mathcal{P}\}$ is the full set of elements in $\mathfrak{R}_{\gamma'}$.

Recall Crawley-Boevey and Schröer's theory of the canonical decomposition explained in section 4.4.

Proposition 7.9 *Let m and n be positive integers and let λ be a partition.*

(i) Any general point T in $I(n)$ is an indecomposable Π*-module.*

(ii) We have $\hom_\Pi(I(m), I(n)) = \operatorname{ext}^1_\Pi(I(m), I(n)) = 0$.

(iii) Writing $\lambda = (\lambda_1, \ldots, \lambda_\ell)$, where ℓ is the number of nonzero parts of λ, the canonical decomposition of $I(\lambda)$ is

$$I(\lambda) = \overline{I(\lambda_1) \oplus \cdots \oplus I(\lambda_\ell)}.$$

(iv) The component $I(\lambda)^*$ is obtained from $I(\lambda)$ by applying the automorphism of Π that exchanges the idempotents e_0 and e_1, the arrows α and $\overline{\alpha}$, and the arrows β and $\overline{\beta}$.

(v) For any general point T in $I(\lambda)$, we have $\dim \operatorname{End}_\Pi(T) = |\lambda|$.

Proof. Let T be a general point in $I(n)$. It provides a Π-module with dimension-vector $n\delta$. If this module were indecomposable, then the matrix $T_{\overline{\beta}}T_\alpha$ would represent an endomorphism of the vector space K^n which stabilizes two complementary subspaces. This is however impossible, since $T_{\overline{\beta}}T_\alpha$ is a nilpotent matrix of Jordan type (n). This shows (i).

The group $G(\delta)$ acts on $I(1)$ with infinitely many orbits, and a general point in $I(1)$ is a γ'-stable Π-module. So if S' and S'' are points in two different orbits, then they are non-isomorphic simple objects in the category $\mathscr{R}_{\gamma'}$, whence $\operatorname{Hom}_\Pi(S', S'') = \operatorname{Hom}_\Pi(S'', S') = 0$ by Schur's lemma. Now let (T', T'') be general in $I(m) \times I(n)$. Then there is a general point $(S', S'') \in I(1)^2$ such that T' is an m-th iterated extension of S' and T'' is an n-th iterated extension of S''. Therefore $\operatorname{Hom}_\Pi(T', T'') = \operatorname{Hom}_\Pi(T'', T') = 0$, and using Crawley-Boevey's formula (4.2), we get $\operatorname{Ext}^1_\Pi(T', T'') = 0$. Item (ii) is proved.

From (ii) and from Crawley-Boevey and Schröer's theory, it follows that for any partition $\lambda = (\lambda_1, \ldots, \lambda_\ell)$, the closure $Z(\lambda) = \overline{I(\lambda_1) \oplus \cdots \oplus I(\lambda_\ell)}$ is an element in $\mathfrak{R}_{\gamma'}$. There thus exists a partition μ such that $Z(\lambda) = I(\mu)$. Let T be a general point in this irreducible component. Looking at the Jordan type of the nilpotent matrix of $T_{\overline{\beta}}T_\alpha$, we conclude that $\lambda = \mu$, whence (iii).

In order to establish (iv), it suffices by (iii) to consider the case $\lambda = (n)$. Certainly, $I(n)^*$ belongs to $\mathfrak{R}_{\gamma''}$, by Remark 4.1. Applying to it the automorphism defined in the statement, we obtain again an element of $\mathfrak{R}_{\gamma'}$, which we can write $I(\mu)$ for a certain partition μ. Now the general point of this component must be indecomposable, by (i), so μ has only one nonzero part. Looking at the dimension-vector, we conclude that $\mu = n$, as desired. Item (iv) is proved.

Lastly, let T be a general point in $I(n)$. The commutator of the matrix $X = T_{\overline{\beta}}T_\alpha$ in $M_n(K)$ is the algebra of polynomials on X, because X is a nilpotent matrix of Jordan type (n), so is a commutative algebra of dimension n. Therefore the stabilizer of $f(T)$ in $\mathrm{GL}_n(K)$ coincides with the group of invertible elements in this algebra; in particular, its dimension is equal to n. The stabilizer of T in $\Pi(n\delta)$ is isomorphic to this group, and is the group of invertible elements in the endomorphism algebra $\operatorname{End}_\Pi(T)$. We conclude that $\dim \operatorname{End}_\Pi(T) = n$. This shows (v) in the particular case where λ has just one nonzero part; the general case then follows from (ii) and (iii). $\square$

Assertion (iv) in this proposition implies that $\{I(\lambda)^* \mid \lambda \in \mathcal{P}\}$ is the full set of elements in $\mathfrak{R}_{\gamma''}$.

7.4 Cores (proofs of Theorems 1.1 and 1.2)

Let $\gamma \in \Gamma$ be a spherical chamber coweight. By equation (7.4), we have

$$\bigcap_{\substack{F \text{ face} \\ \gamma \in \overline{F}}} \mathscr{R}_F = \bigcap_{\substack{C \text{ Weyl chamber} \\ \gamma \in \overline{C}}} \mathscr{R}_C.$$

Objects in this intersection are called γ-cores. They form an abelian, closed under extensions, subcategory of Λ-mod. The dimension-vector of a γ-core is a multiple of δ.

The following proposition provides an alternative definition of γ-cores.

Proposition 7.10 *A object in $\mathscr{R}_\gamma$ is a γ-core if and only if the dimension-vectors of all its Jordan-Hölder components are multiples of δ.*

Proof. By Theorem 7.7 (i), a simple $\mathscr{R}_\gamma$-module whose dimension-vector is a multiple of δ is necessarily a γ-core. The sufficiency of the condition then follows from the fact that the category of γ-cores is closed under extensions.

Conversely, let $T \in \mathscr{R}_\gamma$ be such that the dimension-vector of at least one Jordan-Hölder component of T is not a multiple of δ. We want to show that T is not a γ-core. It suffices to show that T has a direct summand which is not a γ-core, so without loss of generality, we can assume that T is indecomposable. Then all the Jordan-Hölder components of T, regarded as an object of $\mathscr{R}_\gamma$, belong to the same linkage class; by Theorem 7.7, none of these components have a dimension-vector multiple of δ. Now consider a simple object X of $\mathscr{R}_\gamma$ contained in T. Since X belongs to $\mathscr{R}_\gamma$, its dimension-vector $\underline{\dim}\, X$ is in the kernel of γ, without being a multiple of δ, so there exists θ near γ, in a Weyl chamber, such that $\langle \theta, \underline{\dim}\, X \rangle > 0$. It follows that T is not in $\mathscr{R}_\theta$, and therefore is not a γ-core. This proves the necessity of the condition given in the lemma. □

Obviously, the $*$-dual of a γ-core is a $(-\gamma)$-core. More interesting is the following compatibility between cores and reflection functors.

Proposition 7.11 *Let $\gamma \in \Gamma$, let $i \in I$, and let T be a γ-core. If $\langle \gamma, \alpha_i \rangle < 0$, then $T \in \mathscr{T}^{s_i}$ and $\Sigma_i T$ is a $s_i\gamma$-core. If $\langle \gamma, \alpha_i \rangle > 0$, then $T \in \mathscr{F}_{s_i}$ and $\Sigma_i^* T$ is a $s_i\gamma$-core. If $\langle \gamma, \alpha_i \rangle = 0$, then $T \cong \Sigma_i T \cong \Sigma_i^* T$.*

Proof. The first two claims immediately follow from Theorem 5.17, so let us consider the case where $\langle \gamma, \alpha_i \rangle = 0$. Then there exists $\theta \in \mathfrak{t}$ close to γ such that $\langle \theta, \alpha_i \rangle > 0$, which forbids S_i to appear as a submodule of T, and there exists $\eta \in \mathfrak{t}$ close to γ such that $\langle \eta, \alpha_i \rangle < 0$, which forbids S_i to appear as a quotient of T. Therefore the i-socle and the i-head of T are both trivial. Now recall the diagram (4.1). We rewrite it as

$$T_i \xrightarrow{T_{\mathrm{out}(i)}} \widetilde{T}_i \xrightarrow{T_{\mathrm{in}(i)}} T_i.$$

This is a complex, $T_{\mathrm{out}(i)}$ is injective, and $T_{\mathrm{in}(i)}$ is surjective. We have $(\alpha_i, \underline{\dim}\, T) = 0$, for $\underline{\dim}\, T$ is a multiple of δ, so the dimension of $\widetilde{T}_i$ is twice the dimension of T_i. Our complex is therefore acyclic. The isomorphisms $T \cong \Sigma_i T \cong \Sigma_i^* T$ then follow from Proposition 5.1. $\square$

We will produce γ-cores by means of Hall functors. According to section 6, we need to construct pairs (S, R) of Λ-modules that satisfy (6.1). We proceed as follows.

We call a pair (C, F) a flag, if it consists of a spherical Weyl chamber C and a facet F contained in the closure of C. Such a pair determines a spherical chamber coweight $\gamma_{C/F}$ by the equation $C = F + \mathbb{R}_{>0}\, \gamma_{C/F}$. The spherical Weyl group W_0 acts on the set of flags.

Take a flag (C, F) and pick $\theta \in F$. Then $\Phi \cap (\ker \theta)$ is an affine root system of type $\widetilde{A}_1$ and $\Phi^s \cap (\ker \theta)$ is a root system of type A_1, consisting of two opposite roots $\pm\alpha$. By Theorem 7.7, the dimension-vectors of the simple objects in $\mathscr{R}_F$ belong to $\{\iota(\alpha), \iota(-\alpha)\} \cup \big(\mathbb{Z}_{>0}\delta\big)$. This implies that an object in $\mathscr{R}_F$ of dimension-vector $\iota(\pm\alpha)$ is necessarily simple.

Now Corollary 7.6 guarantees the existence of a whole irreducible component of $\Lambda(\iota(\pm\alpha))$ whose general point is in $\mathscr{R}_F$. Therefore there are objects in $\mathscr{R}_F$ of dimension-vector $\iota(\pm\alpha)$. These objects are necessarily simple, and therefore unique up to isomorphism, by Theorem 7.7 (ii). We denote them by $S_{C,F}$ and $R_{C,F}$, the labels being adjusted so that

$$\langle \gamma_{C/F}, \underline{\dim}\, S_{C,F} \rangle = 1 \quad \text{and} \quad \langle \gamma_{C/F}, \underline{\dim}\, R_{C,F} \rangle = -1. \tag{7.5}$$

Examples 7.12. (i) One fashion to produce $S_{C,F}$ and $R_{C,F}$ is to use Proposition 7.4: if we set $A = A_\theta^{\min} = \{\alpha \in \Phi_+ \mid \langle \theta, \alpha \rangle > 0\}$, $B' = A \sqcup \{\iota(\alpha)\}$ and $B'' = A \sqcup \{\iota(-\alpha)\}$, then both modules $L(A, B')$ and $L(A, B'')$ belong to $\mathscr{R}_\theta$ and have the correct dimension-vectors to be $S_{C,F}$ and $R_{C,F}$. Alternatively, we can consider $L(A', B)$ and $L(A'', B)$, where $B = A_\theta^{\max} = \{\alpha \in \Phi_+ \mid \langle \theta, \alpha \rangle \geq 0\}$, $A' = B \setminus \{\iota(\alpha)\}$ and $A'' = B \setminus \{\iota(-\alpha)\}$.

(ii) A facet F separates two chambers, say C' and C''. We then have $(S_{C',F}, R_{C',F}) = (R_{C'',F}, S_{C'',F})$.

(iii) If (C, F) is a flag, then $(-C, -F)$ is also a flag, and we have $(S_{-C,-F}, R_{-C,-F}) = ((R_{C,F})^*, (S_{C,F})^*)$.

(iv) Let us fix an extending vertex $0 \in I$, as explained at the end of section 2.3. Then each $i \in I_0$ provides a flag $(C_0^s, F_{\{i\}})$, where

$$F_{\{i\}} = \{\theta \in \mathfrak{t} \mid \langle \theta, \alpha_i \rangle = 0,\ \langle \theta, \alpha_j \rangle > 0 \text{ for all } j \in I_0 \setminus \{i\}\}.$$

We have $\gamma_{C_0^s/F_{\{i\}}} = \varpi_i$ and $\{\iota(\pm\alpha)\} = \{\alpha_i, \delta - \alpha_i\}$. The dimension-vector of $S_{C_0^s, F_{\{i\}}}$ is α_i, hence $S_{C_0^s, F_{\{i\}}} = S_i$. We define $R_i = R_{C_0^s, F_{\{i\}}}$. Since R_i is in $\mathscr{R}_{F_{\{i\}}}$, it cannot contain any submodule isomorphic to S_j, with $j \in I_0 \setminus \{i\}$, and since it is simple, it cannot contain any submodule isomorphic to S_i either. Therefore R_i has dimension-vector $\delta - \alpha_i$ and its socle is S_0. By Lemma 2 (2) of [17], these two conditions characterize R_i.

(v) Keeping our extending vertex $0 \in I$, let $(i, w) \in I_0 \times W_0$ be such that $\ell(ws_i) > \ell(w)$ and consider $(C, F) = (wC_0^s, wF_{\{i\}})$. By Theorem 5.18 (i), we have equivalences of categories

$$\mathscr{R}_{F_{\{i\}}} \underset{\operatorname{Hom}_\Lambda(I_w, ?)}{\overset{I_w \otimes_\Lambda ?}{\rightleftarrows}} \mathscr{R}_{wF_{\{i\}}},$$

which carries (S_i, R_i) to $(S_{C,F}, R_{C,F})$.

Lemma 7.13 *The modules $S_{C,F}$ and $R_{C,F}$ satisfy the conditions* (6.1).

Proof. The modules $S_{C,F}$ and $R_{C,F}$ are simple objects in $\mathscr{R}_F$, so by Schur's lemma, they are orthogonal bricks:

$$\operatorname{End}_\Lambda(S_{C,F}) = \operatorname{End}_\Lambda(R_{C,F}) = K, \qquad \operatorname{Hom}_\Lambda(S_{C,F}, R_{C,F}) = \operatorname{Hom}_\Lambda(R_{C,F}, S_{C,F}) = 0.$$

The remaining equations

$$\operatorname{Ext}^1_\Lambda(S_{C,F}, S_{C,F}) = \operatorname{Ext}^1_\Lambda(R_{C,F}, R_{C,F}) = 0, \qquad \dim \operatorname{Ext}^1_\Lambda(S_{C,F}, R_{C,F}) = 2$$

follow from Crawley-Boevey's formula (4.2). □

We can thus apply the results of section 6 to the modules $S_{C,F}$ and $R_{C,F}$. We get a Hall functor $\mathscr{H}_{C,F} : \Pi\text{-mod} \to \Lambda\text{-mod}$, which is an equivalence of categories between Π-mod and the subcategory $\langle S_{C,F}, R_{C,F} \rangle$ of $\mathscr{R}_F$.

We denote by $\mathfrak{A} = \bigsqcup_{\mathbf{d} \in \mathbb{N}^2} \operatorname{Irr} \Pi(\mathbf{d})$ the analog of the crystal $\mathfrak{B}$, but for the Kronecker quiver. Proposition 6.4 claims that $\mathscr{H}_{C,F}$ induces a injection $\mathfrak{H}_{C,F} : \mathfrak{A} \to \mathfrak{B}$, whose image consists of those components whose general points lie in $\langle S_{C,F}, R_{C,F} \rangle$.

For λ a partition, recall the irreducible component $I(\lambda) \in \mathfrak{A}$ defined in section 7.3. We will see that the general point of $\mathfrak{H}_{C,F}(I(\lambda))$ is a $\gamma_{C,F}$-core, and moreover that $\mathfrak{H}_{C,F}(I(\lambda))$ depends only on $\gamma_{C,F}$ and on λ.

To establish these results, our strategy is to first focus on the particular case described in Example 7.12 (iv). So now let us fix an extending vertex $0 \in I$, let us take $i \in I_0$, and let us set $I(\varpi_i, 1) = \mathfrak{H}_{C_0^s, F_{\{i\}}}(I(1))$. As we saw earlier in this section, there is a unique irreducible component of $\Lambda(\delta - \alpha_i)$ whose general point is in $\mathscr{R}_{F_{\{i\}}}$. Further, this component is the closure in $\Lambda(\delta - \alpha_i)$ of the set of all modules isomorphic to R_i. Denoting it by Z_i, we thus have $I(\varpi_i, 1) = \tilde{e}_i Z_i$. The following proposition is essentially a reformulation of [17], Theorem 2.

Proposition 7.14 *(i) One has* $\mathfrak{R}_{C_0^s}(\delta) = \{I(\varpi_i, 1) \mid i \in I_0\}$.

(ii) For each $i \in I_0$*, a general point in* $I(\varpi_i, 1)$ *has socle* S_0 *and head* S_i*.*

(iii) For each $i \in I_0$*, a general point in* $I(\varpi_i, 1)$ *is a* ϖ_i*-core. Conversely, any* ϖ_i*-core of dimension-vector* δ *belongs to* $I(\varpi_i, 1)$*.*

Proof. We first note that the socle of a module $T \in \mathscr{R}_{C_0^s}$ is necessarily concentrated at the vertex 0. If in addition $\underline{\dim}\, T = \delta$, then $\operatorname{soc} T \cong S_0$.

Let $\Lambda_0 = \{T \in \Lambda(\delta) \mid T \in \mathscr{R}_{C_0^s}\}$, an open subset of $\Lambda(\delta)$ by Proposition 4.2 (iii). Then $\mathfrak{R}_{C_0^s}(\delta)$ can be identified with the set of irreducible components of Λ_0.

Let $T \in \Lambda_0$ and let S_i in the head of T. There is then a surjective morphism $T \to S_i$. Its kernel has socle S_0 and dimension-vector $\delta - \alpha_i$, so is isomorphic to R_i. We thus have a short exact sequence $0 \to R_i \to T \to S_i \to 0$, which shows that T belongs to $\tilde{e}_i Z_i = I(\varpi_i, 1)$.

Therefore Λ_0 is covered by the irreducible components $I(\varpi_i, 1)$. Now Corollary 7.6 asserts that Λ_0 has r irreducible components. Assertion (i) follows.

If a point in Λ_0 has two different simple modules S_i and S_j in its head, then it belongs to both $I(\varpi_i, 1)$ and $I(\varpi_j, 1)$, and therefore is not general. Therefore the head of a general point in $I(\varpi_i, 1)$ is concentrated at the vertex i. Further, R_i and S_i are non-isomorphic simple objects in $\mathscr{R}_{F_{\{i\}}}$, therefore $\operatorname{Hom}_\Lambda(R_i, S_i) = 0$, and so the i-head of R_i is trivial. It follows that the i-head of a general point T in $I(\varpi_i, 1)$ is one-dimensional. All this shows (ii).

Now fix $i \in I_0$. Let T be a general point in $I(\varpi_i, 1)$. Then $\langle \varpi_i, \underline{\dim}\, T\rangle = 0$ and (ii) implies that $\langle \varpi_i, \underline{\dim}\, X\rangle < 0$ for any proper submodule $X \subseteq T$. The module T is thus a simple object in $\mathscr{R}_{\varpi_i}$. By Theorem 7.7 (i), T is a ϖ_i-core.

Conversely, if T is a ϖ_i-core of dimension-vector δ, then its socle must be S_0 and only S_i can show up in the head of T. The reasoning used to prove (i) applies anew, and we conclude that T belongs to $I(\varpi_i, 1)$. Assertion (iii) is proved. □

Remark 7.15. With the notation of the proof, a point T in $\Lambda_0 \cap I(\varpi_i, 1)$ is always the middle term of a non-split extension $0 \to R_i \to T \to S_i \to 0$. Using the fact that $\operatorname{End}_\Lambda(S_i) = \operatorname{End}_\Lambda(R_i) = K$, one shows that the datum of the isomorphism class of T is equivalent to the datum of the class of the extension, up to scalar. In other words, $G(\delta)$-orbits in $\Lambda_0 \cap I(\varpi_i, 1)$ are in bijection with points in the projective line $\mathbb{P}(\operatorname{Ext}^1_\Lambda(S_i, R_i))$. In [17], Crawley-Boevey shows that in the moduli space $\Lambda_0/\!/G(\delta)$ (the quotient here should be understood in the GIT sense w.r.t. a character $\theta \in C_0^s$, see [40], Definition 2.1), these projective lines intersect as displayed by the edges of the Dynkin diagram.

For $i \in I_0$ and $\lambda \in \mathcal{P}$, let us set $I(\varpi_i, \lambda) = \mathfrak{H}_{C_0^s, F_{\{i\}}}(I(\lambda))$. As before, we simplify the notation $I(\varpi_i, (n))$ into $I(\varpi_i, n)$ when λ has just one nonzero part, here n. It follows from Theorem 6.3 and Propositions 6.4 and 7.9 (iii) that if $\lambda = (\lambda_1, \ldots, \lambda_\ell)$, then

$$I(\varpi_i, \lambda) = \overline{I(\varpi_i, \lambda_1) \oplus \cdots \oplus I(\varpi_i, \lambda_\ell)}. \tag{7.6}$$

Lemma 7.16 *(i) Let $i \in I_0$. For any partition λ, a general point in $I(\varpi_i, \lambda)$ is a ϖ_i-core.*

(ii) Let $(i, n) \in I_0 \times \mathbb{N}$. A general point in $I(\varpi_i, n)$ is indecomposable in Λ-mod.

(iii) For (i, m) and (j, n) in $I_0 \times \mathbb{N}$, we have

$$\hom_\Lambda(I(\varpi_i, m), I(\varpi_j, n)) = \operatorname{ext}^1_\Lambda(I(\varpi_i, m), I(\varpi_j, n)) = 0.$$

In addition, if $(i, m) \neq (j, n)$, then $I(\varpi_i, m) \neq I(\varpi_j, n)$.

Proof. Let n be a positive integer. Let us take a general point in $I(\varpi_i, n)$. It is of the form $\mathscr{H}_{C_0^s, F_{\{i\}}}(T)$, where T is a general point in $I(n)$, and there is a general point $S \in I(1)$ such that T is an n-th iterated extension of S. By Proposition 7.14 (iii), $\mathscr{H}_{C_0^s, F_{\{i\}}}(S)$ is a ϖ_i-core. Thus $\mathscr{H}_{C_0^s, F_{\{i\}}}(T)$, being an iterated extension of ϖ_i-cores, is itself a ϖ_i-core. This shows (i) in the case where λ has just one nonzero part; the general case follows then from (7.6).

The functor $\mathscr{H}_{C_0^s, F_{\{i\}}}$ preserves the indecomposability, because it is fully faithful and because direct summands are kernels of idempotent endomorphisms. Assertion (ii) is thus a consequence of Proposition 7.9 (i).

Let $(i, j, m, n) \in (I_0)^2 \times \mathbb{N}^2$, with $i \neq j$. Let (T', T'') be a general point in $I(\varpi_i, m) \times I(\varpi_j, n)$. Then there is a general point $(S', S'') \in I(\varpi_i, 1) \times I(\varpi_j, 1)$ such that T' is an m-th iterated extension of S' and T'' is an n-iterated extension of S''. By Proposition 7.14 (ii), the head of S' is concentrated at vertex i and the head of S'' is concentrated at vertex j. Therefore the simple objects S' and S'' of $\mathscr{R}_{C_0^s}$ are not isomorphic. By Schur's lemma, we then have $\operatorname{Hom}_\Lambda(S', S'') = \operatorname{Hom}_\Lambda(S'', S') = 0$, whence $\operatorname{Hom}_\Lambda(T', T'') = \operatorname{Hom}_\Lambda(T'', T') = 0$. Using

Crawley-Boevey's formula (4.2), we get $\mathrm{Ext}^1_\Lambda(T', T'') = 0$. Moreover, this argument show that the head of a general point in $I(\varpi_i, m)$ (respectively, $I(\varpi_j, n)$) is concentrated at vertex i (respectively, j), whence $I(\varpi_i, m) \neq I(\varpi_j, n)$. All this shows (iii) in the case $i \neq j$. When $i = j$, item (iii) follows from Proposition 7.9 (ii). □

Any spherical chamber coweight γ can be written as $w\varpi_i$, with $(i, w) \in I_0 \times W_0$. We denote by $I(\gamma, \lambda)$ the image of $I(\varpi_i, \lambda)$ by the functor $I_w \otimes ?$, viewed as operating on irreducible components as in section 5.5. Proposition 7.11 shows that $I(\gamma, \lambda)$ does not depend on the choice of w, which justifies the notation; it also shows that a general point of $I(\gamma, \lambda)$ is a γ-core. Moreover,

$$I(\gamma, \lambda) = \overline{I(\gamma, \lambda_1) \oplus \cdots \oplus I(\gamma, \lambda_\ell)}.$$

Theorem 7.17 *Let C be a spherical Weyl chamber. Then the map*

$$(\lambda_\gamma) \mapsto \overline{\bigoplus_{\gamma \in \Gamma \cap \overline{C}} I(\gamma, \lambda_\gamma)}$$

is a bijection from $\mathcal{P}^{\Gamma \cap \overline{C}}$ onto $\mathfrak{R}_C$.

Proof. Using the reflection functors (specifically, Theorem 5.18 (i) and Propositions 5.7 and 5.23), we can reduce to the case of the dominant chamber C_0^s.

Lemma 7.16 shows that each $\overline{\bigoplus_{\gamma \in \Gamma \cap \overline{C_0^s}} I(\gamma, \lambda_\gamma)}$ is an irreducible component whose general point belongs to $\mathscr{R}_{C_0^s}$. The canonical decompositions of these components show that they are pairwise distinct. The map described in the statement is thus well defined and injective. Its surjectivity follows from Corollary 7.6. □

Theorem 7.17 provides the following intrinsic characterization of $I(\gamma, n)$.

Corollary 7.18 *For each $\gamma \in \Gamma$ and each $n \geq 1$, $I(\gamma, n)$ is the unique irreducible component of $\Lambda(n\delta)$ whose general point is an indecomposable γ-core.*

We have now established Theorem 1.1 and 1.1: Theorem 1.2 is a special case of Corollary 7.18, and with this in hand Theorem 1.2 is exactly Theorem 7.17.

The characterization in Corollary 7.18 also proves that the components $I(\gamma, n)$, and thus also the components $I(\gamma, \lambda)$, do not depend on the various choices made to construct them,

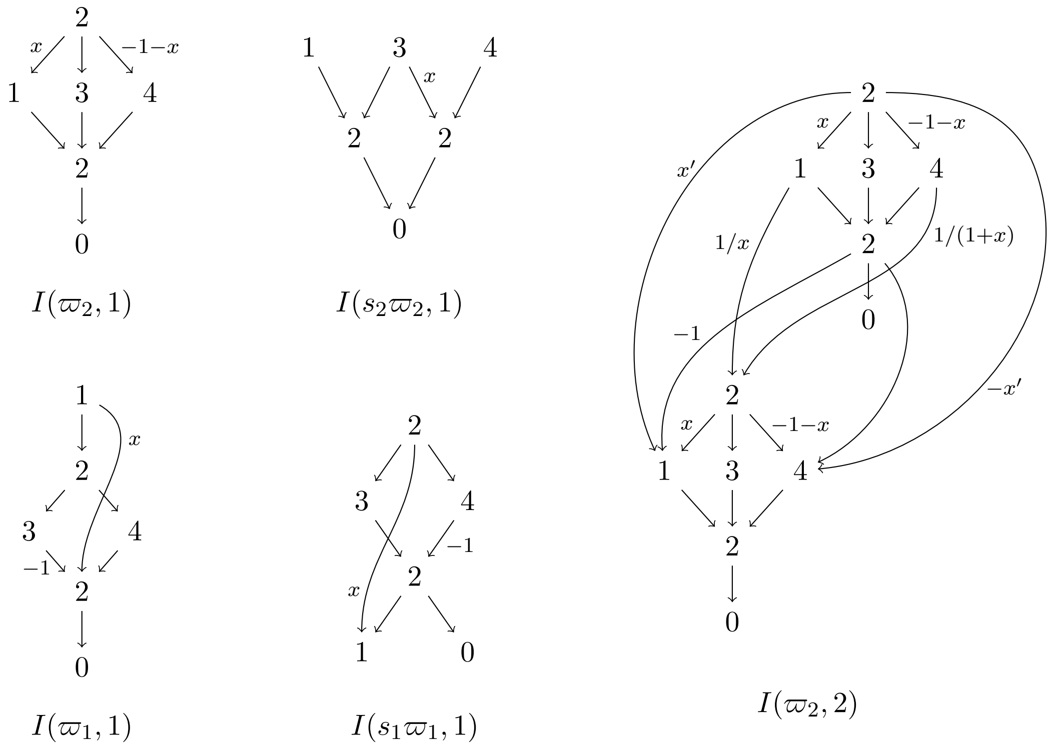

Figure 5: Examples of cores in type D_4 (convention: the central node of the Dynkin diagram is 2). The isomorphism class of the general point of $I(\gamma, n)$ depends on n parameters; here these parameters are called x and x'.

notably in the construction of the functors $\mathscr{H}_{C_0^s, F_{\{i\}}}$ used to define $I(\varpi_i, \lambda)$. It also implies that $I(-\gamma, n) = I(\gamma, n)^*$, whence

$$I(-\gamma, \lambda) = I(\gamma, \lambda)^* \tag{7.7}$$

for any $(\gamma, \lambda) \in \Gamma \times \mathcal{P}$.

This independence on the choice of the Hall functors leads to the following result.

Corollary 7.19 *For all flags (C, F) and all partitions λ, we have $\mathfrak{H}_{C,F}(I(\lambda)) = I(\gamma_{C,F}, \lambda)$.*

Proof. Suppose that we are in the situation of Example 7.12 (v) with $(C, F) = (wC_0^s, wF_{\{i\}})$. Then $\mathrm{Hom}_\Lambda(I_w, ?) \circ \mathscr{H}_{C,F}$ is a Hall functor built from the datum of S_i and R_i, so it can play the role of $\mathscr{H}_{C_0^s,F_{\{i\}}}$. Thus by the above observation, the map on components defined by the functor $\operatorname{Hom}_\Lambda(I_w, ?)$ sends $\mathfrak{H}_{C,F}(I(\lambda))$ to $I(\varpi_i, \lambda)$, for each partition λ. This immediately yields $\mathfrak{H}_{C,F}(I(\lambda)) = I(\gamma_{C,F}, \lambda)$.

This analysis applies to exactly half of the flags. The remaining flags are of the form $(C, F) = (-wC_0^s, -wF_{\{i\}})$; their cases are deduced from the previous situation by duality. □

Combining the above result with Proposition 7.9 (iv) and Example 7.12 (ii), we obtain:

Corollary 7.20 *Let F be a facet and let C' and C'' be the two Weyl chambers that F separates. For any partition λ, we have $\mathfrak{H}_{C',F}(I(\lambda)^*) = I(\gamma_{C''/F}, \lambda)$.*

Our last corollary to Theorem 7.17 rounds off Theorem 7.7 (i).

Corollary 7.21 *Let F be a face of the spherical Weyl fan. Let T be a simple object of $\mathscr{R}_F$. If $\underline{\dim}\, T$ is a multiple of δ, then in fact $\underline{\dim}\, T = \delta$.*

Proof. Using the reflection functors, we can reduce to the case where F is contained in the dominant spherical Weyl chamber C_0^s.

Let n be a positive integer. Take an irreducible component $Z \in \mathfrak{R}_{C_0^s}$ in dimension-vector $n\delta$. This Z can be written in the form provided by Theorem 7.17. Combining Proposition 7.9 (v) with Lemma 7.16, we obtain that $\dim \operatorname{End}_\Lambda(X) = n$ for any general point $X \in Z$. Since the function $X \mapsto \dim \operatorname{End}_\Lambda(X)$ is upper semicontinuous on $\Lambda(n\delta)$, we have $\dim \operatorname{End}_\Lambda(X) \geq n$ for any $X \in Z$. Since Z was arbitrary in $\mathfrak{R}_{C_0^s}$, the latter inequality holds for any $X \in \mathscr{R}_{C_0^s}$ with dimension-vector $n\delta$.

Now take a simple object $T \in \mathscr{R}_F$ with dimension-vector $n\delta$. Then $T \in \mathscr{R}_{C_0^s}$, by Theorem 7.7 (i), and $\dim \operatorname{End}_\Lambda(T) = 1$, by Schur's lemma. The result just established then implies that $1 \geq n$. □

Theorem 7.22 *For any flag (C, F), the map*

$$(Z, (\lambda_\gamma)) \mapsto \overline{\mathfrak{H}_{C,F}(Z) \oplus \bigoplus_{\gamma \in \Gamma \cap \overline{F}} I(\gamma, \lambda_\gamma)}$$

is a bijection from $\mathfrak{A} \times \mathcal{P}^{\Gamma \cap \overline{F}}$ onto $\mathfrak{R}_F$.

Proof. Let $\gamma \in \Gamma \cap \overline{F}$ and let n be a positive integer. Let T be a general point in $I(\gamma, n)$. Then T is the n-th iterated extension of a general point $X \in I(\gamma, 1)$. By Corollary 7.18, X is a γ-core of dimension-vector δ, so X is simple in $\mathscr{R}_\gamma$ by Proposition 7.10, so X is simple in $\mathscr{R}_F$. By Theorem 7.7, $\{X\}$ is a linkage class in $\operatorname{Irr} \mathscr{R}_F$, so X is not linked to the modules $S_{C,F}$ and $R_{C,F}$. It follows that $\operatorname{Ext}^1_\Lambda(Y, T) = 0$ for any module Y in the subcategory $\langle S_{C,F}, R_{C,F}\rangle$, and therefore that $\operatorname{ext}^1_\Lambda(\mathfrak{H}_{C,F}(Z), I(\gamma, n)) = 0$ for any $Z \in \mathfrak{A}$. In view of Crawley-Boevey and Schröer's theory, this implies that our map is well defined.

The argument above also shows that the general point of the component $I(\gamma, n)$ does not belong to the category $\langle S_{C,F}, R_{C,F}\rangle$, so $I(\gamma, n)$ cannot occur in the canonical decomposition of $\mathfrak{H}_{C,F}(Z)$. The set of indecomposable irreducible components that arise from the $\mathfrak{H}_{C,F}(Z)$ is thus disjoint from $\left\{I(\gamma, n) \,\middle|\, \gamma \in \Gamma\cap\overline{F},\ n \geq 1\right\}$. The uniqueness of the canonical decomposition of an element in $\mathfrak{R}_F$ implies then that our map is injective.

Finally, we use a counting argument to prove the surjectivity of our map. Pick $\theta \in F$. If $Z \in \mathfrak{A}$ has weight $\mu = (\mu_0, \mu_1)$, then $\mathfrak{H}_{C,F}(Z) \in \mathfrak{B}$ has weight $\mathbf{K}(\mathscr{H}_{C,F})(\mu) = \mu_0 \underline{\dim}\, S_{C,F} + \mu_1 \underline{\dim}\, R_{C,F}$. We here note that $\mathbf{K}(\mathscr{H}_{C,F})$ maps the imaginary root $\delta \in \Delta_+$ to the imaginary root in $\delta \in \Phi_+$; thus the use of the same notation δ for both root systems does not lead to any confusion. Plugging this information into the character series for $\mathfrak{A}$, given by the analog for Δ_+ of (7.2), and adding the contribution of the components $I(\gamma, \lambda_\gamma)$, we get the character of the image of our map:

$$\prod_{n\in\mathbb{N}} \left(\frac{1}{1 - t^{\underline{\dim}\, S_{C,F} + n\delta}} \times \frac{1}{1 - t^{\underline{\dim}\, R_{C,F} + n\delta}} \times \frac{1}{1 - t^{(n+1)\delta}} \right) \times \prod_{\gamma\in\Gamma\cap\overline{F}} \left(\sum_{\lambda_\gamma\in\mathcal{P}} t^{|\lambda_\gamma|\delta} \right).$$

This is equal to

$$P_{\mathfrak{R}_\theta} = \left(\prod_{\substack{\alpha\in\Phi_+^{\mathrm{re}}\\ \langle\theta,\alpha\rangle=0}} \frac{1}{1 - t^\alpha} \right) \left(\prod_{n\geq 1} \frac{1}{1 - t^{n\delta}} \right)^r$$

(see Corollary 7.6), which ensures that our map is surjective. □

Let $(Z, (\lambda_\gamma)) \in \mathfrak{A} \times \mathcal{P}^{\Gamma\cap\overline{F}}$ and denote by $\widetilde{Z} \in \mathfrak{R}_F$ its image under the map in Theorem 7.22. Write each partition involved as $\lambda_\gamma = (\lambda_{\gamma,1}, \lambda_{\gamma,2}, \ldots)$ and denote by $\ell(\lambda_\gamma)$ the number of nonzero parts of λ_γ. Pick general points $X \in Z$ and $Y_{\gamma,p} \in I(\gamma, \lambda_{\gamma,p})$, for each $\gamma \in \Gamma \cap \overline{F}$ and each $p \in \{1, \ldots, \ell(\lambda_\gamma)\}$. Then

$$\widetilde{X} = \mathscr{H}_{C,F}(X) \oplus \left(\bigoplus_{\gamma\in\Gamma\cap F} \left(\bigoplus_{1\leq p\leq\ell(\lambda_\gamma)} Y_{\gamma,p} \right) \right) \tag{7.8}$$

is a general point of $\widetilde{Z}$. By Remark 3.5 (ii) applied to the category $\mathscr{R}_F$, the HN polytope of $\widetilde{X}$ is the Minkowski sum of the HN polytopes of its summands, that is, the Minkowski sum of $\mathbf{K}(\mathscr{H}_{C,F})_{\mathbb{R}}(\mathrm{Pol}(X))$ and of segments that join 0 to each $[Y_{\gamma,p}]$. Further, if $i : \mathscr{R}_F \subseteq \Lambda$-mod denotes the inclusion functor, then $\mathbf{K}(i)_{\mathbb{R}}([Y_{\gamma,p}]) = p\delta$, so

$$\mathbf{K}(i)_{\mathbb{R}}\big(\mathrm{Pol}\big(\widetilde{X}\big)\big) = \mathbf{K}(i \circ \mathscr{H}_{C,F})_{\mathbb{R}}(\mathrm{Pol}(X)) + \sum_{\gamma \in \Gamma \cap \overline{F}} \big[\, 0,\ |\lambda_\gamma|\,\delta\,\big]. \tag{7.9}$$

Pick now $\theta \in F$ and $\Lambda_b \in \mathfrak{B}$. The bijection Ξ_θ^{-1} maps Λ_b to say $(\Lambda_{b'}, \Lambda_{b''}, \Lambda_{b'''}) \in \mathfrak{I}_F \times \mathfrak{R}_F \times \mathfrak{P}_F$. If $T \in \Lambda_b$ is a general point, then $T_\theta^{\max}/T_\theta^{\min}$ is a general point in $\Lambda_{b''}$. Corollary 3.3 then says that the HN polytope of $T_\theta^{\max}/T_\theta^{\min}$, viewed as an object of $\mathscr{R}_F$, is the 2-face of the HN polytope of T defined by θ. Putting $\widetilde{Z} = \Lambda_{b''}$ in the previous paragraph, we then see that this 2-face can be written as a Minkowski sum (7.9). This sum involves $|\Gamma \cap \overline{F}| = r - 1$ one-dimensional polytopes, which all point in the direction δ, so we can intuitively regard it as an MV polytope of type $\widetilde{A}_1 \times \widetilde{A}_0^{r-1}$ if we equip it with adequate partitions.

We now need to look at these partitions. In particular, we need to show that the partitions used in Theorem 7.22 (including those that decorate the polytope $\mathrm{Pol}(X)$) are the same as the partitions provided by Theorem 7.17, which decorate $\mathrm{Pol}(T)$.

7.5 The MV polytope of a component (proof of Theorem 1.3)

For any spherical Weyl chamber C, we have bijections

$$\mathfrak{I}_C \times \mathfrak{R}_C \times \mathfrak{P}_C \to \mathfrak{B} \quad \text{and} \quad \mathcal{P}^{\Gamma \cap \overline{C}} \to \mathfrak{R}_C,$$

by Proposition 4.5 and Theorem 7.17. Therefore each $\Lambda_b \in \mathfrak{B}$ provides a tuple of partitions $(\lambda_\gamma)_{\gamma \in \Gamma \cap \overline{C}}$. Concretely, for any $\eta \in C$ and any general point T in Λ_b, in the Krull-Schmidt decomposition of $T_\eta^{\max}/T_\eta^{\min}$, there are as many γ-cores of dimension-vector $n\delta$ as parts equal to n in λ_γ.

A priori, λ_γ depends on b, on γ and on C, but in fact it only depends on b and γ. Our aim now is to show this fact. To this end we fix b and we indicate the potential dependence on C by writing $\lambda_\gamma(C)$.

Let us study what happens around a facet F. We thus consider a flag (C', F), whence a functor $\mathscr{H}_{C',F}$ to which we apply Theorem 7.22. Tracing $\Lambda_b \in \mathfrak{B}$ through the bijections

$$\mathfrak{I}_F \times \mathfrak{R}_F \times \mathfrak{P}_F \to \mathfrak{B} \quad \text{and} \quad \mathfrak{A} \times \mathcal{P}^{\Gamma \cap \overline{F}} \to \mathfrak{R}_F$$

given by Proposition 4.5 and Theorem 7.22, we get $(Z, (\lambda_\gamma(F))) \in \mathfrak{A} \times \mathcal{P}^{\Gamma \cap \overline{F}}$. Concretely, for any $\theta \in F$ and any general point T in Λ_b, we can write

$$T_\theta^{\max}/T_\theta^{\min} \cong \mathscr{H}_{C,F}(X) \oplus \left(\bigoplus_{\gamma \in \Gamma \cap F} \left(\bigoplus_{1 \leq p \leq \ell(\lambda_\gamma)} Y_{\gamma,p} \right) \right)$$

where the notations are as in equation (7.8). In addition, with the notation of section 7.3, we can look at the Π-module $Y_{\gamma'}^{\max}/Y_{\gamma'}^{\min}$; this is the general point of an irreducible component $I(\lambda')$. Likewise, the Π-module $Y_{\gamma''}^{\max}/Y_{\gamma''}^{\min}$ is the general point of an irreducible component $I(\lambda'')^*$. Finally, let C'' be the other spherical Weyl chamber bordered by F.

Proposition 7.23 *In the context above,*

$$\lambda_{\gamma_{C'/F}}(C') = \lambda', \quad \lambda_{\gamma_{C''/F}}(C'') = \lambda'',$$

and for each $\gamma \in \Gamma \cap \overline{F}$,

$$\lambda_\gamma(C') = \lambda_\gamma(C'') = \lambda_\gamma(F).$$

Proof. Take m large enough and consider $\eta = m\theta + \gamma_{C'/F}$.

We first note that $\gamma' = \eta \circ \mathbf{K}(\mathscr{H}_{C',F})$. In fact, $\mathbf{K}(\mathscr{H}_{C',F})$ maps a dimension-vector μ in $\mathbf{K}(\Pi\text{-mod})$ to the element $\mu_0 \, \underline{\dim}\, S_{C',F} + \mu_1 \, \underline{\dim}\, R_{C',F}$ of $\mathbf{K}(\Lambda\text{-mod})$. Using (7.5), we then see that $\eta \circ \mathbf{K}(\mathscr{H}_{C',F})$ maps μ to $\mu_0 - \mu_1$, as does γ' (see section 7.3).

Now η is an element of C', so the module $T_\eta^{\max}/T_\eta^{\min}$ bears the information about the partitions $\lambda_\gamma(C')$. Proposition 3.4 explains how this module can be obtained from $T_\theta^{\max}/T_\theta^{\min}$. In the process, the summands $Y_{\gamma,p}$ stay unchanged, because they already belong to $\mathscr{R}_{C'}$. By contrast, $\mathscr{H}_{C',F}(X)$ is truncated to its subquotient $X' = \mathscr{H}_{C',F}(X_{\gamma'}^{\max}/X_{\gamma'}^{\min})$, which is a general point of the component $\mathfrak{H}_{C',F}(I(\lambda')) = I(\gamma_{C'/F}, \lambda')$. We thus have

$$\lambda_\gamma(C') = \begin{cases} \lambda_\gamma(F) & \text{if } \gamma \in \Gamma \cap \overline{F}, \\ \lambda' & \text{if } \gamma = \gamma_{C'/F}. \end{cases}$$

The partitions $\lambda_\gamma(C'')$ are computed in a similar fashion, using Corollary 7.20 at the last step.
□

Proposition 7.23 asserts in particular that $\lambda_\gamma(C') = \lambda_\gamma(C'')$ if C' and C'' are two adjacent spherical Weyl chambers whose closures contain γ. This implies that $\lambda_\gamma(C)$ is independent of C, assuming of course that $\gamma \in \overline{C}$.

To an irreducible component $\Lambda_b \in \mathfrak{B}$, we may thus associate a family of partitions $(\lambda_\gamma) \in \mathcal{P}^\Gamma$. In addition, by Proposition 4.2 (ii), Λ_b contains a dense open subset on which the map $T \mapsto \mathrm{Pol}(T)$ is constant. As in the introduction, we denote by $\widetilde{\mathrm{Pol}}(b)$ the datum of this constant value $\mathrm{Pol}(T)$ and of the family of partitions (λ_γ).

We have shown that $\widetilde{\mathrm{Pol}}(b)$ belongs to the set $\mathcal{MV}$ of decorated lattice polytopes defined in section 1.7:

- The polytope is GGMS, by Corollary 5.21, so its normal fan is coarser than $\mathscr{W}$.
- Given a spherical Weyl chamber C, the partitions (λ_γ) for $\gamma \in \overline{C}$ describe the Krull-Schmidt decomposition of $T_\theta^{\max}/T_\theta^{\min}$, where T is general in Λ_b and $\theta \in C$. Therefore
$$\left(\sum_{\gamma \in \Gamma \cap \overline{C}} |\lambda_\gamma|\right) \delta = \underline{\dim}\, T_\theta^{\max}/T_\theta^{\min}.$$
- Each 2-face of finite type is constrained by the relations in Propositions 5.26 and 5.27, and each 2-face of affine type is an MV polytope of type $\widetilde{A}_1 \times \widetilde{A}_0^{r-1}$, as explained at the end of section 7.4 and in the discussion following Proposition 7.23.

The last point in the list above proves Theorem 1.3.

7.6 Lusztig data (proof of Theorem 1.5)

We are now ready to prove Theorem 1.5. However, first we need to precisely define the map $\Omega_\preceq$ (i.e. the map taking an element of $B(-\infty)$ to its Lusztig data) in affine type.

Fix a convex order $\preccurlyeq$. Given a weight $\nu \in \mathbb{N}I$, we set $E_\nu = \{\alpha \in \Phi_+^{\mathrm{re}} \sqcup \{\delta\} \mid \mathrm{ht}\,\alpha \le \mathrm{ht}\,\nu\}$. Enumerate the elements in E_ν in decreasing order: $\beta_1 \succ \beta_2 \succ \cdots \succ \beta_\ell$. For $1 \le k \le \ell$, set $A_k = \{\alpha \in \Phi_+ \mid \alpha \succ \beta_k\}$ and $B_k = \{\alpha \in \Phi_+ \mid \alpha \succcurlyeq \beta_k\}$. These biconvex subsets provide a nested family of torsion pairs (here we write only the torsion classes):

$$\{0\} \subseteq \mathscr{T}(A_1) \subseteq \mathscr{T}(B_1) \subseteq \mathscr{T}(A_2) \subseteq \cdots \subseteq \mathscr{T}(A_\ell) \subseteq \mathscr{T}(B_\ell) \subseteq \Lambda\text{-mod}.$$

On a Λ-module T, this induces a filtration

$$0 \subseteq T_1 \subseteq \overline{T}_1 \subseteq T_2 \subseteq \cdots \subseteq T_\ell \subseteq \overline{T}_\ell \subseteq T, \tag{7.10}$$

with $\overline{T}_k/T_k \in \mathscr{F}(A_k) \cap \mathscr{T}(B_k)$. If $\underline{\dim}\, T = \nu$, then by Lemma 7.3 the only jumps in the filtration (7.10) occur between a T_k and the corresponding $\overline{T}_k$. If moreover T is a general point in an irreducible component $Z \in \mathfrak{B}(\nu)$, then by Proposition 4.5 each subquotient $\overline{T}_k/T_k$ will be a general point in an irreducible component $Z_k \in \mathfrak{F}(A_k) \cap \mathfrak{T}(B_k)$, and furthermore the map

$$\mathfrak{B}(\nu) \to \left\{ (Z_1, \dots, Z_\ell) \in \prod_{k=1}^{\ell} \Big(\mathfrak{F}(A_k) \cap \mathfrak{T}(B_k) \Big) \;\middle|\; \sum_{k=1}^{\ell} \operatorname{wt} Z_k = \nu \right\} \tag{7.11}$$

is a bijection.

By Proposition 7.4, if β_k is real, then $\mathscr{F}(A_k) \cap \mathscr{T}(B_k) = \operatorname{add} L(A_k, B_k)$, where $L(A_k, B_k)$ is a rigid Λ-module of dimension-vector β_k. Therefore $\mathfrak{F}(A_k) \cap \mathfrak{T}(B_k)$ is in one-to-one correspondence with $\mathbb{N}$: to a natural number n corresponds the closure in $\Lambda(n\beta_k)$ of the orbit that represents $L(A_k, B_k)^{\oplus n}$.

On the other hand, if $\beta_k = \delta$, then we can find a spherical Weyl chamber C such that $(A_k, B_k) = (A_\theta^{\min}, A_\theta^{\max})$ for $\theta \in C$ (Lemma 2.10 (i)), and then $\mathscr{F}(A_k) \cap \mathscr{T}(B_k) = \mathscr{R}_C$ (Proposition 7.2), whence $\mathfrak{F}(A_k) \cap \mathfrak{T}(B_k) = \mathfrak{R}_C$. Further, Theorem 7.17 provides a bijection between $\mathfrak{R}_C$ and $\mathcal{P}^{\Gamma \cap \overline{C}}$.

The bijection (7.11) can thus be rewritten as

$$\mathfrak{B}(\nu) \to \left\{ ((n_\beta), (\lambda_\gamma)) \in \mathbb{N}^{E_\nu \cap \Phi_+^{\mathrm{re}}} \times \mathcal{P}^{\Gamma \cap \overline{C}} \;\middle|\; \sum_{\beta \in E_\nu \cap \Phi_+^{\mathrm{re}}} n_\beta \beta + \left(\sum_{\gamma \in \Gamma \cap \overline{C}} |\lambda_\gamma| \right) \delta = \nu \right\}.$$

Letting ν run over $\mathbb{N}I$ and assembling the resulting maps, we get a bijection

$$\Omega_\preccurlyeq : \mathfrak{B} \to \mathbb{N}^{(\Phi_+^{\mathrm{re}})} \times \mathcal{P}^{\Gamma \cap \overline{C}},$$

proving Theorem 1.5.

One may here observe that the map $\Omega_{\mathbf{i}}$ constructed in section 5.5 gives the beginning of $\Omega_\preccurlyeq$ when the smallest roots for $\preccurlyeq$ are, in order

$$\alpha_{i_1},\ s_{i_1}\alpha_{i_2},\ s_{i_1}s_{i_2}\alpha_{i_3},\ \dots$$

In addition, Remark 5.25 (ii) and Propositions 5.26 and 5.27 justify referring to the bijection $\Omega_{\mathbf{i}}$ as the partial Lusztig datum in direction $\mathbf{i}$. We are thus led to regard $\Omega_\preccurlyeq$ as the Lusztig datum in direction $\preccurlyeq$. A further justification of this terminology is the fact that the components of $\Omega_\preccurlyeq(b)$ can be read as the lengths and the decorations of the edges on the path in the 1-skeleton of $\widetilde{\operatorname{Pol}}(b)$ defined by $\preccurlyeq$.

7.7 Proof of Theorem 1.4

Proof of the injectivity of $\widetilde{\mathrm{Pol}}$. Let us choose a convex order $\preccurlyeq$, and let C be the spherical Weyl chamber whose positive root system is $\pi(\{\beta \in \Phi_+^{\mathrm{re}} \mid \beta \succ \delta\})$.

An element $\widetilde{P} \in \mathcal{MV}$ is the datum of a GGMS polytope P and of a family of partitions $(\lambda_\gamma) \in \mathcal{P}^\Gamma$, subject to certain conditions. To P and $\preccurlyeq$, the construction at the end of section 2.6 associates a finitely supported family of non-negative integers (n_α), indexed by Φ_+^{re}; specifically, n_α is the length of the edge parallel to α on the path in the 1-skeleton of P defined by $\preccurlyeq$. Adding to this datum the partitions λ_γ for all γ in $\Gamma \cap \overline{C}$, we get an element $\widehat{\Omega}_\preccurlyeq(\widetilde{P}) \in \mathbb{N}^{(\Phi_+^{\mathrm{re}})} \times \mathcal{P}^{\Gamma \cap \overline{C}}$.

Proposition 7.1 shows that the resulting map $\widehat{\Omega}_\preccurlyeq$ is compatible with the map $\Omega_\preccurlyeq$ constructed in the previous section, in the sense that the diagram

$$\begin{tikzcd} \mathfrak{B} \arrow[rr, "\widetilde{\mathrm{Pol}}"] \arrow[dr, "\Omega_\preccurlyeq"'] & & \mathcal{MV} \arrow[dl, "\widehat{\Omega}_\preccurlyeq"] \\ & \mathbb{N}^{(\Phi_+^{\mathrm{re}})} \times \mathcal{P}^{\Gamma \cap \overline{C}} & \end{tikzcd} \tag{7.12}$$

commutes. The injectivity of $\widetilde{\mathrm{Pol}}$ then follows from the injectivity of $\Omega_\preccurlyeq$. □

Proof of the surjectivity of $\widetilde{\mathrm{Pol}}$. As explained in Example 2.14 (ii), a sufficiently general $\theta \in (\mathbb{R}I)^*$ defines a convex order. The precise condition is that θ avoids the countably many hyperplanes

$$H_{\alpha,\beta} = \{\theta \in (\mathbb{R}I)^* \mid \theta(\alpha)/\operatorname{ht}(\alpha) = \theta(\beta)/\operatorname{ht}(\beta)\},$$

where (α, β) runs over pairs of non-proportional roots. We denote the collection of all these hyperplanes by $\mathscr{X}$. For simplicity, we will denote by $\widehat{\Omega}_\theta$ the map $\widehat{\Omega}_\preccurlyeq$ relative to the order $\preccurlyeq$ defined by such a θ.

Consider an MV polytope $\widetilde{P} = (P, (\lambda_\gamma))$. Every vertex (respectively, every imaginary edge) of P is a face P_θ, and $\theta \in (\mathbb{R}I)^*$ can certainly be chosen outside all the hyperplanes in the collection $\mathscr{X}$. The position of this vertex (respectively, the partitions relative to this imaginary edge) is then determined by $\widehat{\Omega}_\theta(\widetilde{P})$. Thus $\widetilde{P}$ is determined by the datum of $\widehat{\Omega}_\theta(\widetilde{P})$ for all possible θ.

Now fix an element $\eta_0 \in (\mathbb{R}I)^*$ outside the hyperplanes in the collection $\mathscr{X}$. We will show that for any general $\eta_1 \in (\mathbb{R}I)^*$ and any MV polytope $\widetilde{P}$, the datum $\widehat{\Omega}_{\eta_1}(\widetilde{P})$ can unambiguously be determined from $\widehat{\Omega}_{\eta_0}(\widetilde{P})$ by a rule which depends on $\widetilde{P}$ only through its weight. Together with the argument in the previous paragraph, this establishes the injectivity of the map $\widehat{\Omega}_{\eta_0}$.

Chasing in the diagram (7.12) relative to the order $\preccurlyeq$ defined by η_0, we can then deduce the desired result from the surjectivity of the map $\Omega_{\preccurlyeq}$.

Let $\widetilde{P} = (P, (\lambda_\gamma))$ be an MV polytope, let ν be the weight of $\widetilde{P}$, and as in section 7.6, let $E_\nu = \left\{\alpha \in \Phi_+^{\mathrm{re}} \sqcup \{\delta\} \;\middle|\; \operatorname{ht}\alpha \leq \operatorname{ht}\nu\right\}$. The datum $\widehat{\Omega}_{\preccurlyeq}(\widetilde{P})$ depends on the choice of the convex order $\preccurlyeq$ only through the restriction of $\preccurlyeq$ to E_ν. Therefore, as a function of θ, the datum $\widehat{\Omega}_\theta(\widetilde{P})$ only changes when θ crosses an hyperplane $H \in \mathscr{X}_\nu$, where $\mathscr{X}_\nu$ denotes the collection of all hyperplanes $H_{\alpha,\beta}$ with $(\alpha, \beta) \in (E_\nu)^2$.

Pick $\eta_1 \in (\mathbb{R}I)^*$ outside all the hyperplanes of the collection $\mathscr{X}$, and let (η_t) be a piecewise linear path in $(\mathbb{R}I)^*$ that connects η_0 to η_1, general enough so that η_t lies on a hyperplane of the collection $\mathscr{X}_\nu$ at only finitely many times $0 < t_1 < \cdots < t_m < 1$, and never lies at once on two such hyperplanes. Set $t_0 = 0$ and $t_{m+1} = 1$. For a fixed $j \in \{0, \ldots, m\}$, the data $\widehat{\Omega}_{\eta_t}(\widetilde{P})$ for $t \in (t_j, t_{j+1})$ are all one and the same. We denote this datum by $\widehat{\Omega}_j$ and observe that $\widehat{\Omega}_{\eta_0}(\widetilde{P}) = \widehat{\Omega}_0$ and $\widehat{\Omega}_{\eta_1}(\widetilde{P}) = \widehat{\Omega}_m$. For $j \in \{1, \ldots, m\}$, the data $\widehat{\Omega}_{j-1}$ and $\widehat{\Omega}_j$ record the lengths and the partitions along two paths in the 1-skeleton of P which coincide almost everywhere, the only difference between these paths being that they may traverse the opposite sides of the 2-face P_θ, where $\theta = \eta_{t_j}$. But $\widetilde{P}$ is an MV polytope, so its 2-faces are constrained: by Propositions 5.26 and 5.27 in the case where P_θ is a 2-face of finite type, and by Theorem 7.22 and Proposition 7.23 in the case where P_θ is a 2-face of affine type, the datum $\widehat{\Omega}_{j-1}$ determines $\widehat{\Omega}_j$. Thus $\widehat{\Omega}_{\eta_1}(\widetilde{P})$ is determined by $\widehat{\Omega}_{\eta_0}(\widetilde{P})$, as announced. □

Appendix: Restriction to the tame quiver

The path algebra KQ of an acyclic quiver Q can be seen as a subalgebra of the completed preprojective algebra Λ_Q of Q. In our present situation of an extended Dynkin diagram, Q is tame, so its representation theory is very well understood, thanks to the work of Dlab and Ringel. In this appendix, we discuss our constructions in terms of the representation theory of Q.

We begin with a refinement to Theorem 7.7 in the case where F is a minimal face, that is, the ray generated by a spherical chamber coweight γ. For $(\mu, \nu) \in (\mathbb{Z}I)^2$, we write $\mu \geq \nu$ if $\mu - \nu \in \mathbb{N}I$.

Proposition A.1 *Let F be a ray of the spherical Weyl fan and let L be a linkage class of simple objects in $\mathscr{R}_F$. Then $\sum_{S \in L} \underline{\dim}\, S \leq \delta$.*

Proof. Theorem 7.7 distinguishes two kinds of simple objects, described in its assertions (i) and (ii). For the objects of the first kind, the desired property is proved in Corollary 7.21. In the sequel of this proof, we consider the other case, when the dimension-vectors of objects in L belong to $\iota(\Phi^s)$.

We choose an extending vertex in I and we set $I_0 = I \setminus \{0\}$. The spherical root system Φ^s is then endowed with a basis, namely $\{\pi(\alpha_i) \mid i \in I_0\}$, whence a positive system Φ^s_+, and a dominant spherical Weyl chamber C^s_0. As in section 2.3, we identify the spherical Weyl group W_0 with the parabolic subgroup $\langle s_i \mid i \in I_0\rangle$ of W.

First consider the case where $F \subseteq \overline{C^s_0}$. Then F is spanned by a certain spherical fundamental coweight ϖ_i, with $i \in I_0$.

Given a connected component J of $I_0 \setminus \{i\}$, we can look at the root system $\Phi_J = \Phi \cap \mathbb{Z}J$. This root system is finite and irreducible and comes with a natural basis, so it has a largest root $\widetilde{\alpha}_J$. By the dual statement of [17], Lemma 2 (2), there is a unique Λ-module with socle S_0 and dimension-vector $\delta - \widetilde{\alpha}_J$; we denote it by R_J.

We claim that the head of R_J is isomorphic to S_i. In fact, S_0 does not occur in the head of R_J; otherwise, S_0 would be a direct factor of R_J (because it occurs in the socle of R_J and its Jordan-Hölder multiplicity in R_J is one), which is ruled out by the fact that R_J is indecomposable (the socle of R_J is simple) of dimension-vector $\neq \alpha_0$. If S_j occurs in the head of R_J, then we can produce a Λ-module X with socle S_0 and dimension-vector $\underline{\dim}\, R_J - \alpha_j$; the latter is then a root, by [17], Lemma 2 (1), and therefore $\widetilde{\alpha}_J + \alpha_j$ is a root; this forces $j = i$. If S_i occurred twice in the head of R_J, then $\underline{\dim}\, R_J - 2\alpha_i$ would be a root, so $\widetilde{\alpha}_J + 2\alpha_i$ would be a root, which is impossible because the root system Φ^s is simply laced.

Next we claim that R_J is a simple object in $\mathscr{R}_F$. To prove that, it suffices to show that R_J is ϖ_i-stable, in other words, that $\langle \varpi_i, \underline{\dim}\, R_J\rangle = 0$ and that $\langle \varpi_i, \underline{\dim}\,(R_J/X)\rangle > 0$ for all proper submodules X of R_J. The first equation comes from the fact that Φ_J is contained in $\ker \varpi_i$. To prove the second equation, we observe that S_0 is not a Jordan-Hölder component of R_J/X (because X contains the unique copy of S_0 in R_J), so the simple components in R_J/X are S_j with $j \in I_0$, and S_i appears at least once in R_J/X.

In addition, the modules S_j, for $j \in I_0 \setminus \{i\}$, are also simple objects in $\mathscr{R}_F$. We now claim that the modules R_J and S_j are all the simple objects in $\mathscr{R}_F$ whose dimension-vectors are in $\iota(\Phi^s)$.

Indeed, let $T \in \operatorname{Irr} \mathscr{R}_F$ such that $\underline{\dim}\, T \in \iota(\Phi^s_+)$. The vector space T_0 attached to the extending vertex is thus zero. The vector space T_i attached to i is then also zero, for $\varpi_i(\underline{\dim}\, T) = 0$. Thus T is an iterated extension of the modules S_j with $j \in I_0 \setminus \{i\}$. Since all these modules belong to $\mathscr{R}_F$, we conclude that T is one of these S_j.

On the other hand, let $T \in \operatorname{Irr} \mathscr{R}_F$ such that $\beta = \underline{\dim}\, T$ belongs to $\iota(\Phi^s_-)$. The simplicity of T forbids any S_j with $j \in I_0 \setminus \{i\}$ to appear in the socle or in the head of T. Further, S_i cannot appear in the socle of T, because $T \in \mathscr{R}_F$ and $S_i \in \mathscr{I}_F$. Therefore $\operatorname{soc} T = S_0$. This condition and β completely determine T, by [17], Lemma 2 (2). With the notations of [5], section 3 (see also Example 5.12), we have $T \cong N(\beta-\omega_0)$. (To apply [5], Theorem 3.1, we note the existence of $w \in W_0$ such that $\beta = w\alpha_0$, which implies $\beta - \omega_0 = -ws_0\omega_0$.) Equation (3.1) in [5] (or the proof of Lemma 2 in [17]) then says that

$$\dim \operatorname{hd}_j T = \max(0, (\beta, \alpha_j) - \langle \omega_0, \alpha_j \rangle),$$

and we have seen that the left-hand side is zero for $j \in I_0 \setminus \{i\}$. Let us write $\beta = \iota(-\alpha)$, with $\alpha \in \Phi^s_+$. Then $(\alpha, \alpha_j) \geq 0$ for each $j \in I_0 \setminus \{i\}$. Now $\langle \varpi_i, \beta \rangle = 0$, so the support of α (a subset of I_0) avoids the node i, and therefore $\alpha \in \Phi^s_J$ for a certain connected component J of $I_0 \setminus \{i\}$. We then conclude that $\alpha = \widetilde{\alpha}_J$, and therefore that $T = R_J$.

We now claim that the simple objects linked to R_J are the S_j with $j \in J$. By [10], chapitre 6, §1, n° 6, Proposition 19, there is a sequence $\beta_1, \ldots, \beta_n$ of elements in $\{\alpha_j \mid j \in J\}$ such that $\beta_1 + \cdots + \beta_k$ is a root for each k and $\beta_1 + \cdots + \beta_n = \widetilde{\alpha}_J$. For $k \in \{1, \ldots, n\}$, let T_k be the simple Λ-module with dimension-vector β_k. Let $N_{n+1} = R_J$, and for $1 < k \leq n$, let N_k be the Λ-module with socle S_0 and dimension-vector

$$\delta - (\beta_1 + \cdots + \beta_{k-1}) = \delta - \widetilde{\alpha}_J + (\beta_n + \cdots + \beta_k).$$

The existence and uniqueness of N_k follows from Lemma 2 (2) in [17]. Inspecting the proof of this result, we see that N_k is the middle term of a non-trivial extension of N_{k+1} by T_k, and thus T_k is certainly linked to at least one of the simple components of N_{k+1}. This conclusion also holds for $k = 1$, since $\operatorname{Ext}^1_\Lambda(N_2, T_1) \neq 0$ by Crawley-Boevey's formula (4.2). Thus, all the T_k with $1 \leq k \leq n$ are linked to R_J. Since each S_j with $j \in J$ shows up among these modules T_k, we conclude that all the S_j with $j \in J$ are linked to R_J.

Now take two different connected components J and K of $I_0 \setminus \{i\}$. By Schur's lemma,

$$\operatorname{Hom}_\Lambda(S_j, S_k) = \operatorname{Hom}_\Lambda(S_j, R_K) = \operatorname{Hom}_\Lambda(R_J, S_k) = \operatorname{Hom}_\Lambda(R_J, R_K) = 0$$

for any $j \in J$ and $k \in K$. An easy calculation based on Crawley-Boevey's formula (4.2) then shows that the Ext^1 between S_j or R_J and S_k or R_K is zero. So J and K give rise to different linkage classes.

To each connected component of $I_0 \setminus \{i\}$ corresponds thus a linkage class, formed by R_J and the S_j with $j \in J$. The sum of the dimension-vectors of these objects is $\delta - \widetilde{\alpha}_J + \sum_{j \in J} \alpha_j$, which is smaller than or equal to δ, with equality if and only if J is of type A.

At this point, we have established the desired property in the case where $F \subseteq \overline{C^s_0}$. It remains to handle the case of a general face F. Let $w \in W_0$ of minimal length such that $w^{-1}F \subseteq \overline{C_0}$. Then F is spanned by a certain spherical chamber coweight $w\,\varpi_i$, where $i \in I_0$ and $w \in W_0$ is

$I_0 \setminus \{i\}$-reduced on the right. By Theorem 5.17 (and Corollary 2.2), we have equivalences of categories

$$\mathscr{R}_{\varpi_i} \underset{\operatorname{Hom}_\Lambda(I_w,?)}{\overset{I_w \otimes_\Lambda ?}{\rightleftarrows}} \mathscr{R}_F \ .$$

We can then transfer to $\mathscr{R}_F$ the information obtained above for $\mathscr{R}_{\varpi_i}$.

What is at stake is the fact that the sum of the dimension-vectors of the simple objects in a linkage class is at most δ. As regards $\mathscr{R}_{\varpi_i}$, this sum has the form $\delta - \beta_J$, where J is a connected component of $I_0 \setminus \{i\}$ and $\beta_J = \widetilde{\alpha}_J - \sum_{j \in J} \alpha_j$. Checking the classification of root systems, we observe that β_J is a root; using Corollary 2.2, we see that $\beta_J \notin N_{w^{-1}}$. Therefore $w\beta_J$ is a positive root and $w(\delta - \beta_J) = \delta - w\beta_J$ is less than or equal to δ, as desired. □

Recall the framework of section 4.1. The graph (I, E) can be endowed with several orientations Ω (we only consider acyclic orientations). The datum of Ω gives a quiver Q, whence an Euler form $\langle \, , \, \rangle_Q$ on $\mathbb{Z}I$, defined as

$$\langle \lambda, \mu \rangle_Q = \sum_{i \in I} \lambda_i \mu_i - \sum_{a \in \Omega} \lambda_{s(a)} \mu_{t(a)}.$$

The symmetric bilinear form $(\, , \,) : \mathbb{Z}I \times \mathbb{Z}I \to \mathbb{Z}$ is then the symmetrization of $\langle \, , \, \rangle_Q$.

The imaginary root δ belongs to the kernel of $(\, , \,)$, so $\langle \delta, ? \rangle_Q$ induces a linear form on $\mathfrak{t}^*$, in other words, an element $\gamma_\Omega \in \mathfrak{t}$. For example, in type $\widetilde{A}_1$, there are two orientations

$$\Omega' : \ 0 \underset{\beta}{\overset{\alpha}{\rightrightarrows}} 1 \quad \text{and} \quad \Omega'' : \ 0 \underset{\overline{\beta}}{\overset{\overline{\alpha}}{\leftleftarrows}} 1.$$

The corresponding linear forms are the spherical chamber coweights $\gamma_{\Omega'} = \gamma'$ and $\gamma_{\Omega''} = \gamma''$ of section 7.3.

Proposition A.2 *The map $\Omega \mapsto \gamma_\Omega$ is an injection from the set of all non-cyclic orientations of (I, E) into Γ. In type $\widetilde{A}$, this map is bijective.*

Sketch of proof. We begin by studying the type $\widetilde{A}_n$. The vertices of the graph (I, E) are numbered consecutively from 0 to n and we have $\delta = \alpha_0 + \cdots + \alpha_n$. Let Ω be a non-cyclic orientation. The number $a_i = \gamma_\Omega(\pi(\alpha_i)) = \langle \delta, \alpha_i \rangle_Q$ is equal to 1, 0 or -1 depending on the number of arrows that terminate at i. When i cyclically runs over $\{0, \ldots, n\}$, a_i alternatively takes the values 1 and -1, with zeros interspersed between these values. The sum of all the a_i is $\langle \delta, \delta \rangle_Q = 0$, and the a_i cannot be all zero because Ω is not cyclic. There is thus a unique sequence of values $b_i \in \{0,1\}$, for $i \in \{0,\dots,n+1\}$, such that $a_i = b_i - b_{i+1}$ and $b_0 = b_{n+1}$. Now $\mathfrak{t}^*$ has a standard realization as a hyperplane of the vector space with basis $\{\varepsilon_i \mid 1 \le i \le n+1\}$. In this context, the b_i are the coordinates of γ_Ω in the basis (ε_i^*) dual to (ε_i). In this basis (ε_i^*), the spherical chamber coweights are the sums $\varepsilon_{i_1}^* + \cdots + \varepsilon_{i_k}^*$ with $1 \le k \le n$ and $1 \le i_1 < \cdots < i_k \le n+1$. Certainly γ_Ω matches this pattern, hence is a spherical chamber coweight. We leave to the reader the routine verifications needed to show the announced bijectivity.

Consider the following orientation in type $\widetilde{D}_n$.

$$\begin{tikzcd}
0 & & & & & & n-1 \\
& 2 \arrow[ul] \arrow[dl] & 3 \arrow[l] & \cdots \arrow[l] & n-3 \arrow[l] & n-2 \arrow[l] \arrow[ur] \arrow[dr] & \\
1 & & & & & & n
\end{tikzcd}$$

A direct calculation shows that the associated coweight is $(s_{n-1}s_n)(s_1\cdots s_{n-2})\varpi_{n-2}$. Since the graph is a tree, any orientation Ω can be obtained from this one by a sequence of reflections at sources. Noting that δ is W-invariant and using [5], Lemma 7.2, we deduce that the coweight γ_Ω is W-conjugate to ϖ_{n-2}. We omit the proof of the injectivity of the map $\Omega \mapsto \gamma_\Omega$, for it requires lengthy (but direct) calculations in coordinates.

The types $\widetilde{E}_6$, $\widetilde{E}_7$ and $\widetilde{E}_8$ are dealt with similarly. One finds that the coweights γ_Ω are all W-conjugate to the fundamental coweight corresponding to the branching point in the Dynkin diagram. For these exceptional types, we used a computer to check the injectivity of the map $\Omega \mapsto \gamma_\Omega$. $\square$

Let us fix an orientation Ω, whence a quiver Q. In [54], Ringel describes Λ-mod in terms of the category KQ-mod of finite dimensional representations of Q. More precisely, let τ denote the Auslander-Reiten translation in KQ-mod and let M be a KQ-module. Then the structures of Λ-module on M that extend the given structure of KQ-module are in natural bijection with a certain subspace $\mathcal{N}^{\tau^-}(M)$ of nilpotent elements in $\mathcal{O}^{\tau^-}(M) = \mathrm{Hom}_{KQ}(\tau^{-1}M, M)$.

Recall that indecomposable KQ-modules are classified into preprojective, preinjective and regular types. Every KQ-module M can then be written as $M = I \oplus R \oplus P$, where I, R and P are the submodules of M obtained by gathering all direct summands in a Krull-Schmidt decomposition of M which are respectively preinjective, regular, and preprojective. (The subspaces R and P depend on the choice of the Krull-Schmidt decomposition, but I and $I \oplus R$ do not; see [16], §7, Remark.)

Proposition A.3 *Let T be a Λ-module and decompose the restriction $M = T|_Q$ as a sum $M = I \oplus R \oplus P$ as above. Then $T_{\gamma_\Omega}^{\min} = I$ and $T_{\gamma_\Omega}^{\max} = I \oplus R$ (as subspaces of T).*

Proof. According to Ringel's construction, the datum of T is equivalent to the datum of M and of $f \in \mathcal{N}^{\tau^-}(M)$. By [16], §7, Lemma 3, f must map $\tau^{-1}I$ to I and $\tau^{-1}(R \oplus I)$ to $R \oplus I$, so I and $R \oplus I$ are Λ-submodules of T.

Any nonzero quotient KQ-module of I is preinjective (otherwise, we would have a nonzero map from a preinjective to a preprojective or a regular module), hence has a positive defect (§7, Lemma 2 in [16]). A fortiori, a nonzero quotient Λ-module X of I satisfies $\langle \gamma_\Omega, \underline{\dim}\, X \rangle > 0$. Therefore the Λ-module I belongs to $\mathscr{I}_{\gamma_\Omega}$.

Similarly, a nonzero KQ-submodule of $P \oplus R$ cannot have a preinjective direct summand, so has a nonpositive defect. Thus a nonzero Λ-submodule Y of T/I satisfies $\langle \gamma_\Omega, \underline{\dim}\, Y \rangle \leq 0$, and so T/I belongs to $\overline{\mathscr{P}}_{\gamma_\Omega}$.

We conclude that I is the torsion submodule of T with respect to the torsion pair $(\mathscr{I}_{\gamma_\Omega}, \overline{\mathscr{P}}_{\gamma_\Omega})$, so $T^{\min}_{\gamma_\Omega} = I$. The proof of the equality $T^{\max}_{\gamma_\Omega} = I \oplus R$ is similar. $\square$

Remarks A.4. (i) This proposition explains our choice of the notation $\mathscr{I}_\theta$, $\mathscr{R}_\theta$ and $\mathscr{P}_\theta$: when $\theta = \gamma_\Omega$, the objects of these categories are the Λ-modules whose restriction to Q are preinjective, regular or preprojective, respectively.

(ii) The abelian category $\mathcal{R}_Q$ of regular KQ-modules is well understood (see [16], §8). Indecomposable objects are grouped into tubes, and there is no nonzero morphism or extension between modules that belong to different tubes. Simple objects in $\mathcal{R}_Q$ lie at the mouth of the tubes; two simple objects are linked if and only if they belong to the same tube. The sum of the dimension-vectors of the simple objects in a tube $\mathcal{T}$ is equal to δ. A tube is called homogeneous if it has only one simple object; all but at most three tubes are homogeneous.

This description fits well with Theorem 7.7 and Proposition A.1, with $F = \mathbb{R}_{>0}\gamma_\Omega$. In fact, using Ringel's description of Λ-modules, one easily shows that the arrows in $\overline{\Omega}$ act by zero on any simple object in $\mathscr{R}_F$. Therefore the map $T \mapsto T\big|_Q$ is a bijection from $\operatorname{Irr} \mathscr{R}_F$ onto $\operatorname{Irr} \mathcal{R}_Q$. In this context, the statements (i) and (ii) in Theorem 7.7 correspond to the cases where $T\big|_Q$ belongs to an homogeneous tube or not.

Two simple objects in $\mathscr{R}_F$ are linked if their restrictions to Q are linked in $\operatorname{Irr} \mathcal{R}_Q$. Using Proposition A.1, we conclude that the bijection $T \mapsto T\big|_Q$ maps linkage classes in $\operatorname{Irr} \mathscr{R}_F$ to linkage classes in $\operatorname{Irr} \mathcal{R}_Q$.

(iii) Let us keep $F = \mathbb{R}_{>0}\gamma_\Omega$, let us choose an extending vertex 0 in I, and let $i \in I_0$ be such that γ_Ω is W_0-conjugated to ϖ_i. By item (ii) above, we have $\sum_{S \in L} \underline{\dim}\, S = \delta$ for each linkage class L in $\mathscr{R}_F$. As we observed during the course of the proof of Proposition A.1, this implies that the connected components J of $I_0 \setminus \{i\}$ are of type A. This is certainly compatible with the fact, noticed in the proof of Proposition A.2, that i is the central node of I_0 when I is of type $\widetilde{D}$ or $\widetilde{E}$.

Let $\nu \in \mathbb{N}I$ be a dimension-vector. The nilpotent variety $\Lambda(\nu)$ is a subvariety of the space of representations of the double quiver $\overline{Q}$, which itself can be identified with the cotangent of the space of representations $\mathrm{Rep}(KQ, \nu)$ of the quiver Q. It turns out that any irreducible component of $\Lambda(\nu)$ is the closure of the conormal bundle of a constructible subset $X \subseteq \mathrm{Rep}(KQ, \nu)$. The relevant subsets X were first described by Lusztig [44] in the case of a bipartite orientation Ω, and by Ringel [55] in the general case of an acyclic orientation. We now explain how this works.

Recall that an indecomposable KQ-module N is regular if and only if the Auslander-Reiten translation acts periodically on N: there is a number $p > 0$ such that $\tau^p N \cong N$. One says that a finite-dimensional KQ-module M is aperiodic if for any indecomposable regular module N in a non-homogeneous tube, the sum $\bigoplus_{i=0}^{p-1} \tau^i N$ is not a direct summand of M, where $p \geq 2$ is the τ-period of N. In addition, recall that a homogeneous tube $\mathcal{T}$ contains exactly one module in each dimension-vector $n\delta$; we denote this module by $J(\mathcal{T}, n)$. Lastly, recall that the set of homogeneous tubes is parameterized by the projective line $\mathbb{P}^1_K$, minus at most three points.

Given $\nu \in \mathbb{N}I$, let $\mathscr{S}(\nu)$ be the set of all pairs (σ, λ), where σ is an isomorphism class of aperiodic modules and λ is a partition, with the further condition $\nu = \underline{\dim}\, \sigma + |\lambda|\delta$ (see [44], §4.13). For $(\sigma, \lambda) \in \mathscr{S}(\nu)$, denote by ℓ the number of nonzero parts of λ, write $\lambda = (\lambda_1, \ldots, \lambda_\ell)$, and let $X(\sigma, \lambda)$ be the set of all points in $\mathrm{Rep}(KQ, \nu)$ isomorphic to a module of the form $M \oplus J(\mathcal{T}_1, \lambda_1) \oplus \cdots \oplus J(\mathcal{T}_\ell, \lambda_\ell)$, where $\mathcal{T}_1, \ldots, \mathcal{T}_\ell$ are distinct homogeneous tubes and M is an aperiodic module in the isomorphism class σ. Let also $\mathcal{N}(\sigma, \lambda)$ be the closure of the conormal bundle of $X(\sigma, \lambda)$. Proposition 4.14 in [44] claims that the map $(\sigma, \lambda) \mapsto \mathcal{N}(\sigma, \lambda)$ is a bijection from $\mathscr{S}(\nu)$ onto $\mathfrak{B}(\nu)$. Thanks to [55], Corollary 5.3, $\mathcal{N}(\sigma, \lambda)$ can also be described as the closure of $\{T \in \Lambda(\nu) \mid T\big|_Q \in X(\sigma, \lambda)\}$.

With all these tools in hand, one can prove that $I(\gamma_\Omega, \lambda) = \mathcal{N}(0, \lambda)$ for each partition λ, where 0 is the isomorphism class of the trivial module. Thanks to Proposition A.2, this construction provides another proof (only valid in type $\widetilde{A}$) for the results presented in section 7.4.

References

[1] C. Amiot, O. Iyama, I. Reiten, G. Todorov, *Preprojective algebras and c-sortable words,* Proc. Lond. Math. Soc. (3) **104** (2012), 513–539.

[2] J. E. Anderson, *A polytope calculus for semisimple groups,* Duke Math. J. **116** (2003), 567–588.

[3] I. Assem, *Tilting theory — an introduction,* Topics in algebra, Part 1 (Warszawa, Poland, 1988), ed. by S. Balcerzyk et al., Banach Center Publ. **26** (1990), 127–180.

[4] P. Baumann, T. Dunlap, J. Kamnitzer, P. Tingley, *Rank 2 affine MV polytopes,* to appear in Represent. Theory.

[5] P. Baumann, J. Kamnitzer, *Preprojective algebras and MV polytopes,* Represent. Theory **16** (2012), 152–188.

[6] J. Beck, *Convex bases of PBW type for quantum affine algebras,* Comm. Math. Phys. **165** (1994), 193–199.

[7] J. Beck, V. Chari, A. Pressley, *An algebraic characterization of the affine canonical basis,* Duke Math. J. **99** (1999), 455–487.

[8] J. Beck, H. Nakajima, *Crystal bases and two-sided cells of quantum affine algebras,* Duke Math. J. **123** (2004), 335–402.

[9] A. Berenstein, A. Zelevinsky, *Tensor product multiplicities, canonical bases and totally positive varieties,* Invent. Math. **143** (2001), 77–128.

[10] N. Bourbaki, *Groupes et algèbres de Lie, chapitres 4, 5 et 6,* Hermann, 1968.

[11] A. Braverman, M. Finkelberg, D. Gaitsgory, *Uhlenbeck Spaces via Affine Lie algebra,* The unity of mathematics, Progress in Mathematics vol. 244, Birkhäuser Boston, 2006, pp. 17–135.

[12] A. Braverman, D. Gaitsgory, *Crystals via the affine Grassmannian,* Duke Math. J. **107** (2001), 561–575.

[13] T. Bridgeland, *Spaces of stability conditions,* Algebraic geometry (Seattle, WA, USA, 2005), ed. by D. Abramovich et al., Proceedings of Symposia in Pure Mathematics vol. 80, part 1, Amer. Math. Soc., 2009, pp. 1–21.

[14] A. B. Buan, O. Iyama, I. Reiten, J. Scott, *Cluster structures for* 2*-Calabi-Yau categories and unipotent groups,* Compos. Math. **145** (2009), 1035–1079.

[15] P. Cellini, P. Papi, *The structure of total reflection orders in affine root systems,* J. Algebra **205** (1998), 207–226.

[16] W. Crawley-Boevey, *Lectures on representations of quivers,* lecture notes for a course given in Oxford in Spring 1992, available at `http://www.amsta.leeds.ac.uk/~pmtwc/`.

[17] W. Crawley-Boevey, *On the exceptional fibres of Kleinian singularities,* Amer. J. Math. **122** (2000), 1027–1037.

[18] W. Crawley-Boevey, J. Schröer, *Irreducible components of varieties of modules,* J. reine angew. Math. **553** (2002), 201–220.

[19] T. Dunlap, *Combinatorial representation theory of affine $\mathfrak{sl}_2$ via polytope calculus,* PhD thesis, Northwestern University, 2010.

[20] I. Frenkel, A. Malkin, M. Vybornov, *Affine Lie algebras and tame quivers,* Selecta Math. (N.S.) **7** (2001), 1–56.

[21] I. Frenkel, A. Savage, *Bases of representations of type A affine Lie algebras via quiver varieties and statistical mechanics,* Int. Math. Res. Not. **2003**, no. 28, 1521–1547.

[22] P. Gabriel, A. V. Roiter, *Representations of finite-dimensional algebras,* with a chapter by B. Keller, Encyclopaedia Math. Sci. vol. 73, Algebra VIII, Springer-Verlag, 1992.

[23] S. Gaussent, P. Littelmann, *LS galleries, the path model, and MV cycles,* Duke Math. J. **127** (2005), 35–88.

[24] I. M. Gelfand, R .M .Goresky, R. D. MacPherson, V. V. Serganova, *Combinatorial geometries, convex polyhedra, and Schubert cells,* Adv. Math. **63** (1987), 301–316.

[25] C. Geiß, B. Leclerc, J. Schröer, *Rigid modules over preprojective algebras,* Invent. Math. **165** (2006), 589–632.

[26] C. Geiß, B. Leclerc, J. Schröer, *Semicanonical basis and preprojecive algebras. II: a multiplication formula,* Compos. Math. **143** (2007), 1313–1334.

[27] C. Geiß, B. Leclerc, J. Schröer, *Kac-Moody groups and cluster algebras,* Adv. Math. **228** (2011), 329–433.

[28] K. Ito, *The classification of convex orders on affine root systems,* Commun. Algebra **29** (2001), 5605–5630.

[29] K. Ito, *Parametrizations of infinite biconvex sets in affine root systems,* Hiroshima Math. J. **35** (2005), 425–451.

[30] K. Ito, *A new description of convex bases of PBW type for untwisted quantum affine algebras,* Hiroshima Math. J. **40** (2010), 133–183.

[31] O. Iyama, I. Reiten, *Fomin-Zelevinsky mutation and tilting modules over Calabi-Yau algebras,* Amer. J. Math. **130** (2008), 1087–1149.

[32] O. Iyama, I. Reiten, 2-*Auslander algebras associated with reduced words in Coxeter groups,* Int. Math. Res. Not. **2011**, no. 8, 1782–1803.

[33] Y. Jiang, *Parametrizations of canonical bases and irreducible components of nilpotent varieties,* to appear in Int. Math. Res. Not.

[34] V. G. Kac, *Infinite dimensional Lie algebras,* 3rd ed., Cambridge University Press, 1990.

[35] J. Kamnitzer, *The crystal structure on the set of Mirković-Vilonen polytopes,* Adv. Math. **215** (2007), 66–93.

[36] J. Kamnitzer, *Mirković-Vilonen cycles and polytopes,* Ann. of Math. **171** (2010), 245–294.

[37] M. Kashiwara, *On crystal bases,* Representations of groups (Banff, AB, 1994), CMS Conf. Proc. vol. 16 (1995), pp. 155–197.

[38] M. Kashiwara, Y. Saito, *Geometric construction of crystal bases,* Duke Math. J. **89** (1997), 9–36.

[39] Y. Kimura, *Affine quivers and crystal bases,* Master thesis, University of Kyoto (Japan), 2007.

[40] A. King, *Moduli of representations of finite dimensional algebras,* Quart. J. Math. Oxford (2) **45** (1994), 515–530.

[41] G. Lusztig, *Canonical bases arising from quantized enveloping algebras,* J. Amer. Math. Soc. **3** (1990), 447–498.

[42] G. Lusztig, *Canonical bases arising from quantized enveloping algebras II,* Progr. Theoret. Phys. Suppl. **102** (1990), 175–201.

[43] G. Lusztig, *Quivers, perverse sheaves, and quantized enveloping algebras,* J. Am. Math. Soc. **4** (1991), 365–421.

[44] G. Lusztig, *Affine quivers and canonical bases,* Inst. Hautes Études Sci. Publ. Math. **76** (1992), 111–163.

[45] G. Lusztig, *Introduction to quantum groups,* Progress in Mathematics vol. 110, Birkhäuser Boston, 1993.

[46] G. Lusztig, *Braid group action and canonical bases,* Adv. Math. **122** (1996), 237–261.

[47] I. Mirković, K. Vilonen, *Geometric Langlands duality and representations of algebraic groups over commutative rings,* Ann. of Math. **166** (2007), 95–143.

[48] D. Mumford, *The red book of varieties and schemes,* Lecture Notes in Mathematics, vol. 1358, Springer-Verlag, 1999.

[49] D. Muthiah, *Double MV cycles and the Naito-Sagaki-Saito crystal,* Adv. Math. **240** (2013), 268–290.

[50] D. Muthiah, P. Tingley, *Affine PBW bases and MV polytopes in rank* 2*,* to appear in Selecta Math.

[51] S. Naito, D. Sagaki, Y. Saito, *Toward Berenstein-Zelevinsky data in affine type A, parts I and II,* Representation theory of algebraic groups and quantum groups '10 (Nagoya, Japan, 2010), ed. by S. Ariki et al., Contemp. Math. vol. 565, Amer. Math. Soc., 2012, pp. 143–216.

[52] S. Naito, D. Sagaki, Y. Saito, *Toward Berenstein-Zelevinsky data in affine type A, part III,* Symmetries, integrable systems and representations (Tokyo, Japan; Lyon, France, 2011), ed. by K. Iohara et al., Springer Proceedings in Mathematics and Statistics vol. 40, Springer-Verlag, 2013, pp. 361–402.

[53] M. Reineke, *The Harder-Narasimhan system in quantum groups and cohomology of quiver moduli,* Invent. Math. **152** (2003), 349–368.

[54] C. M. Ringel, *The preprojective algebra of a quiver,* Algebras and modules II, Eighth international conference on representations of algebras (Geiranger, Norway, 1996), ed. by I. Reiten, CMS conf. Proc. vol. 24, Amer. Math. Soc., 1998, pp. 467–480.

[55] C. M. Ringel, *The preprojective algebra of a tame quiver: the irreducible components of the module varieties,* Trends in the representation theory of finite dimensional algebras (Seattle, WA, USA, 1997), ed. by E. L. Green et al., Contemp. Math. vol. 229, Amer. Math. Soc., 1998, pp. 293–306.

[56] A. Rudakov, *Stability for an abelian category,* J. Algebra **197** (1997), 231–245.

[57] Y. Saito, *PBW basis of quantized universal enveloping algebras,* Publ. Res. Inst. Math. Sci. **30** (1994), 209–232.

[58] S. Shatz, *The decomposition and specilization of algebraic families of vector bundles,* Compos. Math. **35** (1977), 163–187.

[59] Y. Sekiya, K. Yamaura, *Tilting theoretical approach to moduli spaces over preprojective algebras,* to appear in Algebr. Represent. Theory.

Pierre Baumann, Institut de Recherche Mathématique Avancée, Université de Strasbourg et CNRS, 7 rue René Descartes, 67084 Strasbourg Cedex, France.
`p.baumann@unistra.fr`

Joel Kamnitzer, Department of Mathematics, University of Toronto, Toronto, ON, M5S 2E4 Canada.
jkamnitz@math.toronto.edu

Peter Tingley, Department of Mathematics and Statistics, Loyola University Chicago, 1032 W. Sheridan Road, Chicago, IL 60660, U.S.A.
ptingley@luc.edu