Derived Quiver Representations: Reflections, Stability and Filtrations

Isabel Longbottom

October 2021

A thesis submitted for the degree of Bachelor of Philosophy
of the Australian National University

Declaration

The work in this thesis is my own except where otherwise stated.

Isabel Longbottom

Acknowledgements

First and foremost, this project would not have been possible without the support of my supervisor, Asilata. Her dedication and the gift of her time, especially during such an unusual year, has been invaluable. Her infectious enthusiasm made even the most difficult parts of this thesis enjoyable. It would be impossible to overstate how much I have enjoyed working with her.

Thanks also go to my family. They have always been supportive of my choice to pursue mathematics, and I appreciate the conversations we have had about the ideas that excite me. Your time, and efforts to be interested, did not go unnoticed.

To my partner Martin: you made the good times better, and the bad times easier. You also helped me write the rest of these acknowledgements, so I guess it's no surprise I can't find the right words to thank you. But seriously, this year would have been so much harder without your support and encouragement every step of the way. Thanks for the proofreading, and the coffee, and for pushing me when I needed it.

Last but not least, I would like to thank the staff of the MSI for fostering such a fantastic mathematical community at the ANU. Because of you all, I have been introduced to new and interesting ideas every day for the past four years. I hope that continues long into the future.

Abstract

We study the derived category of quiver representations for an acyclic quiver Q. This can be described very concretely because the abelian category of Q-representations is hereditary, and our first occupation is to do so. We then use reflection functors to construct explicit equivalences between the derived categories of quivers with the same underlying graph, but whose representation categories are distinct. Passing to the derived category thus makes the representation theory of quivers whose underlying graphs are acyclic entirely uniform under changes of orientation.

Finally, we discuss the iterated weight filtration, which can be computed for any artinian lattice, given a weight function. This gives a refinement of the Harder–Narasimhan filtration coming from any stability condition on $\mathcal{D}^b(\operatorname{Rep} Q)$. The refinement depends only on a choice of positive weights at each vertex of Q, and is thus orientation-agnostic.

Contents

Acknowledgements **v**

Abstract **vii**

Notation and terminology **xi**

Introduction **1**

0 Projective and injective modules **3**

1 The bounded derived category **7**
1.1 Construction and basic properties . . . 8
1.2 Derived functors, with a focus on Ext . . . 14
1.3 The hereditary case . . . 21

2 Quiver representations **25**
2.1 Basic definitions . . . 25
2.2 Classification of simples and indecomposables . . . 27
2.3 The path algebra, and representations as modules . . . 32
2.4 Projective and injective indecomposables . . . 36
2.5 The Euler form . . . 43

3 (Derived) reflection functors **49**
3.1 Reflections on Rep Q . . . 50
3.2 Defining derived reflection functors . . . 60
3.3 C_x^+ is representable . . . 61
3.4 Derived reflections are equivalences . . . 65

4 Stability and filtrations **69**
4.1 Stability conditions . . . 69

4.2 The weight filtration 74
4.3 Three examples 79
4.4 Further directions 86

A Lie algebras, Dynkin diagrams and root systems 87

Bibliography 94

Notation and terminology

In the following, k is a field, and Q is a quiver, typically acyclic.

Notation

Q_0	the vertex set of Q
Q_1	the set of arrows for Q
$\operatorname{Rep} Q$	the abelian category of Q-representations
kQ	the path algebra of Q
A	a finite-dimensional k-algebra
e_x	the path in kQ of length 0 at $x \in Q_0$
S_x	the 1-dimensional simple Q-representation at $x \in Q_0$
P_x	kQe_x, the indecomposable projective Q-representation associated with $x \in Q_0$
I_x	$(e_x kQ)^*$, the indecomposable injective Q-representation associated with $x \in Q_0$
$\widetilde{M}$	the right A-module $\operatorname{Hom}_A(M, A)$, where M is a left A-module
M^*	the right A-module $\operatorname{Hom}_k(M, k)$, where M is a left A-module
Q^{op}	the quiver obtained from Q by reversing all the arrows
$\mathcal{A}$	an abelian category
$\mathcal{D}$	a triangulated category

$\mathcal{D}^b(\mathcal{A})$	the bounded derived category of $\mathcal{A}$
ι	the fully faithful functor $\mathcal{A} \to \mathcal{D}^b(\mathcal{A})$ given by inclusion in degree 0
$\mathcal{A}^\#$	a full abelian subcategory of $\mathcal{D}$, typically different from $\mathcal{A}$
$\sigma_x Q$	the quiver obtained from Q by reversing arrows incident at $x \in Q_0$, when x is a sink or source
C_x^+	the reflection functor $\operatorname{Rep} Q \to \operatorname{Rep} \sigma_x Q$ at x, when x is a sink of Q
C_x^-	the reflection functor $\operatorname{Rep} \sigma_x Q \to \operatorname{Rep} Q$ at x, when x is a sink of Q
RC_x^+	the right-derived functor of C_x^+
LC_x^-	the left-derived functor of C_x^-
RHom	the right-derived Hom functor
Ext^n	the nth cohomology of RHom
$\otimes^L$	the left-derived tensor product

Terminology

acyclic	of a quiver, containing no directed cycles
simple	having no nontrivial proper subobjects
semisimple	being a direct sum of simple subobjects
HN	Harder–Narasimhan

Introduction

Consider the bounded derived category $\mathcal{D}^b(\mathcal{A})$ of some abelian category $\mathcal{A}$. This derived category can variously be viewed as a framework for organising the homological algebra of $\mathcal{A}$, or a natural category on which to define derived functors. For us, the most important perspective is that two non-equivalent abelian categories can have equivalent derived categories — indeed, we will construct many such equivalences. Passing to the derived category thus allows us to study the relationships between abelian categories which arise from similar contexts, but are not obviously comparable. It also allows translation between geometric and algebraic situations. The canonical example of this is that the abelian category $\operatorname{Coh} X$ for a smooth projective scheme X is often derived-equivalent, but not equivalent, to the category $\operatorname{Rep}(Q, R)$ of representations of some quiver with relations.

In this thesis, we will focus on derived equivalences between the representation categories of non-isomorphic quivers. A quiver is a directed graph. We would like to know to what extent the representation theory of a quiver can be determined from its underlying graph, and how sensitive this is to the choice of edge orientations. Although some aspects of the theory, such as the simple representations, can be shown to depend only on the underlying graph, two non-isomorphic quivers never have equivalent representation categories. In particular, this applies to two different orientations of the same graph. However, if the underlying graph is acyclic then *any* two orientations are derived-equivalent, and moreover the corresponding representation categories are related by a series of tilts in their common derived category.

In Chapter 1, we introduce and develop the properties of the bounded derived category, $\mathcal{D}^b(\mathcal{A})$. We discuss the general construction of derived functors, with a specific focus on derived Hom functors, since these can be used to describe morphisms in the derived category. We then give a concrete description of $\mathcal{D}^b(\mathcal{A})$ for a class of abelian categories which includes $\mathcal{A} = \operatorname{Rep} Q$ for any acyclic quiver Q.

In Chapter 2, we present the theory of quiver representations. We give particular emphasis to the projective and injective objects in $\operatorname{Rep} Q$, since these are key to the construction of derived functors. We also prove that a quiver can be recovered from its representation category.

In Chapter 3, we use reflection functors to construct explicit derived equivalences between quivers with the same acyclic underlying graph. On the abelian level, these functors are a tensor-Hom adjoint pair, and are not equivalences. We prove that the right adjoint is representable, and give a specific representing object. We then show that the corresponding derived functors are equivalences of triangulated categories, which allows the abelian categories corresponding to all possible orientations to be viewed simultaneously as full subcategories of a common derived category.

Finally, in Chapter 4, we introduce stability conditions on the derived category of a quiver, and look at the iterated weight filtration, which can be used to give a canonical refinement of the Harder–Narasimhan filtration arising from any stability condition. This weight filtration displays similar wall-crossing behaviour as is exhibited when deforming a stability condition, with more than one iteration of the process often occurring along walls. We conclude with several examples, illustrating interesting phenomena meriting further study.

Chapter 0

Projective and injective modules

This chapter is not intended to be read sequentially. Rather, it is provided as a reference, collecting the properties of projective and injective modules which are used in proofs throughout the thesis. The reader can and should skip to Chapter 1 at this point, returning to this section as necessary.

We give a brief survey of some important properties of projective and injective modules, focusing on descriptions in terms of maps to and from such modules. Throughout, A is a unital algebra over a fixed base field k, and is not assumed to be commutative. We will from time to time require the additional hypothesis that A is hereditary. This is always true for the path algebra of an acyclic quiver. We will use, without comment, basic terminology from the theory of module categories over an algebra. Most proofs are omitted since they are standard.

Definition 0.1. A k-algebra A is *hereditary* if for any projective A-module P, every submodule of P is also projective. Equivalently, every quotient of an injective module is injective.

Lemma 0.2 (See Def 2.1.1 in [DW17])**.** *The following are equivalent for an A-module P. If one, and therefore all, of these hold then we call P projective.*

(a) The functor $\mathrm{Hom}_A(P, -)$ *is exact.*

(b) Maps out of P lift over epimorphisms. That is, for every epimorphism $f : M \to N$ and every morphism $g : P \to N$, there exists a lift $\tilde{g} : P \to M$ with $f \circ \tilde{g} = g$.

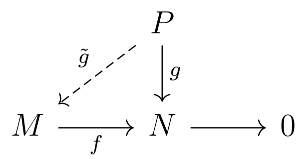

(c) Every short exact sequence with P as the last term splits.

(d) P has a complement Q such that $P \oplus Q$ is free.

The following are equivalent for an A-module I. If one, and therefore all, of these hold we call I injective. Note that these conditions are dual to those for a projective module.

(a) The functor $\mathrm{Hom}_A(-, I)$ is exact.

(b) Maps into I factor through monomorphisms. That is, for every monomorphism $f : M \to N$ and every morphism $g : M \to I$, there exists a factoring $\tilde{g} : N \to I$ with $\tilde{g} \circ f = g$.

$$\begin{array}{ccccc} 0 & \longrightarrow & M & \xrightarrow{\ f\ } & N \\ & & {\scriptstyle g}\big\downarrow & \swarrow {\scriptstyle \exists\, \tilde{g}} & \\ & & I & & \end{array}$$

(c) Every short exact sequence with I as the first term splits.

Properties (a) and (b) are equivalent in any abelian category $\mathcal{A}$, so we take these to define projective and injective objects when $\mathcal{A}$ is not necessarily a module category.

Lemma 0.3 (Lemma 6.2.1 in [DW17])**.** *If P is a projective left A-module, then $\widetilde{P} := \mathrm{Hom}_A(P, A)$ is a projective right A-module, where the action is by precomposition with the action on P. Dually, if I is an injective left-module then $\widetilde{I}$ is an injective right-module.*

Lemma 0.4. *P is a projective left A-module if and only if $P^* := \mathrm{Hom}_k(P, k)$ is an injective right A-module. Again, the action on P^* is via precomposition with the A-action on P.*

Lemma 0.5 (Lemma 6.3.2 in [DW17])**.** *Suppose A is hereditary. If P is a projective A-module and V is an A-module with no projective direct summands, then* $\mathrm{Hom}_A(V, P) = 0$.

Proof. Let $\varphi : V \to P$ be an A-module homomorphism, and $P' := \mathrm{im}\,\varphi$. Then P' is projective since A is hereditary. The induced surjection $V \twoheadrightarrow P'$ therefore splits by Lemma 0.2. Hence P' is a direct summand of V. But we assumed V has no projective summands, so $P' = 0$ and hence $\varphi = 0$. □

Lemma 0.6 (Tensor-Hom adjunction)**.** *Let R, S be rings, and M, N left modules over R, S respectively. Suppose T is an (S, R)-bimodule. Then we have an isomorphism*

$$\mathrm{Hom}_S(T \otimes_R M, N) \cong \mathrm{Hom}_R(M, \mathrm{Hom}_S(T, N))$$

natural in M and N. That is, the functors $T \otimes_R - : R\mathrm{-mod} \to S\mathrm{-mod}$ and $\mathrm{Hom}_S(T, -) : S\mathrm{-mod} \to R\mathrm{-mod}$ are an adjoint pair $(T \otimes_R -) \dashv \mathrm{Hom}_S(T, -)$.

Corollary 0.7. *For a ring R and R-module M, the covariant functor*

$$\mathrm{Hom}_R(M, -) : R\mathrm{-mod} \to \mathrm{Ab}$$

is left-exact, the contravariant functor $\mathrm{Hom}_R(-, M)$ is right-exact, and $- \otimes M$ is right-exact on $\mathrm{mod}-R$.

Proof. Right-adjoints preserve limits, and kernels are limits. Hence $\mathrm{Hom}_R(M, -)$ preserves kernels (and is additive) so is left-exact. Dually, left-adjoints preserve colimits, so the tensor product preserves cokernels, and is thus right-exact. □

Chapter 1

The bounded derived category

In this chapter, we discuss the construction and properties of the bounded derived category $\mathcal{D}^b(\mathcal{A})$, which can be formed for any abelian category $\mathcal{A}$. We then give an overview of derived functors, developing the key example of the derived Hom functor alongside the general theory. Finally, we give an explicit and constructive description of morphisms and irreducible objects in the derived category in the case where the initial abelian category is hereditary, using the derived Hom functor. We are particularly interested in the hereditary case because quiver categories are hereditary, and this is the application we will focus on in later chapters.

This chapter is quite technical, and many of the definitions are difficult to motivate without the context of chapters 2 and 3. Motivation for defining the derived category comes chiefly from two sources: first, a systematic construction of derived functors, which are a useful tool in homological algebra; and second, the derived category allows us to unify the theory of some related constructions which look quite different on the abelian level. We will see concrete examples of this in Example 1.22, Remark 2.31, and Chapter 3 — indeed, this is the purpose of introducing reflections in Chapter 3 — but the essential idea is as follows. Within $\mathcal{D}^b(\mathcal{A})$, one often finds many distinct abelian categories occurring as full subcategories, in addition to $\mathcal{A}$. Such categories may seem unrelated on the abelian level, but using the structure of the derived category we can translate between them.

1.1 Construction and basic properties

In this section, we construct the bounded derived category $\mathcal{D}^b(\mathcal{A})$ for an abelian category $\mathcal{A}$, and discuss its elementary properties. We assume basic knowledge of abelian categories. The reader unfamiliar with the general definitions can choose to think of $\mathcal{A}$ as the category of modules over a finite-dimensional algebra, or see Chapter VIII of [Mac88]. Our treatment is loosely based on Chapter 1 of [KS13] and Sections 1-2 of [Kra07], although this construction is standard and many sources give a similar account. We fix an abelian category $\mathcal{A}$ throughout. First consider the following two categories.

- $\mathcal{C}^b(\mathcal{A})$ is the category of bounded cochain complexes, with objects

  $$\dots \xrightarrow{d^{i-1}} X^i \xrightarrow{d^i} X^{i+1} \xrightarrow{d^{i+1}} X^{i+2} \rightarrow \dots$$

  whose differentials satisfy $d^{i+1} \circ d^i = 0$ for every $i \in \mathbb{Z}$, and such that $H^i(X) = 0$ for all but finitely many i. The cohomology groups $H^i(X) = \ker(d^i)/\operatorname{im}(d^{i-1})$ are defined in the usual way. A morphism $f : X \rightarrow Y$ is a cochain map, consisting of $f_i : X^i \rightarrow Y^i$ such that $f_{i+1} \circ d^i_X = d^i_Y \circ f_i$ for each $i \in \mathbb{Z}$.

- $\mathcal{K}^b(\mathcal{A})$ is the bounded homotopy category. Its objects are the same as those of $\mathcal{C}^b(\mathcal{A})$, and its Hom sets are obtained from those of $\mathcal{C}^b(\mathcal{A})$ by identifying morphisms which are chain-homotopic to each other. That is, $f \sim g : X \rightarrow Y$ if there exists a chain homotopy, $h^i : X^i \rightarrow Y^{i-1}$ such that $f_i - g_i = h^{i+1} \circ d^i_X + d^{i-1}_Y \circ h^i$ for all $i \in \mathbb{Z}$. Note that in the following diagram, only the solid arrows commute.

  $$\begin{array}{ccccc} X^{i-1} & \xrightarrow{d_X^{i-1}} & X^i & \xrightarrow{d_X^i} & X^{i+1} \\ {\scriptstyle f^{i-1}-g^{i-1}}\downarrow & \swarrow{\scriptstyle h^i} & \downarrow{\scriptstyle f^i-g^i} & \swarrow{\scriptstyle h^{i+1}} & \downarrow{\scriptstyle f^{i+1}-g^{i+1}} \\ Y^{i-1} & \xrightarrow[d_Y^{i-1}]{} & Y^i & \xrightarrow[d_Y^i]{} & Y^{i+1} \end{array}$$

  Then $\operatorname{Hom}_{\mathcal{K}^b(\mathcal{A})}(X,Y) := \operatorname{Hom}_{\mathcal{C}^b(\mathcal{A})}(X,Y)/\sim$.

There are unbounded versions of these two categories, given by omitting the requirement that any object have finitely many nonzero cohomology pieces. The bounded versions are full subcategories, and we will work only with these.

To avoid further technicalities, we assume that infinite resolutions or other unbounded complexes do not arise. This will be true in all situations that we study concretely.

There is a natural autoequivalence of $\mathcal{C}^b(\mathcal{A})$ given by translation by an integer, denoted $[n]$. This is defined on objects by $(X[n])^i := X^{i+n}$, $d^i_{X[n]} := (-1)^n d^{i+n}_X$ and similarly on morphisms. For $n > 0$, $[n]$ is a left shift, and for $n < 0$ a right shift. The translation functors descend to $\mathcal{K}^b(\mathcal{A})$ where they remain autoequivalences. Cohomology is functorial on $\mathcal{C}^b(\mathcal{A})$, and similarly descends to $\mathcal{K}^b(\mathcal{A})$. Moreover, cohomology and translation are compatible, with $H^i \circ [n] = H^{i+n}$.

The category $\mathcal{C}^b(\mathcal{A})$ is abelian, but $\mathcal{K}^b(\mathcal{A})$ is generally not because not all kernels and cokernels exist[1]. An important construction on $\mathcal{K}^b(\mathcal{A})$ is the following.

Definition 1.1 (Mapping cone)**.** Any $f \in \mathrm{Hom}_{\mathcal{C}^b(\mathcal{A})}(X, Y)$ has a corresponding *mapping cone* complex $M(f)$, defined by

$$(M(f))^n := X^{n+1} \oplus Y^n$$

with differentials

$$d^n : X^{n+1} \oplus Y^n \to X^{n+2} \oplus Y^{n+1} := \begin{bmatrix} d^n_{X[1]} & 0 \\ f_{n+1} & d^n_Y \end{bmatrix}.$$

The pointwise inclusion $Y \hookrightarrow M(f)$ and pointwise projection $M(f) \twoheadrightarrow X[1]$ are cochain maps.

The mapping cone descends to $\mathcal{K}^b(\mathcal{A})$. For any morphism $f \in \mathrm{Hom}_{\mathcal{K}^b(\mathcal{A})}(X, Y)$ we then have a sequence

$$Y \xrightarrow{g} M(f) \xrightarrow{h} X[1] \xrightarrow{f[1]} Y[1] \tag{1.1}$$

which continues in either direction via translations of these maps. In $\mathcal{K}^b(\mathcal{A})$, this sequence has the property that the composition of any two adjacent maps vanishes. We call a sequence in $\mathcal{K}^b(\mathcal{A})$ that is isomorphic to one of this form a *distinguished triangle*, and the collection of such triangles makes $\mathcal{K}^b(\mathcal{A})$ a triangulated category (see Sections 1.4-1.5 of [KS13] for the axioms (TR1)-(TR5)).

Theorem 1.2. *Given a sequence as in* (1.1)*, there is an isomorphism* $\varphi : X[1] \to M(g)$ *such that the diagram*

[1]This is because a morphism in a triangulated category has a kernel if and only if it is a split kernel, and similarly for cokernels. So, to find a counterexample, we need only construct a map in $\mathcal{K}^b(\mathcal{A})$ whose kernel cannot be split.

$$\begin{array}{ccccccc} Y & \xrightarrow{g} & M(f) & \xrightarrow{h} & X[1] & \xrightarrow{-f[1]} & Y[1] \\ \downarrow{\scriptstyle \mathrm{id}} & & \downarrow{\scriptstyle \mathrm{id}} & & \downarrow{\scriptstyle \varphi} & & \downarrow{\scriptstyle \mathrm{id}} \\ Y & \xrightarrow{g} & M(f) & \xrightarrow{g'} & M(g) & \xrightarrow{h'} & Y[1] \end{array}$$

commutes, where g', h' are the inclusion and projection respectively for $M(g)$.

The theorem implies that any rotation of a distinguished triangle in $\mathcal{K}^b(\mathcal{A})$ is a distinguished triangle, that is if $X \to Y \to Z \to X[1]$ is distinguished, then so is $Y \to Z \to X[1] \to Y[1]$. This is (TR3), one of the axioms for a triangulated category.

Distinguished triangles in a triangulated category play a similar role to short exact sequences in an abelian category. For example, the snake lemma and the splitting lemma for abelian categories both have analogues[2] for distinguished triangles in a triangulated category. In particular if any of the maps in a distinguished triangle is 0 then the triangle splits.

Definition 1.3. A morphism f in $\mathcal{K}^b(\mathcal{A})$ is called a *quasi-isomorphism* if the induced maps on cohomology are isomorphisms. That is, $H^i(f)$ is an isomorphism for each $i \in \mathbb{Z}$.

One motivation for constructing the derived category is as follows. Any short exact sequence in $\mathcal{A}$ induces a long exact sequence on cohomology by the snake lemma for abelian categories. Similarly, any distinguished triangle in $\mathcal{K}^b(\mathcal{A})$ induces a long exact sequence on cohomology. Moreover, a short exact sequence $0 \to X \xrightarrow{f} X \to Z \to 0$ in $\mathcal{C}^b(\mathcal{A})$ induces a triangle (not necessarily distinguished!) of the form

$$X \to Y \to Z \to X[1] \in \mathcal{K}^b(\mathcal{A}). \tag{1.2}$$

By triangle, we mean a sequence of three maps, not necessarily arising from a mapping cone. Comparing (1.2) to the distinguished triangle $X \to Y \to M(f) \to X[1]$, one can construct a natural map $q : M(f) \to Z$ which gives a morphism between the two triangles (taking the other maps to be the identity). Moreover, q is a quasi-isomorphism. In this sense, the triangle (1.2) fails to be distinguished to the extent that q fails to be an isomorphism.

The derived category is built from the homotopy category by inverting quasi-isomorphisms. This in particular means that any short exact sequence in $\mathcal{C}^b(\mathcal{A})$ induces a corresponding distinguished triangle in the derived category.

[2]See Section 2.5 of [Kra07] for the snake lemma on $\mathcal{K}(\mathcal{A})$, and Section 1.2 of [Nee14] for the splitting lemma on a general triangulated category.

Definition 1.4 (The bounded derived category)**.** The *bounded derived category* $\mathcal{D}^b(\mathcal{A})$ has the same objects as $\mathcal{C}^b(\mathcal{A})$ and $\mathcal{K}^b(\mathcal{A})$, namely bounded cochain complexes. A morphism (or *roof*) $X \to Y$ in $\mathcal{D}^b(\mathcal{A})$ consists of a pair of morphisms $(q : Z \to X, f : Z \to Y)$ in $\mathcal{K}^b(\mathcal{A})$ where q is a quasi-isomorphism, considered up to equivalence. The equivalence relation on morphisms identifies $(q' : B \to X, f' : B \to Y)$ and $(q'' : C \to X, f'' : C \to Y)$ if there exists a third morphism $(q : A \to B, f : A \to C)$ such that the diagram

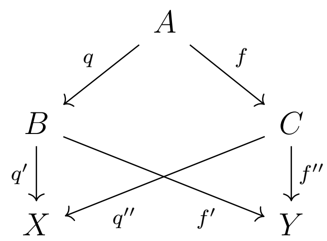

commutes. We will often denote morphisms in the derived category by $[q, f]$.

With morphisms as above, it is not immediately clear how they are composed. Essentially, this is done by building a larger roof using the mapping cone construction.

Definition 1.5 (Definition/Proposition)**.** Let $[q : X \to A, f : X \to B]$ and $[r : Y \to B, g : Y \to C]$ be morphisms in $\mathcal{D}^b(\mathcal{A})$. The maps f, r together define a map $X \oplus Y \to B$ in $\mathcal{K}^b(\mathcal{A})$. Taking the mapping cone of this morphism, we get a distinguished triangle

$$M(f,r)[-1] \overset{h}{\to} X \oplus Y \overset{(f,r)}{\to} B \to M(f,r).$$

Let $Z := M(f,r)[-1]$. We have maps $h_1 := \mathrm{pr}_1 \circ h, h_2 := -\mathrm{pr}_2 \circ h$ from Z to X and Y respectively. Then we get a roof

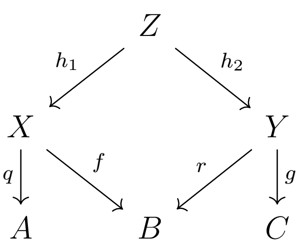

This diagram commutes, h_1 is a quasi-isomorphism, and we define the composition $[r, g] \circ [q, f] := [q \circ h_1, g \circ h_2]$. This is well-defined under the equivalence relation on morphisms.

We think of the morphism $[q, f]$ in $\mathcal{D}^b(\mathcal{A})$ as "$f \circ q^{-1}$". In particular for a quasi-isomorphism $q : X \to Y$ in $\mathcal{K}^b(\mathcal{A})$, the roof $[q, \mathrm{id}_X]$ is a morphism $Y \to X$ representing q^{-1}. Indeed, $[q, \mathrm{id}_X]$ and $[\mathrm{id}_X, q]$ are inverse isomorphisms in $\mathcal{D}^b(\mathcal{A})$ is such case. As for $\mathcal{C}^b(\mathcal{A})$ and $\mathcal{K}^b(\mathcal{A})$, translation is an autoequivalence of $\mathcal{D}^b(\mathcal{A})$. There is also a functor $\iota : \mathcal{A} \to \mathcal{D}^b(\mathcal{A})$ which sends an object $A \in \mathcal{A}$ to the complex with A in degree zero and zeroes elsewhere. In Theorem 1.20 we will show that ι is fully faithful, realising $\mathcal{A}$ as a full subcategory of $\mathcal{D}^b(\mathcal{A})$.

Remark 1.6 (The derived category is a localisation)**.** The definition of a morphism in the derived category may, on first reading, appear strange and unmotivated. For the reader comfortable with the notion of ring localisation, the derived category should be viewed as the localisation of the homotopy category at the set of quasi-isomorphisms. That is, $\mathcal{D}^b(\mathcal{A})$ is obtained from $\mathcal{K}^b(\mathcal{A})$ merely by formally inverting all quasi-isomorphisms. From this perspective, the equivalence relation on morphisms in the derived category can be re-interpreted as the familiar equivalence relation on the localisation of a ring. Indeed, there is a universal property for the derived category analogous to that for ring localisation.

Practically, what this means is that any construction on $\mathcal{K}^b(\mathcal{A})$ that respects quasi-isomorphisms descends to the derived category, and any construction on $\mathcal{A}$ that respects chain homotopy and quasi-isomorphisms descends to the derived category.

One can extend the definition of cohomology functors to the derived category by setting $H^i_{\mathcal{D}^b(\mathcal{A})}([q, f]) = H^i_{\mathcal{K}^b(\mathcal{A})}(f) \circ H^i_{\mathcal{K}^b(\mathcal{A})}(q)^{-1}$ for morphisms. This is motivated by the fact that we are thinking of $[q, f]$ as $f \circ q^{-1}$.

An object of $\mathcal{D}^b(\mathcal{A})$ whose cohomology pieces are all trivial is necessarily the zero object. However, the same is not true for morphisms — there can exist nonzero morphisms f in $\mathcal{D}^b(\mathcal{A})$ with $H^i(f) = 0$ for all $i \in \mathbb{Z}$.

In certain special cases, morphisms in $\mathcal{D}^b(\mathcal{A})$ are the same as morphisms in $\mathcal{K}^b(\mathcal{A})$. The dual to the following Proposition also holds, for a complex of injectives in the second coordinate.

Proposition 1.7 (see [Kra07] Section 1.5)**.** *Let P be a bounded complex, whose terms are projective in $\mathcal{A}$. Let X be a bounded complex.*

(1) Any quasi-isomorphism $q : P' \to P$ has a right-inverse $q' : P \to P'$ such that $q \circ q' = \mathrm{id}_P$ in $\mathcal{K}^b(\mathcal{A})$.

(2) The map $\mathrm{Hom}_{\mathcal{K}^b(\mathcal{A})}(P, X) \to \mathrm{Hom}_{\mathcal{D}^b(\mathcal{A})}(P, X), f \mapsto [\mathrm{id}_P, f]$ is an isomorphism.

Proof. For (1), complete q to a distinguished triangle by taking the mapping cone $M(q)$, and consider the induced long exact sequence in cohomology. All the maps q_i^* in this sequence are isomorphisms, so exactness implies that $H^i(M(q)) = 0$ for every i, and thus $M(q) = 0$. Then the distinguished triangle $P' \xrightarrow{q} P \to M(q) = 0 \to P'[1]$ splits, giving a right-inverse for q.

Now for (2), given a roof $[q : P' \to P, f : P' \to X]$, we know q has a right-inverse q'. Then $[q, f]$ is equivalent to $[q \circ q', f \circ q']$ which is in the image of the map since $q \circ q' = \mathrm{id}_P$.

If $f_1, f_2 \in \mathrm{Hom}_{\mathcal{K}(\mathcal{A})}(P, X)$ are such that $[\mathrm{id}_P, f_1]$ and $[\mathrm{id}_P, f_2]$ are equivalent roofs, then there is a quasi-isomorphism $q : P' \to P$ with $f_1 \circ q = f_2 \circ q$. Since q has a right-inverse, it is right-cancellative, and hence $f_1 = f_2$. □

The following results summarise a few basic facts about morphisms in the derived category.

Proposition 1.8. *Let $X \in \mathcal{D}^b(\mathcal{A})$ with i_0 minimal and i_1 maximal such that $H^i(X) \neq 0$. Then there is a morphism $f_0 : H^{i_0}(X)[-i_0] \to X$ and a morphism $f_1 : X \to H^{i_1}(X)[-i_1]$, each inducing the identity map on the appropriate cohomology grade. In particular, $f_0, f_1 \neq 0$.*

Proposition 1.9. *Finite products and coproducts in $\mathcal{D}(\mathcal{A})$ and $\mathcal{D}^b(\mathcal{A})$ are inherited from $\mathcal{A}$. That is, for complexes B and C, the complex $B \oplus C$ defined termwise satisfies the universal property of both the product and coproduct in $\mathcal{D}^b(\mathcal{A})$.*

Proof sketch. From the localisation perspective, this follows from the fact that the direct sum of two quasi-isomorphisms is a quasi-isomorphism. □

Proposition 1.10. *Suppose C, D are objects of $\mathcal{D}^b(\mathcal{A})$ such that $H^i(C) = 0$ for all $i > 0$ and $H^i(D) = 0$ for all $i < 1$. Then*

$$\mathrm{Hom}_{\mathcal{D}(\mathcal{A})}(C, D) = 0.$$

That is, there are no maps in $\mathcal{D}^b(\mathcal{A})$ that go up at least one grade, between complexes whose cohomology is concentrated in distinct grades.

Since translation is an autoequivalence, the same holds for any two complexes C, D where all the cohomology of C lives in strictly lower grades than the cohomology of D.

Proof sketch. Since C satisfies $H^i(C) = 0$ for all $i > 0$, we can find a chain complex C_* that is isomorphic to C in $\mathcal{D}^b(\mathcal{A})$ and such that $C_*^i = 0$ for all $i > 0$. We do this by taking what is called a truncation. Define

$$C_*^i = \begin{cases} C^i & \text{for } i < 0, \\ 0 & \text{for } i > 0, \\ \ker(d^0) & \text{for } i = 0 \end{cases}$$

This preserves the cohomology of C, and the natural chain map $C_* \hookrightarrow C$ is a quasi-isomorphism. Similarly, we can truncate D from below at degree 1 to get the complex $D_* \simeq D$ defined by

$$D_*^i = \begin{cases} D^i & \text{for } i > 1, \\ 0 & \text{for } i < 1, \\ \operatorname{coker}(d^0) & \text{for } i = 1. \end{cases}$$

Then a morphism $\varphi : C \to D$ in $\mathcal{D}^b(\mathcal{A})$ can be represented by a roof $[q : X \to C, f : X \to D]$. Since q is a quasi-isomorphism, X also has no cohomology above zero, so can be truncated to give a representative complex X_* with $X_*^i = 0$ for $i > 0$ as well. The map f induces $f_* : X_* \to D_*$, and we obtain a new roof $[q : X_* \to C_*, f_* : X_* \to D_*]$ representing φ. But then f_* is a map in the homotopy category between chain complexes which have no nonzero terms in common grades. Hence $f_* = 0$ and so $\varphi = 0$. □

We next shift our focus to discuss (right- and left-) derived functors, which are simultaneously motivation for constructing the derived category and key tools for its study. We will give a framework for the general construction of a derived functor, while using the right-derived Hom functor, commonly denoted Ext, as a worked example.

1.2 Derived functors, with a focus on Ext

The motivation for derived functors is homological in nature. Suppose we have abelian categories $\mathcal{A}, \mathcal{B}$ and an additive functor $F : \mathcal{A} \to \mathcal{B}$. Let us assume for simplicity that F is covariant. For a short exact sequence $0 \to A \to B \to C \to 0$ in $\mathcal{A}$, we are interested in the sequence

$$0 \to FA \to FB \to FC \to 0. \tag{1.3}$$

Unfortunately, we cannot hope that this sequence will be exact for every F. Indeed, in the general case we can put essentially no restrictions on the possible cohomology groups of this sequence. Let us now restrict to the case where F is left (or right) exact. Then (1.3) is exact if we remove the zero on the right (or left respectively). So we have a sequence

$$0 \to FA \to FB \to FC$$

which we would like to continue to a long exact sequence by adding terms on the right. Derived functors give a canonical, functorial method for doing so. Specifically, if $RF : \mathcal{D}^b(\mathcal{A}) \to \mathcal{D}^b(\mathcal{B})$ is the right-derived functor of F then we have an exact sequence

$$0 \to R^0FA \to R^0FB \to R^0FC \to R^1FA \to R^1FB \to R^1FC \to \dots \tag{1.4}$$

where $R^iF := H^i \circ RF \circ \iota$. Since $R^0F = F$, this is as desired. Similar sequences arise when F is instead right-exact. Moreover, RF is compatible with translation.

While developing the general theory of derived functors in this section, we will pay particular attention to the right-derived Hom functor, which characterises maps in the derived category.

If the functor F were exact on $\mathcal{A}$, then we would immediately obtain a new functor on $\mathcal{D}^b(\mathcal{A})$. This is because if F is exact on $\mathcal{A}$, it preserves quasi-isomorphisms, and so the functor on $\mathcal{D}^b(\mathcal{A})$ given by applying F to each object and map in a chain complex is well-defined. However, if F is not exact, this is not well-defined. To see this explicitly, pick a short exact sequence $0 \to A \to B \to C \to 0$ in $\mathcal{A}$ so that the sequence obtained by applying F is not exact. Then the complex $0 \to A \to B \to C \to 0$ in $\mathcal{D}^b(\mathcal{A})$ is isomorphic to the zero complex, since it has trivial cohomology. Hence F takes different values on these two representative complexes for the zero object.

From this perspective, we can think of derived functors as a construction allowing a functor F that is not exact to descend to the derived category. The key to this construction is finding a special class of objects on which F is exact, and in terms of which a resolution of any object of A exists. We can then replace an object of $\mathcal{A}$ by a resolution in terms of these special objects, and define RF by applying F to this resolution. This is well-defined in the same way that an exact functor descends immediately to $\mathcal{D}^b(\mathcal{A})$.

Definition 1.11. Given a covariant, additive, left or right exact functor F, an *adapted class* for F is a collection Γ of objects in $\mathcal{A}$ such that:

(i) $0 \in \Gamma$, and Γ is closed under finite products (which are the same as finite coproducts, since $\mathcal{A}$ is abelian), and closed under isomorphisms;

(ii) if F is right exact, then we require that for every object $A \in \mathcal{A}$ there exists an epimorphism from some element of Γ to A, and if F is left exact we dually require existence of a monomorphism from any object of $\mathcal{A}$ to some element of Γ;

(iii) the functor F is exact on Γ, meaning it is exact on any short exact sequence all of whose terms come from Γ.

The conditions corresponding to left and right exactness in (ii) are swapped if F is contravariant. Note that this is not a completely standardised definition; conditions vary in the literature.

If (i) and (ii) hold for a collection Γ without the context of a functor F, we say Γ is a *sufficient class* for $\mathcal{A}$, or we say there are *enough* objects in Γ. For example, $R-$mod always has enough projectives and enough injectives.

Let us unpack the above definition in the context of the Hom functor. We recall that $\mathrm{Hom}_{\mathcal{A}}(X, -)$ is covariant left-exact while $\mathrm{Hom}_{\mathcal{A}}(-, X)$ is contravariant left-exact. Our adapted class for the covariant Hom functor will be the injective objects of $\mathcal{A}$, and for the contravariant Hom will be projective objects. These are common choices of adapted class because any short exact sequence ending with a projective or beginning with an injective object automatically splits. This means (iii) holds for any additive functor on projective or injective objects. Similarly, (i) is clear for injectives and projectives. The only condition which is not always satisfied for projectives is (ii). We will restrict ourselves to the case where $\mathcal{A} = R-$mod for some ring R, or similarly a category of right-modules, so that this condition holds for both injectives and projectives.

Remark 1.12. If $\mathcal{A}$ has enough projectives, then they are an adapted class for any (covariant right-exact or contravariant left-exact) additive functor. A dual statement holds for injectives. So, we need only look for more exotic adapted classes when this is not the case.

However, not all abelian categories of interest have enough injectives and projectives. For example, $\mathrm{Coh}\, X$ generally has very few injective objects, and when X is a smooth projective variety $Q\,\mathrm{Coh}\, X$ never has enough projectives (unless X is a collection of points). For a proof when $X = \mathbb{P}^1_R$, see [EEGRO04], Corollary 2.3. The general case is similar.

Proposition 1.13. *Let Γ be a sufficient class for $\mathcal{A}$. Then for any object $A \in \mathcal{A}$, there exists a (possibly infinite) resolution of A by objects in Γ, called a Γ-resolution for A.*

Proof. Consider a left-handed adapted class, so there exists a surjection $P_0 \to A$ for some $P_0 \in \Gamma$. We want to find an exact sequence of the form

$$\ldots \to P_i \to P_{i-1} \to \ldots \to P_1 \to P_0 \to A \to 0$$

where each $P_i \in \Gamma$. Such a sequence exists by induction. Once we have constructed the resolution up to $d_i : P_i \to P_{i-1}$, we know by assumption there exists a surjection $\widetilde{d}_{i+1} : P_{i+1} \twoheadrightarrow \ker(d_i)$ with $P_{i+1} \in \Gamma$, so let d_{i+1} be the composition of this map with the inclusion $\ker(d_i) \hookrightarrow P_i$. □

We define the *length* of such a resolution to be the minimum index i such that $P_{i+1} = 0$, that is the index of the last non-zero term in the resolution. This could be infinite; however, in many interesting cases there are known finite resolutions which we can describe explicitly. We will see in Chapter 2 that for the category of representations of an acyclic quiver, there is a canonical length one projective resolution and a dual length one injective resolution, for any object.

Proposition 1.13 is not quite enough to define RF in general, since we only have Γ-resolutions of objects in $\mathcal{A}$. However, such resolutions can be combined to give a Γ-resolution of any bounded complex representing an object of $\mathcal{D}^b(\mathcal{A})$.

Proposition 1.14. *If Γ is a sufficient class for $\mathcal{A}$, then for any object C of $\mathcal{D}^b(\mathcal{A})$ there exists a complex $\mathcal{I}$ whose terms lie in Γ and such that $\mathcal{I} \cong C$ in $\mathcal{D}^b(\mathcal{A})$.*

This is proved in greater generality in Section 4.2 of [Mur06]. We omit the details of the proof, since we can avoid this construction in the case of hereditary abelian categories, including quiver categories. The idea is to take an adapted resolution of each term of the complex, and lift the differentials to give a double complex. Then the desired resolution is formed by taking diagonal sums.

If F has an adapted class, we can hope to define its derived functor.

Definition 1.15. Let $F : \mathcal{A} \to \mathcal{B}$ be additive, covariant, and left exact, with an adapted class Γ. Given $C \in \mathcal{D}^b(\mathcal{A})$, find a bounded complex $\mathcal{I}$ whose terms lie in Γ, and such that $C \simeq \mathcal{I}$ in $\mathcal{D}^b(\mathcal{A})$. Then $RF(C) := F\mathcal{I}$ and $R^iF(C) := H^i(RF(C)) = H^i(F\mathcal{I})$.

There is of course an analogous version of this definition when F is right-exact, or contravariant. Moreover, if $A \in \mathcal{A}$ then this is equivalent to taking a Γ-resolution $0 \to A \to \mathcal{I} \to 0$. By left-exactness of F, such a sequence remains exact at A after applying F. In particular, this means that $H^0 RF\iota(A) = H^0(F\mathcal{I}) \cong FA$ when $A \in \mathcal{A}$.

There is an important check we must conduct to determine that derived functors are well-defined — namely, we must show that if $\mathcal{I} \simeq \mathcal{I}'$ in $\mathcal{D}^b(\mathcal{A})$ with $\mathcal{I}, \mathcal{I}'$ both Γ-complexes, then $F\mathcal{I} \simeq F\mathcal{I}'$ in $\mathcal{D}^b(\mathcal{A})$. Now, a morphism in $\mathcal{D}^b(\mathcal{A})$ is a quasi-isomorphism if and only if its mapping cone is 0 in $\mathcal{D}^b(\mathcal{A})$. Additive functors preserve mapping cones, so given a quasi-isomorphism $q : \mathcal{I} \to \mathcal{I}'$ we know $F(M(q)) \simeq M(F(q))$. Moreover, $M(q)$ is a Γ-complex, since $\mathcal{I}, \mathcal{I}'$ were Γ-complexes and Γ is closed under finite direct sums. Then because $M(q)$ has trivial cohomology and F is exact on Γ, it follows that $0 \simeq F(M(q)) \simeq M(F(q))$ in $\mathcal{D}^b(\mathcal{B})$, so $F(q) : F\mathcal{I} \to F\mathcal{I}'$ is a quasi-isomorphism. In particular, derived functors may be computed on any choice of adapted resolution.

Proposition 1.16. *Let $F : \mathcal{A} \to \mathcal{B}$ have right-derived functor RF, constructed as in Definition 1.15. Then we have a long exact sequence*

$$0 \to R^0FA \to R^0FB \to R^0FC \to R^1FA \to R^1FB \to R^1FC \to \dots$$

for any short exact sequence $0 \to A \to B \to C \to 0$ in $\mathcal{A}$.

Proof. Take Γ-resolutions $\mathcal{I}_A, \mathcal{I}_b, \mathcal{I}_C$ of A, B, C respectively. A diagram chase shows that the maps in the short exact sequence extend along these resolutions, giving a short exact sequence of chain complexes $0 \to \mathcal{I}_A \to \mathcal{I}_B \to \mathcal{I}_C \to 0$.

Since Γ is an adapted class, F is exact on sequences in Γ, and so the rows remain exact after applying F to this double complex. Hence $0 \to F(\mathcal{I}_A) \to F(\mathcal{I}_B) \to F(\mathcal{I}_C) \to 0$ is a short exact sequence of chain complexes. Then the snake lemma gives the desired long exact sequence in cohomology. □

In the case of Hom, we have two different sets of right derived functors $R^i \operatorname{Hom}(A, B)$ for any pair of objects $A, B \in \mathcal{A}$, one arising from taking a projective resolution of A and the other from an injective resolution of B. We know that these are both well-defined, but we have not yet shown that they agree. We will do this only in the case where $\mathcal{A}$ has enough injectives and enough projectives, in Theorem 1.20. A more direct proof is given in [Wei95], Theorem 2.7.6.

Definition 1.17. We use the notation $\operatorname{Ext}^i(A, B) := R^i \operatorname{Hom}(A, B)$ for the cohomology pieces of $\operatorname{RHom}(A, B)$. They can be computed by taking either a projective resolution in the first coordinate, or an injective resolution in the second.

The fact that two complexes with the same cohomology are not necessarily isomorphic in $\mathcal{D}^b(\mathcal{A})$ seems to be a source of considerable complication. For example, for $A \in \mathcal{A}$ the complex

$$\bigoplus_{i \geq 0} R^i F(A)[-i]$$

has cohomology isomorphic to that of $RF(A)$, but is not necessarily an isomorphic object in $\mathcal{D}^b(\mathcal{A})$. If it were, we could give a simpler definition of RF by splitting a complex up into its cohomology pieces (which live in translations of $\mathcal{A}$) and only defining the functors R^iF. This would circumvent the need for Proposition 1.14 when performing computations. The hypothesis on $\mathcal{A}$ required to allow this simplification is that it be hereditary, meaning all the higher Ext groups vanish. This assumption has convenient consequences, including that every object of $\mathcal{D}^b(\mathcal{A})$ is simply isomorphic to the direct sum of its cohomology pieces, and will allow us to describe $\mathcal{D}^b(\mathcal{A})$ very concretely. However, even when $\mathcal{A}$ is not hereditary, we have a cohomology filtration for any object of $\mathcal{D}^b(\mathcal{A})$.

Proposition 1.18. *Given a complex $X \in \mathcal{D}^b(\mathcal{A})$ there exists a sequence of distinguished triangles*

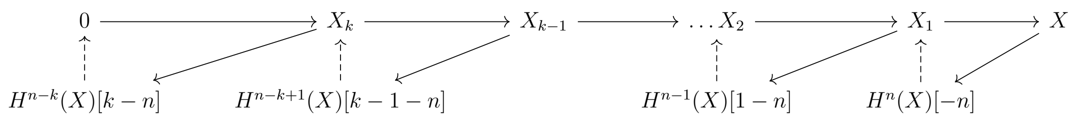

That is, X may be a nontrivial extension of its cohomology, rather than the direct sum.

The proof is very similar to the proof of Theorem 1.23, except that in the nonhereditary case the distinguished triangles that arise need not split. Here $n > n-k$ are the maximum and minimum nonzero cohomology degrees respectively. The map $X \to H^k(X)[-k]$ is f_1 from Proposition 1.8, and we proceed by induction. After removing some triangles from the right of the diagram, what remains is the cohomology filtration of the rightmost object.

Before discussing the hereditary case in detail, we finish this section by giving two new perspectives on Ext^i.

Definition 1.19. An *n-extension* of A by B in $\mathcal{A}$ is an exact sequence

$$0 \to B \to E_n \to \ldots \to E_2 \to E_1 \to A \to 0.$$

Two n-extensions are equivalent if there exists a chain map[3] between them which is the identity on A and B. Note that a 1-extension is a short exact sequence.

In the next proof, we will need a result for $\mathcal{K}^b(\mathcal{A})$. If $X, Y \in \mathcal{C}^b(\mathcal{A})$, there is a natural isomorphism $H^n(\mathrm{Hom}^\bullet(X,Y)) = \mathrm{Hom}_{\mathcal{K}^b(\mathcal{A})}(X, Y[n])$, where $\mathrm{Hom}^\bullet(X,Y)$ is the total Hom-complex defined by $\mathrm{Hom}^n(X,Y) = \prod_{k\in\mathbb{Z}} \mathrm{Hom}_{\mathcal{A}}(X^k, Y^{k+n})$. This is elementary to check, because $f = (f_k)_{k\in\mathbb{Z}}$ is a cocycle in $\mathrm{Hom}^n(X,Y)$ if and only if $f : X \to Y[n]$ is a chain map, and being a coboundary means f is homotopic to 0.

Theorem 1.20. *The collection of n-extensions of A by B up to equivalence is isomorphic to $\mathrm{Ext}^n(A,B)$. Moreover, there are natural isomorphisms*

$$\mathrm{Ext}^n(A,B) \cong \mathrm{Hom}_{\mathcal{D}^b(\mathcal{A})}(A, B[n]) \cong \mathrm{Hom}_{\mathcal{D}^b(\mathcal{A})}(A[-n], B).$$

That is, extensions correspond to maps in $\mathcal{D}^b(\mathcal{A})$ between complexes each concentrated in a single degree. Hence the functor $\iota : \mathcal{A} \to \mathcal{D}^b(\mathcal{A})$ is fully faithful.

Proof. Since translation is an autoequivalence, we have

$$\mathrm{Hom}_{\mathcal{D}^b(\mathcal{A})}(A, B[n]) \cong \mathrm{Hom}_{\mathcal{D}^b(\mathcal{A})}(A[-n], B).$$

To show that these agree with $\mathrm{Ext}^n(A,B)$, we use Proposition 1.7. We give a proof only in the case $\mathcal{A}$ has enough projectives (or enough injectives). Take a projective resolution P of A, so in particular $P \simeq A$ in $\mathcal{D}^b(\mathcal{A})$. Then

$$\begin{aligned} \mathrm{Ext}^n_{\mathcal{A}}(A,B) &= H^n \mathrm{Hom}_{\mathcal{A}}(P,B) = \mathrm{Hom}_{\mathcal{K}^b(\mathcal{A})}(P, B[n]) \\ &\cong \mathrm{Hom}_{\mathcal{D}^b(\mathcal{A})}(P, B[n]) = \mathrm{Hom}_{\mathcal{D}^b(\mathcal{A})}(A, B[n]). \end{aligned}$$

This proves that the right-derived functor Ext^n computed by taking a projective resolution in the first coordinate corresponds to morphisms in the derived category. But a dual proof using an injective resolution I of B shows that

$$H^n \mathrm{Hom}_{\mathcal{A}}(A, I) \cong \mathrm{Hom}_{\mathcal{D}^b(\mathcal{A})}(A[-n], B) = \mathrm{Hom}_{\mathcal{D}^b(\mathcal{A})}(A, B[n])$$

also. So, if $\mathcal{A}$ has enough injectives and enough projectives, then the two definitions of Ext^n concur.

As a consequence, ι is fully faithful since for $A, B \in \mathcal{A}$ we have

$$\mathrm{Hom}_{\mathcal{D}^b(\mathcal{A})}(\iota(A), \iota(B)) \cong \mathrm{Ext}^0_{\mathcal{A}}(A,B) = \mathrm{Hom}_{\mathcal{A}}(A,B).$$

[3] We do not require a chain isomorphism. In the case of 1-extensions, any chain map will a posteriori be an isomorphism, but this is not true of higher extensions.

Finally, we outline the argument relating n-extensions in $\mathcal{A}$ to morphisms in $\mathcal{D}^b(\mathcal{A})$ only in the case $n = 1$.

When constructing the derived category, we observed that any 1-extension

$$0 \to B \to M \to A \to 0$$

in $\mathcal{A}$ gives rise to a distinguished triangle

$$B \to M \to A \to B[1]$$

in $\mathcal{D}^b(\mathcal{A})$. In particular we obtain a morphism $A \to B[1]$. One can check that the equivalence relation on morphisms in $\mathcal{D}^b(\mathcal{A})$ corresponds to equivalence of extensions. In the other direction, a map $A \to B[1]$ in $\mathcal{D}^b(\mathcal{A})$ can be completed to a distinguished triangle via the mapping cone construction. After rotating, we get the distinguished triangle

$$B[1] \to C \to A[1] \to B[2]$$

which then induces a long exact sequence in cohomology. Since $A[1], B[1]$ only have nonzero cohomology in degree -1, the same holds for C. Thus the rotation

$$B \to C[-1] \to A \to B[1]$$

is a distinguished triangle whose terms $B, C[-1], A$ lie in $\iota(\mathcal{A})$, and thus the maps $B \to C[-1], C[-1] \to A$ are morphisms in $\mathcal{A}$ since ι is fully faithful. One checks that these maps form a short exact sequence. □

This gives three different characterisations of the extension groups Ext^i, and allows us to describe all maps in the derived category by computing extensions.

1.3 The hereditary case

Let $\mathcal{A}$ be a hereditary abelian category, by which we mean $\text{Ext}^i(-,-) = 0$ for $i > 1$. This is in particular the case if $\mathcal{A} = A\text{-mod}$ for a hereditary[4] algebra A, because every object M of $\mathcal{A}$ has a length one projective resolution. Such a resolution is given by finding any surjection $P \twoheadrightarrow M$ with P projective, since the kernel is a submodule of P and thus automatically projective. We can use this resolution to compute Ext^i, so $\text{Ext}^i = 0$ for $i \neq 0, 1$. In particular, we get the following exact sequence.

$$0 \to \text{Hom}_{\mathcal{A}}(A, B) \to \text{Hom}_{\mathcal{A}}(P_0, B) \to \text{Hom}_{\mathcal{A}}(P_1, B) \to \text{Ext}^1(A, B) \to 0$$

[4] Recall Definition 0.1.

Notation 1.21. In a hereditary category we write $\operatorname{Ext} := \operatorname{Ext}^1$ for the single nontrivial derived Hom functor.

Here is an initial example of two distinct abelian categories whose derived categories are equivalent. For an introduction to the relevant quiver theory, see the beginning of Section 2.1.

Example 1.22. Let Q_{Kr} be the directed graph with two vertices x, y and two arrows a, b from x to y. This is called the Kronecker quiver.

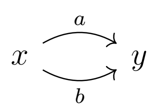

A representation of Q_{Kr} consists of finite-dimensional $\mathbb{C}$-vector spaces V_x, V_y and linear maps $L_a, L_b : V_x \to V_y$. Such representations form an abelian category, $\operatorname{Rep}_{\mathbb{C}} Q_{\mathrm{Kr}}$. If V_x, V_y are 1-dimensional, then the linear maps $L_a, L_b \in \mathbb{C}$ are simply scalars, and define a point $[L_x : L_y] \in \mathbb{P}^1_{\mathbb{C}}$. It is easy to check that two such representations are isomorphic if and only if they correspond to the same point in projective space.

But in fact, the connection to $\mathbb{P}^1_{\mathbb{C}}$ runs much deeper. The abelian category $\operatorname{Coh} \mathbb{P}^1_{\mathbb{C}}$ has objects coherent sheaves on $\mathbb{P}^1_{\mathbb{C}}$. Consider the object $T = \mathscr{O} \oplus \mathscr{O}(1)$. We can define the functor

$$\operatorname{Hom}_{\mathbb{P}^1}(T, -) : \operatorname{Coh}(\mathbb{P}^1) \to \operatorname{mod-} \operatorname{End}_{\mathbb{P}^1}(T)$$

since there is a right-action of $\operatorname{End}_{\mathbb{P}^1}(T)$ on $\operatorname{Hom}_{\mathbb{P}^1}(T, -)$ by precomposition. This functor is not an equivalence of abelian categories because for example $\operatorname{Hom}_{\mathbb{P}^1}(T, \mathscr{O}(-n)) = 0$ for all $n > 0$. However, the derived functor

$$\operatorname{RHom}_{\mathbb{P}^1}(T, -) : \mathcal{D}^b \operatorname{Coh}(\mathbb{P}^1) \to \mathcal{D}^b(\operatorname{End}_{\mathbb{P}^1}(T)^{\mathrm{op}}\text{–mod})$$

is an equivalence of triangulated categories. Moreover, $\operatorname{End}_{\mathbb{P}^1}(T)^{\mathrm{op}} \cong \mathbb{C}Q_{\mathrm{Kr}}$ can be identified with the path algebra of the Kronecker quiver, via $\mathrm{id}_{\mathscr{O}} \leftrightarrow e_y$, $\mathrm{id}_{\mathscr{O}(1)} \leftrightarrow e_x$, $\operatorname{Hom}_{\mathbb{P}^1}(\mathscr{O}, \mathscr{O}(1)) \leftrightarrow \operatorname{span}_{\mathbb{C}}\{a, b\}$, and $\operatorname{Hom}_{\mathbb{P}^1}(\mathscr{O}(1), \mathscr{O}) = 0$. Hence $\operatorname{RHom}(T, -)$ gives a derived equivalence $\mathcal{D}^b \operatorname{Coh}(\mathbb{P}^1) \xrightarrow{\simeq} \mathcal{D}^b(\operatorname{Rep}_{\mathbb{C}} Q_{\mathrm{Kr}})$. Concretely, for a point $[c_a : c_b] \in \mathbb{P}^1$, the skyscraper sheaf $\mathscr{O}_{[c_a:c_b]}$ is sent to the $(1,1)$-dimensional representation with maps c_a, c_b under this equivalence.

We will return to this example in Remark 2.31, once we have developed the language to appreciate the quiver side of the story.

For us, the most important result regarding hereditary categories is that any object of the derived category is isomorphic to the direct sum of its cohomology.

Theorem 1.23. *If $\mathcal{A}$ is hereditary then any object in $\mathcal{D}^b(\mathcal{A})$ is the direct sum of its cohomology pieces. That is,*

$$X \cong \bigoplus_{i\in\mathbb{Z}} H^i(X)[-i]$$

for any object X of $\mathcal{D}^b(\mathcal{A})$.

The following argument is based on the proof of Corollary 3.15 in [Huy06].

Proof. X has finitely many nonzero cohomology groups. We induct on the length k of X, with the $k = 0$ case immediate. For $k > 0$, let i_0 be minimal such that $H^{i_0}(X) \neq 0$. By Proposition 1.8, we have a map $H^{i_0}(X)[-i_0] \to X$ which we complete to a distinguished triangle

$$H^{i_0}(X)[-i_0] \to X \to X_1 \to H^{i_0}(X)[1 - i_0] \tag{1.5}$$

From this distinguished triangle we get a long exact sequence in cohomology, and we find that $H^i(X_1) \cong H^i(X)$ for $i > i_0$, and $H^i(X_1) = 0$ for $i \leq i_0$. In particular X_1 has length $k - 1$. It is enough to show (1.5) splits, since then $X \cong X_1 \oplus H^{i_0}(X)[-i_0]$ and by induction X_1 decomposes as desired.

To show (1.5) splits, we will prove that $\mathrm{Hom}(X_1, H^{i_0}(X)[-i_0+1]) = 0$. Hence the third map in the distinguished triangle is zero, and by the splitting lemma it splits. This Hom-vanishing is a consequence of the higher extension groups being trivial.

Our inductive hypothesis allows us to write:

$$X_1 \cong \bigoplus_{i>i_0} H^i(X_1)[-i] \cong \bigoplus_{i>i_0} H^i(X)[-i]$$

and then using Proposition 1.9 and Theorem 1.20, we have

$$\begin{aligned}
\mathrm{Hom}_{\mathcal{D}^b(\mathcal{A})}(X_1, H^{i_0}(X)[-i_0+1]) &= \bigoplus_{i>i_0} \mathrm{Hom}_{\mathcal{D}^b(\mathcal{A})}(H^i(X)[-i], H^{i_0}(X)[-i_0+1]) \\
&= \bigoplus_{i>i_0} \mathrm{Ext}_{\mathcal{A}}^{i-i_0+1}(H^i(X), H^{i_0}(X)) \\
&= 0
\end{aligned}$$

which vanishes because $i > i_0$ so $i - i_0 + 1 \geq 2$. □

Corollary 1.24. *For any pair of objects X, Y of $\mathcal{D}^b(\mathcal{A})$, we have*

$$\mathrm{Hom}_{\mathcal{D}^b(\mathcal{A})}(X, Y) = \bigoplus_{i,j\in\mathbb{Z}} \mathrm{Ext}_{\mathcal{A}}^{i-j}(H^i(X), H^j(Y)).$$

Proof.

$$\begin{aligned}
\mathrm{Hom}_{\mathcal{D}^b(\mathcal{A})}(X,Y) &= \mathrm{Hom}_{\mathcal{D}^b(\mathcal{A})}\left(\bigoplus_{i\in\mathbb{Z}} H^i(X)[-i], \bigoplus_{j\in\mathbb{Z}} H^j(Y)[-j]\right) && \text{Theorem 1.23}\\
&\cong \bigoplus_{i,j\in\mathbb{Z}} \mathrm{Hom}_{\mathcal{D}^b(\mathcal{A})}(H^i(X)[-i], H^j(Y)[-j]) && \text{Proposition 1.9}\\
&\cong \bigoplus_{i,j\in\mathbb{Z}} \mathrm{Ext}^{i-j}_{\mathcal{A}}(H^i(X), H^j(Y)) && \text{Theorem 1.20}
\end{aligned}$$

□

As a corollary, we now have a very explicit description of $\mathcal{D}^b(\mathcal{A})$ when $\mathcal{A}$ is hereditary.

Theorem 1.25. *Let $\mathcal{A}$ be hereditary. The objects of $\mathcal{D}^b(\mathcal{A})$ are bounded complexes up to cohomology, and morphisms are additively generated by elements of $\mathrm{Ext}_{\mathcal{A}}(A,B)$ for $A,B \in \mathcal{A}$, which correspond to maps $\iota(A) \to \iota(B)[1]$ in $\mathcal{D}^b(\mathcal{A})$, together with elements of $\mathrm{Hom}_{\mathcal{A}}(A,B)$, corresponding to maps $\iota(A) \to \iota(B) \in \mathcal{D}^b(\mathcal{A})$. We also have shifts of such maps.*

In particular, every morphism is the direct sum of morphisms between complexes each concentrated in a single degree, and these degrees can differ by at most 1.

Another corollary is that any indecomposable object in $\mathcal{D}^b(\mathcal{A})$ must have a representative which is concentrated in a single degree. Hence indecomposable objects in $\mathcal{D}^b(\mathcal{A})$ are just shifts of indecomposables in $\mathcal{A}$.

Corollary 1.26. *Suppose $\mathcal{A}$ is hereditary. Then an object X of $\mathcal{D}^b(\mathcal{A})$ is indecomposable iff $X \cong \iota(A)[n]$ for some $n \in \mathbb{Z}$, where A is an indecomposable object of $\mathcal{A}$.*

Proof. We use the fact that the product (or coproduct) in $\mathcal{D}^b(\mathcal{A})$ is the one inherited from $\mathcal{A}$.

Suppose X is indecomposable. Since X is isomorphic to the direct sum of its cohomology, it must only have one nonzero cohomology piece, and then $X \cong H^n(X)[-n]$ where $H^n(X) \in \mathcal{A}$. The inclusion of $\mathcal{A}$ into $\mathcal{D}^b(\mathcal{A})$ in degree $-n$ is fully faithful, so $H^n(X)$ must be indecomposable in $\mathcal{A}$.

Conversely, since the inclusion of $\mathcal{A}$ in degree zero is fully faithful, any indecomposable object of $\mathcal{A}$ remains indecomposable under this inclusion. Translations are autoequivalences of $\mathcal{D}^b(\mathcal{A})$ so preserve indecomposables. □

Chapter 2

Quiver representations

In this chapter, we develop the theory of quiver representations. We study the abelian category $\operatorname{Rep} Q$ of representations alongside its derived companion $\mathcal{D}^b(\operatorname{Rep} Q)$. Our goal is to arm ourselves with all the standard tools we will need to study (derived) reflection functors in Chapter 3. One of the first questions we will be interested in is determining when two quivers have equivalent representation categories. In later chapters we will be able to address the same question for the corresponding derived category.

2.1 Basic definitions

A quiver is a directed graph, which we assume for simplicity is connected. We allow loops and multiple edges. Formally, we have the following definition.

Definition 2.1. A *quiver* Q consists of a finite set of vertices Q_0, a finite set of arrows Q_1, and two functions $h, t : Q_1 \to Q_0$. For each arrow $\alpha \in Q_1$ we call $h(\alpha)$ the *head* and $t(\alpha)$ the *tail*. We think of α as a directed edge from $t(\alpha)$ to $h(\alpha)$. We require for simplicity that Q be connected.

A quiver Q is called *acyclic* if it contains no directed cycles.

Notation 2.2. Write $h\alpha$ for $h(\alpha)$ and $t\alpha$ for $t(\alpha)$. We will often use x, y, z to denote vertices and $\alpha, \beta, \ldots$ to denote arrows.

From now on, we work over a fixed base field k. A representation of a quiver Q associates a finite-dimensional vector space with every vertex and a linear map with every arrow.

Definition 2.3. A *representation* V of a quiver Q consists of a finite-dimensional k-vector space $V(x)$ at each vertex $x \in Q_0$, and a k-linear map $V(\alpha) : V(t\alpha) \to V(h\alpha)$ for each arrow $\alpha \in Q_1$.

Definition 2.4. Let V, W be two representations of a quiver V. A *morphism* $\phi : V \to W$ consists of a k-linear map $\phi(x) : V(x) \to W(x)$ at each vertex $x \in Q_0$, such that $W(\alpha) \circ \phi(t\alpha) = \phi(h\alpha) \circ V(\alpha)$ for every arrow $\alpha \in Q_1$. In other words, the diagram

$$\begin{array}{ccc} V(t\alpha) & \xrightarrow{V(\alpha)} & V(h\alpha) \\ {\scriptstyle \phi(t\alpha)}\downarrow & & \downarrow{\scriptstyle \phi(h\alpha)} \\ W(t\alpha) & \xrightarrow{W(\alpha)} & W(h\alpha) \end{array}$$

commutes, for every $\alpha \in Q_1$.

An important invariant for quiver representations is the dimension vector, which fully captures the vector space at each vertex but ignores the maps between vertices.

Definition 2.5. Given a representation V of a quiver Q, we define the *dimension vector* of V to be the vector $\dim_k V := (\dim_k V(x))_{x \in Q_0}$. We think of this vector as belonging to the group $\mathbb{Z}^{Q_0}$ or the vector space $\mathbb{R}^{Q_0}$, as convenience dictates.

Both $\mathbb{Z}^{Q_0}$ and $\mathbb{R}^{Q_0}$ come equipped with a partial order, where $\mathbf{v} \leq \mathbf{w}$ if $v_i \leq w_i$ for every $i \in Q_0$, and $\mathbf{v} < \mathbf{w}$ if both $\mathbf{v} \leq \mathbf{w}$ and $\mathbf{v} \neq \mathbf{w}$.

We next introduce the category $\operatorname{Rep}_k Q$, whose objects are k-representations of Q and whose morphisms are as in Definition 2.4. We often omit the field k from the notation. This category is abelian, with kernels, images and cokernels defined vertex-wise. Its zero object is the representation of Q consisting of the zero vector space at each vertex, often called the *trivial representation* of Q.

Many standard vector space operations have natural analogues for quiver representations, defined vertex-wise. This includes direct sums and duals. Finite direct sums satisfy the universal property of both products and coproducts, as is the case for vector spaces.

We will soon be interested in the category $\mathcal{D}^b(\operatorname{Rep} Q)$, the bounded derived category of representations of Q. This is obtained from $\operatorname{Rep} Q$ in the usual way. Before we discuss this category in more detail, we finish our characterisation of $\operatorname{Rep} Q$.

Lemma 2.6. *A morphism ϕ in* $\operatorname{Rep} Q$ *is mono if and only if each $\phi(x)$ for $x \in Q_0$ is an injective linear map, and epi if and only if each $\phi(x)$ is surjective. Moreover, ϕ is an isomorphism if and only if each $\phi(x)$ is.*

The proof of Lemma 2.6 is omitted; it follows from the same fact in the category of finite-dimensional vector spaces over k. In fact, this holds in any abelian category.

We have the following entirely categorical definitions.

Definition 2.7. Fix a quiver Q. A *subrepresentation* of a representation W is a monomorphism $\phi : V \to W$, often denoted $V \leq W$. A subrepresentation is *proper* if ϕ is not an isomorphism. A *quotient* of W is an epimorphism $\phi : W \to V$. The representation W is *simple* if it has no proper nontrivial subrepresentations, or equivalently no proper nontrivial quotients.

Finally, W is *indecomposable* if it is not isomorphic to a direct sum of two nonzero representations.

Every simple representation is indecomposable, since given a decomposition $W \cong U \oplus V$ with U, V nonzero, both U and V can be realised as proper subrepresentations of W via inclusion composed with the isomorphism. However, the converse is not true, as we will see in the next section. We also observe that given $V \leq W$ we automatically have $\dim_k V \leq \dim_k W$, and if V is a proper subrepresentation then $\dim_k V < \dim_k W$.

Any representation can be written as a finite direct sum of indecomposable representations. In fact, this decomposition is unique up to permuting the summands. See [DW17], Theorem 1.7.4.

2.2 Classification of simples and indecomposables

A standard classification problem in the theory of quiver representations is to determine all isomorphism classes of indecomposable representations. This problem is completely solved for quivers whose underlying graph is an ADE type Dynkin diagram (we will define these shortly), but is very difficult in general. We begin with the relatively easy problem of classifying the simple representations of an acyclic quiver.

Theorem 2.8. *Let Q be an acyclic quiver. Then the isomorphism classes of simple representations of Q are in bijection with the vertices of Q. For each vertex $x \in Q_0$, we define the representation S_x by $S_x(x) = k$ and $S_x(y) = 0$ for all $y \in Q_0 \setminus \{k\}$. Then each S_x is simple, and every simple representation has the form S_x for some vertex x.*

Before proving this, we will need some results for acyclic quivers.

Definition 2.9. A vertex $x \in Q_0$ is a *source* in Q if there are no arrows $\alpha \in Q_1$ such that $h\alpha = x$. Dually, x is a *sink* if there are no arrows α such that $t\alpha = x$.

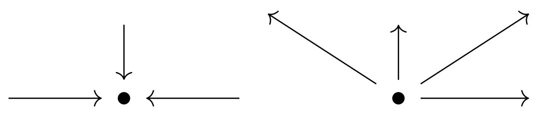

sink source

Given a representation W of Q, a vertex x is a *nonzero source w.r.t.* W if $W(x) \neq 0$, and for every arrow α with $h\alpha = x$, we have $W(t\alpha) = 0$. Dually, x is a *nonzero sink w.r.t.* W if $W(x) \neq 0$ and for every arrow α with $t\alpha = x$, then $W(h\alpha) = 0$. A nonzero source x w.r.t. W is a vertex of Q which becomes a sink when all the vertices with associated vector space 0 in W are removed from Q.

Lemma 2.10. *Every acyclic quiver has a source and a sink. Given a nontrivial representation of an acyclic quiver, there exists a nonzero source and a nonzero sink w.r.t. the representation.*

Moreover, given an acyclic quiver on n vertices, there exists a labelling of the vertices $1, \ldots, n$ such that each vertex is a source in the subquiver obtained by discarding all vertices with higher numbers. Equivalently, there exists a labelling of the vertices such that every $\alpha \in Q_1$ has $h\alpha < t\alpha$.

Proof. Let Q be an acyclic quiver. Pick a vertex x. If x is a source, we are done. If not, there exists an arrow α with $h\alpha = x$. Replace x by $t\alpha$ and repeat. We cannot encounter the same vertex twice in this process, otherwise the arrows we traversed in between would form a directed cycle. Hence the process terminates, since Q has finitely many vertices. It must terminate at a source. A symmetric argument shows Q must have a sink.

Now let W be a nontrivial representation of Q. Let Q' be the nonempty subquiver of Q consisting of all vertices $x \in Q_0$ such that $W(x) \neq 0$, and arrows between these vertices. Restricting W to Q' gives a representation W' which is nonzero at every vertex. Since Q' is still acyclic, by the previous argument Q' has at least one source, $x \in Q_0$. Then x is a nonzero source w.r.t. W, in the quiver Q. Symmetrically, Q has a nonzero sink w.r.t. W.

Now let n denote the number of vertices in Q. Pick a source, and label it n. Assume for an induction argument that we can successfully label vertices $i, i+1, \ldots, n$. Consider the subquiver Q' consisting of all the unlabelled vertices. Then Q' has a source, which we may label $i-1$ (in Q). Continuing in this way, we label all the vertices of Q.

Consider an arrow $\alpha \in Q_1$. Since $h\alpha$ is a source when all vertices with higher numbers are removed from Q, we know that $t\alpha < h\alpha$ is impossible. We cannot have $t\alpha = h\alpha$, as then Q would have a loop at this vertex. Thus $t\alpha > h\alpha$, as claimed. □

Proof of Theorem 2.8. First we claim that each S_x is simple. Let W be a proper subrepresentation of S_x. Then we have $\dim_k W < \dim_k S_x$. But the dimension vector of S_x consists of all zeroes except for a single 1, so the only possible strictly smaller dimension vector is the zero vector. Therefore W is the trivial representation.

Conversely, let W be a simple representation of acyclic Q. By Lemma 2.10, there exists a vertex $x \in Q_0$ which is a nonzero source w.r.t. W. Consider the simple representation S_x. Since $W(x) \neq 0$, there exists a linear inclusion $\phi_x : S_x(x) \to W(x)$. One can check that the map ϕ obtained from ϕ_x together with the zero map at every other vertex is a morphism of representations, because x is a nonzero source w.r.t. W. By simplicity of W, ϕ must be an isomorphism so $W \cong S_x$ has the desired form. □

Theorem 2.8 does not hold in general for quivers containing cycles. Indeed for any choice of directed cycle in a quiver, we obtain a simple representation by taking a 1-dimensional space at each vertex in the cycle and the zero vector space at other vertices, with the linear maps involved in the cycle all isomorphisms. The composition of these isomorphisms gives some scalar $c \in k^\times$, and two such representations cannot be isomorphic if they have different associated scalars c. When k is an infinite field, this gives an infinite family of non-isomorphic simple representations.

We progress now to the problem of determining the indecomposable representations of an ADE quiver, that is one whose underlying graph (obtained by forgetting the orientations of the arrows) is one of the following

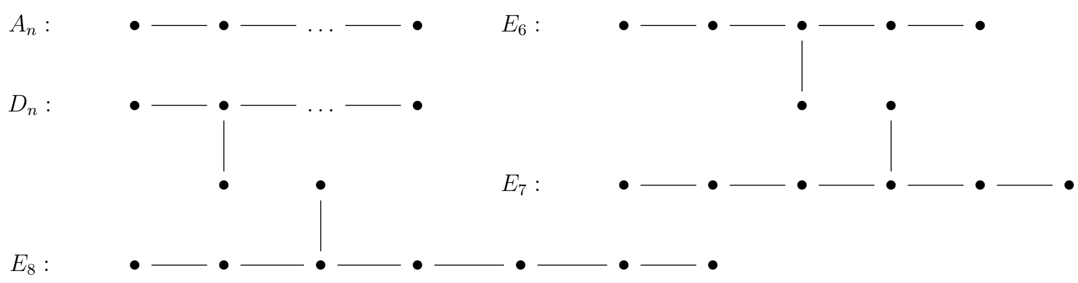

where subscripts denote the number of vertices. These graphs are the ADE Dynkin graphs, and they also appear in the classification of semisimple Lie algebras. We find that ADE quivers are precisely the ones for which the category of representations contains finitely many isomorphism classes of indecomposable objects. Moreover, for ADE quivers the indecomposable representations are in bijection with the positive roots in the root system of the corresponding Dynkin graph. Practically, this means that we can give a complete list of indecomposable representations of ADE quivers. See Appendix A for a brief exploration of the relationship between ADE Dynkin diagrams, root systems, and semisimple Lie algebras.

Definition 2.11. A quiver Q has *finite (representation) type* if the category $\operatorname{Rep} Q$ contains finitely many isomorphism classes of indecomposable objects, meaning that it is finitely generated under the direct sum operation.

We saw that if Q is cyclic then it has infinitely many isomorphism classes of simple objects, so in particular cannot be finite representation type. However, being acyclic is not sufficient to give Q finite representation type.

Theorem 2.12 (Gabriel)**.** *A quiver Q has finite type if and only if its underlying graph is a disjoint union of ADE-type Dynkin graphs. If one (and therefore both) of these conditions holds, then isomorphism classes of indecomposable representations of Q are in bijection with positive roots of the Dynkin graph, via the map sending a representation to its dimension vector.*

Gabriel's original proof can be found in [Gab72], while an alternative proof due to Bernstein, Gelfand and Ponomarev is given in [DW17], Theorem 4.4.13.

Corollary 2.13. *In all of the following constructions, maps between spaces of the same dimension are taken to be isomorphisms.*

Label the vertices of A_n from 1 to n in order. For $1 \leq i \leq j \leq n$, define the representation $E_{i,j}$ by

$$\dim_k E_{i,j}(x) := \begin{cases} 1 & \textit{if } i \leq x \leq j \\ 0 & \textit{else} \end{cases}$$

This is indecomposable, and for A_n all isomorphism classes of indecomposable representations have this form.

Label the vertices of D_n as in the diagram, $n \geq 4$.

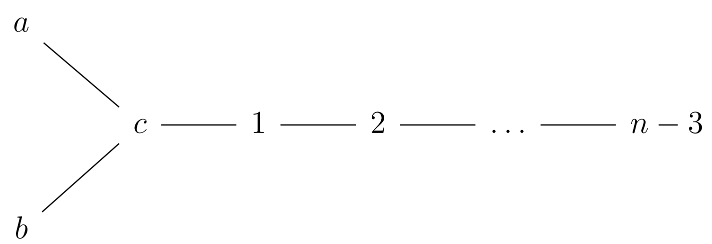

There are two families of indecomposable representations of D_n. The first is given by choosing a connected subgraph of D_n, and taking the representation with 1-dimensional spaces only at vertices in this subgraph. The second is given by choosing vertices $1 \leq i \leq j \leq n$, and defining

$$\dim_k F_{i,j}(x) := \begin{cases} 1 & \textit{if } x = a, b \textit{ or } i \leq x \leq j \\ 2 & \textit{if } x = c \textit{ or } x < i \\ 0 & \textit{if } x > j. \end{cases}$$

The three exceptional maps between a 1-dimensional space and a 2-dimensional space are the same for all $F_{i,j}$, and they are given by A, B, C where $A + B = C$ when all are directed from the 1-dimensional space to the 2-dimensional space.

Even when the full classification of indecomposables is not required, we will often use the bijection result to characterise an indecomposable by its much simpler dimension vector. We note also that the number and type of indecomposable representations of an ADE quiver Q depends only on the underlying graph, not on the chosen edge orientations. This gives a bijection between the objects of $\operatorname{Rep} Q$ and $\operatorname{Rep} Q'$ whenever Q and Q' are ADE quivers which differ only in the orientations of some edges. Unfortunately this bijection is not functorial. However, in Chapter 3 we will construct a functor which (almost) gives a bijection between the indecomposable representations of two such quivers, and which descends to a derived equivalence.

We finish this section with a warning: in general, the classification of indecomposable representations for a quiver Q is a wild problem, and is only completely solved in the ADE case and a few other related special cases. However, there has been progress in understanding the possible dimension vectors of indecomposable representations (see Kac's Theorem, Chapter 8 in [DW17]).

2.3 The path algebra, and representations as modules

There is an alternate framework for studying quiver representations, in which each representation over a fixed field k is viewed as a (left or right) module over a particular associative k-algebra, called the path algebra. In fact, we have an equivalence of categories between $kQ-\text{mod}$, the category of finitely-generated left modules over the path algebra, and $\text{Rep}_k Q$. This viewpoint allows us to bring to bear some useful results regarding module categories. In particular, we obtain a canonical length one projective resolution for any quiver representation, and a dual length one injective resolution.

Definition 2.14. A *path* in a quiver Q is a nonempty finite sequence $p = \alpha_\ell \alpha_{\ell-1} \dots \alpha_1$ of arrows $\alpha_i \in Q_1$ such that $h\alpha_i = t\alpha_{i+1}$ for each $1 \leq i \leq \ell - 1$. The length of such a path is $|p| := \ell$, the number of arrows, and we define $t(p) = tp = t\alpha_1$ and $h(p) = hp = h\alpha_\ell$ to be the tail and head respectively of the path. In addition, we introduce the path e_x of length zero for each $x \in Q_0$. This is the empty path at the vertex x, with $he_x = x = te_x$.

Given a representation V of Q and a path p, we define

$$V(p) = V(\alpha_\ell) \dots V(\alpha_1).$$

For $x \in Q_0$, we also define

$$V(e_x) = \mathbb{1}_{V(x)}.$$

Then V has a linear map $V(p) : V_{tp} \to V_{hp}$ associated to each path p of Q, not just each arrow. These satisfy $V(pq) = V(p)V(q)$ whenever $hq = tp$, so p and q compose.

Remark 2.15. We can think of a quiver Q as a category, whose objects are vertices and whose morphisms are paths. Then the above discussion shows that a representation of Q is the same thing as a functor from this category to the category $\mathrm{Vec}_f(k)$ of finite-dimensional k-vector spaces. From this perspective, morphisms between quiver representations become natural transformations between their respective functors.

Definition 2.16. The *path algebra* kQ corresponding to a quiver Q has a k-basis labelled by paths in Q. The multiplication is defined by composition of paths where they share an endpoint, so

$$p \cdot q = \begin{cases} pq & \text{if } hq = tp \\ 0 & \text{else.} \end{cases}$$

This extends linearly to a multiplication law on kQ, which is associative because path composition is. In particular, note that $e_x p = p$ if $hp = x$ and 0 otherwise, and $pe_x = p$ if $tp = x$, 0 otherwise.

The elements $e_x, x \in Q_0$ satisfy $e_x^2 = e_x$ for all $x \in Q_0$, and $e_x e_y = 0$ when $x \neq y$. Hence they form a set of pairwise orthogonal idempotents. Moreover, for every path p we have $e_x p = p$ exactly when $x = hp$ and this product is zero otherwise, so $\sum_{x \in Q_0} e_x = \mathbb{1}_{kQ}$ is the identity element of the path algebra.

Proposition 2.17 (Lemma 1.5.3 in [DW17]). *kQ is finite-dimensional over k if and only if Q is acyclic.*

Proof. If Q has a directed cycle, this gives a path p with $hp = tp$. Then $p, p^2, p^3, \ldots$ are distinct paths, so kQ is infinite-dimensional.

Conversely, suppose Q is acyclic. By Lemma 2.10 there exists a labelling of the vertices of Q such that every arrow $\alpha \in Q_1$ has $h\alpha < t\alpha$. Then as a path p is traversed from tail to head, vertices lying on the path must be encountered in strictly decreasing order. In particular, this means that $|p| < |Q_0|$ for every path p. Since there is a global bound on path length and Q has finitely many arrows from which path segments can be drawn, there are finitely many distinct paths, including the $|Q_0|$ many length zero paths. Thus kQ has a finite k-basis. □

Proposition 2.18 (Proposition 1.8 in [Wey]). *The categories* $kQ-\mathrm{mod}$ *and* $\mathrm{Rep}_k\, Q$ *are equivalent.*

Proof. We explicitly define the equivalences. Let $F : kQ-\mathrm{mod} \to \mathrm{Rep}_k\, Q$ send a module M to the collection of vector spaces $FM(x) := e_x M$ for $x \in Q_0$. Each $\alpha \in Q_1$ acts k-linearly on M via left multiplication, and we take $FM(\alpha)$ to be the restriction of this action to the vector subspace $FM(t\alpha) = e_{t\alpha}M$. This defines a linear map $FM(\alpha) : e_{t\alpha}M \to e_{h\alpha}M$ since

$$\alpha e_{t\alpha}M = \alpha M = e_{h\alpha}\alpha M \subseteq e_{h\alpha}M.$$

Given a homomorphism of kQ-modules $\phi : M \to N$, we define $F\phi(x) = \phi|_{e_xM}$, for each $x \in Q_0$. Then

$$\phi(e_xM) = e_x\phi(M) \subseteq e_xN$$

so with this definition, $F\phi(x)$ is a map $FM(x) \to FN(x)$ as desired. This collection of linear maps forms a morphism of quiver representations because the map ϕ commutes with the kQ-action.

In the other direction, let $G : \operatorname{Rep}_k Q \to kQ-\text{mod}$ send a representation V to the module

$$GV = \bigoplus_{x \in Q_0} V(x)$$

where this equality is as k-vector spaces. For any path p, left-multiplication by p is defined to be the linear map $V(p)$ on the summand $V(tp)$, and zero on all other summands. This makes GV a kQ-module.

For a morphism of representations $\phi : V \to W$, we define $G\phi$ to be the direct sum of all the maps $\phi(x), x \in Q_0$. This commutes with the module action because it commutes with the action by arrows in Q_1, since ϕ is a morphism of representations.

One can easily check that F, G are functors, and that $F \circ G$ is the identity endomorphism on $\operatorname{Rep}_k Q$. The composition $G \circ F$ is naturally isomorphic to the identity functor on $kQ-\text{mod}$ via the collection of isomorphisms

$$M \xrightarrow{\sim} \bigoplus_{x \in Q_0} e_xM$$

$$m \mapsto (e_xm)_{x \in Q_0}$$

where left-multiplication by a path p in the codomain is the same as in the domain when restricted to the subspace $e_{tp}M$, and zero on the other summands. Using that the elements $e_x, x \in Q_0$ are orthogonal idempotents which sum to the identity, one can show bijectivity of this map. Naturality can be checked similarly. □

Note that in the above proof, it is important that we worked with the category of left-modules rather than the category of right-modules over kQ, so that the multiplication in the module is compatible with path composition. If we wanted an algebra A so that $\operatorname{Rep} Q \simeq \text{mod}-A$, we could take A to be kQ but with the order of multiplication reversed.

$\operatorname{Rep} Q$ has two different kinds of duality, coming from the duality between $kQ-\text{mod}$ and $\text{mod}-kQ$. Formally, we have two functors $\operatorname{Rep} Q \to \operatorname{Rep} Q^{\text{op}}$, one of which preserves injective and projective objects, and one of which interchanges injectives and projectives.

Definition 2.19. Let A be an algebra. The *opposite algebra* A^{op} is equal to A as a set, with the same addition and scalar multiplication, but with the order of multiplication reversed.

Let Q be a quiver. The *opposite quiver* Q^{op} is defined by $Q_0^{\text{op}} = Q_0$, $Q_1^{\text{op}} = Q_1$ and for each arrow $\alpha \in Q_1^{\text{op}}$ we define $h_{\text{op}}\alpha = t\alpha, t_{\text{op}}\alpha = h\alpha$. That is, Q^{op} is the quiver Q with all the arrows reversed.

Lemma 2.20. *The path algebra of the opposite quiver is the opposite algebra, that is $k(Q^{\text{op}})$ is canonically isomorphic to $(kQ)^{\text{op}}$ for any quiver Q.*

Proof. The identification between vertices and arrows of Q and Q^{op} induces the isomorphism. □

With this canonical identification in mind, we will denote the path algebra of the opposite quiver simply by kQ^{op}. Now, there are natural equivalences

$$A^{\text{op}}-\text{mod} \simeq \text{mod}-A \text{ and } \text{mod}-A^{\text{op}} \simeq A-\text{mod}$$

since a left-action by A is the same as a right-action by A^{op}, and vice versa. Hence the natural duality between $kQ-\text{mod}$ and $\text{mod}-kQ$ induces a corresponding duality between $\operatorname{Rep} Q$ and $\operatorname{Rep} Q^{\text{op}}$.

Definition 2.21. Given a quiver Q, we have two duality functors $\operatorname{Rep} Q \to \operatorname{Rep} Q^{\text{op}}$. On objects, these are defined by

$$V^* := \operatorname{Hom}_k(V, k) \text{ and } \widetilde{V} := \operatorname{Hom}_{kQ}(V, kQ).$$

V^* has a natural right-action by kQ given by precomposing with the left-action on V. $\widetilde{V}$ has a natural right-multiplication by kQ arising from the right-action of kQ on itself in the codomain. We will sometimes denote the functor $V \mapsto V^*$ by D, and the corresponding functor $\operatorname{Rep} Q^{\text{op}} \to \operatorname{Rep} Q$ by D^{op}.

The functor V^* is simply an extension of the usual vector space dual. That is, it sends a representation of Q to the corresponding representation of Q^{op} which has the dual vector space at each vertex and the dual linear map for each arrow, and similarly dualises morphisms. $D^{\mathrm{op}} \circ D$ is naturally isomorphic to the identity on $kQ-$mod, with the functors D, D^{op} making $\operatorname{Rep} Q, \operatorname{Rep} Q^{\mathrm{op}}$ into a pair of dual categories.

The same is true of our other dualising functor, on projective objects. Namely, for any projective $P \in \operatorname{Rep} Q$, we have

$$P \cong \operatorname{Hom}_{Q^{\mathrm{op}}}(\widetilde{P}, kQ)$$

and this isomorphism is natural.

Lemma 2.22. *If V is a projective object in $\operatorname{Rep} Q$ then V^* is injective and $\widetilde{V}$ is projective in $\operatorname{Rep} Q^{\mathrm{op}}$. Dually, if V is injective then V^* is projective and $\widetilde{V}$ is injective.*

Proof. We consider only the case where V is projective. Then $\widetilde{V}$ is projective by Lemma 0.3, and V^* is injective by Lemma 0.4. □

Since we have canonical identifications

$$\text{mod}-kQ^{\mathrm{op}} \simeq kQ-\text{mod} \simeq \operatorname{Rep} Q$$

and

$$\text{mod}-kQ \simeq kQ^{\mathrm{op}}-\text{mod} \simeq \operatorname{Rep} Q^{\mathrm{op}}$$

we will henceforth treat these categories as identical and use them interchangeably.

2.4 Projective and injective indecomposables

We now restrict our attention to acyclic quivers. We wish to determine which objects in $\operatorname{Rep}_k Q$ are projective. Using the equivalence of categories established in Proposition 2.18, we can apply standard alternate characterisations of projective modules, recalled in Chapter 0. We will show that indecomposable projective representations are in bijection with the vertices of an acyclic quiver. The projective representation corresponding to a vertex can be cleanly described using the path algebra. Dually, we develop the theory of injective indecomposable representations.

Since direct summands and direct sums of projective modules are projective, to determine all projective representations we need only determine the indecomposable projectives. Then every projective can be written as a direct sum of some indecomposable projectives.

Definition 2.23. Let Q be an acyclic quiver. For each $x \in Q_0$, we define kQ-modules $P_x := kQe_x$ and $I_x := (e_x kQ)^*$.

Lemma 2.24. *There are direct sum decompositions*

$$kQ = \bigoplus_{x \in Q_0} P_x \qquad \text{and} \qquad (kQ)^* = \bigoplus_{x \in Q_0} I_x.$$

Proof. We prove the first decomposition, and the second follows dually (by showing kQ decomposes as a direct sum of right modules of the form $e_x kQ$, and then applying the duality functor D). First, given $v \in kQ$,

$$v = v \cdot 1 = v \sum_{x \in Q_0} e_x \in \bigoplus_{x \in Q_0} kQe_x = \bigoplus_{x \in Q_0} P_x.$$

Hence $kQ = \sum_{x \in Q_0} P_x$ where the sum is internal. To show the sum is direct, suppose $\sum_{x \in Q_0} a_x = 0$ with each $a_x \in P_x$. In particular, this means that $a_x e_y = a_x e_x e_y = 0$ for $y \neq x$. Then

$$0 = \left(\sum_{x \in Q_0} a_x \right) e_y = a_y e_y = a_y$$

for each $y \in Q_0$. □

Corollary 2.25. *The kQ-modules P_x (resp. I_x) are projective (resp. injective).*

Proof. A direct summand of a projective module is projective and kQ is free, therefore projective. Hence $kQe_x = P_x$ is a projective left module, and $e_x kQ$ is a projective right module. Then by Lemma 0.4, I_x is an injective left module. □

Remark 2.26. Viewing P_x as a representation of Q, the vector space $P_x(y) = e_y kQe_x$ has a basis of paths from x to y for each $y \in Q_0$. Then P_x has a k-basis of paths starting at x. Dually, the right module $I_x^* = e_x kQ$ has a k-basis consisting of paths ending at x, and each vector space $I_x(y)^* = e_x kQe_y$ has a basis of paths from y to x.

We will next collect some useful facts regarding the projective objects P_x. In the following series of results, we work in the category $kQ-$mod (or mod$-kQ$ for the dual results).

Lemma 2.27 (Proposition 2.2.2 in [DW17])**.** *There are canonical isomorphisms*

$$\mathrm{Hom}_Q(P_x, V) \cong V(x) \qquad \text{and} \qquad \mathrm{Hom}_Q(V, I_x) \cong V(x)^*.$$

Proof. We define linear maps each way by

$$\begin{aligned}\Psi : \mathrm{Hom}_Q(P_x, V) &\to V(x)\\ \phi &\mapsto \phi(x)(e_x) \in V(x)\\ \Phi : V(x) &\to \mathrm{Hom}_Q(P_x, V)\\ v &\mapsto (p \mapsto V(p)(v)).\end{aligned}$$

These are mutual inverses, giving the desired isomorphism. The result for injectives is dual. □

Proposition 2.28. *The projective objects P_x are indecomposable, and up to isomorphism these are all the indecomposable projectives. Similarly, the injective indecomposables are precisely the objects I_x for $x \in Q_0$.*

Proof. To see that P_x is indecomposable, we note that for a kQ-module M with n indecomposable direct summands, $\dim_k \mathrm{Hom}_Q(M, M) \geq n$. This is because the identity maps on the various indecomposable summands are independent. Hence by Lemma 2.27, P_x must have only one indecomposable summand.

If P is any projective object, then it must be a direct summand of $(kQ)^{\oplus n}$ for some n. But by uniqueness of decompositions and Lemma 2.24, the only indecomposable summands of $(kQ)^{\oplus n}$ are the P_x. Hence any projective indecomposable must be isomorphic to P_x, some $x \in Q_0$. □

Our interest in projective and injective objects arises because injective and projective resolutions can be used to compute the extension groups $\mathrm{Ext}^i(-,-)$, which determine maps in the derived category $\mathcal{D}^b(\mathrm{Rep}\, Q)$. We will construct a canonical length one projective resolution for any object of $\mathrm{Rep}\, Q$, showing that $\mathrm{Rep}\, Q$ is hereditary for acyclic Q. This effectively gives a complete description of $\mathcal{D}^b(\mathrm{Rep}\, Q)$ in the acyclic case, see Theorem 1.23.

Theorem 2.29. *Let Q be an acyclic quiver. Every kQ-module M has a length one projective resolution, given explicitly by the exact complex*

$$\mathcal{C}(M) : 0 \to \bigoplus_{\alpha \in Q_1} P_{h\alpha} \otimes_k e_{t\alpha} M \xrightarrow{d^M} \bigoplus_{x \in Q_0} P_x \otimes_k e_x M \xrightarrow{f^M} M \to 0$$

where the maps are induced by

$$f^M(p \otimes m) = p \cdot m$$

and

$$d^M(p e_{h\alpha} \otimes e_{t\alpha} m) = p \otimes (\alpha \cdot m) - (p\alpha) \otimes m.$$

The maps d^M and f^M are not mysterious. f^M is simply multiplication, taking advantage of the linear action of any path on the vector space at its tail. For d^M, we have a path $p \in P_{h\alpha}$ beginning at $h\alpha$, and a vector $m \in M(t\alpha)$. There are two natural ways we could act by α to obtain a term of the central sum, landing in the summands corresponding to $h\alpha$ and $t\alpha$ respectively. Namely, we could act on m by α to get $p \otimes (\alpha \cdot m) \in P_{h\alpha} \otimes e_{h\alpha} M$, or we could extend the path p by precomposing with α, giving $(p\alpha) \otimes m \in P_{t\alpha} \otimes e_{t\alpha} M$. To define d^M we do both, and take their difference to ensure $\mathcal{C}(M)$ is a chain complex. The global sign is immaterial.

Proof. It is clear that $\mathcal{C}(M)$ is a chain complex, because $f^M \circ d^M$ is zero on pure tensors. Moreover, $e_x M$ is a finite-dimensional k-vector space, so each summand $P_y \otimes_k e_x M \cong P_y^{\oplus \dim_k e_x M}$ is projective. Therefore the two left-most terms in the complex are projective.

For surjectivity of f^M, let $m \in M$. Then

$$f^M \left(\sum_{x \in Q_0} e_x \otimes e_x m \right) = \sum_{x \in Q_0} e_x^2 \cdot m = \left(\sum_{x \in Q_0} e_x \right) \cdot m = 1 \cdot m = m$$

giving surjectivity.

For injectivity of d^M, suppose for a contradiction that $d^M(S) = 0$ and $S \neq 0$. Write

$$S = \sum_{\alpha \in Q_1} p_\alpha \otimes m_\alpha$$

with $m_\alpha \in e_{t\alpha} M, p_\alpha \in P_{h\alpha}$ for all $\alpha \in Q_1$. Label the vertices of Q by $\{1, \ldots, n\}$ as in Lemma 2.10 such that $h\alpha < t\alpha$ for every $\alpha \in Q_1$. Choose a maximal vertex y such that there is an arrow with $t\alpha = y$ and $p_\alpha \otimes m_\alpha \neq 0$. Then for any arrow β with $h\beta = y$, we have $t\alpha > y$ and $p_\beta \otimes m_\beta = 0$. For any $y \in Q_0$, the component of $d^M(S)$ lying in the summand $P_y \otimes e_y M$ is then

$$0 = -\sum_{i=1}^{t} p_{\alpha_i} \alpha_i \otimes m_{\alpha_i}$$

where $\alpha_1, \ldots, \alpha_t$ are precisely the arrows whose tails are y. Now, the paths $p_{\alpha_1}\alpha_1, \ldots, p_{\alpha_t}\alpha_t$ are all distinct since their first segments are different, so they are linearly independent in kQ. Thus $m_{\alpha_i} = 0$ for all $1 \leq i \leq t$. But this contradicts our assumption on y.

For exactness in the middle, we first consider a special case and then use a dimension-counting argument. When $M = S_x$ is simple, our resolution becomes

$$0 \to \bigoplus_{\substack{\alpha \in Q_1 \\ t\alpha = x}} P_{h\alpha} \xrightarrow{d^x} P_x \xrightarrow{f^x} S_x \to 0$$

where $d^x(p) = -p\alpha$ for $p \in P_{h\alpha}$ and $f^x(p) = e_x p \cdot 1$ for $p \in P_x$. If $p \in \ker(f^x) \subset P_x$, then $0 = f^x(p) = e_x p \cdot 1$ and so the coefficient of e_x in p must be zero. Thus p is a linear combination of paths of length at least one, and it is enough to show that any path beginning at x of length at least one appears in the image of d^x. Let p be such a path, and write $p = p'\alpha$ where $tp' = h\alpha$ and $t\alpha = x$. Then $d^x(p') = p$ as desired.

Exactness of this special case yields the following equality of dimensions, for any $x \in Q_0$:

$$\dim_k P_x = \sum_{\substack{\alpha \in Q_1 \\ t\alpha = x}} \dim_k P_{h\alpha} + 1$$

since $\dim_k S_x = 1$. Now in the general case, we have

$$\begin{aligned} \sum_{x \in Q_0} \dim_k P_x \cdot \dim_k e_x M &= \sum_{x \in Q_0} \left(1 + \sum_{\substack{\alpha \in Q_1 \\ t\alpha = x}} \dim_k P_{h\alpha} \right) \cdot \dim_k e_x M \\ &= \sum_{x \in Q_0} e_x M + \sum_{\alpha \in Q_1} \dim_k P_{h\alpha} \cdot \dim_k e_x M \end{aligned}$$

since each arrow has a unique tail. Thus dimension is additive in our complex, giving exactness in the middle. □

For any Q-representation V, we also have a canonical length one injective resolution. This is given by taking the projective resolution of V^* as a Q^{op}-representation, and dualising. This resolution has the form

$$0 \to V \xrightarrow{f_V} \bigoplus_{x \in Q_0} I_x \otimes_k V(x) \xrightarrow{d_V} \bigoplus_{\alpha \in Q_1} I_{t\alpha} \otimes_k W(h\alpha) \to 0$$

where f_V, d_V are the duals of the maps from the projective resolution of V^*.

The existence of such resolutions means that $\operatorname{Rep} Q \simeq kQ-$mod is a hereditary category for any acyclic quiver Q, with enough injectives and enough projectives. In particular, all the results of Sections 1.2 and 1.3 apply to $\mathcal{D}^b(\operatorname{Rep} Q)$. Moreover, it follows that kQ is a hereditary algebra.

Theorem 2.30. *Let Q be an acyclic quiver. Then the path algebra kQ is hereditary, meaning that any submodule of a projective (left or right) module is itself projective.*

Proof. Let $R_1 \subset R_0$ be a subrepresentation of a projective representation, and define $V = R_0/R_1$ so we have an exact sequence

$$0 \to R_1 \xrightarrow{f} R_0 \xrightarrow{r} V \to 0.$$

We also have the canonical projective resolution

$$0 \to P_1 \xrightarrow{d} P_0 \xrightarrow{p} V \to 0$$

from Theorem 2.29. Then we claim that $R_1 \oplus P_0 \cong R_0 \oplus P_1$. In particular, since R_0 and P_1 are projective and summands of projectives are projective, this shows that R_1 is projective.

A short diagram chase using that R_0 is projective shows id_V lifts to a commuting diagram

$$\begin{array}{ccccccccc} 0 & \longrightarrow & R_1 & \xrightarrow{f} & R_0 & \xrightarrow{r} & V & \longrightarrow & 0 \\ & & \downarrow{\scriptstyle h_1} & & \downarrow{\scriptstyle h_0} & & \downarrow{\scriptstyle \mathrm{id}_V} & & \\ 0 & \longrightarrow & P_1 & \xrightarrow{d} & P_0 & \xrightarrow{p} & V & \longrightarrow & 0 \end{array}$$

We now consider the complex

$$\mathcal{C} : 0 \to R_1 \xrightarrow{a_1} P_1 \oplus R_0 \xrightarrow{a_0} P_0 \to 0.$$

It is enough to show that $\mathcal{C}$ is exact, since the fact that P_0 is projective will then imply it splits. Here $a_1 = (-h_1\ f)^T$ and $a_0 = (d\ h_0)$ are the natural maps. The addition of a sign makes this a chain complex.

We know a_1 is injective because f is. To see that a_0 is surjective, let $x_0 \in P_0$. Since $r = ph_0$ is surjective, there exists $y_0 \in R_0$ with $ph_0(y_0) = r(y_0) = p(x_0)$. Then $x_0 - h_0(y_0) \in \ker p = \operatorname{im} d$ so we can find $x_1 \in P_1$ with $x_0 = d(x_1) + h_0(y_0) \in \operatorname{im}(a_0)$.

Finally, a dimension count gives exactness in the middle. Here we are using that kQ is finite-dimensional over k, because Q is acyclic. From exactness of the original two sequences,

$$\dim_k R_0 - \dim_k R_1 = \dim_k V = \dim_k P_0 - \dim_k P_1$$

and hence $\dim_k(P_1 \oplus R_0) - \dim_k R_1 - \dim_k P_0 = 0$. So, the dimensions are additive in $\mathcal{C}$, giving exactness. □

Remark 2.31 (continuing Example 1.22)**.** We can easily check the derived equivalence $\mathcal{D}^b(\operatorname{Coh}\mathbb{P}^1_{\mathbb{C}}) \simeq \mathcal{D}^b(\operatorname{Rep}_{\mathbb{C}} Q_{\mathrm{Kr}})$ by hand, since both abelian categories are hereditary. More generally, given a smooth variety X we may wish to find a quiver Q such that $\mathcal{D}^b(\operatorname{Coh} X) \simeq \mathcal{D}^b(\operatorname{Rep} Q)$. This can be done in two parts, by finding an algebra A such that $\mathcal{D}^b(\operatorname{Coh} X) \simeq \mathcal{D}^b(\text{mod}-A)$, and a quiver Q with $kQ \cong A^{\mathrm{op}}$. We will first address how such an algebra may be constructed.

Let $\mathcal{A}$ be an abelian category over a field k, and $\mathcal{D}^b(\mathcal{A})$ its bounded derived category. An object $E \in \mathcal{D}^b(\mathcal{A})$ is *exceptional* if $\operatorname{Ext}^i(E, E) = 0$ for every $i \neq 0$, and $\operatorname{Hom}(E, E) = k$. An ordered set of exceptional objects is a *strong exceptional collection* if, in addition, $\operatorname{Hom}(E_j, E_k[l]) = 0$ whenever $l \neq 0$ or $j > k$. Bondal ([Bon90], Theorem 6.2) showed that if X is a smooth manifold, and $(E_0, E_1, \ldots, E_n)$ is a strong exceptional collection which generates $\mathcal{D}^b(\operatorname{Coh} X)$, then

$$\operatorname{RHom}(E, -) : \mathcal{D}^b(\operatorname{Coh} X) \to \mathcal{D}^b(\text{mod}-A)$$

is an equivalence of triangulated categories, where $E = \bigoplus_{i=0}^n E_i$ and $A = \operatorname{End}(E)$. For example, $(\mathscr{O}_{\mathbb{P}^1}, \mathscr{O}_{\mathbb{P}^1}(1))$ is a strong exceptional collection generating $\mathcal{D}^b(\mathbb{P}^1_{\mathbb{C}})$, which gives the equivalence from our example. More generally, Beilinson [Bei78] constructed the strong exceptional sequence $(\mathscr{O}_{\mathbb{P}^n}, \mathscr{O}_{\mathbb{P}^n}(1), \ldots, \mathscr{O}_{\mathbb{P}^n}(n))$ generating $\mathcal{D}^b(\operatorname{Coh}\mathbb{P}^n)$, and showed that

$$\operatorname{RHom}(E, -) : \mathcal{D}^b(\mathbb{P}^n) \to \mathcal{D}^b(\text{mod}-\operatorname{End}(E))$$

is an equivalence, where again $\bigoplus_{a=0}^n \mathscr{O}_{\mathbb{P}^n}(a)$.

Does this allow us to find a quiver Q such that $\mathcal{D}^b(\operatorname{Coh}\mathbb{P}^n) \simeq \mathcal{D}^b(\operatorname{Rep} Q)$? The short answer is no, because the algebra $\operatorname{End}(\bigoplus_{a=0}^n \mathscr{O}_{\mathbb{P}^n}(a))$ is not hereditary when $n > 1$, so cannot be the path algebra of a quiver. The solution is to impose a finite subset $R \subset kQ$ of relations on Q, so that $\operatorname{Rep}(Q, R) \simeq (kQ/\langle R\rangle)-\text{mod}$. Then $kQ/\langle R\rangle$ is not, in general, hereditary. The Beilinson quiver

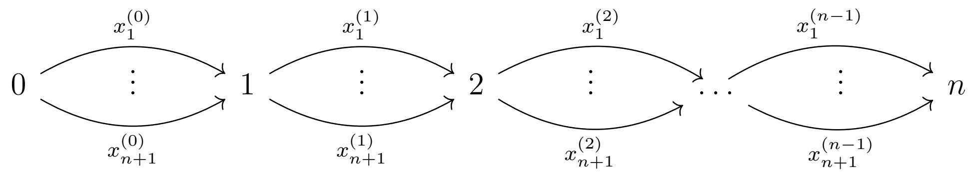

with relations

$$x_i^{(l+1)} x_j^{(l)} = x_j^{(l+1)} x_i^{(l)} \quad \text{for } 0 \leq l \leq n-2 \text{ and } 1 \leq i, j \leq n+1$$

is constructed to have path algebra isomorphic to $\operatorname{End}(\bigoplus_{a=0}^n \mathscr{O}_{\mathbb{P}^n}(a))$, and is therefore derived-equivalent to $\mathbb{P}^n$.

2.5 The Euler form

Let Q be an acyclic quiver. The Euler characteristic of a pair of representations of Q is a homological quantity, defined as the alternating sum of the dimensions of extension groups. As we saw in the previous section, any quiver representation has a length one projective resolution and so the higher extension groups are trivial. (See also Section 1.3.) Hence the Euler characteristic of a pair of quiver representations (V, W) is given by

$$\chi(V, W) := \dim_k \operatorname{Hom}_Q(V, W) - \dim_k \operatorname{Ext}_Q(V, W).$$

In this section, we will show that the Euler characteristic can also be computed via a bilinear form on the space of dimension vectors. This shows that the Euler characteristic depends only on the dimension vectors involved. The Euler form is closely related to the adjacency matrix of the quiver, and gives us a means of determining numerical quantities of the quiver Q from the category $\operatorname{Rep} Q$. This will allow us to show that the path algebra of an acyclic quiver uniquely determines the quiver, as does the category $\operatorname{Rep} Q$.

Definition 2.32. The *Euler form* is a bilinear form on $\mathbb{R}^{Q_0}$ defined by

$$\langle v, w \rangle = \sum_{x \in Q_0} v_x w_x - \sum_{\alpha \in Q_1} v_{t\alpha} w_{h\alpha}.$$

Proposition 2.33. *The Euler form computes the Euler characteristic. That is, for representations V and W of Q,*

$$\chi(V, W) = \langle \dim_k V, \dim_k W \rangle.$$

Proof. We have an exact sequence

$$0 \to \operatorname{Hom}_Q(V, W) \to \bigoplus_{x \in Q_0} \operatorname{Hom}_k(V(x), W(x)) \xrightarrow{d_W^V} \bigoplus_{\alpha \in Q_1} \operatorname{Hom}_k(V(t\alpha), W(h\alpha)) \to \operatorname{Ext}_Q(V, W) \to 0$$

where the map d_W^V is defined by

$$d_W^V((\phi(x))_{x \in Q_0}) = (\phi(h\alpha) \circ V(\alpha) - W(\alpha) \circ \phi(t\alpha))_{\alpha \in Q_1}.$$

The kernel of d_W^V is naturally identified with $\operatorname{Hom}_Q(V, W)$, because a collection ϕ of linear maps at the vertices of Q is a morphism of quiver representations precisely when the commutativity condition checked by d^V_W holds for every arrow. A computation shows

$$d^V_W = \mathrm{Hom}_Q(d^V, W) = \mathrm{Hom}_Q(V, d_W)$$

where d^V, d_W are the maps from the canonical projective resolution for V and injective resolution for W respectively. Hence $\mathrm{coker}(d^V_W)$ is $\mathrm{Ext}_Q(V, W)$. This shows the sequence is exact as claimed.

Then the dimensions are additive, giving

$$\begin{aligned}\chi(V, W) &= \sum_{x\in Q_0} \dim \mathrm{Hom}(V(x), W(x)) - \sum_{\alpha\in Q_1} \dim \mathrm{Hom}(V(t\alpha), W(h\alpha)) \\ &= \sum_{x\in Q_0} \dim V(x) \cdot \dim W(x) - \sum_{\alpha\in Q_1} \dim V(t\alpha) \cdot \dim W(h\alpha) \\ &= \langle \dim_k V, \dim_k W\rangle.\end{aligned}$$

The last equality here is the motivation for the definition of the Euler form. □

An immediate corollary is that the Euler characteristic depends only on the dimension vectors of V and W. Moreover, Proposition 2.36 and Corollary 2.38 tell us that an acyclic quiver is uniquely determined respectively by its representation category and its path algebra. In the next chapter, we will see that the same is not true on the derived level — two non-isomorphic quivers can have equivalent derived categories, indeed we will construct a family of such equivalences.

Lemma 2.34. *Let Q be acyclic. Then for any pair of vertices $i, j \in Q_0$, we have*

$$\chi(S_i, S_j) = \begin{cases} -|\{\alpha \in Q_1 \mid t\alpha = i, h\alpha = j\}| & \text{if } i \neq j \\ 1 & \text{if } i = j. \end{cases}$$

In particular, when $i \neq j$ then $\dim_k \mathrm{Ext}_Q(S_i, S_j) = |\{\alpha \in Q_1 \mid t\alpha = i, h\alpha = j\}|$ counts the number of arrows from i to j in Q. Moreover, $\mathrm{Ext}_Q(S_i, S_i) = 0$.

Proof. We compute the Euler form for $e_i = \dim_k S_i$ and $e_j = \dim_k S_j$. If $i \neq j$, we have

$$\begin{aligned}\langle e_i, e_j\rangle &= \sum_{x\in Q_0} \dim S_i(x) \cdot \dim S_j(x) - \sum_{\alpha\in Q_1} \dim S_i(t\alpha) \cdot \dim S_j(h\alpha) \\ &= 0 - \sum_{\substack{\alpha\in Q_1 \\ t\alpha=i, h\alpha=j}} 1 \\ &= -|\{\alpha \in Q_1 \mid t\alpha = i, h\alpha = j\}|.\end{aligned}$$

In the case $i = j$, the first sum evaluates to 1, while the second sum is zero since there are no loops at the vertex i in the acyclic quiver Q.

When $i \neq j$, the representations S_i, S_j are concentrated at distinct vertices so $\mathrm{Hom}_Q(S_i, S_j) = 0$. When $i = j$ we have $\dim \mathrm{Hom}_Q(S_i, S_i) = 1$ since S_i is one-dimensional concentrated at a single vertex. In either case we compute $\mathrm{Ext}_Q(S_i, S_j)$ using Proposition 2.33 and our above calculation of $\chi(S_i, S_j)$. □

Remark 2.35. A bilinear form on a finite-dimensional vector space can be represented by a matrix. Given a basis $\{e_i\}_{i \in I}$, the entries of the corresponding matrix M are given by $M_{ij} := \langle e_i, e_j \rangle$ and then $\langle v, w \rangle = v^t M w$. In the case of the Euler form, the dimension vectors of the simple representations are the standard basis of $\mathbb{R}^{Q_0}$ so the matrix with entries

$$M_{ij} = \chi(S_i, S_j)$$

represents the Euler form in this basis. By Lemma 2.34, the off-diagonal entries in this matrix are non-positive integers which are exactly the negatives of the corresponding entries in the (directed) adjacency matrix for Q. The diagonal entries are all ones.

By Lemma 2.10, there exists a numbering of the vertices such that $h\alpha < t\alpha$ for every arrow $\alpha \in Q_1$. If we order our basis of simple representations in this way, then the matrix representing the Euler form is upper-triangular.

Proposition 2.36. *Let Q, Q' be acyclic. Then* $\mathrm{Rep}\, Q \simeq \mathrm{Rep}\, Q'$ *if and only if Q and Q' are isomorphic directed graphs.*

Proof. Only the forwards direction is interesting. By Theorem 2.8 the simple objects in each category are in bijection with the vertices of the corresponding quiver. Since equivalences preserve simple objects, the map $\mathrm{Rep}\, Q \xrightarrow{\sim} \mathrm{Rep}\, Q'$ on objects induces a bijection between the vertices of Q and Q'. Use this bijection to give the vertices a common labelling $1, \ldots, n$. To show that we have an isomorphism of directed graphs, it is enough to show that for each pair of vertices (i, j) there are the same number of arrows $i \to j$ in Q and Q'.

For any pair of vertices, by Lemma 2.34 the Euler form $\langle S_i, S_j \rangle_Q$ captures the number of such arrows in Q, and similarly for Q'. But by Proposition 2.33 the Euler form agrees with the Euler characteristic, which is purely homological and so is preserved under equivalences. Hence $Q \cong Q'$. □

Remark 2.37. In fact, $\mathrm{Rep}\, Q$ determines the quiver Q even when Q is not acyclic. If we allow oriented cycles but no loops at any vertex, then Q has infinitely many simple representations, but the 1-dimensional simple representations are still in bijection with vertices of Q. An equivalence of categories $\operatorname{Rep} Q \simeq \operatorname{Rep} Q'$ preserves dimension, so still gives a bijection between vertices of Q and Q'. Extensions between 1-dimensional simple modules again determine the arrows in Q, but the proof of this is more involved since the projective representations P_x are no longer finite-dimensional. One uses the natural grading on kQ by path-length to reduce to a finite-dimensional situation.

When Q has loops there are more 1-dimensional simples, but a similar argument can be made by choosing an appropriate 1-dimensional simple at each vertex.

Proposition 2.36 is somewhat dismaying, as the representation category of a quiver is entirely sensitive to a choice of orientation and cannot be determined from the underlying graph. However, we have already seen that there are many aspects of the theory that do not depend on the choice of orientation, such as the simple representations in an acyclic quiver or the classification of indecomposable representations in an ADE quiver. This leads us to hope that there is some way of unifying the representation theory of various orientations on the same underlying graph. Indeed, this can be done by passing to the derived category. As long as the underlying graph is acyclic, we will be able to construct equivalences between the derived representation categories of any two orientations. In particular, this means that the derived category depends only on the underlying graph, and the abelian categories corresponding to various orientations all arise as full subcategories.

Corollary 2.38. *Two acyclic quivers are isomorphic (as directed graphs) if and only if their path algebras are isomorphic k-algebras.*

Proof. An isomorphism of directed graphs induces an isomorphism of path algebras. Conversely, if $kQ \cong kQ'$ then Proposition 2.18 gives

$$\operatorname{Rep} Q \simeq kQ\text{−mod} \simeq kQ'\text{−mod} \simeq \operatorname{Rep} Q'$$

so the result follows from Proposition 2.36. □

Although acyclic quivers with the same path algebra are necessarily isomorphic, it is not true that every automorphism of the path algebra arises from an automorphism of the quiver. For example, when $n \geq 2$ the equioriented[1] quiver

[1] Equioriented means all arrows point in the same direction.

A_n has no nontrivial automorphisms as a directed graph, but its path algebra does have nontrivial automorphisms. In this sense the path algebra is a richer object of study than the quiver itself.

Example 2.39. Let us explicitly construct a nontrivial automorphism of kA_n^{eq}, where A_n^{eq} denotes the equioriented quiver with underlying graph A_n.

$$1 \overset{a_1}{\to} 2 \overset{a_2}{\to} 3 \overset{a_3}{\to} \dots \overset{a_{n-1}}{\to} n$$

Let $p := a_{n-1}a_{n-2} \dots a_2 a_1$ be the path from the first to the last vertex. Then for $\lambda \in k^\times$, we claim that the linear map sending

$$e_1 \mapsto e_1 + \lambda p$$
$$e_n \mapsto e_n - \lambda p$$

and leaving all other paths fixed is a nontrivial automorphism of the path algebra. The matrix for this map in the basis of paths is lower triangular with ones on the diagonal, so has determinant 1 and is thus an automorphism as k-vector spaces. To see that this map respects the multiplication, we observe that for any path $q \neq e_1, e_n$ in A_n^{eq}, we have $qp = 0 = pq$. So, pairwise products of paths involving at least one such q are automatically preserved. Finally, it is easy to check that the four pairwise products involving e_1 and e_n are preserved.

This construction realises the additive group k as a subgroup of $\mathrm{Aut}(kA_n^{eq})$, although $\mathrm{Aut}(A_n^{eq})$ is trivial.

Given a vertex $x \in Q_0$, we can define a linear operator $\sigma_x : \mathbb{R}^{Q_0} \to \mathbb{R}^{Q_0}$ which is a *reflection*, in the sense that it has order 2 and fixes a subspace of dimension $|Q_0| - 1$. These reflections respect the Euler form. The reflection functors we construct in the next chapter are a categorification of the reflections σ_x, in the sense that they lift the map of dimension vectors to a functor between representation categories.

Definition 2.40. Let $x \in Q_0$. The linear map $\sigma_x : \mathbb{R}^{Q_0} \to \mathbb{R}^{Q_0}$ is defined by

$$\sigma_x\left(\sum_{y \in Q_0} \lambda_y e_y\right) = \sum_{y \in Q_0 \setminus \{x\}} \lambda_y e_y + \left(\sum_{\substack{\alpha \in Q_1 \\ h\alpha = x}} \lambda_{t\alpha} + \sum_{\substack{\alpha \in Q_1 \\ t\alpha = x}} \lambda_{h\alpha} - \lambda_x\right) e_x.$$

That is, in coordinates w.r.t. the standard basis $\{e_y\}_{y \in Q_0}$, the map σ_x modifies the e_x coordinate using the arrows incident at x.

A calculation shows that $\sigma_x^2 = \mathrm{id}$ because Q is acyclic. Moreover, although σ_x does not preserve the Euler form itself, it does preserve the symmetric bilinear form

$$(v, w) := \langle v, w \rangle + \langle w, v \rangle.$$

Remark 2.41. With respect to the above symmetric bilinear form, the automorphism σ_x is a reflection[2] in the sense that

$$\sigma_x(v) = v - 2\frac{(v, e_x)}{(e_x, e_x)} e_x = v - (v, e_x) e_x$$

since one can compute $(e_x, e_x) = 2$. It has an eigenbasis with eigenvalues ± 1. To find eigenvectors, we note that

$$\sigma_x(e_x) = \left(2 \sum_{\substack{\alpha \in Q_1 \\ h\alpha = x, t\alpha = x}} 1 - 1 \right) e_x = -e_x$$

while for $y \neq x \in Q_0$,

$$\sigma_x(e_y) = e_y + \left(\sum_{\substack{\alpha \in Q_1 \\ h\alpha = x, t\alpha = y}} 1 + \sum_{\substack{\alpha \in Q_1 \\ t\alpha = x, h\alpha = y}} 1 - 0 \right) e_x = e_y + n_{xy} e_x$$

where n_{xy} is the nonnegative integer which counts the numbers of arrows between x and y in Q, in either direction. In particular, this means that $e_y + \frac{n_{xy}}{2} e_x$ is fixed by σ_x, so the set

$$\{e_x\} \sqcup \left\{ e_y - \frac{n_{xy}}{2} e_x \;\middle|\; y \neq x \in Q_0 \right\}$$

is an eigenbasis.

In the next chapter, we will construct functors between the representation categories of two quivers which are related by changing the orientation of all the edges incident at a vertex $x \in Q_0$, and such that the induced map on dimension vectors is precisely σ_x. Such functors will give equivalences between the derived categories of quivers with the same underlying graph but distinct edge orientations.

[2] See Appendix A for the definition of a reflection, and some context on root systems. From this perspective, the symmetric bilinear form $(-, -)$ is the real inner product that comes with the root system.

Chapter 3

(Derived) reflection functors

In this chapter, we develop the theory of reflection functors and their counterparts on the derived category. Reflection functors were introduced by Bernstein, Gelfand and Ponomarev in [BGP73], to give a proof of Gabriel's Theorem which exploited the connection to root systems and used similar techniques to the classification of simple Lie algebras. They exist between the abelian categories $\operatorname{Rep} Q$ and $\operatorname{Rep} Q'$ when Q and Q' have the same underlying graph, but different choices of edge orientation, where Q' is obtained from Q by changing all the edge orientations incident at a single source or sink. They are not equivalences on the level of abelian categories, but in this chapter we will prove that they are equivalences on the level of the derived category. In the previous chapter we saw that the abelian categories $\operatorname{Rep} Q$ and $\operatorname{Rep} Q'$ are equivalent if and only if Q and Q' are isomorphic directed graphs. However, this is not true on the derived level, and reflection functors will give a partial answer to the question of when two non-isomorphic quivers are derived-equivalent. We will also show that a reflection C_x^+ is representable, while C_x^- can be written as a tensor. This will allow us to describe derived reflections several different ways.

In the last section, we discuss the relationship between the abelian categories coming from different orientations of the same underlying graph. We use the fact that all such abelian categories occur as full subcategories of a common derived category to relate them via tilting at a torsion pair on the derived category. Throughout this chapter, we consider only acyclic quivers.

3.1 Reflections on Rep Q

Let $Q = (Q_0, Q_1)$ be an acyclic quiver. Recall (see Definition 2.9) that a vertex $x \in Q_0$ is a sink if all the arrows incident at x are directed towards x, and a source if all arrows incident at x are directed away from x. When $x \in Q_0$ is a source or a sink of Q, we let $\sigma_x Q$ be the quiver obtained from Q by reversing each of the arrows incident at x. In particular, note that if x is a source in Q, then it becomes a sink in $\sigma_x Q$, and vice versa.

This operation on Q would make sense even if the vertex x were not a sink or source, but in the case where it is, we will construct a corresponding functor $C_x^{\pm} : \operatorname{Rep} Q \to \operatorname{Rep} \sigma_x Q$. This functor is defined only when x is a source or sink, so we will restrict the notation $\sigma_x Q$ to this case as well to avoid confusion.

Definition 3.1 (Reflection functors)**.** Let Q be an acyclic quiver.

(i) Given a sink $x \in Q_0$, we define the functor $C_x^+ : \operatorname{Rep} Q \to \operatorname{Rep} \sigma_x Q$ as follows. For a representation V of Q, we set $C_x^+ V(y) = V(y)$ at each vertex $y \neq x$, while $C_x^+ V(x)$ is defined so that the following is (left) exact:

$$0 \to C_x^+ V(x) \xrightarrow{\tau} \bigoplus_{\substack{\alpha \in Q_1 \\ h\alpha = x}} V(t\alpha) \xrightarrow{\xi} V(x) \tag{3.1}$$

where $\xi = \sum_{\substack{\alpha \in Q_1 \\ h\alpha = x}} V(\alpha)$ is the natural map. That is, $C_x^+ V(x) = \ker \xi$, thought of as a subspace of $\bigoplus_{h\alpha = x} V(t\alpha)$.

For arrows $\alpha \in (\sigma_x Q)_1$, let $C_x^+ V(\alpha) = V(\alpha)$ when $t\alpha \neq x$. If $t\alpha = x$, then the map $C_x^+ V(\alpha) : \ker \xi \to V(h\alpha)$ is defined to be τ post-composed with the canonical projection onto the summand $V(h\alpha)$.

Finally, for a morphism $f : V \to W$, we again let $C_x^+ f(y) = f(y)$ for vertices $y \neq x$, while $C_x^+ f(x)$ is given by restricting the natural map

$$\bigoplus_{\substack{\alpha \in Q_1 \\ h\alpha = x}} f(t\alpha) : \bigoplus_{\substack{\alpha \in Q_1 \\ h\alpha = x}} V(t\alpha) \to \bigoplus_{\substack{\alpha \in Q_1 \\ h\alpha = x}} W(t\alpha).$$

(ii) Given a source $x \in Q_0$, we define the functor $C_x^- : \operatorname{Rep} Q \to \operatorname{Rep} \sigma_x Q$ dually to the above. For a representation V of Q, set $C_x^- V(y) = V(y)$ for $y \neq x$ and define $C_x^- V(x)$ so the following is (right) exact:

$$V(x) \xrightarrow{\zeta} \bigoplus_{\substack{\alpha \in Q_1 \\ t\alpha = x}} V(h\alpha) \xrightarrow{\rho} C_x^- V(x) \to 0. \tag{3.2}$$

That is, $C_x^- V(x)$ is the cokernel of the natural map $\zeta = (V(\alpha))_{\alpha \in Q_1, t\alpha = x}$.

For an arrow $\alpha \in (\sigma_x Q)_1$ with $h\alpha \neq x$, set $C_x^- V(\alpha) = V(\alpha)$, and if $h\alpha = x$ then we let $C_x^- V(\alpha) : V(t\alpha) \to \operatorname{coker} \zeta$ be the restriction of ρ to the summand $V(t\alpha)$.

For a morphism $f : V \to W$, let $C_x^- f(y) = f(y)$ at a vertex $y \neq x$, and let $C_x^- f(x) : \operatorname{coker} \zeta_V \to \operatorname{coker} \zeta_W$ be the map between the cokernels induced from

$$\bigoplus_{\substack{\alpha \in Q_1 \\ h\alpha = x}} f(t\alpha) : \bigoplus_{\substack{\alpha \in Q_1 \\ h\alpha = x}} V(t\alpha) \to \bigoplus_{\substack{\alpha \in Q_1 \\ h\alpha = x}} W(t\alpha).$$

Notation 3.2. Throughout this chapter, we will take the convention that x is a sink of Q (and thus a source of $\sigma_x Q$), and consider the functors

$$C_x^+ : \operatorname{Rep} Q \to \operatorname{Rep} \sigma_x Q, \quad C_x^- : \operatorname{Rep} \sigma_x Q \to \operatorname{Rep} Q.$$

We make a few initial observations regarding the definition. First, C_x^+, C_x^- are well-defined functors. Second, reflecting twice at the same vertex gives a functor $\operatorname{Rep} Q \to \operatorname{Rep} Q$. This functor preserves $V(y)$ for any vertex $y \neq x$, and similarly preserves $V(\alpha)$ when α is not incident at x, but usually does not preserve $V(x)$ or the linear maps corresponding to arrows incident at x. This is because the sequences (3.1) and (3.2) are not in general exact. In fact, $C_x^- C_x^+ V \cong V$ if and only if the sequence (3.1) is exact on the right, and dually for the other composition. We have a natural transformation ι_x comparing $C_x^- C_x^+$ to the identity, defined by the collection of monomorphisms

$$\iota_x V : C_x^- C_x^+ V \to V$$

where $\iota_x V(y) = \mathrm{id}_{V(y)}$ for $y \neq x$, and $\iota_x V(x)$ is the canonical inclusion

$$C_x^- C_x^+ V(x) = \operatorname{coker} \ker \xi \cong \operatorname{im} \xi \hookrightarrow V(x).$$

Similarly, there is a natural transformation π_x defined by the epimorphisms

$$\pi_x V : V \to C_x^+ C_x^- V$$

where $\pi_x V(y) = \mathrm{id}_{V(y)}$ for $y \neq x$, and $\pi_x V(x)$ is the canonical quotient map

$$V(x) \twoheadrightarrow \operatorname{im} \zeta \cong \ker \operatorname{coker} \zeta = C_x^+ C_x^- V(x).$$

Then $C_x^- C_x^+ V \cong V$ if and only if $\iota_x V$ is an isomorphism, and similarly for $C_x^+ C_x^-$ and π_x.

The reflection functors C_x^+, C_x^- are never equivalences between the categories $\operatorname{Rep} Q$ and $\operatorname{Rep} \sigma_x Q$, but they are close to being equivalences. We will make this precise in Theorem 3.6. Consider the result of reflecting the simple representation S_x at the vertex x. Since all the spaces $S_x(y)$ for $y \neq x$ are zero, the maps τ and ξ when x is a sink or ζ and ρ when x is a source are zero, because the middle term in (3.1) or (3.2) is zero. Hence $C_x^+ S_x^Q = 0 = C_x^- S_x^{\sigma_x Q}$. If C_x^+, C_x^- were equivalences then they would send indecomposable objects to indecomposable objects (and simples to simples), but reflecting at a vertex annihilates the simple representation at that vertex.

However, this is essentially the only way in which reflection functors fail to be equivalences. Every other indecomposable representation of Q reflects to an indecomposable representation of $\sigma_x Q$, and the two reflections $C_x^+ : \operatorname{Rep} Q \to \operatorname{Rep} \sigma_x Q$ and $C_x^- : \operatorname{Rep} \sigma_x Q \to \operatorname{Rep} Q$ give inverse bijections between the indecomposable objects of the two categories, excluding the simple representation at x in each. Moreover, they are an adjoint pair.

Proposition 3.3. *The functors (C_x^-, C_x^+) are an adjoint pair. That is, for a Q-representation V and a $\sigma_x Q$-representation W, we have*

$$\operatorname{Hom}_Q(C_x^- W, V) \cong \operatorname{Hom}_{\sigma_x Q}(W, C_x^+ V)$$

and these isomorphisms are natural.

Proof. We prove this using the unit-counit characterisation of an adjunction, see [Mac88], Chapter IV, Theorem 2. The unit is ι_x and the counit is π_x. We must check that the compositions

$$\begin{aligned} \mathrm{id}_{C_x^- W} &= \iota_x(C_x^- W) \circ C_x^-(\pi_x W) \\ \mathrm{id}_{C_x^+ V} &= C_x^+(\iota_x V) \circ \pi_x(C_x^+ V) \end{aligned}$$

are satisfied, for any $V \in \operatorname{Rep} Q$ and $W \in \operatorname{Rep} \sigma_x Q$. We show only the first equality, as the second is similar. Checking at each vertex of Q separately, the result is immediate for any vertex $y \neq x$ since both reflections are the identity there. At x, we have the following commuting diagram:

$$\begin{array}{ccccccccc}
 & & W(x) & \longrightarrow & \bigoplus_{t\alpha=x} W(h\alpha) & \longrightarrow & C_x^- W(x) & \longrightarrow & 0 \\
 & & \downarrow{\scriptstyle \pi_x W(x)} & & \downarrow{\scriptstyle \mathrm{id}} & & \downarrow{\scriptstyle C_x^-(\pi_x W)(x)} & & \\
0 & \longrightarrow & C_x^+ C_x^- W(x) & \longrightarrow & \bigoplus_{t\alpha=x} W(h\alpha) & \longrightarrow & C_x^- C_x^+ C_x^- W(x) & \longrightarrow & 0 \\
 & & \downarrow{\scriptstyle \mathrm{id}} & & \downarrow{\scriptstyle \mathrm{id}} & & \downarrow{\scriptstyle \iota_x(C_x^- W)(x)} & & \\
0 & \longrightarrow & C_x^+ C_x^- W(x) & \longrightarrow & \bigoplus_{t\alpha=x} W(h\alpha) & \longrightarrow & C_x^- W(x) & \longrightarrow & 0
\end{array}$$

with exact rows which compute various reflections of W. The composition we are interested in is the last column; call this map φ. If we consider just the rectangle on the right, we have

$$\begin{array}{ccc} \bigoplus_{t\alpha=x} W(h\alpha) & \twoheadrightarrow & C_x^- W(x) \\ \downarrow \mathrm{id} & & \downarrow \varphi \\ \bigoplus_{t\alpha=x} W(h\alpha) & \twoheadrightarrow & C_x^- W(x) \end{array}$$

so clearly $\varphi = \mathrm{id}_{C_x^- W(x)}$ as desired. □

More informally, this is true essentially because $C_x^- C_x^+ C_x^- W$ can be canonically identified with $C_x^- W$ and then $\iota_x(C_x^- W), C_x^-(\pi_x W)$ are each the identity map between these two representations.

Corollary 3.4 (Lemma 7.4.1 in [Kra08]). *C_x^- is right-exact, C_x^+ is left exact, and both preserve direct sums.*

Proof. Right adjoints preserve limits, and the direct sum is a product, which is a limit. Hence C_x^+ is additive. Kernels are also limits, so C_x^+ preserves kernels and is left-exact. Similarly, left adjoints preserve colimits, and the direct sum is a coproduct, which is a colimit. Cokernels are also colimits. So, C_x^- is additive and right-exact. □

We now give a description of the key properties of reflection functors on the abelian level.

Lemma 3.5. *(i) $V \in \mathrm{Rep}\, Q$ has a summand isomorphic to S_x if and only if $\iota_x V(x)$ is not surjective.*

(ii) $V \in \mathrm{Rep}\, \sigma_x Q$ has a summand isomorphic to S_x if and only if $\pi_x V(x)$ is not injective.

Proof. Since C_x^+, C_x^- are both additive, ι_x and π_x split across direct sums. So, it is enough to check that for V indecomposable, $\iota_x V(x)$ is surjective iff $V \not\cong S_x$, and similarly for (ii).

Let V be any indecomposable. Then $\iota_x V(x)$ is surjective iff $\mathrm{im}(\xi) = V(x)$, that is iff

$$\xi : \bigoplus_{\substack{\alpha \in Q_1 \\ h\alpha = x}} V(t\alpha) \to V(x)$$

is surjective. Let W be the Q-representation obtained from V by replacing $V(x)$ with $\mathrm{im}(\xi)$. Then V splits as $V \cong W \oplus S_x^n$ where $n = \dim V(x) - \dim \mathrm{im}(\xi)$, since each map $V(\alpha)$ for $h\alpha = x$ splits as a direct sum in this way. In particular, because V is indecomposable, either $W = 0$ is surjective or $n = 0$. If $V \ncong S_x$ then $n = 0$ and ξ is surjective. If $V \cong S_x$ then $n = 1$ so $W = 0$ and ξ is not surjective.

A similar argument shows that if x is a source, then V splits as a direct sum $V \cong W \oplus S_x^n$ where $n = \dim(\ker \zeta)$ and so $\pi_x V(x)$ is injective iff V does not contain S_x as a summand. □

In the following theorem, σ_x denotes the map $\mathbb{Z}^{Q_0} \to \mathbb{Z}^{Q_0} = \mathbb{Z}^{\sigma_x Q_0}$ defined in Section 2.5, which is known to respect the symmetrisation of the Euler form.

Theorem 3.6 (Bernstein-Gelfand-Ponomarev, Thm 4.3.9 in [DW17])**.**

(i) Let $V \in \operatorname{Rep} Q$ be indecomposable. Then $V \cong S_x$ if and only if $C_x^+ V = 0$. If $V \ncong S_x$ then $C_x^+ V$ is indecomposable of dimension vector $\sigma_x(\dim_k V)$, and moreover $C_x^- C_x^+ V \cong V$.

(ii) Let $V \in \operatorname{Rep} \sigma_x Q$ be indecomposable. Then $V \cong S_x$ if and only if $C_x^- V = 0$. If $V \ncong S_x$ then C_x^- is indecomposable of dimension vector $\sigma_x(\dim_k V)$, and moreover $C_x^+ C_x^- V \cong V$.

Proof. We prove only (i) since (ii) is dual. We have already seen that $C_x^+ S_x = 0$. If $C_x^+ V = 0$ then $V(y) = 0$ for every $y \neq x$, so V is a direct sum of copies of S_x. But V is indecomposable, so $V \cong S_x$.

Now suppose $V \ncong S_x$. By the Lemma, $\iota_x V(x)$ is surjective and hence an isomorphism. But this implies $\iota_x V$ is an isomorphism $C_x^- C_x^+ V \cong V$, since $\iota_x V(y)$ for $y \neq x$ is always bijective. In particular, the sequence (3.1) is exact. Let $D := \dim_k V$; then since dimensions are additive in short exact sequences, we have

$$\dim_k(C_x^+ V)(x) = \sum_{\substack{\alpha \in Q_1 \\ h\alpha = x}} \dim V(t\alpha) - \dim V(x) = \sum_{h\alpha = x} D(t\alpha) - D(x) = \sigma_x(D)(x).$$

But also, $\dim_k(C_x^+ V)(y) = \dim_k V(y) = \sigma_x(D)(y)$ for every $y \neq x$, so $\dim_k C_x^+ V = \sigma_x(\dim_k V)$ as desired. A similar proof shows that when x is a source, we have $V \cong S_x$ if and only if $C_x^- V = 0$ and the same dimension calculation goes through.

Let $V \ncong S_x$. To show $C_x^+ V$ is indecomposable, write it in terms of its indecomposable summands as

$$C_x^+ V \cong W_1 \oplus \ldots \oplus W_r$$

where $r \geq 1$ because $C_x^+ V \neq 0$. By additivity of C_x^-, we have

$$V \cong C_x^- C_x^+ V \cong C_x^- W_1 \oplus \ldots \oplus C_x^- W_r.$$

The sequences (3.1) for V and (3.2) for $C_x^+ V$ are the same, so both of them are exact. That is, since $\iota_x V$ is an isomorphism, $\pi_x(C_x^+ V)$ is also an isomorphism. Hence by the Lemma, $C_x^+ V$ has no summands isomorphic to S_x. Then the first part of (ii) implies that $C^- W_i \neq 0$ for each i, and so each summand $C_x^- W_i$ of V decomposes into at least one indecomposable piece. Thus V has at least i indecomposable summands. But V is indecomposable, so $i = 1$ and therefore $C_x^+ V$ is indecomposable as claimed. □

Corollary 3.7. *The functors C_x^+ and C_x^- are mutually inverse bijections between the isomorphism classes of indecomposable representations of Q and of $\sigma_x Q$, with the exception of the simple representation S_x which is annihilated by both functors.*

Proof. We saw in the proof of the theorem that for $V \ncong S_x$ indecomposable, we have $C_x^+ V \ncong S_x$ indecomposable also. Hence C_x^+ gives a map

$$\{\text{indecomposables of } Q\} \setminus \{S_x\} \to \{\text{indecomposables of } \sigma_x Q\} \setminus \{S_x\}.$$

Similarly, C_x^- gives a map the other way. But again by the theorem, both compositions are the identity, so these are mutually inverse bijections. □

Since (C_x^-, C_x^+) are an adjoint pair, if either were an equivalence then they would be mutual inverse equivalences. This is almost the case; they are mutual inverses on all but one of the isomorphism classes of indecomposables. Moreover, C_x^+, C_x^- send (indecomposable) projectives to projectives and injectives to injectives, with the exception of the injective and projective indecomposable at the vertex x.

Let P_y^Q denote the indecomposable projective Q-representation at the vertex $y \in Q_0$, and $P_y^{\sigma_x Q}$ the indecomposable $\sigma_x Q$-representation at $y \in Q_0$. Similarly for injectives.

Proposition 3.8. *For each $y \neq x$, we have $C_x^+(I_y^Q) = I_y^{\sigma_x Q}$ and $C_x^-(P_y^{\sigma_x Q}) = P_y^Q$. Hence C_x^-, C_x^+ both preserve the injective and projective indecomposables at vertices $y \neq x$.*

Proof. This follows directly from Definition 3.1, but Theorem 3.6 will reduce the number of calculations. We will show that $C_x^-(P_y^{\sigma_x Q}) = P_y^Q$, the argument in the injective case being dual.

Recall x is a source in $\sigma_x Q$, and $P_y^{\sigma_x Q} = k\sigma_x Q e_y$. For any vertex $z \neq x$, we have

$$C_x^-(P_y^{\sigma_x Q})(z) = P_y^{\sigma_x Q}(z) = e_z k\sigma_x Q e_y = e_z kQe_y = P_y^Q(z)$$

and for any arrow $\alpha \in Q_1$ with $h\alpha \neq x$, similarly

$$C_x^-(P_y^{\sigma_x Q})(\alpha) = P_y^{\sigma_x Q}(\alpha) = P_y^Q(\alpha)$$

since both maps are left-multiplication by α. At x, we have an exact sequence

$$e_x k\sigma_x Q e_y \to \bigoplus_{h\alpha = x} e_{t\alpha} k\sigma_x Q e_y \xrightarrow{\rho} C^-(P_y^{\sigma_x Q})(x) \to 0$$

and since x is a source in $\sigma_x Q$, $e_x k\sigma_x Q e_y = 0$ so ρ is an isomorphism. Then for $\alpha \in Q_1$ with $h\alpha = x$, the map $C_x^-(P_y^{\sigma_x Q})(\alpha) = \rho \circ \iota_\alpha$ is by definition the inclusion of the summand $e_{t\alpha} k\sigma_x Q e_x$ followed by ρ. We also have an identification

$$P_y^Q(x) = e_x kQe_y \xrightarrow{f} \bigoplus_{h\alpha = x} e_{t\alpha} kQe_y = \bigoplus_{h\alpha = x} e_{t\alpha} k\sigma_x Q e_y$$

where the isomorphism f is given by removing the last path segment α from a path ending at x. Then for $\alpha \in Q_1$ with $h\alpha = x$, $P_y^Q(\alpha) = f^{-1} \circ \iota_\alpha$. Thus the isomorphism of vector spaces

$$\rho^{-1} \circ f : P_y^Q(x) \to C_x^-(P_y^{\sigma_x Q})(x)$$

identifies $P_y^Q(\alpha)$ with $C_x^-(P_y^{\sigma_x Q})(\alpha)$ so defines an isomorphism $C_x^-(P_y^{\sigma_x Q}) \cong P_y^Q$ as Q-representations.

After performing the same calculation for $C_x^+(I_y^Q)$, we get by Theorem 3.6 that

$$C_x^-(I_y^{\sigma_x Q}) = C_x^-(C_x^+(I_y^Q)) = I_y^Q$$

and similarly

$$C_x^+(P_y^Q) = C_x^+(C_x^-(P_y^{\sigma_x Q})) = P_y^{\sigma_x Q}$$

so C_x^+, C_x^- preserve both projective and injective indecomposables at vertices $y \neq x$. □

We have now fully described C_x^+ and C_x^- on projective and injective indecomposables, apart from the exceptional values

$$C_x^+(P_x^Q), C_x^+(I_x^Q), C_x^-(P_x^{\sigma_x Q}), \text{ and } C_x^-(I_x^{\sigma_x Q}).$$

But since x is a sink in Q, there are no paths in Q from x to any other vertex and we have $P_x^Q = kQe_x = S_x$. Similarly, $I_x^{\sigma_x Q} = (e_x kQ)^* = S_x$ since x is a source in $\sigma_x Q$. So we get $C_x^-(I_x^{\sigma_x Q}) = 0 = C_x^+(P_x^Q)$. When we introduce the derived reflections LC_x^- and LC_x^+ in the next section, we will see that they map the simple representation S_x to its translation, so identify $I_x^{\sigma_x Q}$ with P_x^Q up to a shift.

But what about the remaining exceptional values $C_x^-(P_x^{\sigma_x Q})$ and $C_x^+(I_x^Q)$? We know by Theorem 3.6 that these are indecomposable, so we might guess that they are interchanged by reflection. Unfortunately, this need not be the case. For example, $C_x^-(P_x^{\sigma_x Q})$ is never projective, and in general is not injective either. We present two examples.

Example 3.9. (i) Let $Q = A_2$.

$$Q : v_1 \to v_2, \qquad \sigma_i Q : v_1 \leftarrow v_2$$

There are 3 indecomposable representations, corresponding to dimension vectors $(1,0), (1,1)$ and $(0,1)$. The projective indecomposables are given by $\dim P_1^Q = (1,1)$ and $\dim P_2^Q = (0,1)$. The injective indecomposables are $\dim I_1^Q = (1,0)$ and $\dim I_2^Q = (1,1)$. In particular the indecomposable with dimension vector $(1,1)$ is both injective and projective. Reflecting at either vertex gives an isomorphic quiver, where the isomorphism swaps the numbers of the two vertices. So we have

$$\dim P_1^{\sigma_i Q} = (0,1), \dim P_2^{\sigma_i Q} = (1,1), \dim I_1^{\sigma_i Q} = (1,1), \dim I_2^{\sigma_i Q} = (1,0)$$

for $i = 1, 2$.

All the indecomposables are either projective or injective, so $C_i^+(I_i^Q) = P_i^{\sigma_i Q}$ by process of elimination.

(ii) Let $Q = D_4$, oriented so the central vertex x is a sink. Then I_x^Q has dimension vector $(1,1,1,1)$ with all the maps isomorphisms.

$$Q = D_4 : \quad \begin{array}{ccccc} & & b & & \\ & & \downarrow & & \\ a & \longrightarrow & x & \longleftarrow & c \end{array} \qquad\qquad I_x : \quad \begin{array}{ccccc} & & k & & \\ & & \downarrow{\scriptstyle \sim} & & \\ k & \xrightarrow{\sim} & k & \xleftarrow{\sim} & k \end{array}$$

The vector space $C_x^+(I_x^Q)(x)$ is the kernel of the sum of the three maps in I_x^Q. In particular, it is the kernel of a map from a 3-dimensional space to a 1-dimensional space, so is 2-dimensional. Thus $\dim C_x^+(I_x^Q) = (1,1,1,2)$. But the representation $P_x^{\sigma_x Q}$ has dimension vector $(1,1,1,1)$, so $C_x^+(I_x^Q)$ is neither injective nor projective.

More generally, if the underlying graph of Q is acyclic then any two vertices have at most one path between them, so the dimension vectors of projective and injective indecomposable representations contain only 0 and 1 (see Remark 2.26). Let us now compute $\dim C_x^+(I_x^Q)(x) = \sigma_x(\dim I_x^Q)(x)$ for a sink x of degree n. We get

$$\sigma_x(\dim I_x^Q)(x) = -1 + \sum_{h\alpha=x} 1 = n - 1.$$

If x has degree at least 3, then the dimension vector of $C_x^+(I_x^Q)$ contains an entry greater than 1, so it cannot be projective.

Remark 3.10. Theorem 3.6 tells us that the map σ_x determines the dimension vector of the reflection C_x^+V or C_x^-V. In general, this does not fully determine the reflection C_x^+, even on indecomposables, since $\sigma_x Q$ might have two distinct indecomposables with the same dimension vector. However, if Q is an ADE-type quiver, then it is part of the statement of Gabriel's Theorem 2.12 that reflection functors on Q are fully determined by σ_x.

Corollary 3.11 (to Thm 3.6)**.** *If Q is an ADE-type quiver then the dimension vectors of the indecomposable Q- and $\sigma_x Q$-representations are the same, and we can view a reflection functor C_x^+, C_x^- as a permutation of these dimension vectors, that fixes the exceptional element* $\dim_k S_x$*. This permutation has order 2.*

Proof. By the Theorem, C_x^+ and C_x^- induce mutually inverse permutations of the dimension vectors. But also, they induce the same permutation, since both are computed by σ_x on all the dimension vectors except $\dim_k S_x$. □

Reflection functors allow us to pass between the representation categories of different quivers with the same underlying graph but different choices of edge orientation. For a given underlying graph, we might ask whether any orientation defining an acyclic quiver can be achieved from any other such orientation by a sequence of reflections. We call a graph *reflection-transitive* if this is the case. Unfortunately, not every graph is reflection-transitive; however, the next proposition gives a characterisation of reflection-transitive graphs, which in particular implies that ADE graphs are reflection-transitive.

Proposition 3.12. *Let G be an undirected graph without loops. Then G is reflection-transitive if and only if G is acyclic, but with multiple edges allowed. That is, G has no cycles that involve more than two vertices, but need not be simple.*

Proof. One direction follows from the fact that for any cycle in G and any choice of orientation on G, the number of clockwise (or equivalently anticlockwise) arrows in the cycle is preserved under reflection. Hence if G has a cycle involving more than 2 vertices, we can find orientations with no sequence of reflections between them by making all the arrows in the cycle except one clockwise in one orientation, and all except one anticlockwise in the other.

The converse is by induction on the number of vertices. First reduce to the case where G is simple. Any orientation of G giving an acyclic quiver must have all arrows between a given pair of vertices oriented in the same direction, and any reflection that changes one of these arrows changes all of them. So we may treat them as a single arrow.

Suppose G is simple and acyclic, that is a tree, and consider two quivers Q, Q' with underlying graph G. If $|V(G)| = 1$ the result is immediate. For $n = |V(G)| > 1$, choose a leaf $v \in V(G)$. Let $w \in V(G)$ be adjacent to z, via the arrow α, which may have distinct orientations in Q and Q'. By induction, if we delete z and α from G, Q and Q', then the resulting quivers $\overline{Q}$ and $\overline{Q}'$ differ by a sequence of reflections. We now consider two cases.

If this sequence does not involve a reflection at w, then the same sequence of reflections can be applied to Q. This will give the quiver Q', except that the arrow α may be oriented the wrong way. We reflect at z if necessary to correct the orientation of α.

If w does appear in the sequence of reflections, we modify the sequence. Each time a reflection at w appears in the sequence, add a reflection at z first if necessary so that after the reflection at z, the orientation of α matches that of all other arrows incident at w. Again, after performing the sequence of reflections we obtain Q', except possibly for the orientation of α, which can be corrected by reflecting at z. $\square$

In fact, for any two acyclic quivers Q, Q' on a reflection-transitive graph G, there exists a sequence of positive reflections (that is, reflections at sinks) transforming Q into Q', and also such a sequence of negative reflections. The proof is much the same, but the sequence of reflections must be modified more carefully in the inductive step.

3.2 Defining derived reflection functors

In this section, we explain why reflection functors give rise to derived reflections, and obtain an initial definition of the functors

$$RC_x^+ : \mathcal{D}^b(\operatorname{Rep} Q) \to \mathcal{D}^b(\operatorname{Rep} \sigma_x Q)$$

and

$$LC_x^- : \mathcal{D}^b(\operatorname{Rep} \sigma_x Q) \to \mathcal{D}^b(\operatorname{Rep} Q).$$

Our discussion relies heavily on Section 1.2.

The following Lemma proves that injectives are an adapted class for C_x^+ and projectives are an adapted class for C_x^-, when taken together with Theorem 2.29 and Corollary 3.4.

Lemma 3.13. *The functors C_x^+, C_x^- are both exact on projectives and exact on injectives.*

Proof. Since $\operatorname{Rep} Q \simeq kQ-$mod is a module category, we use the fact that an exact sequence whose first term is injective or whose last term is projective must be split. Hence any short exact sequence of injectives or projectives is split. But C_x^+, C_x^- are both additive, so they certainly preserve split short exact sequences. □

Note that this result uses no specific properties of the reflection functors at all, except that they are additive. We have therefore essentially shown that any additive, left- or right- exact functor on $\operatorname{Rep} Q$ has corresponding right- or left-derived functors.

Since C_x^- is right-exact, exact on projectives, and $\operatorname{Rep} \sigma_x Q$ has enough projectives, we can define LC_x^-. With the canonical projective resolution

$$0 \to P_1(V) \to P_0(V) \to V \to 0$$

of the representation V, let $LC_x^- V$ be the complex obtained by applying C_x^- to the complex $0 \to P_1(V) \to P_0(V) \to 0$, with $P_0(V)$ in degree 0. Then $L_i C_x^- V := H^i(LC_x^- V)$ are the cohomology pieces. Similarly, to compute RC_x^+ we take the canonical injective resolution

$$0 \to V \to I_0(V) \to I_1(V) \to 0$$

and obtain RC_x^+V by applying C_x^+ to the complex $0 \to I_0(V) \to I_1(V) \to 0$ with $I_0(V)$ in degree 0. By essentially the same proof as the proof that Ext is well-defined, this does not depend on the choice of projective or injective resolution, although it is usually convenient to use the canonical one in calculations.

To define LC_x^-, RC_x^+ on complexes which are not concentrated in a single degree, we have two choices. We can either decompose the complex as a direct sum of its cohomology pieces using Theorem 1.23 and use the above to compute the image of each summand separately, or we can find a complex of projectives or injectives respectively isomorphic to our starting complex, as in Proposition 1.14, and apply C_x^-, C_x^+ to this complex directly. These two approaches give the same result.

We have omitted a variety of well-definedness checks in the above discussion. Rather than giving these proofs in detail, we come back to the question of well-definedness in the next section. We will soon show that C_x^+ is representable, that is it can be written in terms of Hom_Q for a specific representing object. We already saw in Section 1.2 that Ext is well-defined over all appropriate choices of resolution (projective in the first coordinate or injective in the second), so any representable functor is too. Then the tensor-Hom adjunction means LC_x^- is similarly well-defined.

3.3 C_x^+ is representable

In this section we use the adjunction $C_x^- \dashv C_x^+$ to show that C_x^+ is representable, and compute the representing object. As a corollary, C_x^- can be written as a tensor product via the tensor-Hom adjunction. Note that since C_x^- is right-exact covariant (and not left-exact), it cannot be equivalent to a Hom functor, so cannot itself be representable. We will later use our representing object to rewrite derived reflections from several different perspectives. In this section, x denotes a sink of Q that is not an isolated vertex. (If x were an isolated vertex, the reflection at x would be trivial so this case is not interesting.)

Our treatment is inspired by Chapter 6 of [DW17], although we give a more streamlined proof by exploiting the adjunction.

We would like to reinterpret reflections in the language of modules. Consider the object of $\mathrm{Rep}\, Q$

$$T_x := C_x^-(k\sigma_x Q)$$

which we will show represents the reflection C_x^+. Recall that for an R-module M, the functor $\mathrm{Hom}_R(M,-)$ naturally returns an abelian group. If we instead want to consider $\mathrm{Hom}_R(M,-)$ as a functor $R-\mathrm{mod} \to S-\mathrm{mod}$, we need M to have the structure of a right module over S, compatible with its R-module structure. This gives $\mathrm{Hom}_R(M,-)$ the structure of a left S-module by precomposition with the S-action on M. So M should be an (R,S)-bimodule.

In our setting, we need T_x to be a $(kQ, k\sigma_x Q)$-bimodule. The definition of T_x endows it with the structure of a kQ-module, and the following Proposition will give us a compatible right-action by $k\sigma_x Q$.

Proposition 3.14. *There is a natural isomorphism of k-algebras*

$$\mathrm{Hom}_Q(T_x, T_x) \cong k\sigma_x Q^{\mathrm{op}}.$$

Proof. We have via the adjunction $C_x^- \dashv C_x^+$ that

$$\begin{aligned}\mathrm{Hom}_Q(T_x, T_x) &= \mathrm{Hom}_Q(C_x^-(k\sigma_x Q), C_x^-(k\sigma_x Q)) \\ &\cong \mathrm{Hom}_{\sigma_x Q}(k\sigma_x Q, C_x^+ C_x^-(k\sigma_x Q)) \\ &= \mathrm{Hom}_{\sigma_x Q}(k\sigma_x Q, k\sigma_x Q) \\ &\cong (k\sigma_x Q)^{\mathrm{op}}.\end{aligned}$$

Here we are using the fact that $k\sigma_x Q$ does not have $I_x^{\sigma_x Q} = S_x$ as a summand, so by Theorem 3.6 reflecting twice recovers $k\sigma_x Q$. The adjunction is a priori only an isomorphism of vector spaces, so to upgrade to an algebra isomorphism we need to check that the bijection $\mathrm{Hom}_Q(C_x^-(k\sigma_x Q), C_x^-(k\sigma_x Q)) \to \mathrm{Hom}_{\sigma_x Q}(k\sigma_x Q, k\sigma_x Q)$ preserves composition. But this is clear because the map in the other direction is given by applying the functor C_x^-.

□

Any object V of $\mathrm{Rep}\, Q$ is naturally a $(kQ, \mathrm{End}_Q(V)^{\mathrm{op}})$-bimodule, since the endomorphism action is by definition compatible with multiplication by elements of kQ. Here $\mathrm{End}_Q(V)$ is a k-algebra under composition. Since $\mathrm{End}_Q(T_x)$ can be identified with $k\sigma_x Q^{\mathrm{op}}$, this gives T_x the structure of a $(kQ, k\sigma_x Q)$-bimodule. Hence we can view $\mathrm{Hom}_Q(T_x,-)$ as a functor $kQ-\mathrm{mod} \to k\sigma_x Q-\mathrm{mod}$ and similarly the tensor $T_x \otimes -$ defines a functor $k\sigma_x Q-\mathrm{mod} \to kQ-\mathrm{mod}$. We will prove that these functors are isomorphic to C_x^+, C_x^- respectively.

From Lemma 0.6 we have an adjunction $(T_x \otimes -) \dashv \mathrm{Hom}_Q(T_x,-)$, since T_x is a bimodule. This, together with $C_x^- \dashv C_x^-$, will allow us to show that the bimodule T_x represents C_x^+.

Theorem 3.15. *T_x represents C_x^+. That is,* $\mathrm{Hom}_Q(T_x, -) \simeq C_x^+$*. Also* $(T_x \otimes -) \simeq C_x^-$*.*

Proof. Let W be a kQ-module and V a $k\sigma_x Q$-module. We have a series of adjunctions

$$\begin{aligned}\mathrm{Hom}_Q(T_x \otimes_{\sigma_x Q} V, W) &\cong \mathrm{Hom}_{\sigma_x Q}(V, \mathrm{Hom}_Q(T_x, W)) \\ &= \mathrm{Hom}_{\sigma_x Q}(V, \mathrm{Hom}_Q(C_x^-(k\sigma_x Q), W)) \\ &\cong \mathrm{Hom}_{\sigma_x Q}(V, \mathrm{Hom}_{\sigma_x Q}(k\sigma_x Q, C_x^+(W))) \\ &= \mathrm{Hom}_{\sigma_x Q}(V, C_x^+(W))\end{aligned}$$

since $\mathrm{Hom}_A(A, M)$ is naturally identified with M for any algebra A and A-module M. These isomorphisms come from adjunctions so are natural, and hence their composition is also natural. Thus we have an adjunction $(T_x \otimes_{\sigma_x Q} -) \dashv C_x^+$. This means that C_x^- and $T_x \otimes_{\sigma_x Q} -$ are both left-adjoint to C_x^+, so by uniqueness of adjoints $(T_x \otimes -) \simeq C_x^-$. Also C_x^+ and $\mathrm{Hom}_Q(T_x, -)$ are both right-adjoint to $T_x \otimes_{\sigma_x Q} -$, so $\mathrm{Hom}_Q(T_x, -) \simeq C_x^+$.

□

Corollary 3.16. *C_x^+ has a right-derived functor. That is, we have functors*

$$RC_x^+ := R\,\mathrm{Hom}_Q(T_x, -), R^0C_x^+ := \mathrm{Hom}_Q(T_x, -) = C_x^+, R^1C_x^+ := \mathrm{Ext}_Q(T_x, -)$$

such that for any short exact sequence $0 \to A \to B \to C \to 0$ *in* $\mathrm{Rep}\, Q$*, we get a long exact sequence*

$$0 \to R^0C_x^+A \to R^0C_x^+B \to R^0C_x^+C \to R^1C_x^+A \to R^1C_x^+B \to R^1C_x^+C \to 0.$$

Moreover, RC_x^+ defines a functor $\mathcal{D}^b(\mathrm{Rep}\, Q) \to \mathcal{D}^b(\mathrm{Rep}\, \sigma_x Q)$ *which can be computed by finding a complex of injectives representing any object in* $\mathcal{D}^b(\mathrm{Rep}\, Q)$ *and applying* $\mathrm{Hom}_Q(T_x, -)$ *termwise.*

Proof. This follows from Theorem 3.15 and the discussion in Section 1.2 where we showed that Ext is well-defined. □

Projective modules are an adapted class for the tensor product and hence for C_x^-. We get a left-derived functor

$$LC_x^- \simeq T_x \otimes^L - : \mathcal{D}^b(\mathrm{Rep}\, \sigma_x Q) \to \mathcal{D}^b(\mathrm{Rep}\, Q)$$

where $\otimes^L$ denotes the left-derived tensor. This functor can be computed by taking a complex of projectives that is isomorphic in $\mathcal{D}^b(\mathrm{Rep}\, \sigma_x Q)$ to a given complex, and applying C_x^- or equivalently $T_x \otimes -$ to this complex of projectives. A similar discussion to that regarding Ext in Section 1.2 applies to show that the derived tensor does not depend on the choice of projective resolution. Unlike $\mathrm{Hom}(-,-)$, the tensor product is covariant in both coordinates, so the derived tensor can be computed by taking a projective resolution in either coordinate.

Let us use the properties of T_x to compute its decomposition into indecomposable summands as both a left- and right- module. As a kQ-module, we have

$$T_x = C_x^-(k\sigma_x Q) = C_x^-\left(\bigoplus_{y\in Q_0} P_y^{\sigma_x Q}\right) \cong C_x^-(P_x^{\sigma_x Q}) \oplus \bigoplus_{y\neq x} P_y^Q$$

using additivity of C_x^- and Proposition 3.8. This is a decomposition into indecomposable summands by Theorem 3.6. Note the appearance of the exceptional term $C_x^-(P_x^{\sigma_x Q})$. In some sense, this indecomposable representation characterises the reflection C_x^-. Also observe that all the summands except $C_x^-(P_x^{\sigma_x Q})$ are projective, while this exceptional summand is never projective.

To compute the decomposition of T_x as a right-module, we note that because T_x is a bimodule, the right-action by $k\sigma_x Q$ respects the vector space decomposition

$$T_x = 1 \cdot T_x = \left(\sum_{y\in Q_0} e_y\right) \cdot T_x = \bigoplus_{y\in Q_0} e_y T_x$$

so this is a decomposition as right $k\sigma_x Q$ modules. We can compute the summands using Propositions 2.27 and 3.8, by calculating the reflection $C_x^+(I_y^Q)$ two different ways. For $y \neq x \in Q_0$, we have

$$I_y^{\sigma_x Q} \cong C_x^+(I_y^Q) \cong \mathrm{Hom}_Q(T_x, I_y^Q) \cong (e_y T_x)^*$$

as $k\sigma_x Q$-modules. So $e_y T_x \cong (I_y^{\sigma_x Q})^* = e_y k\sigma_x Q$ as right-modules. Similarly, we get

$$C_x^+(I_x^Q) \cong \mathrm{Hom}_Q(T_x, I_x^Q) \cong (e_y T_x)^*.$$

Then we have a decomposition

$$T_x^* \cong C_x^+(I_x^Q) \oplus \bigoplus_{y\neq x} I_y^{\sigma_x Q} \cong C_x^+\left(\bigoplus_{y\in Q_0} I_y^Q\right) = C_x^+(kQ^*).$$

The summands $(I_y^{\sigma_x Q})^*$ of T_x are projective right-modules while the exceptional indecomposable $C_x^+(I_y^Q)$ cannot be injective, so $C_x^+(I_y^Q)^*$ cannot be projective. In particular:

Corollary 3.17. *$T_x \cong C_x^-(k\sigma_x Q)$ in $kQ-\mathrm{mod}$ and similarly $T_x \cong C_x^+(kQ^*)^*$ in $\mathrm{mod}-k\sigma_x Q$.*

3.4 Derived reflections are equivalences

In this section, we prove that derived reflection functors are equivalences. In particular, this means that two quivers with the same underlying graph which are related by a sequence of reflections have the same derived category. We previously showed that if the underlying graph of a quiver is acyclic, then any two choices of orientation are related by a sequence of reflections, so in this case the derived category depends only on the underlying graph. Finally, we interpret the equivalences RC_x^+, LC_x^- as tilting at a torsion pair. This means that if we take two acyclic quivers with the same underlying graph, and the underlying graph is reflection-transitive, then their abelian categories are related by a sequence of tilts in the common derived category.

Proposition 3.18. *Let Q be acyclic, and $x \in Q_0$ a sink that is not an isolated vertex. For $V \in \operatorname{Rep} Q$ indecomposable, we have $R^1C_x^+V = 0$ unless $V \cong S_x = P_x$, and $R^1C_x^+S_x \cong S_x$.*

Similarly, for $V \in \operatorname{Rep} \sigma_x Q$ indecomposable, we have $L_1C_x^-V = 0$ unless $V \cong S_x = I_x$, and $L_1C_x^-S_x \cong S_x$.

Proof. We prove this for C_x^-, since this is computed on a projective resolution. The other proof is dual. For $V \not\cong S_x$, take the canonical projective resolution

$$0 \to P_1 \to P_0 \to V \to 0.$$

The claim that $L_1C_x^-V = 0$ is equivalent to exactness of the sequence

$$0 \to C_x^-P_1 \to C_x^-P_0 \to C_x^-V \to 0$$

since a priori this is right-exact, and $L_1C_x^-V = 0$ is the kernel on the left. Since we already have right-exactness, the sequence is exact iff the dimensions are balanced. So it is enough to show

$$\dim_k(C_x^-P_1) + \dim_k(C_x^-V) = \dim_k(C_x^-P_0) \in \mathbb{R}^{Q_0}.$$

Now, each of V, P_0, P_1 is a direct sum of indecomposables, and none of the summands is S_x. This is because x is a source in $\sigma_x Q$, so $S_x = I_x$ is injective but not projective (see Proposition 2.28), and hence cannot occur as a summand of the projectives P_0, P_1. Then by Theorem 3.6

$$\dim_k(C_x^-P_i) = \sigma_x(\dim_k P_i) \text{ and } \dim_k(C_x^-V) = \sigma_x(\dim_k V)$$

so (3.4) holds by exactness of the original resolution.

Now we calculate $L_1 C_x^- S_x$. Take the canonical projective resolution and reflect, so

$$0 \to L_1 C_x^- S_x \to C_x^- P_1 \to C_x^- P_0 \to 0$$

is exact. Again, S_x cannot occur as a summand of P_0 or P_1, so

$$\begin{aligned}\dim_k(L_1 C_x^- S_x) &= \dim_k(C_x^- P_1) - \dim_k(C_x^- P_0) = \sigma_x(\dim_k P_1 - \dim_k P_0)\\ &= -\sigma_x(\dim_k S_x) = -\sigma_x(e_x) = e_x\end{aligned}$$

computing $\sigma_x(e_x) = -e_x$ directly. Hence $L_1 C_x^- S_x = S_x$ since S_x is the unique Q-representation of its dimension vector. □

Theorem 3.19. *The derived reflection functors RC_x^+ and LC_x^- are inverse equivalences between $\mathcal{D}^b(\operatorname{Rep} Q)$ and $\mathcal{D}^b(\operatorname{Rep} \sigma_x Q)$.*

Proof. We already know LC^-, RC^+ are additive, and in $\mathcal{D}^b(\operatorname{Rep} Q)$ any object is the direct sum of its cohomology pieces, so it is enough to check this on objects concentrated in a single grade. Moreover, all derived functors commute with translation (essentially by definition) so it is sufficient to check on indecomposable objects of $\operatorname{Rep} Q$ as complexes concentrated in degree 0. We will consider the composition $RC_x^+ LC_x^-$; the other is similar.

First, since $L_0 C_x^- S_x = 0 = R^0 C_x^+ S_x$ we have

$$RC_x^+(LC_x^- S_x)) \cong RC_x^+(S_x[1]) \cong (RC_x^+ S_x)[1] \cong S_x[-1][1] \cong S_x$$

so the statement holds for the indecomposable S_x that is annihilated by $C_x^{\pm}$. Now let $V \ncong S_x$ be indecomposable. By the previous proposition, $LC_x^- V = C_x^- V \ncong S_x$ is an indecomposable object of $\operatorname{Rep} Q$, and hence

$$RC_x^+(LC_x^- V) \cong C_x^+ C_x^- V \cong V$$

by Theorem 3.6. □

We can give a complete description of LC_x^- and RC_x^+ on indecomposable objects of $\mathcal{A}$, which fully determines both functors on the derived category because $\operatorname{Rep} Q$ is hereditary.

Corollary 3.20. *$RC_x^+(V) \cong C_x^+(V)$ for indecomposables $V \ncong S_x \in \operatorname{Rep} Q$, and $RC_x^+(S_x) \cong S_x[-1]$. Similarly, $LC_x^-(W) \cong C_x^-(W)$ for $W \ncong S_x \in \operatorname{Rep} \sigma_x Q$ and $C_x^-(S_x) \cong S_x[1]$.*

Remark 3.21 (Tilting at a torsion pair)**.** Derived reflection functors are an example of a more general phenomenon called tilting at a torsion pair, where we start with a full abelian subcategory of a derived category and essentially apply a translation to some objects, obtaining a new abelian category which is usually not isomorphic to the one we started with.

Take the abelian category $\mathcal{A} \hookrightarrow \mathcal{D}^b(\mathcal{A})$. A torsion pair is a pair $(\mathcal{T}, \mathcal{F})$ of full additive subcategories of $\mathcal{A}$ such that $\mathrm{Hom}_{\mathcal{A}}(\mathcal{T}, \mathcal{F}) = 0$ and any object $E \in \mathcal{A}$ can be resolved $0 \to T \to E \to F \to 0$ with $T \in \mathcal{T}, F \in \mathcal{F}$. That is, $\mathcal{T}$ and $\mathcal{F}$ generate $\mathcal{A}$ under extensions.

From a torsion pair, we obtain a new full subcategory of $\mathcal{D}^b(\mathcal{A})$ given by

$$\mathcal{A}^{\#} := \left\{ E \in \mathcal{D}^b(\mathcal{A}) \;\middle|\; H^0(E) \in \mathcal{T}, H^{-1}(E) \in \mathcal{F}, H^i(E) = 0 \text{ for } i \neq 0, -1 \right\}$$

which is always abelian. We proved that an object is isomorphic to the direct sum of its cohomology in $\mathcal{D}^b(\mathrm{Rep}\, Q)$, so in our case

$$\mathcal{A}^{\#} = \mathcal{T} \oplus \mathcal{F}[1].$$

Take $\mathcal{F}$ to be the subcategory of $\mathrm{Rep}\, Q$ generated additively by S_x, and $\mathcal{T}$ to be the image of C_x^-, generated by all the indecomposables except S_x. Then $\mathrm{Hom}_Q(\mathcal{T}, \mathcal{F}) = 0$ via the adjunction

$$\mathrm{Hom}_Q(C_x^- V, S_x) \cong \mathrm{Hom}_{\sigma_x Q}(V, C_x^+ S_x) = 0.$$

We get a resolution $0 \to C_x^- C_x^+ V \to V \to S_x^{\oplus n} \to 0$ for any $V \in \mathrm{Rep}\, Q$ using $\iota_x V$, whose cokernel is concentrated at x. Then the new abelian category $\mathcal{A}^{\#}$ is precisely the image of $\mathrm{Rep}(\sigma_x Q)$ under the equivalence LC_x^-.

Chapter 4

Stability and filtrations

In this chapter we discuss two different filtrations of objects in a triangulated category $\mathcal{D}$. The first is the Harder–Narasimhan filtration arising from a stability condition. The second is the (iterated) weight filtration, which can be performed on any object in an abelian subcategory of $\mathcal{D}$, and which in particular can be used to give a canonical refinement of the Harder–Narasimhan filtration.

We use the language of triangulated categories, but we refrain from giving the technical definition. The unfamiliar reader should think of the bounded derived category $\mathcal{D}^b(\mathcal{A})$ in place of a generic triangulated category $\mathcal{D}$. See [Nee14] for the general theory of triangulated categories.

4.1 Stability conditions

In this section we define stability conditions on a triangulated category, following the treatment in Section 4 of [Bay11]. The reader may also be interested in the more technical paper of Bridgeland [Bri07] which introduced stability conditions. We present two equivalent definitions of a stability condition on $\mathcal{D}^b(\mathcal{A})$.

Recall that the Grothendieck group of an abelian category $\mathcal{A}$, denoted $K(\mathcal{A})$, is the abelian group generated by isomorphism classes $[A]$ of objects $A \in \mathcal{A}$, subject to relations $[A] - [B] + [C] = 0$ whenever $0 \to A \to B \to C \to 0$ is exact.

Similarly, the Grothendieck group $K(\mathscr{C})$ of a triangulated category $\mathscr{C}$ is the abelian group with generators corresponding to isomorphism classes of objects, and relations $[A] - [B] + [C] = 0$ whenever $A \to B \to C \to A[1]$ is a distinguished triangle. This requires that $[A[1]] = -[A]$ for any object A, because there is a distinguished triangle $A \to 0 \to A[1] \xrightarrow{\mathrm{id}_{A[1]}} A[1]$. The natural map $K(\mathcal{A}) \to K(\mathcal{D}^b(\mathcal{A}))$ is an isomorphism for any abelian category $\mathcal{A}$. This is a consequence of the discussion relating the three definitions of Ext in Section 1.2 and the cohomology filtration, Proposition 1.18.

Let $\mathcal{D}$ be a triangulated category. A stability condition on $\mathcal{D}$ consists of two components, a slicing and a central charge homomorphism.

Definition 4.1 (Slicing)**.** A slicing $\mathcal{P}$ of $\mathcal{D}$ is a collection of full additive subcategories $\mathcal{P}(\phi) \subset \mathcal{D}$ for $\phi \in \mathbb{R}$ such that

(i) The slicing is compatible with translation, that is $\mathcal{P}(\phi + 1) = \mathcal{P}(\phi)[1]$.

(ii) For all $\phi_1 > \phi_2$, $\mathrm{Hom}_{\mathcal{D}}(\mathcal{P}(\phi_1), \mathcal{P}(\phi_2)) = 0$.

(iii) Each object $0 \neq E \in \mathcal{D}$ has sequence $\phi_1 > \phi_2 > \ldots > \phi_n \in \mathbb{R}$ and a sequence of distinguished triangles

$$\begin{array}{ccccccccc} 0 = E_0 & \longrightarrow & E_1 & \longrightarrow & E_2 & \longrightarrow & \ldots E_{n-1} & \longrightarrow & E_n = E \\ \uparrow & \swarrow & \uparrow & \swarrow & & & \uparrow & \swarrow & \\ A_1 & & A_2 & & & & A_n & & \end{array}$$

with $a_i \in \mathcal{P}(\phi_i)$, called the Harder–Narasimhan filtration. The A_i are filtration quotients. Set $\phi^-(E) = \phi_n$ and $\phi^+(E) = \phi_1$.

The objects in $\mathcal{P}(\phi)$ are called semistable of phase ϕ.

Remark 4.2. (a) Given a slicing, the sequence ϕ_i and the objects A_i are unique for any $E \in \mathcal{D}$.

(b) It follows from (ii) that for any two objects X, Y with $\phi^-(X) > \phi^+(Y)$, $\mathrm{Hom}(X, Y) = 0$. This is because a nonzero map $f \in \mathrm{Hom}(X, Y)$ induces a nonzero map from some filtration quotient of X to some filtration quotient of Y, but if $\phi^-(X) > \phi^+(Y)$ then every $\phi_i(X) > \phi_j(Y)$ so no such map can exist (cf Proposition 1.8, which can be used to prove a similar result for the cohomology filtration).

(c) Taking $\mathcal{P}(0) = \mathcal{A}$ and $\mathcal{P}(\phi) = 0$ for $\phi \notin \mathbb{Z}$ gives a (fairly trivial) slicing on $\mathcal{D}^b(\mathcal{A})$, for which the Harder–Narasimhan filtration is the cohomology filtration.

(d) More generally, given a slicing $\mathcal{P}$, take $\mathcal{A}^\# = \mathcal{P}((0, 1])$ to be the full subcategory generated as extensions of semistable objects in $\mathcal{P}(\phi)$ for $0 < \phi \leq 1$. Equivalently, $\mathcal{A}^\#$ consists of objects E satisfying $0 < \phi^-(E) \leq \phi^+(E) \leq 1$.

Then $\mathcal{A}^{\#}$ is an abelian category, and any object of $\mathcal{D}$ has a cohomology filtration with respect to $\mathcal{A}^{\#}$. The Harder–Narasimhan filtration with respect to the slicing is a refinement of this cohomology filtration. In this setting, the pair $\mathcal{T} = \mathcal{P}((\phi, 1]), \mathcal{F} = \mathcal{P}((0, \phi])$ when $\phi \in (0, 1)$ is a torsion pair for $\mathcal{A}^{\#}$, and tilting at $(\mathcal{T}, \mathcal{F})$ gives the abelian category $\mathcal{P}((\phi, \phi + 1])$.

Definition 4.3. A stability condition on a triangulated category $\mathcal{D}$ is a pair $(Z, \mathcal{P})$ where $\mathcal{P}$ is a slicing and the *central charge* $Z : K(\mathcal{D}) \to \mathbb{C}$ is a group homomorphism compatible with $\mathcal{P}$. That is, for every $0 \neq E \in \mathcal{P}(\phi)$,

$$Z(E) = m_E \cdot e^{i\pi\phi}$$

where $m_E \in \mathbb{R}^+$.

For any stability condition $(Z, \mathcal{P})$ on a triangulated category $\mathcal{D}$, the slices $\mathcal{P}(\phi)$ are necessarily abelian. This is Lemma 5.2 in [Bri07], and will allow us to use the weight filtration in the next section to refine the Harder–Narasimhan filtration arising from any stability condition.

If we start with a specific full abelian subcategory $\mathcal{A}^{\#} \subset \mathcal{D}$, and want to determine a stability condition $(Z, \mathcal{P})$ such that $\mathcal{P}((0, 1]) = \mathcal{A}^{\#}$, what do we need to specify? Roughly speaking, we must give an appropriate stability function $Z_{\mathcal{A}^{\#}}$ that will be used to define Z, and this automatically determines the slicing $\mathcal{P}$. We make this precise in Theorem 4.5.

Definition 4.4. A stability function for an abelian category $\mathcal{A}$ is a homomorphism $Z : K(\mathcal{A}) \to \mathbb{C}$ such that $Z(E) \in \mathbb{H}$ for every nonzero $E \in \mathcal{A}$. Here $\mathbb{H}$ is the semi-closed upper half-plane

$$\mathbb{H} = \left\{ z = m \cdot e^{i\pi\phi} \;\middle|\; m > 0, \phi \in (0, 1] \right\}.$$

Given such a function, any nonzero $E \in \mathcal{A}$ has a phase $\phi(E) \in (0, 1]$ such that $Z(E) = m \cdot e^{i\pi\phi(E)}$ for some $m > 0$. A nonzero $E \in \mathcal{A}$ is called Z-semistable if $\phi(A) \leq \phi(E)$ whenever $A \hookrightarrow E$ is a subobject.

We could instead define semistability by requiring that $\phi(E) \leq \phi(B)$ whenever $E \twoheadrightarrow B$ is a quotient of B. This is equivalent because any inclusion or quotient can be extended to a short exact sequence, on which Z is additive by definition of $K(\mathcal{A})$. Then $\phi(Z(A)) \leq \phi(Z(A) + Z(B))$ if and only if $\phi(Z(A) + Z(B)) \leq \phi(B)$, because the sum of two complex numbers of distinct phases has a phase strictly between the two.

Theorem 4.5. *Giving a stability condition on a triangulated category $\mathcal{D}$ is equivalent to specifying*

(1) a full abelian subcategory $\mathcal{A}^\# \subset \mathcal{D}$, and

(2) a stability function $Z : K(\mathcal{A}^\#) \to \mathbb{C}$

such that:

(a) If $n_1 > n_2$ then $\mathrm{Hom}(\mathcal{A}^\#[n_1], \mathcal{A}^\#[n_2]) = 0$.

(b) Every object $E \in \mathcal{D}$ has a cohomology filtration by objects in translations of $\mathcal{A}^\#$.

(c) Every object in $\mathcal{A}^\#$ has a filtration by Z-semistable objects.

The filtration in (c) takes place in the abelian category $\mathcal{A}^\#$, so we require a sequence of inclusions

$$0 = A_0 \subset A_1 \subset \ldots \subset A_n = A$$

where all the quotients A_i/A_{i-1} are Z-semistable. The translations in the filtration (b) must be ordered as in 1.18.

Proof. Starting from a stability condition $(W, \mathcal{P})$ we take $\mathcal{A}^\# = \mathcal{P}((0,1])$, or more generally $\mathcal{A}^\# = \mathcal{P}((\phi, \phi+1])$ for any $\phi \in \mathbb{R}$. Then let $Z : K(\mathcal{A}^\#) \to \mathbb{C}$ be given by $e^{-i\pi\phi} \cdot W$ so that Z maps $\mathcal{A}^\#$ to the upper half-plane. The $\mathcal{A}^\#$-cohomology filtration of $E \in \mathcal{D}$ is given by coarsening the HN filtration, combining the semistable objects with phase in $(\phi+n, \phi+n+1]$ to give a single object in $\mathcal{A}^\#[n]$ for each $n \in \mathbb{Z}$. The filtration (c) comes from the HN filtration of any object in $\mathcal{A}$. An object of $\mathcal{A}$ is Z-semistable iff W-semistable with this construction.

Conversely, we first check that $K(\mathcal{A}^\#) = \mathcal{K}(\mathcal{D})$ so the central charge W is determined by Z. For each $\phi \in (0,1]$, take $\mathcal{P}(\phi)$ to be the Z-semistable objects in $\mathcal{A}^\#$ of phase ϕ. Then we must have $\mathcal{P}(\phi+n) = \mathcal{P}(\phi)[n] \subset \mathcal{A}^\#[n]$ so this extends to all $\phi \in \mathbb{R}$ by translation. The compatibility condition in Definition 4.3 is satisfied by construction, so we need only verify the conditions in Definition 4.1 for $\mathcal{P}$. (i) is immediate. The Hom-vanishing condition (ii) follows from (a) when $\mathcal{P}(\phi_1), \mathcal{P}(\phi_2)$ are contained in distinct integer translations of $\mathcal{A}^\#$. If they lie in the same translation of $\mathcal{A}^\#$, without loss of generality we may assume $\mathcal{P}(\phi_1), \mathcal{P}(\phi_2) \subset \mathcal{A}^\#$. Given any nonzero morphism $f : E \to E' \in \mathcal{A}^\#$, we have a sequence

$$E \twoheadrightarrow \operatorname{im} f \hookrightarrow E'$$

exact in the middle. If E, E' are Z-semistable then by definition of semistability, $\phi(E) \leq \phi(\operatorname{im} f) \leq \phi(E')$. In particular, there are no such nonzero morphisms if $\phi(E) = \phi_1 > \phi_2 = \phi(E')$.

Finally, the Harder–Narasimhan filtration (iii) of any $E \in \mathcal{D}$ is obtained by first taking the cohomology filtration (b) of E, and then refining it using the Harder–Narasimhan filtration (c) in $\mathcal{A}^{\#}$ of each cohomology piece, with Z-semistable filtration quotients. □

The fact that filtrations in $\mathcal{D}$ can be combined, as employed at the end of the previous proof, does not rely on any special properties of the filtration itself. Given a filtration of an object $E \in \mathcal{D}$, together with further filtrations of the filtration quotients A_i, we can always construct a composite filtration of E which refines the original filtration, and whose quotients are the quotients of the filtrations[1] of the A_i.

Theorem 4.5 gives an easier method for constructing stability conditions, since to do so on an abelian category we need only define a stability function, and the semistable objects are determined. Defining a stability condition on a triangulated category $\mathcal{D}$ is more difficult, since we must specify the semistable objects and then check various compatibility properties. When $\mathcal{D} = \mathcal{D}^b(\operatorname{Rep} Q)$ for an acyclic quiver, we have many appropriate abelian categories $\mathcal{A}^{\#}$, given by $\operatorname{Rep} Q'$ for various orientations Q' of the underlying graph of Q. Moreover, reflection functors give embeddings of these categories into $\mathcal{D}^b(\operatorname{Rep} Q)$, which automatically satisfy (a) and (b) since they come from equivalences $\mathcal{D}^b(\operatorname{Rep} Q') \simeq \mathcal{D}^b(\operatorname{Rep} Q)$. This means a variety of stability conditions can be constructed by finding stability functions on these abelian hearts.

However, this simplification does not make construction of stability conditions easy, as it is often difficult to satisfy the condition that Z sends nonzero objects of $\mathcal{A}^{\#}$ to the upper half-plane $\mathbb{H}$. Bridgeland's deformation result (Theorem 7.1 in [Bri07]) offers a solution to this difficulty. Essentially, this result tells us that with a mild finiteness hypothesis on the stability condition, if we deform the central charge Z to Z' there is a unique deformed slicing $\mathcal{P}'$ giving a stability condition $(Z', \mathcal{P}')$. Then $\mathcal{P}'$ determines a new abelian category $\mathcal{A}'$ for us.

Definition 4.6. Given a stability condition and with $\mathcal{A}^{\#}$ as in Theorem 4.5, every object $E \in \mathcal{A}^{\#}$ has a *mass* $m(E)$ coming from its Harder–Narasimhan

[1] The author wonders whether we might have fit more instances of 'filtration' into this paragraph. It seems doubtful.

filtration:

$$m(E) := \sum_{i \in I} |Z(A_i)| = \sum_{i \in I} m(A_i) \in \mathbb{R}_{\geq 0}$$

where $A_i, i \in I$ are the Z-semistable objects in the HN filtration of E. The mass of each Z-semistable object A_i is $|Z(A_i)|$, because such an object has trivial HN filtration.

4.2 The weight filtration

The remaining sections of this chapter are chiefly based on [HKKP20], Sections 2 and 4. This paper introduces a new filtration, the weight filtration 4.11, which can be used to give a canonical refinement of the Harder–Narasimhan filtration coming from a stability condition. We present this as a filtration of an artinian lattice, but we will chiefly apply it to the lattice of subobjects of a given object E in an abelian category, thus giving a filtration of E.

Definition 4.7. A *lattice* L is a partially ordered set closed under gcd and lcm. We use the notation $a \wedge b := gcd(a, b)$ and $a \vee b := lcm(a, b)$. This notation is inspired by the fact that for $A, B \subset E$ subobjects in an abelian category, $gcd(A, B) = A \cap B$ and $lcm(A, B) = A + B$ is the smallest subobject of E containing $A \cup B$.

A lattice is *bound* if it has a global minimum, denoted 0, and global maximum, denoted 1. For $a \leq b \in L$, the interval

$$[a, b] := \{x \in L \mid a \leq x \leq b\}$$

is a bound sublattice of L. Given $x \in L$, there are two possible ways to project x to the interval $[a, b]$, which satisfy an inequality

$$(x \wedge b) \vee a \leq (x \vee a) \wedge b.$$

We call the lattice L *modular* if this inequality is always an equality. A lattice L has *finite length* if there is a global upper bound n on the length of any chain

$$a_0 < a_1 < \ldots < a_n$$

in L. Any (possibly infinite) collection of elements in a finite length lattice has a gcd and lcm. In particular, this means a finite length lattice is necessarily bound. A lattice that is both finite length and modular is called *artinian*.

Example 4.8. A natural example of a modular lattice is the lattice of subobjects of a given object E in an abelian category. Such a lattice is bound, with maximum E and minimum 0. If $A \subseteq B \subseteq E$, then for any $X \subseteq E$ we have

$$(X \wedge B) \vee A = (X \cap B) + A = (X + A) \cap B = (X \vee A) \wedge B$$

so such a lattice is modular.

Remark 4.9. Recall that a simple object in an abelian category is one whose only subobject is 0, and a semisimple object is a direct sum of simple objects.

If $\mathcal{A}$ is a finite length abelian category, then every lattice of subobjects in $\mathcal{A}$ is finite length. Recall that $\mathcal{A}$ is finite length if every object $X \in \mathcal{A}$ has a (finite) Jordan–Hölder filtration

$$0 = X_0 \hookrightarrow X_1 \hookrightarrow X_2 \hookrightarrow \dots \hookrightarrow X_n = X$$

where each quotient X_i/X_{i-1} is simple. If such a filtration exists, its length n only depends on the object X, and moreover the filtration quotients are uniquely determined up to permutation.

In the abelian category $\operatorname{Rep} Q \simeq kQ-$mod of an acyclic quiver, every simple object has total dimension 1, so the length of any module $X \in kQ-$mod is merely its dimension over k. In particular, $\operatorname{Rep} Q$ is finite length if Q is acyclic. Uniqueness of Jordan–Hölder subquotients is also clear from the classification of simples, although it is true more generally in any abelian category. The lattice of subobjects of a representation of an acyclic quiver will be our key example of an artinian lattice.

We have an analogue of the Grothendieck group for an artinian lattice L. $K(L)$ is the abelian group generated by intervals $[a, b]$ for $a \leq b \in L$, with relations

$$[a, b] + [b, c] = [a, c] \tag{4.1}$$

$$[a, a \vee b] = [a \wedge b, b] \tag{4.2}$$

whenever $a \leq b \leq c$. It follows from modularity of L that the map

$$[a, a \vee b] \to [a \wedge b, b], x \mapsto x \wedge b$$

with inverse $x \mapsto x \vee a$ is an isomorphism, so the relation (4.2) identifies generators corresponding to isomorphic sublattices. We let $K^+(L)$ be the sub-semigroup of $K(L)$ generated by intervals $[a, b]$ for $a < b$.

For $A, B \subset E$ in an abelian category, we can think of the interval $[A, B]$ as the quotient $B/A \in \mathcal{A}$. In this case $K(L)$ is the subgroup of $K(\mathcal{A})$ generated by subquotients of E, and the relation (4.1) corresponds to the extension

$$0 \to B/A \to C/A \to C/B \to 0 \in \mathcal{A},$$

while the relation (4.2) corresponds to the isomorphism $(A+B)/A \cong B/(A \cap B)$.

Definition 4.4 as well as the notions of mass, phase and semistability generalise immediately from the setting of an abelian category to an artinian lattice. A *polarization* on an artinian lattice L is a homomorphism $Z : K(L) \to \mathbb{C}$ such that $Z(K^+(L))$ is contained in some half-plane. That is, we fix a half-open interval $I \subset \mathbb{R}$ of length 1, and require

$$Z(K^+(L)) \subset \{re^{i\pi\phi} \mid r \in \mathbb{R}^+, \phi \in I\}.$$

This is a translation of the notion of a stability function on an abelian category, up to a rotation of the half-plane. Each interval $[a, b]$ with $a < b$ has a phase $\phi([a, b]) \in I$, and a mass defined analogously to Definition 4.6. We call a polarized lattice semistable if $\phi([0, x]) \leq \phi(L)$ for any $x \in L$. This agrees with the notion for an object of an abelian category, see Definition 4.4.

The Harder–Narasimhan filtration of a polarized lattice is unique, see for example Theorem 4.2 in [HKKP20]. This also proves uniqueness of the HN filtration on an abelian category, cf Theorem 4.5 (c). In fact, the Harder–Narasimhan filtration of a polarized lattice L is the unique mass-minimising filtration, meaning that if $0 = a_0 < a_1 < \ldots < a_n = 1$ is any chain in L, then

$$m(L) \leq \sum_{k=1}^{n} m([a_{k-1}, a_k])$$

with equality if and only if the chain is the HN filtration of L. This is Theorem 4.3 in [HKKP20]. The weight filtration can also be viewed as a mass-minimising filtration, in a slightly different way.

Definition 4.10. An artinian lattice L is *complemented* if any $a \in L$ has a complement $b \in l$ with $a \wedge b = 0$ and $a \vee b = 1$.

The lattice of subobjects of $E \in \mathcal{A}$ is complemented if and only if E is semisimple. If E is semisimple, then any subobject A is the direct sum of some simple summands of E, and its complement is the sum of the remaining summands. Conversely, if the lattice is complemented then each subobject of E occurs as a direct summand.

The reader may wish to consider the first example in Section 4.3 alongside the following definition.

Theorem/Definition 4.11 (see Sections 1.2 and 1.3 of [HKKP20])**.** *Let $L \neq \emptyset$ be an artinian lattice, and $X : K^+(L) \to \mathbb{R}^+$ an additive map. Then there exists a unique filtration*

$$0 = a_0 < a_1 < \ldots < a_n = 1$$

with intervals $[a_{k-1}, a_k] \neq 0$ labelled by real numbers $\lambda_1 < \ldots < \lambda_n \in \mathbb{R}$, satisfying the following:

(i) $[a_{k-1}, a_l]$ is complemented for $1 \leq k \leq l \leq n$ whenever $\lambda_l - \lambda_k < 1$. In particular, we always require $[a_{k-1}, a_k]$ to be complemented.

(ii) The balancing condition

$$\sum_{k=1}^{n} \lambda_k X([a_{k-1}, a_k]) = 0$$

holds.

(iii) For any collection of objects $b_k \in [a_{k-1}, a_k]$ for $1 \leq k \leq n$, either the balancing condition

$$\sum_{k=1}^{n} \lambda_k X([a_{k-1}, b_k]) \leq 0$$

holds, or there exists some $1 \leq k < l \leq n$ with $\lambda_l - \lambda_k \leq 1$ but $[b_k, b_l]$ not complemented.

The uniquely defined filtration is called the weight filtration of L, and depends on X. Uniqueness of the filtration includes the labels λ_k.

Remark 4.12. (a) If L is the lattice of subobjects of $E \in \mathcal{A}$, then we obtain a filtration of E, where (i) requires that appropriate subquotients of the filtration be semisimple.

(b) For any choice of X, the weight filtration of a lattice L is trivial iff L is complemented. Assuming uniqueness, this can be verified by checking that the trivial filtration (with $\lambda_1 = 0$) satisfies (1)-(3). In particular, an object of an abelian category $\mathcal{A}$ has trivial weight filtration iff it is semisimple.

(c) Given a stability condition $(Z, \mathcal{P})$ on a triangulated category $\mathcal{D}$, for any $\phi \in \mathbb{R}$ the image of $\mathcal{P}(\phi)$ under Z is contained in a single ray in $\mathbb{C}$. Z thus induces a map $X : K(\mathcal{P}(\phi)) \to \mathbb{R}$ which is positive on classes of nonzero objects. This allows the weight filtration of any object in $\mathcal{P}(\phi)$ to be computed. Doing so for each $\phi \in \mathbb{R}$ gives a canonical refinement of the HN filtration for the stability condition.

Condition (iii) implicitly defines a new artinian lattice L' whose elements $b \in L'$ are given by collections $b_k \in [a_{k-1}, a_k], 1 \leq k \leq n$ such that $[b_k, b_l]$ is complemented whenever $1 \leq k < l \leq n$ with $\lambda_l - \lambda_k \leq 1$, and the balancing condition

$$\sum_{k=1}^{n} \lambda_k X([a_{k-1}, b_k]) = 0 \tag{4.3}$$

is satisfied. This is a sublattice of the product $\prod_{1 \leq k \leq n} [a_{k-1}, a_k]$, and L' inherits a homomorphism X' defined by

$$X'([b, c]) = \sum_{k=1}^{n} X([b_k, c_k]). \tag{4.4}$$

Then we can take the weight filtration of (L', X'). Projecting the weight filtration of L' onto the factors $[a_{k-1}, a_k]$ of the product, we obtain a (possibly trivial) filtration of each interval $[a_{k-1}, a_k]$ in the weight filtration of L. This refined filtration of L has labels in $\mathbb{R}^2$, ordered lexicographically. We may repeat this process inductively until we reach a lattice which is complemented and hence has trivial weight filtration, at which point no further refinement is possible. The filtration of L indexed over some $\mathbb{R}^n$ obtained in this way is called the *iterated weight filtration*, where the number of iterations n is the *depth*. This process is finite because L' can be shown to have strictly smaller length than L.

The proof of Theorem 4.11 is given in Sections 4.3 and 4.4 of [HKKP20]. We will not give the details here, but we outline some key ideas. To prove existence of the weight filtration, we rephrase this as a minimisation problem. Let $\mathcal{B}(L)$ denote the space of labelled filtrations of L which satisfy (i), topologised using the coefficients $\lambda_k \in \mathbb{R}$. For each candidate filtration $a \in \mathcal{B}(L)$, we construct a new artinian lattice $\Lambda(a)$. This is a sublattice of the product

$$\prod_{k=1}^{\text{length}(a)} [a_{k-1}, a_k]$$

and is defined such that when a is the weight filtration of L, condition (ii) is equivalent to $\phi(\Lambda(a)) = 0$ and condition (iii) is equivalent to $\Lambda(a)$ being semistable.

The weight filtration of L minimises the function $a \mapsto m(\Lambda(a))$ over $\mathcal{B}(L)$. It is easy to show that for any $a \in \mathcal{B}(L)$,

$$m(\Lambda(a)) \geq X(L)$$

with equality if and only if $\Lambda(a)$ is semistable of phase 0, that is a is a weight filtration. One then proves that any local minimum of this function must be a weight filtration, and that a local minimum must exist using topological properties of $\mathcal{B}(L)$.

For any $a \in \mathcal{B}(L)$, the associated $\Lambda(a)$ has a canonical polarization defined by

$$Z([x, y]) = \sum_{k=1}^{\text{length}(a)} (1 + \lambda_k i) X([x_{\lambda_k}, y_{\lambda_k}]).$$

With a the weight filtration of L, the lattice L' used to construct the iterated weight filtration is the sublattice of $\Lambda(a)$ consisting of $x \in \Lambda(a)$ with $\operatorname{Im} Z([0, x]) = 0$, where $0 \in \Lambda(a)$ is defined by $0_{\lambda_k} = a_{k-1}$, and similarly $1_{\lambda_k} = a_k$ is the maximum of $\Lambda(a)$. This is the source of the balancing condition (4.3) for L'. The homomorphism X' defined in (4.4) is then the restriction of Z to L'. Note that Z is a polarization in the right half-plane, not the upper half-plane. In particular, there may exist $x \in \Lambda(a)$ with $\operatorname{Im} Z([0, x]) < 0$, and such x do not appear in the sublattice L'.

4.3 Three examples

In this section, we present three examples which demonstrate how the weight filtration may be computed and its features. The first will familiarise the reader with Definition 4.11 in practice, the second outlines a class of weight filtrations which turn out to be trivial, while the third is a small example which exhibits an iterated weight filtration.

Let $\mathcal{A} = \operatorname{Rep} Q$, for Q acyclic. To construct an appropriate function X, we take a weight vector $\mathbf{X} \in \mathbb{R}_+^{Q_0}$ and define

$$X(V) := \mathbf{X} \cdot \dim V = \sum_{i \in Q_0} X_i \dim_k V(i)$$

to be the dimension of V weighted by $\mathbf{X}$, for any $V \in \operatorname{Rep} Q$. This gives a weight function $X : K^+(L(W)) \to \mathbb{R}^+$ on the lattice $L(W)$ of subobjects of any $W \in \operatorname{Rep} Q$, by identifying an interval $[A, B] \in L(W)$ with the quotient object $B/A \in \operatorname{Rep} Q$. In the following examples, references in Roman numerals are to Definition 4.11.

Example 4.13. Consider $Q = A_3$ with the orientation:

$$1 \to 2 \leftarrow 3$$

and take a weight vector $\mathbf{X} = (X_1, X_2, X_3) \in \mathbb{R}^3_+$. Consider the object W with dimension vector $(1,1,1)$ and both maps the identity. Subobjects of W can be identified by their dimension vectors. We use the notation 2 for the subobject of dimension vector $(0,1,0)$, and so on. The lattice $L(W)$ is

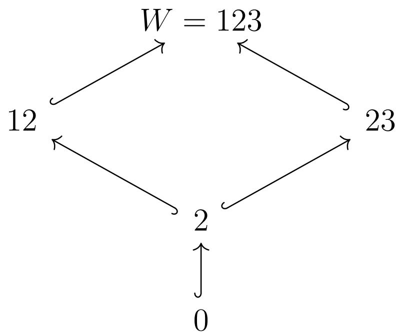

Each subrepresentation is indecomposable in this case.

We start by considering a maximal length filtration, say

$$0 < 2 < 23 < W$$

with labels $\lambda_1, \lambda_2, \lambda_3$. Note that all intervals of length 1 are complemented, and $[2, W]$ is complemented, but $[0, 23]$ is not complemented. So $\lambda_2 \geq \lambda_1 + 1$. Letting $\lambda_2 = \lambda_1 + b$ and $\lambda_3 = \lambda_2 + c$, the balancing condition gives

$$0 = \lambda_1 X_2 + \lambda_2 X_3 + \lambda_3 X_1 = \lambda_1(X_1 + X_2 + X_3) + b(X_1 + X_3) + cX_1. \tag{4.5}$$

Now consider the chain $0 < 2 < W$ in (iii). The balancing condition

$$0 \geq \lambda_1 \cdot 0 + \lambda_2 \cdot 0 + \lambda_3 \cdot X_1 = (\lambda_1 + b + c)X_1$$

cannot hold for this chain, since multiplying both sides by $X_1 + X_2 + X_3 \in \mathbb{R}^+$ we have

$$0 \geq ((b+c)(X_1 + X_2 + X_3) - b(X_1 + X_3) - cX_1)X_1 = (bX_2 + cX_2 + cX_3)X_1$$

using (4.5), which is impossible since all terms on the RHS are positive. The only interval in b that is not complemented is $[0, W]$, so we must have $\lambda_3 - \lambda_1 \leq 1$. But then

$$1 \geq \lambda_3 - \lambda_2 \geq \lambda_2 - \lambda_1 \geq 1$$

so these are all equalities, and $\lambda_3 = \lambda_2$. Hence we should combine the intervals corresponding to λ_3 and λ_2 in the filtration, and consider the chain

$$0 < 2 < W.$$

The balancing condition for this chain is still given by (4.5), and simplifies to

$$\lambda_1(X_1 + X_2 + X_3) + b(X_1 + X_3) = 0$$

since we now have $c = 0$. Consider the trivial chain $0 < W$ in (iii). The balancing condition is

$$0 \geq \lambda_1 \cdot 0 + \lambda_2(X_1 + X_3) = (\lambda_1 + b)(X_1 + X_3)$$

and again multiplying by $X_1 + X_2 + X_3 \in \mathbb{R}^+$, we see that this condition cannot hold since the RHS simplifies to

$$(X_1+X_2+X_3)(\lambda_1+b)(X_1+X_3) = (b(X_1+X_2+X_3)-b(X_1+X_3))X_1 = bX_2X_1 > 0.$$

The interval $[0, W]$ is not complemented, and for (iii) to hold we must have $b = \lambda_2 - \lambda_1 = 1$. Thus

$$\lambda_1 = -\frac{X_1 + X_3}{X_1 + X_2 + X_3} \in (-1, 0)$$

and $\lambda_2 = \lambda_1 + 1 \in (0, 1)$. We claim that this gives the weight filtration. To check, we must verify (iii).

There are 8 possible pairs $(a_1, a_2) \in [0, 2] \times [2, W]$. Of these, 3 are not complemented, and thus need not be balanced. Moreover, the chain $(0, 2)$ has trivial balancing condition. The remaining 4 are $(2, 2), (2, 23), (2, 12)$, and $(2, W)$.

The balancing condition for $(2, W)$ evaluates to 0, because it is the same expression as in (ii). The condition for $(2, 2)$ evaluates to $\lambda_1 X_2 < 0$.

For $(2, 23)$ the balancing condition evaluates to $\frac{-X_1X_2}{X_1+X_2+X_3} < 0$, and for $(2, 12)$ to $\frac{-X_2X_3}{X_1+X_2+X_3} < 0$. This verifies that we have found the weight filtration.

The lattice $L(W)'$ obtained from $L(W)$ has two elements, since of the 8 elements of the product $[0, 2] \times [2, W]$, 3 are not complemented and the balancing conditions of 3 others are strictly negative. Then $L(W)' = \{(0, 2) < (2, W)\}$ is complemented and there are no iterations.

A key feature illustrated in the above example is that if we start with a maximal length filtration and assign not necessarily distinct labels to each interval, we can solve for the labels and the weight filtration at the same time by combining intervals when the corresponding labels must be equal. Another feature of note is that the sequence $0 < (0, 1, 0) < W$ which caused the maximal-length filtration to fail to be the weight filtration was, itself, the weight filtration.

Example 4.14. Let $Q = A_n^{eq}$ be the equioriented A_n quiver, with vertex 1 the source and n the sink. Consider the indecomposable representation $E_{1,n}$ with a 1-dimensional space at each vertex. The lattice of subobjects is a single line of length n.

$$L(E_{1,n}) : 0 \hookrightarrow E_n \hookrightarrow E_{n-1,n} \hookrightarrow E_{n-2,n} \hookrightarrow \cdots \hookrightarrow E_{2,n} \hookrightarrow E_{1,n}.$$

The only intervals in this lattice which are complemented are intervals of length 1, so the weight filtration must include every subobject. We claim that the labels are $\lambda_1 = \lambda, \lambda_{i+1} = \lambda_1 + i$ for $1 \leq i \leq n-1$, where the constant $\lambda < 0$ is determined by (ii).

To see that (iii) holds for these labels, note that since any pair of consecutive labels differ by 1, the balancing inequality need only hold for n-tuples b where $[b_{k-1}, b_k]$ is complemented for every k. Since the only complemented intervals in $L(E_{1,n})$ are length 1, all such b either have the form

$$b^i : E_n \hookrightarrow E_{n-1,n} \hookrightarrow \ldots \hookrightarrow E_{n-i,n} = E_{n-i,n} \hookrightarrow E_{n-i-1,n} \hookrightarrow \ldots \hookrightarrow E_{2,n}$$

for some $0 \leq i \leq n-2$, or are one of b^-, b^+ defined by $b_0^- = 0, b_{j+1}^- = E_{n-j+1,n}$ and $b_j^+ = E_{n-j+1,n}$ respectively. The balancing inequality for b^- is trivial, and the expression for b^+ is precisely (ii) so also evaluates to zero, given by

$$0 = \sum_{j=0}^{n-1} X_{n-j}\lambda_{j+1} = \sum_{j=0}^{n-1} X_{n-j}(\lambda_1 + j). \tag{4.6}$$

We claim that the inequality in (iii) holds strictly for each b^i, so these labels give the weight filtration. Moreover, this means $L(E_{1,n})' = \{0 = b^- < b^+ = 1\}$ is complemented, so the iterated weight filtration is the same as the weight filtration. To justify this claim, note that the balancing condition for b^i is the partial sum

$$S_i := \sum_{j=0}^{i} X_{n-j}\lambda_{j+1}$$

of (4.6). We have $S_0 = \lambda_1 X_n < 0$, and the λ_j are strictly increasing. Let J be the minimal index j such that λ_{j+1} is nonnegative. Then all the partial sums up to S_{J-1} are strictly negative since they are a sum of negative terms. In particular, such an index $J \leq n-1$ must exist because $S_{n-1} = 0$. For $J \leq j < n-1$, we have $S_{J-1} \leq S_j < S_{n-1} = 0$ since S_j is obtained from S_{J-1} by adding nonnegative terms. The strict inequality on the right comes from the fact that $\lambda_n > 0$. This shows every partial sum except $S_{n-1} = 0$ is negative, as claimed.

So, in the equioriented quiver A_n, the weight filtration for $E_{1,n}$ is maximal. The same argument applies to any indecomposable representation of A_n^{eq}, since we can view any such as the representation $E_{1,n'}$ for some $n' < n$ by discarding vertices at which the representation is 0-dimensional. To find a weight filtration that has iterations or that depends on the choice of weights, we will need to consider a less uniform orientation.

Even on A_n, we can obtain interesting weight filtrations by choosing a different orientation, so that the lattice of subobjects is nondegenerate. In the previous two examples the weight filtration did not depend on the weight vector $\mathbf{X}$, however this is not a general feature, as we will see in the next example. See also the example detailed in Section 2.1 of [HKKP20].

Example 4.15. Consider $Q = A_5$ with the orientation:

$$1 \to 2 \leftarrow 3 \to 4 \to 5$$

and take the representation Y with dimension vector $(1, 1, 1, 1, 1)$ and the identity map at each arrow. We will denote each subrepresentation by listing the vertices at which there is a 1-dimensional space, so for example 245 denotes the representation with dimension vector $(0, 1, 0, 1, 1)$. Then $L(Y)$ is:

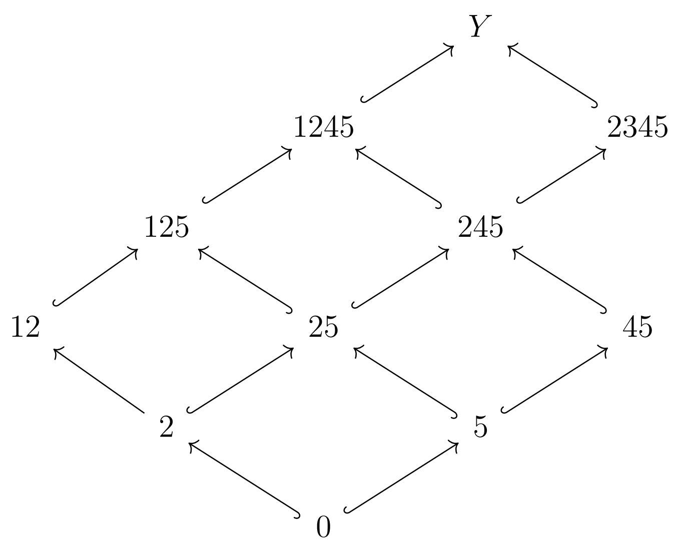

We omit the computations and summarise the results. The weight filtration depends on the signs of the two expressions $f_{15} = X_1X_5-(X_1X_3+2X_2X_3+X_2X_4)$ and $f_{23} = X_2X_3 - (X_1X_4 + 2X_1X_5 + X_2X_5)$, with an iterated weight filtration occurring only along the wall $X_2X_3 = X_1X_4 + 2X_1X_5 + X_2X_5$.

- $f_{23} > 0$: The weight filtration is $0 < 5 < 245 < Y$ with labels
$$\lambda = -\frac{2X_1 + X_2 + 2X_3 + X_4}{X_1 + X_2 + X_3 + X_4 + X_5} < \lambda + 1 < \lambda + 2.$$
The lattice $L' = \{0 < 1\}$ is trivial (and hence complemented) so there is no refinement.

- $f_{23} = 0$: The weight filtration is again $0 < 5 < 245 < Y$, with the same labels as in the previous case. However, L' now has one additional element $(0, 25, 1245)$, corresponding to the balancing condition $\lambda_2 X_2 + \lambda_3 X_1 = 0$ which is equivalent to $f_{23} = 0$. In particular, this means L' is not complemented, and it has maximal weight filtration $0 < (0, 25, 1245) < 1$ with labels
$$\mu = -\frac{X_3 + X_4 + X_5}{X_1 + X_2 + X_3 + X_4 + X_5} \in (-1, 0)$$
and $\mu + 1$. L'' is trivial, giving no further refinement, so the iterated weight filtration of $L(Y)$ is $0 < 5 < 25 < 245 < 1245 < Y$ with labels
$$(\lambda, \mu + 1) < (\lambda + 1, \mu) < (\lambda + 1, \mu + 1) < (\lambda + 2, \mu) < (\lambda + 2, \mu + 1) \in \mathbb{R}^2_+.$$

- $f_{23} < 0$ *and* $f_{15} < 0$: The weight filtration is $0 < 5 < 25 < 245 < 1245 < Y$ with labels
$$\lambda = -\frac{2X_3 + X_4}{X_3 + X_4 + X_5} < \mu = -\frac{X_1}{X_1 + X_2} < \lambda + 1 < \mu + 1 < \lambda + 2.$$
The conditions $f_{23} < 0$ and $f_{15} < 0$ are equivalent to $\mu < \lambda + 1$ and $\lambda < \mu$ respectively. The lattice L' is complemented with 4 terms, where the two intermediate terms correspond to the balancing conditions $\mu X_2 + (\mu + 1)X_1 = 0$ and $\lambda X_5 + (\lambda + 1)X_4 + (\lambda + 2)X_3 = 0$. Thus the iterated weight filtration is the same as the weight filtration.

- $f_{15} = 0$: The weight filtration is $0 < 25 < 1245 < Y$ with labels
$$\lambda = -\frac{X_1}{X_1 + X_2} < \lambda + 1 < \lambda + 2.$$
The lattice L' is complemented with 4 entries, so the iterated weight filtration is trivial along this wall.

- $f_{15} > 0$: The weight filtration is $0 < 2 < 25 < 125 < 1245 < Y$ with labels
$$\mu = -\frac{X_1}{X_1 + X_2} < \lambda = -\frac{2X_3 + X_4}{X_3 + X_4 + X_5} < \mu + 1 < \lambda + 1 < \lambda + 2.$$

These are the same as the labels we had when $f_{23}, f_{15} < 0$, but reordered since now $f_{15} > 0$ so $\mu < \lambda$. L' is again complemented with 4 entries, so the iterated weight filtration is the same as the weight filtration.

Note that the above cases are disjoint, because f_{23} and f_{15} cannot be simultaneously nonnegative.

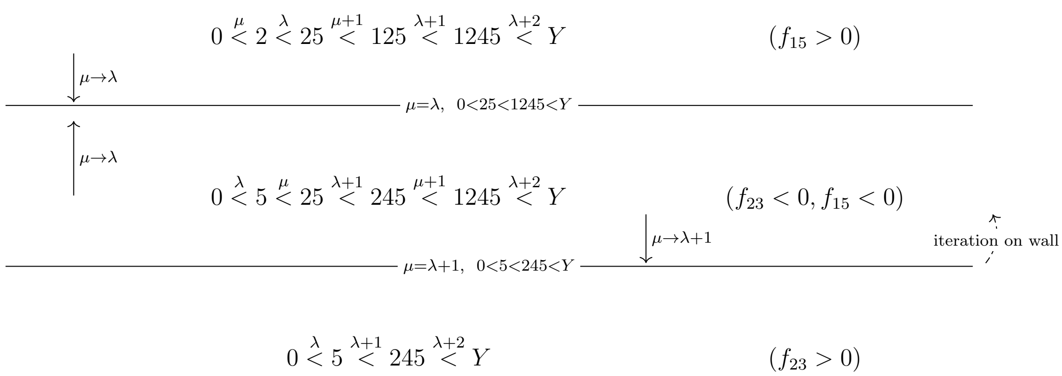

Figure 4.1: A schematic of the parameter space for the weight filtration of Y, showing the two walls and 3 chambers.

This example has several important features. First, we have two disjoint walls, $X_1X_5 = X_1X_3 + 2X_2X_3 + X_2X_4$ and $X_2X_3 = X_1X_4 + 2X_1X_5 + X_2X_5$, which together divide the parameter space $\mathbb{R}^5_+$ into 3 chambers. The weight filtration is constant across each open chamber with trivial iterated weight filtration, while the two walls exhibit distinct behaviours.

Along the wall $f_{23} = 0$, the weight filtration agrees with that of one of the two adjacent chambers, while the other adjacent chamber has a strictly finer weight filtration. In this case the wall has an iteration, which passes from the coarser adjacent filtration to the finer one. If we approach the wall from the chamber $f_{23}, f_{15} < 0$ containing the finer filtration, we find that $\mu \to \lambda+1$ as $f_{23} \to 0$. Then on the wall $f_{23} = 0$, we obtain the weight filtration by combining the intervals corresponding to μ and $\lambda + 1$, which are now equal, and similarly combining the intervals corresponding to $\mu + 1$ and $\lambda + 2$. The filtration along the wall is thus the limit of the filtrations in the two adjacent chambers.

Along the wall $f_{15} = 0$, the weight filtration is coarser than either of the weight filtrations in adjacent chambers, being their intersection. Approaching the wall from either direction, we find that $\mu \to \lambda$. We may obtain the weight filtration on the wall from the weight filtration of either open chamber, by combining any intervals whose corresponding labels now agree.

In particular, the weight filtration is continuous on the parameter space (although we have not made this precise). Walls are analogous to stationary points, where we think of $f_{23} = 0$ as a saddle and $f_{15} = 0$ as a local maximum. In this example, an iteration occurs only at the saddle.

Remark 4.16. Given a simple subobject $S \hookrightarrow V$ in an abelian category $\mathcal{A}$, we may obtain the weight filtration for V/S from the weight filtration for V by discarding S everywhere it appears. (It is not hard to check that the conditions in 4.11 hold for the induced filtration on the quotient.) In Example 4.15, for instance, we obtain the weight filtration for Y/S_5 by setting $X_5 = 0$, and deleting all intervals in the various filtrations and lattices which now have weight 0. Doing so recovers the example in Section 2.1 of [HKKP20]. Effectively, we are restricting to the hyperplane $X_5 = 0$ along the boundary of the parameter space $\mathbb{R}^5_+$. The wall $f_{15} = 0$ and the chamber $f_{15} > 0$ intersect this hyperplane only at the origin, so we lose these two possible weight filtrations upon passing to the quotient.

4.4 Further directions

Our third example raises some interesting general questions, deserving of further study. Here are some of them:

1. Does the wall-and-chamber structure we observed always arise? Can iterated weight filtrations of depth greater than one occur in open chambers, rather than just along walls? Given an iterated weight filtration along a wall, are the successive refinements determined by the weight filtrations in adjacent open chambers?

2. Starting with a stability condition, what kinds of refinements to the HN filtration are given by the iterated weight filtration? Given a trivial stability condition (where $\mathcal{P}(\phi) = 0$ for $\phi \notin \mathbb{Z}$) does this refinement correspond to the HN filtration for a nearby nontrivial stability condition?

3. How does the weight filtration interact with reflections? Can we bound the weight filtration of a reflected representation using the initial filtration?

Section 3 of [HKKP20] gives a simplified criterion for the weight filtration of a Q-representation with 1-dimensional spaces only, which might profitably be used to address some of these questions. Indeed, the authors of that paper use this criterion to construct iterated weight filtrations of arbitrary depth. The rest, we must postpone for later exploration.

Appendix A

Lie algebras, Dynkin diagrams and root systems

This appendix is provided to give context for the classification of quivers of finite representation type. In particular, it is aimed at readers who have not encountered the ADE Dynkin diagrams elsewhere. These diagrams can seem arbitrary when one is first introduced to them, but in reality they appear in classification results in several not-obviously-related areas, often giving a first indication of deeper connections. We will attempt to give some indication of the relationship between root systems, Dynkin diagrams and semisimple Lie algebras here, although our treatment is by no means comprehensive. The reader is directed to [Hum12] as the standard reference for Lie algebras. In the following let V be a real inner product space, with inner product denoted $(-,-)$.

Each nonzero $\alpha \in V$ has a corresponding linear operator called a reflection,

$$s_\alpha : V \to V, \quad s_\alpha(v) := v - 2\frac{(v,\alpha)}{(\alpha,\alpha)}\alpha.$$

Such a reflection has a $(\dim V - 1)$-dimensional eigenspace with eigenvalue 1, namely $\{\alpha\}^\perp$, and $s_\alpha(\alpha) = -\alpha$. In particular s_α is orthogonal, order 2, and diagonalisable with an orthonormal eigenbasis.

Definition A.1. A finite subset $\Phi \subset V\backslash\{0\}$ is a *root system* for V if the following are satisfied:

(R1) Φ spans V.

(R2) If $\alpha \in \Phi$ then $\mathrm{span}_{\mathbb{R}}\{\alpha\} \cap \Phi = \{\pm\alpha\}$.

(R3) For each $\alpha \in \Phi$, $s_\alpha(\Phi) = \Phi$.

(R4) The value $2\frac{(\beta,\alpha)}{(\alpha,\alpha)}$ is an integer for any two roots $\alpha, \beta \in \Phi$.

A root system has an associated Weyl group $W(\Phi)$, defined to be the subgroup of $\mathrm{GL}(V)$ generated by the reflections s_α for $\alpha \in \Phi$.

A *simple system* $\Delta \subset \Phi$ for Φ is a basis of V such that each root $\alpha \in \Phi$ can be written as either a purely nonnegative or purely nonpositive linear combination of elements of Δ. The *positive system* Π corresponding to a simple system consists of elements of Φ which are a nonnegative combination of the simple roots. Then we have

$$\Phi = \Pi \sqcup -\Pi.$$

The requirement (R4) means that when written with respect to any basis of simple roots, the matrix for a reflection s_α with $\alpha \in \Phi$ has integer entries. Moreover, this property restricts the possible angles between roots to a short finite list, namely: $\pi, \pi/2, \pi/3, 2\pi/3, \pi/4, 3\pi/4, \pi/6, 5\pi/6$. One can show that for any two roots $\alpha, \beta \in \Phi$, the order of $s_\alpha s_\beta$ in $W(\Phi)$ is the denominator of the angle between them, written in lowest terms. Denote these orders $m(\alpha\beta) \in \{0, 1, 2, 3, 4, 6\}$. If $\alpha \neq \beta \in \Delta$ then $m(\alpha\beta) \neq 0, 1$.

Theorem A.2. *Let Φ be a root system. The Weyl group W acts transitively on the simple systems for Φ, and for any simple system $\Delta \in \Phi$, the reflections corresponding to simple roots generate $W(\Phi)$. Moreover, $W(\Phi)$ has a presentation*

$$W(\Phi) = \langle s_\alpha, \alpha \in \Delta \mid s_\alpha^2 = 1 = (s_\alpha s_\beta)^{m(\alpha\beta)} \rangle$$

Definition A.3. Given a root system Φ with choice of simple system $\Delta \subset \Phi$, the *Dynkin diagram* for Φ has vertices corresponding to simple roots, with an edge of weight $m(\alpha, \beta) - 2$ between the vertices corresponding to $\alpha, \beta \in \Delta$. If $m(\alpha, \beta) > 3$ this edge is decorated with an arrow pointing towards the incident vertex corresponding to the root of greater magnitude.

The diagrams that arise in this way are called finite Dynkin diagrams.

The Dynkin diagram for a root system determines the presentation of the Weyl group in A.2 which in turn determines the angle between any pair of roots. The magnitudes of the roots are also determined, up to a global scalar, since the value of $m(\alpha\beta)$ determines the ratio of the magnitudes of α and β, and arrows on the Dynkin diagram determine which is larger. As a consequence, the Dynkin diagram of a root system determines the root system up to isomorphism.

If a Dynkin diagram is disconnected, the corresponding root system decomposes into summands corresponding to the connected components. So, indecomposable root systems are classified by connected (finite) Dynkin diagrams, and vice versa.

Theorem A.4. *The following is a complete list of connected finite Dynkin diagrams, that is Dynkin diagrams corresponding to irreducible root systems as in Definition A.1. In each case, n denotes the number of vertices, and edges are marked with $m(\alpha\beta)$, with 3s implicit.*

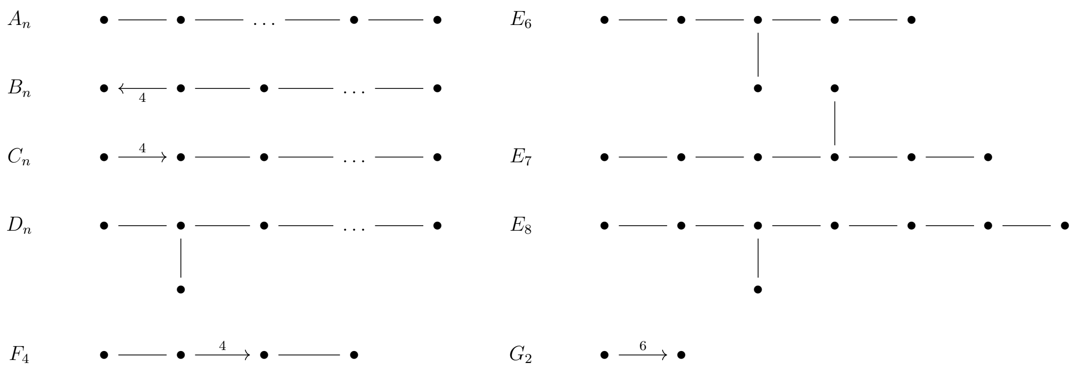

The Dynkin diagrams in Theorem A.4 (or irreducible root systems) also classify the finite-dimensional simple Lie algebras over an algebraically closed field of characteristic zero, by taking the root system corresponding to the action of the Lie algebra on a Cartan subalgebra. In particular, simple Lie algebras having the same root system are isomorphic. Finite-dimensional semisimple Lie algebras are then classified by Dynkin diagrams whose connected components are as above.

The finite Dynkin diagrams also appear in the classification of quivers of finite representation type. Such quivers cannot have multiple edges, so the Dynkin diagrams of type B_n, C_n, F_4, and G_2 do not appear. Gabriel's Theorem 2.12 states that the quivers of finite representation type are exactly those whose underlying graphs are ADE Dynkin diagrams, and that in such case the indecomposable representations are in bijection with the positive roots in the corresponding root system. There is a marked similarity between the standard proof of the classification of indecomposable root systems, Theorem A.4, and Berstein, Gelfand and Ponomarev's proof of Gabriel's Theorem, see [BGP73]. Both results are proved by considering a certain bilinear form, which is positive definite on quivers of finite representation type (and Dynkin diagrams corresponding to indecomposable root systems). One shows that if the form is positive semidefinite on some quiver, then it must be positive definite on any subquiver. Then we find a sufficient list of quivers on which the bilinear form is not positive definite — and which therefore cannot occur as subquivers of any finite type quiver — to rule out all except the ADE quivers. Reflections are then used to show that the dimension vectors of indecomposable representations for ADE quivers do not depend on the orientation, and to classify such.

Let us now explicitly construct the root systems for types A and D. We omit the construction for the three exceptional graphs E_6, E_7, E_8 as it is somewhat involved (the Weyl group for E_8 has order $2^{14}3^5 5^2 7$). The following constructions are taken from Section 12.1 of [Hum12]. We work in $\mathbb{R}^k$ with the standard inner product, and let ε_i for $1 \leq i \leq k$ denote the standard (orthonormal) basis.

Consider the diagram A_n, with vertices numbered from 1 to n in order. The Weyl group is S_{n+1}, where the reflection corresponding to the vertex i is identified with the transposition $(i\ i+1)$. This can be seen from the presentation

$$W(A_n) = \langle s_i, 1 \leq i \leq n \mid s_i s_{i+1} s_i = s_{i+1} s_i s_{i+1}, s_i^2 = 1, s_i s_j = s_j s_i \text{ when } |i-j| > 1 \rangle$$

for the Weyl group.

To construct the root system, we take V to be the n-dimensional subspace of $\mathbb{R}^{n+1}$ orthogonal to the vector $\sum_i \varepsilon_i$, with its inherited inner product. Let Φ be the collection of vectors in V which are $\mathbb{Z}$-linear combinations of the standard basis vectors, and such that $(\alpha, \alpha) = 2$ for all $\alpha \in \Phi$. Explicitly,

$$\Phi = \{\varepsilon_i - \varepsilon_j \mid i \neq j \in [n+1]\}.$$

Note that the reflection $s_{\varepsilon_i - \varepsilon_j}$ interchanges the ε_i and ε_j coordinates of any vector in $\mathbb{R}^{n+1}$ while leaving the other coordinates fixed, acting as the transposition $(ij) \in S_{n+1}$. This can be used to show that the root system Φ has A_n as its Dynkin diagram.

The set $\Delta = \{\alpha_i := \varepsilon_i - \varepsilon_{i+1} \mid 1 \leq i \leq n\}$ is linearly independent and therefore a basis for V. For $i < j$ we have $\varepsilon_i - \varepsilon_j = \sum_{l=i}^{j-1} \alpha_l$, so Δ is a simple system with corresponding positive system

$$\Pi = \{\varepsilon_i - \varepsilon_j \mid i < j\}.$$

Expressing the positive roots in terms of the simple roots, we get coordinate vectors of the form

$$\alpha_i + \alpha_{i+1} + \ldots + \alpha_j$$

for any interval $1 \leq i \leq j \leq n$. These are the dimension vectors of indecomposable representations of A_n. We denote the indecomposable representation of A_n with a 1-dimensional space at the vertices in the interval $[i, j]$ by $E_{i,j}$.

Now consider the diagram D_n, for $n \geq 4$ (if $n \leq 3$ then $A_n = D_n$). Let $V = \mathbb{R}^n$ and again Φ consists of vectors in the $\mathbb{Z}$-span of the standard basis vectors satisfying $(\alpha, \alpha) = 2$, this time without the orthogonality condition. Explicitly,

$$\Phi = \{\pm(\varepsilon_i \pm \varepsilon_j) \mid i \neq j \in [n]\}.$$

The Weyl group for D_n is the group of permutations and sign changes involving only an even number of sign changes, on $\{\varepsilon_1, \ldots, \varepsilon_n\}$. As for A_n, the reflections corresponding to $\varepsilon_i - \varepsilon_j$ are transpositions of the basis elements, so generate the permutations. The reflection $s_{\varepsilon_i+\varepsilon_j}$ permutes the ε_i and ε_j coordinates and changes both signs. So, compositions of the form $s_{\varepsilon_i+\varepsilon_j} s_{\varepsilon_i-\varepsilon_j}$ achieve any pair of sign changes. Hence Φ is a D_n root system.

As a basis, take the set $\Delta = \{\alpha_i := \varepsilon_i - \varepsilon_{i+1} \mid 1 \leq i \leq n-1\} \sqcup \{\beta := \varepsilon_{n-1} + \varepsilon_n\}$. This is a simple system, generating the positive system

$$\Pi = \{\varepsilon_i \pm \varepsilon_j \mid 1 \leq i < j \leq n\}.$$

The following express the positive roots in terms of the simple roots.

$$\varepsilon_i - \varepsilon_j = \sum_{\ell=i}^{j-1} \alpha_\ell \qquad \text{for } 1 \leq i < j \leq n \qquad \text{(A.1)}$$

$$\varepsilon_i + \varepsilon_n = \sum_{\ell=i}^{n-2} \alpha_\ell + \beta \qquad \text{for } 1 \leq i \leq n-1 \qquad \text{(A.2)}$$

$$\varepsilon_i + \varepsilon_j = \sum_{\ell=i}^{j-1} \alpha_\ell + 2\sum_{\ell=j}^{n-2} \alpha_\ell + \alpha_{n-1} + \beta \qquad \text{for } 1 \leq i < j \leq n-1 \qquad \text{(A.3)}$$

The simple roots are associated to the vertices of D_n in the following way. The roots α_i for $1 \leq i \leq n-1$ generate a subgroup of the Weyl group isomorphic to S_n, and correspond to an A_{n-1} subgraph of D_n. The remaining simple root *beta* is orthogonal to all the α_i except α_{n-2}. Hence α_{n-2} corresponds to the degree 3 vertex in D_n, and β, α_{n-1} are the degree 1 vertices adjacent to α_{n-2}.

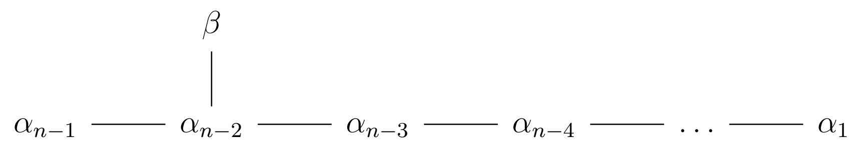

Then there are three families of dimension vectors giving indecomposable representations of D_n. The first two (A.1) and (A.2) are the case where at least one of β and α_{n-1} has dimension 0. This case reduces to the classification for A_{n-1}, so we simply have an interval of adjacent vertices, all with dimension 1.

For the third family (A.3), both β and a_{n-1} have dimension 1. Then vertices in a (possibly empty) interval of the form $[\alpha_{n-2}, \alpha_{n-1-i}]$ have dimension 2, and vertices in a nonempty interval of the form $[\alpha_{n-2-i}, \alpha_{n-2-j}]$ have dimension 1.

Those indecomposable representations in the third family which have at least one dimension 2 vertex are obtained from the unique such representation of D_4 by adding in some extra vertices of dimension 2 and maps between them that are all isomorphisms. This means indecomposable representations of D_n with at least one dimension 2 vertex all have the following form.

$$
\begin{array}{ccccccccccccccc}
 & & k & & & & & & & & & & & & \\
 & & \downarrow{\scriptstyle A} & & & & & & & & & & & & \\
k & \xrightarrow{B} & k^2 & \xrightarrow{\sim} & k^2 & \xrightarrow{\sim} & \dots & \xleftarrow{\sim} & k^2 & \xleftarrow{C} & k & \xrightarrow{\sim} & k & \xrightarrow{\sim} & \dots \longrightarrow 0
\end{array}
$$

$$
\begin{array}{ccccc}
k & \xrightarrow{C} & k^2 & \xleftarrow{A} & k \\
 & & \uparrow{\scriptstyle B} & & \\
 & & k & &
\end{array}
\qquad \text{where } A + B = C.
$$

Bibliography

[Bay11] Arend Bayer. A tour to stability conditions on derived categories. *notes*, 2011.

[Bei78] Alexander A Beilinson. Coherent sheaves on $\mathbb{P}^n$ and problems of linear algebra. *Functional Analysis and Its Applications*, 12(3):214–216, 1978.

[BGP73] I N Bernstein, I M Gel'fand, and V A Ponomarev. Coxeter functors and Gabriel's theorem. *Russian mathematical surveys*, 28(2):17–32, 1973.

[Bon90] Alexei I Bondal. Representation of associative algebras and coherent sheaves. *Mathematics of the USSR-Izvestiya*, 34(1):23–42, 1990.

[Bri07] Tom Bridgeland. Stability conditions on triangulated categories. *Annals of Mathematics*, 166(2):317–345, 2007.

[DW17] Harm Derksen and Jerzy Weyman. *An introduction to quiver representations*, volume 184. American Mathematical Society, 2017.

[EEGRO04] Edgar Enochs, Sergio Estrada, Juan R García-Rozas, and Luis Oyonarte. Flat and cotorsion quasi-coherent sheaves. applications. *Algebras and representation theory*, 7(4):441–456, 2004.

[Gab72] Peter Gabriel. Unzerlegbare darstellungen I. *Manuscripta mathematica*, 6(1):71–103, 1972.

[HKKP20] Fabian Haiden, Ludmil Katzarkov, Maxim Kontsevich, and Pranav Pandit. Semistability, modular lattices, and iterated logarithms. *arXiv preprint arXiv:1706.01073v2*, 2020.

[Hum12] James E Humphreys. *Introduction to Lie algebras and representation theory*, volume 9. Springer Science & Business Media, 2012.

[Huy06] Daniel Huybrechts. *Fourier-Mukai transforms in algebraic geometry.* Oxford University Press, 2006.

[Kra07] Henning Krause. Derived categories, resolutions, and Brown representability. *Contemporary Mathematics*, 436:101, 2007.

[Kra08] Henning Krause. Representations of quivers via reflection functors. *arXiv preprint arXiv:0804.1428*, 2008.

[KS13] Masaki Kashiwara and Pierre Schapira. *Sheaves on Manifolds*, volume 292. Springer Science & Business Media, 2013.

[Mac88] Saunders MacLane. Categories for the working mathematician. *Graduate texts in mathematics*, 5, 1988.

[Mur06] Daniel Murfet. Derived categories part I. `http://therisingsea.org/notes/DerivedCategories.pdf`, 2006.

[Nee14] Amnon Neeman. *Triangulated Categories*, volume 148 of *Annals of Mathematics Studies*. Princeton University Press, 2014.

[Wei95] Charles Weibel. *An introduction to homological algebra.* Number 38 in Cambridge studies in advanced mathematics. Cambridge university press, 1995.

[Wey] Jerzy Weyman. Representations of quivers and their applications. Lecture notes, `http://mathserver.neu.edu/~todorov/Weyman.pdf`.