A deterministic algorithm for Harder–Narasimhan filtrations for representations of acyclic quivers

Chi-Yu Cheng

msp

A deterministic algorithm for Harder–Narasimhan filtrations for representations of acyclic quivers

Chi-Yu Cheng

Let M be a representation of an acyclic quiver Q over an infinite field k. We establish a deterministic algorithm for computing the Harder–Narasimhan filtration of M. The algorithm is polynomial in the dimensions of M, the weights that induce the Harder–Narasimhan filtration of M, and the number of paths in Q. As a direct application, we also show that when k is algebraically closed and when M is unstable, the same algorithm produces Kempf's maximally destabilizing one parameter subgroups for M.

1. Introduction

Our goal in this paper is to provide a constructive approach to Harder–Narasimhan filtrations of representations of acyclic quivers. Specifically, we first establish in Theorem A a link between the Harder–Narasimhan filtration and the *discrepancy* of a representation. The later notion, introduced in [Chindris and Kline 2021], allows us to apply algebraic complexity tools to the study of Harder–Narasimhan filtrations of representations of acyclic quivers. The main result, Theorem 2 of [Chindris and Kline 2021] is a deterministic, polynomial time algorithm for finding witnesses to the discrepancies of representations of bipartite quivers. Huszar [2021, Proposition 3.5] then extended the result to acyclic quivers in Theorem 3.5 in our paper. Our first main result Theorem A affirms that for any representation, its Harder–Narasimhan filtration has a term that witnesses its discrepancy. So the next natural question to ask is that is there also a systematic way to compute Harder–Narasimhan filtrations? Combining Theorem 3.5 and Theorem A, we are able to establish a deterministic algorithm, Algorithm 1, in Theorem B that computes the Harder–Narasimhan filtration of a representation of an acyclic quiver. We would like to point out that Algorithm 1

is not polynomial in the number of edges of the quiver, but in the number of paths (see Theorem 3.7). When the quiver is bipartite where arrows only go from one partition to the other, the number of arrows equals that of the paths. In this case Algorithm 1 indeed has polynomial time complexity.

This paper was originally motivated by the relations between "maximally destabilizing subobjects" under different stability conditions in the case of representations of quivers. Stability plays an important role in algebraic geometry for constructing moduli spaces of algebro-geometric objects. In the case of representations of quivers, two commonly used stability conditions are weight stability and slope stability. We now briefly introduce the two stability conditions, and we shall see that failing either one of them would require some subobject that contradicts the condition. We let Q be an acyclic quiver.[†] That is, Q does not have an oriented cycle, and we let Q_0 denote the set of vertices of Q.

Weight stability was formulated in [King 1994] to construct moduli of representations of finite dimensional algebras. King considered representations of Q of a fixed dimension $\boldsymbol{d} \in \mathbb{N}^{Q_0}$, and a $\mathbb{Z}$-linear weight function $\theta : \mathbb{Z}^{Q_0} \to \mathbb{Z}$ such that $\theta(\boldsymbol{d}) = 0$. For a representation M of Q, we let $\dim M \in \mathbb{N}^{Q_0}$ be its dimension vector. By $\theta(M)$ we mean $\theta(\dim M)$. King defined that for a representation M of dimension $\boldsymbol{d}$, M is θ-*semistable* if and only if $\theta(M') \leq \theta(M) = 0$ for every subrepresentation M' of M. Otherwise, M is θ-*unstable*. Namely, there is a subrepresentation M' of M such that $\theta(M') > 0$. In [Chindris and Kline 2021], the *discrepancy* of any representation M *with respect to* θ is defined to be the number

$$\operatorname{disc}(M, \theta) = \max_{M' \subseteq M} \theta(M'),$$

where $M' \subseteq M$ means M' is a subrepresentation of M. We see that if M is θ-unstable, then the subrepresentation that witnesses $\operatorname{disc}(M, \theta)$ contradicts θ-semistability of M the most.

On the other hand, weight stability is in fact a reinterpretation of the Hilbert–Mumford criterion Theorem 4.3. When M is θ-unstable, the original Hilbert–Mumford criterion states that there is a one parameter subgroup (of the group to be defined in Section 4B) that contradicts θ-semistability for M. Kempf [1978] showed that among all destabilizing one parameter subgroups, there is a unique indivisible one, up to conjugation by a parabolic subgroup, that maximally contradicts θ-semistability in the numerical sense that will be made precise in Section 4. The maximally destabilizing one parameter subgroups of an unstable representation induce a unique filtration, which was referred to as Kempf filtration in [Zamora 2014]. The other type of filtration, known as the Harder–Narasimhan filtration, comes from slope instability.

Slope stability depends on two weights $\Theta, \kappa : \mathbb{Z}^{Q_0} \to \mathbb{Z}$, where $\kappa((\mathbb{Z}^+)^{Q_0}) > 0$. The slope μ of a nonzero representation M is defined as

$$\mu(M) = \frac{\Theta(M)}{\kappa(M)}.$$

[†]The assumption that the quiver is acyclic is not required to define either stability condition. Nevertheless, we will stick to acyclic quivers throughout the paper for consistency, as the algorithm deals with acyclic quivers only.

Unlike weight stability, slope stability is defined for any nonzero representation, whether its slope is zero or not. In [Hille and de la Peña 2002], a nonzero representation M is said to be μ-semistable if $\mu(M') \leq \mu(M)$ for all $0 \neq M' \subseteq M$. In Section 2, we recall in Lemma 2.7 a way to go between weight and slope stability. In any case, we shall stick to the following convention throughout the paper, whose convenience will become apparent later: If a statement only involves weight stability, we shall use the lower case θ as the weight. When slope stability is involved, we use the upper case Θ as the numerator of the slope.

The Harder–Narasimhan filtration of a representation M with respect to μ is constructed inductively in [Hille and de la Peña 2002]. As the first step, set M_1 as the unique subrepresentation that is maximal among all subrepresentations having the highest slope. Such a subrepresentation is called the *strongly contradicting semistability* (abbreviated as scss) *subrepresentation of M*. One then goes on to set M_{i+1}/M_i as the scss subrepresentation of M/M_i until M/M_i is μ-semistable. The resulting filtration $0 \subsetneq M_1 \subsetneq \cdots \subsetneq M_r = M$ is called the *Harder–Narasimhan filtration* of M.

We now have the following maximally destabilizing subobjects associated to an unstable representation M:

(1) The discrepancy of M, and its witnessing subrepresentations (not necessarily unique).

(2) The scss subrepresentation of M, and the Harder–Narasimhan filtration of M.

(3) The one parameter subgroups that maximally contradict the semistability of M, and the Kempf filtration induced by those one parameter subgroups.

The relation between (2) and (3) is already unraveled by the paper [Zamora 2014]. The main result is that for an unstable representation, its Harder–Narasimhan filtration coincides with its Kempf filtration. The first of our two main results Theorem A establishes a link between (1) and (2). The second main result Theorem B establishes a deterministic algorithm Algorithm 1 for computing the Harder–Narasimhan filtration and therefore the maximally destabilizing one parameter subgroups of any unstable representation.

1A. *The main results.* We let Q be an acyclic quiver.

Theorem A (Lemma 3.1, Theorem 3.3, Corollary 3.4). *Let Θ be a weight and let M be a representation of Q with $\Theta(M) = 0$. For any weight κ with $\kappa((\mathbb{Z}^+)^{Q_0}) > 0$, if $\mu = \Theta/\kappa$ is the slope function, then any subrepresentation M' with $\Theta(M') = \text{disc}(M, \Theta)$ contains the scss subrepresentation of M (with respect to μ). Moreover, if*

$$0 \subsetneq M_1 \subsetneq \cdots \subsetneq M_{r-1} \subsetneq M_r = M$$

is the Harder–Narasimhan filtration of M (with respect to μ), then there is an M_l in the filtration such that $\Theta(M_l) = \text{disc}(M, \Theta)$.

It is noteworthy that the discrepancy depends only on Θ, but the Harder–Narasimhan filtration depends on μ, which has an extra piece κ. While the Harder–Narasimhan filtration of M may change due to different choices of κ, the results of Theorem A are independent of κ for a fixed Θ.

Theorem B (Theorems 3.7 and 4.6). *Let M be a representation of Q over an infinite field. There exists a deterministic algorithm to compute the Harder–Narasimhan filtration of M. When the ground field is algebraically closed and when M is unstable, a maximally destabilizing one parameter subgroup for M can be constructed based on its Harder–Narasimhan filtration. Hence the algorithm also produces a maximally destabilizing one parameter subgroup for M. Moreover, in the case that Q is bipartite where all arrows go from one partite to the other, the algorithm has polynomial time complexity.*

Theorem B is inviting for several purposes. First, finding Kempf's maximally destabilizing one parameter subgroups in general is very hard. Roughly speaking, suppose V is a representation of a reductive group G. If G is a torus, then finding Kempf's one parameter subgroup for an unstable point $v \in V$ is done by linear programming, whose constraints come from the states of v. For a general reductive group, one can first restrict the action to a maximal torus, then conduct linear programming on the entire orbit $G \cdot v$. The set of states of all points in the orbit can be as large as the power set of $\{1, 2, \ldots, \dim V\}$, minus the empty subset. Hence the time complexity to compute Kempf's one parameter subgroups using brute force is exponential. Theorem B provides an alternative to brute force for finding Kempf's one parameter subgroups in the context of representations of acyclic quivers. In addition, if the quiver is bipartite as stated in Theorem B, the algorithm has polynomial time complexity.

Second, Hoskins [2014] shows that the stratification of the space of representations of a quiver of a fixed dimension by Harder–Narasimhan types coincides with the stratification by Kempf's one parameter subgroups. Hence Theorem B sheds new light on computing the stratification by either type.

1B. *Outline of the paper.* We briefly recall slope stability Definition 2.1, weight stability Definition 2.5, and a way to go between the two Lemma 2.7 in Section 2. In Section 3, we establish Theorem A (Lemma 3.1, Theorem 3.3, Corollary 3.4) and the first part of Theorem B (Theorem 3.7). Specifically the first part of Theorem B constructs Algorithm 1, and calculates its complexity. The second part of Theorem B constructs maximally destabilizing one parameter subgroups for unstable representations from their Harder–Narasimhan filtrations. Due to the technicality involved in invariant theory, we postpone the second part of Theorem B to Section 4.

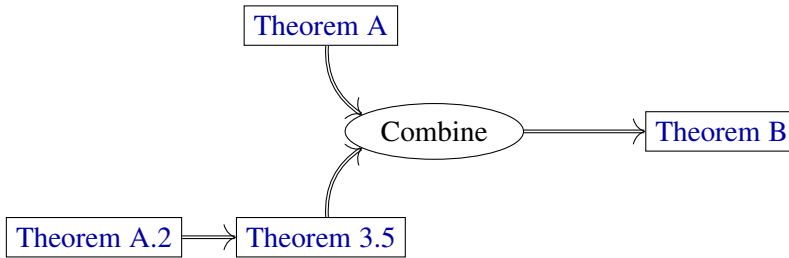In Section 4A, we recall the Hilbert–Mumford criterion Theorem 4.3, and define maximally destabilizing one parameter subgroups Definition 4.4. We also present Kempf's main result, our Theorem 4.5, from the paper [Kempf 1978]. Theorem 4.5 describes the existence and the uniqueness of maximally destabilizing one parameter subgroups. We then apply the machinery to the case of representation of quivers at Section 4B. Finally, we are able to finish the second part of Theorem B in Theorem 4.6.

In Section 5, we construct an unstable representation to showcase that not all subrepresentations that witness the discrepancy occur in the Harder–Narasimhan filtration. We also provide two approaches to compute its maximally destabilizing one parameter subgroups. The first one uses Theorem B and the second one is the brute force procedure that can be applied to any representation of any reductive group. Although the second approach works in more generality, the cost will be apparent in the example. The

point is to give the readers an idea of the hardness of finding maximally destabilizing one parameter subgroups and the value of Theorem B.

We note that our Algorithm 1 relies on Theorem 3.5, which is an algorithm that computes discrepancies and was established in [Huszar 2021]. Theorem 3.5 in turn is a consequence of [Ivanyos et al. 2018, Theorem 1.5], our Theorem A.2. We supply an almost self-contained account of how Theorem 3.5 is established based on Theorem A.2 in the Appendix. Although this was the work of [2021], we intend to fill in some details not supplied in the original paper. Here is a visualization of the logical dependency of the results in our paper:

$$
\begin{array}{ccc}
\boxed{\text{Theorem A}} & & \\
& \searrow & \\
& \big(\text{Combine}\big) \longrightarrow \boxed{\text{Theorem B}} \\
& \nearrow & \\
\boxed{\text{Theorem A.2}} \Longrightarrow \boxed{\text{Theorem 3.5}} &
\end{array}
$$

.

Finally, we will work with an arbitrary ground field throughout Sections 2 and 3 until Theorem 3.5. Originally, the machinery of Theorem 3.5 only requires the ground field to have *enough* elements. For simplicity we will work with an infinite field, starting from Theorem 3.5. In Section 4, we will assume further that the ground field is algebraically closed for the formulation of Kempf's maximally destabilizing one parameter subgroups.

2. Stability conditions for representations of quivers

In this section we recall slope stability (Definition 2.1) and weight stability (Definition 2.5) for representations of quivers.

Weight stability defined in this section is originally the Hilbert–Mumford criterion (Theorem 4.3) formulated in [King 1994] in the context of representations of quivers. For the purpose of the exposition of this paper, we postpone a more detailed account of weight stability to Section 4. At the end of this section we recall in Lemma 2.7 a way to translate between slope stability and weight stability.

2A. *Set up.* Let k be a field, not necessarily algebraically closed. A quiver $Q = (Q_0, Q_1, t, h)$ consists of Q_0 (vertices) and Q_1 (arrows) together with two maps $t, h : Q_1 \to Q_0$. The image of an arrow under t (resp. h) is the tail (resp. head) of the arrow. We represent Q as a directed graph with vertices Q_0 and directed arrows $a : ta \to ha$ for $a \in Q_1$. A *path* in Q is a composition of arrows in Q. We give each vertex v the trivial arrow $e_v : v \to v$. For any path $p : x \to v$ (resp. $q : v \to y$), we set $e_v p = p$ (resp. $q e_v = q$). We let Q be a finite acyclic quiver throughout the paper. That is, both Q_0 and Q_1 are finite, and Q does not have any oriented cycle other than the self-loops $\{e_v\}_{v \in Q_0}$. In this way, the number of paths in Q is finite.

A representation M of Q consists of a collection of finite dimensional k-vector spaces M_v, for each $v \in Q_0$ together with a collection of k-linear maps $M(a) : M_{ta} \to M_{ha}$, for each $a \in Q_1$. For each $v \in Q_0$, we set $M(e_v) : M_v \to M_v$ to be the identity map.

If M, N are two representations of Q, then a morphism from M to N is a collection of k-linear maps $\varphi_v : M_v \to N_v$ for each $v \in Q_0$, such that $N(a)\varphi_{ta} = \varphi_{ha}M(a)$ for each $a \in Q_1$. In particular, M' is a subrepresentation of M whenever each M'_v is a subspace of M_v and each $M'(a)$ is the restriction of $M(a)$ to M'_{ta}. We will write $M' \subseteq M$ if M' is a subrepresentation of M. We also note that the collection of representations of Q forms an abelian category.

The dimension vector of a representation M of Q is the tuple $\dim M \in \mathbb{Z}^{Q_0}$ where $(\dim M)_v = \dim M_v$ as a k-vector space for each $v \in Q_0$. A *weight* is a $\mathbb{Z}$-linear function $\mathbb{Z}^{Q_0} \to \mathbb{Z}$. If θ is a weight, we write $\theta(\dim M)$ as $\theta(M)$.

For clarity, below is a summary of some important notations that we stick to throughout the paper:

Notation. Let M be a representation of Q. For each vertex $v \in Q_0$, we write M_v as the vector space of M at v. For each $a \in Q_1$, $M(a)$ denotes the map $M_{ta} \to M_{ha}$ that is part of the data of M. For any weight $\theta : \mathbb{Z}^{Q_0} \to \mathbb{Z}$, $\theta(v)$ instead of θ_v is the weight of θ at v. We also decree that $\theta(M)$ means $\sum_{v \in Q_0} \theta(v) \dim M_v$. For a dimension vector $\boldsymbol{d} \in \mathbb{Z}^{Q_0}$, we write $\boldsymbol{d}_v$ instead of $\boldsymbol{d}(v)$ for its component at v.

2B. *Slope stability.* In this section we recall semistability of representations of quivers with respect to a slope, along with some important properties stated in Lemmas 2.2 and 2.3. We then recall the construction from [Hille and de la Peña 2002] of the Harder–Narasimhan filtration of a representation in Theorem 2.4. These results will be widely used at Section 3.

Fix two weights $\Theta, \kappa : \mathbb{Z}^{Q_0} \to \mathbb{Z}$ where we require $\kappa((\mathbb{Z}^+)^{Q_0}) > 0$. We set $\mu(M) = \Theta(M)/\kappa(M)$ for any nonzero representation M. We implicitly assume that a representation is nonzero whenever the slope function is applied to it.

Definition 2.1. We say a representation M of Q is *μ-semistable* if $\mu(N) \leq \mu(M)$ for every subrepresentation N of M.

Lemma 2.2 [Hille and de la Peña 2002, Lemma 2.1]. *Let $0 \to L \to M \to N \to 0$ be a short exact sequence of representations of Q. Then the following conditions are equivalent*:

(1) $\mu(L) \leq \mu(M)$.

(2) $\mu(L) \leq \mu(N)$.

(3) $\mu(M) \leq \mu(N)$.

A fundamental result is the following:

Lemma 2.3 [Hille and de la Peña 2002, Lemma 2.2]. *Let M be a representation of Q. There is a unique subrepresentation N such that*

(1) $\mu(N)$ *is maximal among subrepresentations of M, and*

(2) *if $\mu(N') = \mu(N)$, then $N' \subseteq N$.*

We call such a representation N the *strongly contradicting semistability* (abbreviated as scss) *subrepresentation of M*.

Theorem 2.4 [Hille and de la Peña 2002, Theorem 2.5]. *Let M be a representation of Q. There is a unique filtration $0 = M_0 \subsetneq M_1 \subsetneq \cdots \subsetneq M_{r-1} \subsetneq M_r = M$ such that*

(1) $\mu(M_i/M_{i-1}) > \mu(M_{i+1}/M_i)$ *for* $i = 1, \ldots, r-1$,

(2) *the representation M_i/M_{i-1} is μ-semistable for* $i = 1, \ldots, r$.

The above filtration is called the *Harder–Narasimhan filtration of M*. It is constructed by setting M_1 to be the scss subrepresentation of M, then by inductively setting M_{i+1}/M_i as the scss subrepresentation of M/M_i until M/M_i is μ-semistable.

2C. *Weight stability.* Another stability condition on representations of quivers comes from GIT and was introduced in [King 1994]. Like the slope stability we introduced earlier, GIT stability also depends on a choice of parameters. In the case of representations of quivers the parameter is a single weight. We therefore refer to this stability condition as weight stability. We will relate weight and slope stability in Lemma 2.7 at the end of this section.

Definition 2.5. Let $d \in \mathbb{Z}^{Q_0}$ be a dimension vector and let θ be a weight with $\theta(d) = 0$. A representation M with dimension vector d is θ-*semistable* if $\theta(N) \leq 0$ for all subrepresentations $N \subseteq M$.

The notion of *discrepancy* was introduced in [Chindris and Kline 2021] to describe the subrepresentations that contradict the weight stability the most. This very notion allows us to compute the Harder–Narasimhan filtrations of representations of an acyclic quiver using tools from algebraic complexity, as we shall see in Section 3.

Definition 2.6. Let M be a representation and let θ be a weight. We do not require $\theta(M) = 0$. The *discrepancy of M with respect to θ* is defined as

$$\mathrm{disc}(M, \theta) = \max_{N \subseteq M} \theta(N).$$

We say a subrepresentation N is θ-*optimal in M*, or *N witnesses* $\mathrm{disc}(M, \theta)$ if $\theta(N) = \mathrm{disc}(M, \theta)$.

We recall the following way to go between weight stability and slope stability.

Lemma 2.7. *Let Θ and κ be two weights with $\kappa(N) > 0$ for every nonzero representation N. Let μ be the corresponding slope Θ/κ. For every dimension vector $d \in \mathbb{Z}^{Q_0}$, define the new weight*

$$\theta_d(N) := \kappa(d)\Theta(N) - \Theta(d)\kappa(N) \quad \text{for each } N.$$

Then a representation M of dimension d is μ-semistable if and only if M is θ_d-semistable. Moreover, the new slope $\mu_d = \theta_d/\kappa$ defines the same slope stability condition so that the Harder–Narasimhan filtrations of a representation with respect to μ and μ_d coincide.

Proof. Let M be a representation of dimension $\boldsymbol{d}$ and let $N \subset M$ be any subrepresentation. We then have

$$\mu(N) = \frac{\Theta(N)}{\kappa(N)} \le \mu(M) = \frac{\Theta(\boldsymbol{d})}{\kappa(\boldsymbol{d})} \Leftrightarrow \kappa(\boldsymbol{d})\Theta(N) - \Theta(\boldsymbol{d})\kappa(N) = \theta_{\boldsymbol{d}}(N) \le 0$$

The first statement follows.

For the second statement, simply note that

$$\mu_{\boldsymbol{d}}(N) = \kappa(\boldsymbol{d}) \cdot \mu(N) - \Theta(\boldsymbol{d}).$$

Hence $\mu_{\boldsymbol{d}}$ is a positive scalar multiple of μ, followed by a translation. $\qquad\square$

3. The main results

Fix an acyclic quiver Q. In this section we prove two major results: Theorems 3.3 and 3.7. Theorem 3.3 provides a sufficient condition for the Harder–Narasimhan filtration of a representation of Q to contain a term that witnesses the discrepancy. We then propose Algorithm 1 to compute the Harder–Narasimhan filtrations of representations of Q. The correctness of Algorithm 1 is ensured by Proposition 3.6. We also recall the algebraic complexity machinery Theorem 3.5 derived from [Ivanyos et al. 2018]. With the aid of Theorem 3.5, we establish in Theorem 3.7 the time complexity of Algorithm 1. The derivation of Theorem 3.5 is supplied in the Appendix for interested readers.

We fix two weights Θ, κ on $\mathbb{Z}^{Q_0}$ where $\kappa(N) > 0$ for each nonzero representation N. We let $\mu = \Theta/\kappa$ be the slope. We do not assume anything about the ground field until Theorem 3.5, where we will begin to assume the ground field is infinite.

We now present Lemma 3.1, which is the cornerstone of this paper.

Lemma 3.1. *Let M be a μ-unstable representation, and let M_1 be the scss subrepresentation of M. Suppose $\Theta(M_1) > 0$ and M' is Θ-optimal in M. Then M' contains M_1. Moreover, M' is μ-semistable if and only if $M' = M_1$.*

Proof. Note that $\Theta(M_1) > 0$ implies $\mathrm{disc}(M, \Theta) = \Theta(M') > 0$. Suppose on the contrary that $M_1 \not\subset M'$. We then have a proper inclusion $M' \subsetneq M_1 + M'$. In particular, $\kappa(M') < \kappa(M_1 + M')$. On the other hand, $\Theta(M') \ge \Theta(M_1 + M')$ as M' attains the discrepancy. If $\Theta(M_1 + M') \le 0$, we automatically have $\mu(M') > \mu(M_1 + M')$. If $\Theta(M_1 + M') > 0$, we have

$$\frac{\Theta(M')}{\kappa(M')} \ge \frac{\Theta(M_1 + M')}{\kappa(M')} > \frac{\Theta(M_1 + M')}{\kappa(M_1 + M')}.$$

In any case we deduce that $\mu(M') > \mu(M_1 + M')$. By Lemma 2.2, the following short exact sequence

$$0 \to M' \to M_1 + M' \to (M_1 + M')/M' \simeq M_1/(M_1 \cap M') \to 0$$

implies $\mu(M') > \mu(M_1 + M') > \mu(M_1/(M_1 \cap M'))$.

On the other hand, since M_1 is the scss subrepresentation of M, we have $\mu(M_1 \cap M') \leq \mu(M_1)$. Now Lemma 2.2 applied to the following short exact sequence

$$0 \to M_1 \cap M' \to M_1 \to M_1/(M_1 \cap M') \to 0$$

implies $\mu(M_1/(M_1 \cap M')) \geq \mu(M_1)$. In sum, we obtained

$$\mu(M') > \mu(M_1 + M') > \mu(M_1/(M_1 \cap M')) \geq \mu(M_1).$$

However, M_1 is the scss subrepresentation of M. We arrive at a contradiction.

For the second implication of the lemma, if M' is the scss subrepresentation of M, it is obviously μ-semistable. Conversely, if M' is μ-semistable, since it contains M_1, we must have $\mu(M_1) \leq \mu(M')$. This implies $\mu(M_1) = \mu(M')$ so that $M' = M_1$. $\qquad \square$

Now we can make sense of the following proposition.

Proposition 3.2. *Let M be a μ-unstable representation, and let M_1 be the scss subrepresentation of M with $\Theta(M_1) > 0$. If M' is Θ-optimal in M, we then have*

$$\Theta(M'/M_1) = \mathrm{disc}(M/M_1, \Theta).$$

Namely, M'/M_1 is Θ-optimal in M/M_1.

Proof. Suppose N/M_1 is Θ-optimal in M/M_1. We then obviously have

$$\Theta(M'/M_1) \leq \mathrm{disc}(M/M_1, \Theta) = \Theta(N/M_1).$$

This is equivalent to $\Theta(M') \leq \Theta(N)$. However, M' is Θ-optimal in M by assumption. So we must have $\Theta(M') = \Theta(N)$, resulting in $\Theta(M'/M_1) = \Theta(N/M_1) = \mathrm{disc}(M/M_1, \Theta)$. The proposition is proved. $\quad \square$

Theorem 3.3. *Let M be a μ-unstable representation and let*

$$0 \subsetneq M_1 \subsetneq \cdots \subsetneq M_{r-1} \subsetneq M_r = M$$

be its Harder–Narasimhan filtration. Suppose there is an integer l such that $\mu(M_l/M_{l-1}) > 0$ but $\mu(M_{l+1}/M_l) \leq 0$. Then the term M_l in the filtration is Θ-optimal in M. In addition, M_l has the highest slope among all Θ-optimal subrepresentations.

Proof. We will prove the theorem by induction on l. Let $l = 1$. We need to show that $\Theta(M_1) \geq \Theta(M')$ for every $M' \subset M$. We break this down into several steps:

(1) For any $M_1' \subset M_1$, we show that $\Theta(M_1') \leq \Theta(M_1)$.

(2) Next, assuming on the contrary that there is an $M' \subset M$ with $\Theta(M') > \Theta(M_1)$, we prove that such an M' induces the proper inclusion $M_1 \subsetneq M_1 + M'$.

(3) Finally, we deduce from steps (1) and (2) that $\mu((M_1 + M')/M_1) > \mu(M_2/M_1)$, contradicting the fact that M_2/M_1 is the scss subrepresentation of M/M_1.

For step (1), suppose on the contrary that there is an $M_1' \subsetneq M_1$ such that $\Theta(M_1') > \Theta(M_1)$. We would then have

$$\mu(M_1') = \frac{\Theta(M_1')}{\kappa(M_1')} > \frac{\Theta(M_1)}{\kappa(M_1')} > \frac{\Theta(M_1)}{\kappa(M_1)} = \mu(M_1).$$

Note that again we are using the assumption that $\Theta(M_1)$ is positive. Since M_1 is the scss subrepresentation of M, we arrive at a contradiction. Step (1) is completed.

For step (2), since M_1 is scss in M, we get $\mu(M') \leq \mu(M_1)$. This together with the assumption that $\Theta(M') > \Theta(M_1)$ imply $\kappa(M') > \kappa(M_1)$. This clearly implies that M' is not contained in M_1 so that $M_1 \subsetneq M_1 + M'$, finishing step (2).

For step (3), simply note that $\Theta(M') > \Theta(M_1) \geq \Theta(M_1 \cap M')$ by step (1). We then have

$$\mu((M_1 + M')/M_1) = \mu(M'/(M_1 \cap M')) = \frac{\Theta(M') - \Theta(M_1 \cap M')}{\kappa(M') - \kappa(M_1 \cap M')} > 0 \geq \mu(M_2/M_1).$$

These complete all three steps and therefore the base case of the induction.

If $\mu(M_{l+1}/M_l) > 0$ but $\mu(M_{l+2}/M_{l+1}) \leq 0$, consider the Harder–Narasimhan filtration for M/M_1:

$$0 \subsetneq M_2/M_1 \subsetneq \cdots \subsetneq M/M_1.$$

By the induction hypothesis, $\Theta(M_{l+1}/M_1) = \mathrm{disc}(M/M_1, \theta)$. Moreover, Proposition 3.2 implies $\mathrm{disc}(M/M_1, \Theta) = \mathrm{disc}(M, \Theta) - \Theta(M_1)$. Combining these two, we get $\Theta(M_{l+1}) = \mathrm{disc}(M, \Theta)$, as desired.

Finally, the maximality of slope of M_l can also be proved by induction on l, together with Lemma 3.1. $\square$

Corollary 3.4. *Suppose M is a μ-unstable representation of Q, and $\Theta(M) = 0$. Then there exists a term M_i in the Harder–Narasimhan filtration of M such that $\Theta(M_i) = \mathrm{disc}(M, \Theta)$, and such that M_i has the highest slope among all subrepresentations of M that witness $\mathrm{disc}(M, \Theta)$.*

Proof. Let $0 \subsetneq M_1 \subsetneq \cdots \subsetneq M_r = M$ be the Harder–Narasimhan filtration of M. The integer l described in Theorem 3.3 must exist. For if this is not the case, we then have

$$0 = \Theta(M) > \Theta(M_{r-1}) > \cdots > \Theta(M_1) > 0,$$

which is absurd. $\square$

We now turn our attention to the computation of Harder–Narasimhan filtrations for representations of Q. This comes down to computing the scss subrepresentation. The main ideas come from Lemma 3.1, and the following machinery that works for arbitrary infinite fields.

Theorem 3.5. *Let Q be an acylcic quiver and let M be a representation of Q over an infinite field. Fix a weight $\theta : \mathbb{Z}^{Q_0} \to \mathbb{Z}$ with $\theta(M) = 0$. There is a deterministic algorithm that finds the discrepancy $\mathrm{disc}(M, \theta)$, together with a subrepresentation M' so that $\theta(M') = \mathrm{disc}(M, \theta)$. If we set $\Omega = \sum_{v \in Q_0} |\theta(v)|$, $K = \sum_{v \in Q_0} \dim M_v$, and P as the number of paths in Q, then the algorithm has run time complexity that is polynomial in Ω, K, P.*

We refer interested readers to the Appendix for a proof of the theorem. It is presented at Corollary A.8. We now explain how Lemma 3.1, together with the algorithm of Theorem 3.5 can be applied for deriving the scss subrepresentation. For convenience, we adopt the following notations.

Notation. For any representation M of Q of dimension $\boldsymbol{d}$, $F(M)$ denotes the output subrepresentation of the algorithm in Theorem 3.5 applied to M and to the weight $\theta_{\boldsymbol{d}}$. We let $G(M) = \mathrm{disc}(M, \theta_{\boldsymbol{d}}) = \theta_{\boldsymbol{d}}(F(M))$.

As a first instance of the convenience of notations just introduced, we can say a representation M is μ-semistable (resp. μ-unstable) if and only if $G(M) = 0$ (resp. $G(M) > 0$) (see Lemma 2.7).

Now let $M = M^0$ be a representation of Q. If $G(M^0) = 0$, then M^0 itself is the scss subrepresentation. Otherwise, let $M^1 = F(M^0)$. Lemma 3.1 then dictates that $G(M^1) = 0$ if and only if M^1 is the scss subrepresentation of M. If $G(M^1) > 0$, then $M^2 = F(M^1)$ is a proper subrepresentation of M^1. We continue to produce $M^i = F(M^{i-1})$ if $G(M^{i-1}) > 0$. The procedure must end in a finite number of steps as M is finite dimensional at each vertex, and each M^i is a proper subrepresentation of M^{i-1}. At the end, we have a filtration

$$0 \subsetneq M^r \subsetneq M^{r-1} \subsetneq \cdots \subsetneq M^1 \subsetneq M^0,$$

where

(1) $\theta_{\boldsymbol{d}_{i-1}}(M^i) > 0$,

(2) M^i is $\theta_{\boldsymbol{d}_{i-1}}$-optimal in M^{i-1}, and

(3) M^r is μ-semistable.

The following proposition ensures that in this case M^r is the scss subrepresentation of M.

Proposition 3.6. *Let* $0 \subsetneq M^r \subsetneq \ldots \subsetneq M^0$ *be a filtration of representations of Q and let* $\dim M^i = \boldsymbol{d}_i \in \mathbb{Z}^{Q_0}$. *Suppose for each i,*

$$\theta_{\boldsymbol{d}_{i-1}}(M^i) = \mathrm{disc}(M^{i-1}, \theta_{\boldsymbol{d}_{i-1}}) > 0.$$

Then M^r is μ-semistable if and only if M^r is the scss subrepresentation of M^i for $i = 0, \ldots, r-1$.

Proof. The if part is trivial. The proof for the only if part can be carried out by induction. We let M_1^i be the scss subrepresentation of M^i. Suppose M^r is μ-semistable. Then Lemma 3.1 (applied to $\theta_{\boldsymbol{d}_{r-1}}$) implies $M^r = M_1^{r-1}$. Now suppose $M^r = M_1^i$, we will show that $M^r = M_1^{i-1}$. For this, since M^i is $\theta_{\boldsymbol{d}_{i-1}}$-optimal in M^{i-1}, Lemma 3.1 implies M^i contains M_1^{i-1}. This immediately implies $M_1^{i-1} = M_1^i$. By the induction hypothesis, $M^r = M_1^i$ so that $M^r = M_1^{i-1}$, finishing the induction. $\square$

We therefore propose the following algorithm to compute the Harder–Narasimhan filtration of a representation of an acyclic quiver. The outer while loop tests at each i-th step if M/M_i is μ-semistable. If not, the inner while loop then computes the scss subrepresentation M_{i+1}/M_i of M/M_i. The correctness of the inner loop is established by Proposition 3.6. In the end of the algorithm we have M_1, $M_2/M_1, \ldots, M/M_{r-1}$ for some $r \geq 1$, where each M_{i+1}/M_i is the scss subrepresentation of M/M_{i+1}, and M/M_{r-1} is μ-semistable. From these the Harder–Narasimhan filtration for M is immediate.

We now show that the above algorithm satisfies the complexity bound given in Theorem 3.5.

Algorithm 1: An algorithm for Harder–Narasimhan filtrations

Input : A Representation M of Q
Output: The Harder–Narasimhan filtration of M

if $G(M) = 0$ **then**
    return M /* In this case M is semistable, so the filtration is M itself. */
end if
while $G(M) > 0$ **do**
    $N \leftarrow F(M)$
    while $G(N) > 0$ **do**
        $N \leftarrow F(N)$
    end while
    Record N /* N is the scss subrepresentation of M */
    Compute M/N
    $M \leftarrow M/N$
end while

Theorem 3.7. *Let Q be an acyclic quiver and let M be a representation of Q over an infinite field. Let $\Theta, \kappa : \mathbb{Z}^{Q_0} \to \mathbb{Z}$ be two weights that define the slope $\mu = \Theta/\kappa$, P be the number of paths in Q, $\Omega = \sum_v |\Theta(v)|$, and let $K = \kappa(M)$. Algorithm 1 constructs the Harder–Narasimhan filtration of M within time complexity that is polynomial in Ω, K, and P.*

Note that this is only the first part of Theorem B. We postpone the precise statement and the proof of the second part to Theorem 4.6 at the end of the next section.

Proof of Theorem 3.7. For starters, in the i-th outer while loop of Algorithm 1 (i starts from 0), the first line uses the algorithm of Theorem 3.5 to test μ-stability of M/M_i. Let $\boldsymbol{d}^i$ be the dimension vector for the quotient M/M_i. Theorem 3.5 states that the time complexity to determine stability of M/M_i is polynomial in $\sum_v |\theta_{\boldsymbol{d}^i}(v)|$, $\sum_v \dim(M/M_i)_v$, and P. Obviously the second term is bounded by K. To bound the first term by a polynomial in Ω and K, we first note that

$$|\Theta(M/M_i)| \leq \sum_v |\Theta(v)| \dim(M/M_i)_v \leq \Omega K.$$

Therefore

$$\sum_v |\theta_{\boldsymbol{d}^i}(v)| = \sum_v |\kappa(M/M_i)\Theta(v) - \Theta(M/M_i)\kappa(v)|$$

$$\leq \sum_v \kappa(M/M_i)|\Theta(v)| + |\Theta(M/M_i)|\kappa(v)$$

$$\leq \Omega K + \Omega K^2.$$

Hence, the first line of Algorithm 1 in the outer while loop has time complexity polynomial in Ω, K, P.

Next, in the j-th inner while loop, we are testing μ-stability of some subrepresentation M_i^j/M_i of M/M_i using Theorem 3.5. As before, the time complexity is bounded by a polynomial in Ω, K, P. Since

there are at most $\sum_v \dim M_v \leq K$ many inner while loops, completing the inner while loop has again time complexity polynomial in Ω, K, P.

After the inner loop, we arrive at the scss subrepresentation M_{i+1}/M_i of M/M_i. To compute a basis for the quotient $(M/M_{i+1})_v$ at each vertex v, we may compute a basis for a complement of $(M_{i+1}/M_i)_v$ in $(M/M_i)_v$. This can be done by computing the null space of the basis matrix for $(M_{i+1}/M_i)_v$, which has complexity $O(\dim(M_{i+1}/M_i)_v^2 \dim(M/M_i)_v)$. The complexity to compute bases at all vertices can therefore be bounded by $O(K^3)$. Next, for each arrow $a \in Q_1$, the map $(M/M_{i+1})(a)$ can be obtained via the following change of bases to the original map

$$(M/M_i)(a) : (M_{i+1}/M_i)_{ta} \oplus (M/M_{i+1})_{ta} \to (M_{i+1}/M_i)_{ha} \oplus (M/M_{i+1})_{ha},$$

which has time complexity bounded by $O(K^3)$. Therefore, the time complexity to compute the quotient M/M_{i+1} is bounded by a polynomial in K and P.

In sum, we showed that each outer while loop has time complexity polynomial in Ω, K, P. Since there can be at most $\sum_v \dim M_v \leq K$ outer loops, the total time complexity is still polynomial in Ω, K, P. $\square$

4. Completing Theorem B

The major goal of this section is to complete Theorem B. That is, we are going to construct a maximally destabilizing one parameter subgroup from the Harder–Narasimhan filtration of an unstable representation. This is where GIT comes into play.

GIT stability and its numerical criterion, known as the Hilbert–Mumford criterion (Theorem 4.3), were established by Mumford in [Mumford et al. 1994]. The Hilbert–Mumford criterion reduces testing stability for reductive group actions to one dimensional torus actions. This is done by restricting the action to one parameter subgroups of the reductive group. Mumford first conjectured the existence of the one parameter subgroups that fail the numerical criterion maximally. The conjecture was then resolved by Kempf [1978] in his famous paper. Beyond existence Kempf actually established the uniqueness of maximally destabilizing one parameter subgroups up to conjugacy by some parabolic subgroup (Theorem 4.5).

In Section 4A, we will first recall Mumford's definition of stability (Definition 4.1) and the Hilbert–Mumford criterion (Theorem 4.3) in the affine setting. We then make precise the meaning of maximally destabilizing one parameter subgroups (Definition 4.4), and present Kempf's theorem (Theorem 4.5).

We then apply these machinery in the context of representations of quivers in Section 4B. Finally in Section 4C, we recall in Theorem 4.6 a way due to Zamora [2014] to construct a maximally destabilizing one parameter subgroup from the Harder–Narasimhan filtration of an unstable representation. This completes Theorem B.

4A. *Instability in affine geometric invariant theory.* Here we recall necessary notions and theorems from affine GIT. We work with a fixed algebraically closed field k. Let G be a reductive group acting on an affine variety X. We let $\mathbf{\Gamma}(G)$ denote the set of one parameter subgroups of G. We fix a norm $\|-\|$ on $\mathbf{\Gamma}(G)$. The norm satisfies the following two properties:

(1) $\|-\|$ is invariant under conjugation. Namely, for any $g \in G$ and any $\lambda \in \mathbf{\Gamma}(G)$, we have

$$\|\lambda\| = \|g\lambda g^{-1}\|.$$

(2) For any maximal torus $T \subset G$, the restriction of $\|-\|$ to the lattice $\mathbf{\Gamma}(T)$ is induced by an inner product on the vector space $\mathbf{\Gamma}(T) \otimes \mathbb{R}$ that is integral on $\mathbf{\Gamma}(T) \times \mathbf{\Gamma}(T)$.

For the precise notion of a norm on the set of one parameter subgroups, we refer the readers to page 58 of [Mumford et al. 1994].

Stability in GIT depends on the choice of a linearized line bundle. In the affine setting, a character $\chi : G \to k^\times$ of G induces a linearization of the trivial line bundle $\mathscr{O}_X$. Fix a character χ.

For any one parameter subgroup $\lambda : k^\times \to G$ of G, we let $\langle \chi, \lambda \rangle$ be the integer that satisfies $(\chi \circ \lambda)(t) = t^{\langle \chi, \lambda \rangle}$ for all $t \in k^\times$.

Definition 4.1. An element $f \in k[X]$ is χ-*invariant of weight d* if $f(g \cdot x) = \chi^d(g^{-1}) f(x)$ for all $g \in G$ and for all $x \in X$. We say a point $x \in X$ is χ-*semistable* if there is a χ-invariant f of positive weight such that $f(x) \neq 0$. We say a point $x \in X$ is χ-*unstable* if x is not χ-semistable. We write $X^{\mathrm{ss}}(\chi)$ as the set of χ-semistable points in X and $X^{\mathrm{us}}(\chi)$ as the complement $X - X^{\mathrm{ss}}(\chi)$.

It follows from the definition that $X^{\mathrm{ss}}(\chi)$ is a G-invariant open subvariety and that $X^{\mathrm{us}}(\chi)$ is a G-invariant closed subvariety. Moreover, $X^{\mathrm{ss}}(\chi) = X^{\mathrm{ss}}(\chi^d)$ for any $d > 0$.

Remark 4.2 [Mumford et al. 1994]. Let $k[X]_{\chi,d}$ be the space of χ-invariant elements of weight d. The space $\bigoplus_{d \geq 0} k[X]_{\chi,d}$ has a natural graded ring structure. Let

$$X /\!\!/_\chi G := \mathrm{Proj}\left(\bigoplus_{d \geq 0} k[X]_{\chi,d} \right).$$

Then there is a map $X^{\mathrm{ss}}(\chi) \to X /\!\!/_\chi G$ that is constant on G-orbits, submersive and induces a bijection between points in $X /\!\!/_\chi G$ and closed orbits in $X^{\mathrm{ss}}(\chi)$. Moreover, $X /\!\!/_\chi G$ is a quasiprojective variety that is known as *the GIT quotient of X by G with respect to χ*.

We now introduce the Hilbert–Mumford criterion, which tests stability by restricting the group action to certain one parameter subgroups. Let $\lambda : k^\times \to G$ be a one parameter subgroup and let $x \in X$ be a point. We say $\lim_{t \to 0} \lambda(t) \cdot x$ exists if the domain of the map $\lambda_x : k^\times \to X$ defined by $t \mapsto \lambda(t) \cdot x$ can be extended to the entire affine line.

Theorem 4.3 (Hilbert–Mumford criterion [Mumford et al. 1994; King 1994]). *A point $x \in X$ is χ-semistable if and only if for each one parameter subgroup $\lambda : k^\times \to G$ such that $\lim_{t \to 0} \lambda(t) \cdot x$ exists, we have $\langle \chi, \lambda \rangle \leq 0$.*

Next, we define a numerical measure for instabilities contributed by destabilizing one parameter subgroups. For a point $x \in X$, set

$$C_x = \{\lambda \in \mathbf{\Gamma}(G) \mid \lim_{t \to 0} \lambda(t) \cdot x \text{ exists}\}.$$

According to Theorem 4.3, $x \in X^{\mathrm{us}}(\chi)$ if and only if there is a one parameter subgroup λ such that

(1) $\lambda \in C_x$, and

(2) $\langle \chi, \lambda \rangle > 0$.

Therefore, for an $x \in X^{\mathrm{us}}(\chi)$, it is natural to ask if there is a one parameter subgroup that contributes to the highest instability measured by the quantities $\langle \chi, \lambda \rangle$ among all $\lambda \in C_x$. An immediate problem is that $\langle \chi, \lambda^N \rangle = N \cdot \langle \chi, \lambda \rangle$ for any $N \in \mathbb{N}$. To get rid of the dependency on multiples of one parameter subgroups, we divide the function $\langle \chi, - \rangle : \Gamma(G) \to \mathbb{Z}$ by the norm $\| - \|$ on $\Gamma(G)$.

For $x \in X^{\mathrm{us}}(\chi)$, we set

$$\mathcal{M}^\chi(x) := \sup_{\lambda \in C_x \setminus \{0\}} \frac{\langle \chi, \lambda \rangle}{\|\lambda\|}.$$

Definition 4.4. We say a one parameter subgroup λ is *indivisible* if λ is a primitive lattice in a (and hence in any) maximal torus containing λ. We say a one parameter subgroup λ is χ-*adapted to x* if $\langle \chi, \lambda \rangle / \|\lambda\| = \mathcal{M}^\chi(x)$. We let $\Lambda^\chi(x)$ denote the set of indivisible one parameter subgroups that are χ-adapted to x.

We also recall that for any one parameter subgroup λ, there is the associated parabolic subgroup

$$P(\lambda) = \{g \in G \mid \lim_{t \to 0} \lambda(t) g \lambda(t)^{-1} \text{ exists in } G\}$$

of G. An important property about $P(\lambda)$ we want to mention is that it preserves filtration in the following sense: Let V be a representation of G and let $\bigoplus_{n \in \mathbb{Z}} V^{(n)}$ be the weight decomposition of some one parameter subgroup λ of G. Namely, λ acts on $V^{(n)}$ with weight n. Suppose $p \in P(\lambda)$, we then have

$$p \cdot V^{(\geq n)} = V^{(\geq n)}$$

where $V^{(\geq n)} = \bigoplus_{m \geq n} V^{(m)}$.

We are now ready for Kempf's theorem.

Theorem 4.5 [Kempf 1978]. *Let G be a reductive group, acting on an affine variety X. Let χ be a character of G, and let $x \in X$ be a χ-unstable point. Then:*

(1) *$\Lambda^\chi(x)$ is not empty.*

(2) *For any $g \in G$, $\Lambda^\chi(g \cdot x) = g \Lambda^\chi(x) g^{-1}$, and $\mathcal{M}^\chi(g \cdot x) = \mathcal{M}^\chi(x)$.*

(3) *There is a parabolic subgroup $P(\chi, x)$ of G such that for all $\lambda \in \Lambda^\chi(x)$, $P(\lambda) = P(\chi, x)$, and such that any two elements of $\Lambda^\chi(x)$ are conjugate to each other by an element in $P(\chi, x)$.*

4B. *The quiver setting.* We now unpack what these machineries mean in the case of representation of quivers. For starters, let us fix an acyclic quiver Q, and a dimension vector $\boldsymbol{d} \in \mathbb{Z}^{Q_0}$. By $\mathrm{GL}(n)$ we mean the linear algebraic group of invertible $n \times n$ matrices with entries in k. Let the linear algebraic group

$$\mathrm{GL}(\boldsymbol{d}) := \prod_{v \in Q_0} \mathrm{GL}(\boldsymbol{d}_v)$$

act on the k-vector space

$$R(Q, \boldsymbol{d}) := \prod_{a \in Q_1} \mathrm{Hom}_k(k^{\boldsymbol{d}_{ta}}, k^{\boldsymbol{d}_{ha}})$$

via the following law:

$$(g \cdot \phi)_a = g_{ha} \circ \phi_a \circ (g_{ta})^{-1} \text{ for any } g \in G, \phi \in R(Q, \boldsymbol{d}).$$

In this way, the GL($\boldsymbol{d}$)-orbits in $R(Q, \boldsymbol{d})$ correspond to the isomorphism classes of representations of Q of dimension $\boldsymbol{d}$. A weight $\theta : \mathbb{Z}^{Q_0} \to \mathbb{Z}$ defines a character $\chi_\theta : \mathrm{GL}(\boldsymbol{d}) \to k^\times$ by

$$\chi_\theta(g) = \prod_{v \in Q_0} \det(g_v)^{\theta(v)} \text{ for any } g \in \mathrm{GL}(\boldsymbol{d}).$$

We fix a weight θ with $\theta(\boldsymbol{d}) = 0$.

Let M be a representation of Q of dimension $\boldsymbol{d}$. Since χ_θ-stability is constant on GL($\boldsymbol{d}$)-orbits, different choices of bases for each M_v does not affect the χ_θ-stability of the corresponding elements of M in $R(Q, \boldsymbol{d})$. From now on we fix an identification $M_v \simeq \mathbb{R}^{d_v}$ for each v, and still write the identification of M in $R(Q, \boldsymbol{d})$ as M.

We now recall the following relation between weighted filtrations of M, and the one parameter subgroups having limits at M, namely, the one parameter subgroups in C_M. Any one parameter subgroup $\lambda \in C_M$ induces a weighted filtration

$$0 = M_0 \subsetneq M_1 \subsetneq M_2 \subsetneq \cdots \subsetneq M_s = M,$$

where λ acts on each quotient M_i/M_{i-1} with weight Γ_i. Moreover, the weights enjoy the following property: $\Gamma_1 > \Gamma_2 > \cdots > \Gamma_s$; see [King 1994].

Conversely, given a weighted filtration $0 = M_0 \subsetneq M_1 \subsetneq \cdots \subsetneq M_s = M$ with weights $\Gamma_1 > \cdots > \Gamma_s$, we may choose new bases for M at all vertices that are compatible with the filtration. Namely, choose a basis for M_1, then extends the basis to M_2, M_3, and so on. We then let $g_v \in \mathrm{GL}(\boldsymbol{d}_v)$ be the one parameter subgroup that acts on $(M_i/M_{i-1})_v$ by weight Γ_i. In this way, the one parameter subgroup $\lambda = \prod_{v \in Q_0} g_v$ of GL($\boldsymbol{d}$) induces the original filtration. Moreover, since $\theta(M) = 0$, we have

$$\langle \chi_\theta, \lambda \rangle = \sum_{i=1}^s \Gamma_i \theta(M_i/M_{i-1}) = \sum_{i=1}^{s-1} (\Gamma_i - \Gamma_{i+1}) \theta(M_i). \tag{4-1}$$

In particular, for any proper subrepresentation M' of M, one may form the filtration $0 \subsetneq M' \subsetneq M$, and assign weights Γ_1, Γ_2 to M' and M respectively, where Γ_1 is any integer and $\Gamma_2 = \Gamma_1 - 1$. It then follows from the above discussion and (4-1) that there is a one parameter subgroup λ with $\langle \chi_\theta, \lambda \rangle = \theta(M')$. In view of Theorem 4.3, we see that M is χ_θ-semistable if and only if $\theta(M') \le 0$ for any subrepresentation M' of M. This was the definition, Definition 2.5 we used for weight stability.

4C. *From the Harder–Narasimhan filtration to maximally destabilizing one parameter subgroups.* In this section, we address how to derive a maximally destabilizing one parameter subgroup of an unstable representation M of Q from its Harder–Narasimhan filtration. We need to set up appropriate weights and slope.

Let $\Theta, \kappa : \mathbb{Z}^{Q_0} \to \mathbb{Z}$ be two weights, and let $\mu = \Theta/\kappa$ be the associated slope function. Suppose $\dim M = \mathbf{d}$. Recall that we defined the weight $\theta_{\mathbf{d}}$ where

$$\theta_{\mathbf{d}}(N) := \kappa(\mathbf{d})\Theta(N) - \Theta(\mathbf{d})\kappa(N) \text{ for any representation } N.$$

With these we have that $\theta_{\mathbf{d}}(M) = 0$, and that M is μ-semistable if and only if it is $\chi_{\theta_{\mathbf{d}}}$-semistable Lemma 2.7.

We now define a norm $\|-\|$ on $\mathbf{\Gamma}(\mathrm{GL}(\mathbf{d}))$. This norm will depend on the weight κ. We first build the norm on the diagonal maximal torus $T \subset \mathrm{GL}(\mathbf{d})$. There is a natural identification

$$\mathbf{\Gamma}(T)_{\mathbb{R}} \simeq \bigoplus_{v \in Q_0} \mathbb{R}^{\mathbf{d}_v}.$$

We weight the standard norm on $\mathbb{R}^{\mathbf{d}_v}$ by $\kappa(v)$ for each v. More explicitly, if $\lambda \in \mathbf{\Gamma}(T)$ with $\lambda_v = (\lambda_{1,v}, \ldots, \lambda_{\mathbf{d}_v, v})$, we set

$$\|\lambda\| = \left(\sum_{v \in Q_0} \sum_{i=1}^{\mathbf{d}_v} \kappa(v)(\lambda_{i,v})^2 \right)^{1/2}. \tag{4-2}$$

This norm $\|-\|$ extends to the entire set $\mathbf{\Gamma}(\mathrm{GL}(\mathbf{d}))$ and is invariant under conjugation. We refer the readers to page 58 in [Mumford et al. 1994] for a detailed discussion on the construction of norms on the set of one parameter subgroup of a reductive group.

Therefore, whenever $\lambda \in C_M$ (that is, $\lim_{t \to 0} \lambda(t) \cdot M$ exists), its norm can be expressed in terms of its weights and the filtration it induces: If $0 = M_0 \subsetneq M_1 \subsetneq \cdots \subsetneq M_s = M$ is the filtration induced by λ with weights $\Gamma_1 > \cdots > \Gamma_s$, then

$$\|\lambda\| = \left(\sum_{i=1}^{s} \sum_{v \in Q_0} \kappa(v)\Gamma_i^2 \dim(M_i/M_{i-1})_v \right)^{1/2} = \left(\sum_{i=1}^{s} \Gamma_i^2 \kappa(M_i/M_{i-1}) \right)^{1/2}. \tag{4-3}$$

We adopt the above norm to define (using Definition 4.4) one parameter subgroups that are $\chi_{\theta_{\mathbf{d}}}$-adapted to M.

Since all one parameter subgroups that are $\chi_{\theta_{\mathbf{d}}}$-adapted to M are in a full conjugacy class of their parabolic subgroup (Theorem 4.5), the filtrations induced by those one parameter subgroups are the same. In [Zamora 2014], such a filtration is named as the Kempf filtration of M. By Theorem 5.3 from [loc. cit.], it is also the Harder–Narasimhan filtration for M with respect to the slope μ. In addition, there is the following recipe for reverse engineering a maximally destabilizing one parameter subgroup from the Harder–Narasimhan filtration.

Theorem 4.6 [Zamora 2014, Theorem 4.1, Lemma 5.2]. *Let* $\Theta, \kappa : \mathbb{Z}^{Q_0} \to \mathbb{Z}$ *be two weights and let* $\mu = \Theta/\kappa$ *be the slope. Let* M *be a* μ-*unstable representation of dimension* $\boldsymbol{d}$, *and let*

$$0 = M_0 \subsetneq M_1 \subsetneq \cdots \subsetneq M_r = M$$

be the Harder–Narasimhan filtration. Set

$$u_i = \kappa(M)\mu(M_i/M_{i-1}) - \Theta(M) \; for \; i = 1, \ldots, r.$$

Choose a basis for M_v *for each vertex* v *compatible with the filtration, and let* $T \subset \mathrm{GL}(\boldsymbol{d})$ *be the maximal torus diagonal with respect to these bases. Define* $\tilde{\lambda} \in \Gamma(T)_{\mathbb{Q}} \simeq \bigoplus_{v \in Q_0} \mathbb{Q}^{d_v}$ *where for each* i *and* v, *the entries of* $\tilde{\lambda}_v$ *that correspond to* $(M_i/M_{i-1})_v$ *are* u_i. *Then the lattice points on the ray* $\mathbb{Q}_{>0} \cdot \tilde{\lambda}$ *are the one parameter subgroups in* T *that are* χ_{θ_d}-*adapted to* M. *Moreover,*

$$\mathcal{M}^{\chi_{\theta_d}}(M) = \left(\sum_{i=1}^{r} (\kappa(M)\mu(M_i/M_{i-1}) - \Theta(M))^2 \kappa(M_i/M_{i-1}) \right)^{1/2}.$$

Remark 4.7. In the case that $\Theta(M) = 0$ in the first place, one can take

$$u_i = \mu(M_i/M_{i-1}),$$

and

$$\mathcal{M}^{\chi_{\Theta}}(M) = \left(\sum_i \mu(M_i/M_{i-1})^2 \kappa(M_i/M_{i-1}) \right)^{1/2} = \left(\sum_i \Theta(M_i/M_{i-1})\mu(M_i/M_{i-1}) \right)^{1/2}.$$

Therefore, knowing the Harder–Narasimhan filtration of an unstable representation is equivalent to knowing its maximally destabilizing one parameter subgroups. A full treatment of Theorem B is now complete. An account of the correctness of the construction in Theorem 4.6 is given in the next section for interested readers.

4D. *Maximizing Kempf function.* Let M be a μ-unstable representation of Q of dimension $\boldsymbol{d}$ with the Harder–Narasimhan filtration

$$0 = M_0 \subsetneq M_1 \subsetneq \cdots \subsetneq M_r = M.$$

By Theorem 5.3 from [Zamora 2014], the filtration is the same as the Kempf filtration. Recall that C_M denotes the set of one parameter subgroups having limits at M. Let

$$f : C_M \backslash \{\boldsymbol{0}\} \to \mathbb{R}$$

be the function defined by

$$f(\lambda) = \frac{\langle \chi_{\theta_d}, \lambda \rangle}{\|\lambda\|}.$$

To maximize this function, we can restrict the search range to the set of one parameter subgroups that induce the Kempf filtration, which is also the Harder–Narasimhan filtration.

Consider the inner product $(\,-\,,\,-\,) : \mathbb{R}^r \times \mathbb{R}^r \to \mathbb{R}$ defined by the matrix

$$\begin{pmatrix} \kappa(M_1) & & & 0 \\ & \kappa(M_2/M_1) & & \\ & & \ddots & \\ 0 & & & \kappa(M/M_{r-1}) \end{pmatrix}.$$

If λ is a one parameter subgroup that induces the Harder–Narasimhan filtration of M with the weight vector $\Gamma = (\Gamma_1, \ldots, \Gamma_s)$, then upon using θ_d as θ in (4-1), we derive that

$$\langle \chi_{\theta_d}, \lambda \rangle = \sum_i \Gamma_i \big(\kappa(M)\Theta(M_i/M_{i-1}) - \Theta(M)\kappa(M_i/M_{i-1}) \big)$$

$$= \sum_i \Gamma_i \kappa(M_i/M_{i-1})(\kappa(M)\mu(M_i/M_{i-1}) - \Theta(M))$$

$$= (\Gamma, u), \tag{4-4}$$

where $u_i = \kappa(M)\mu(M_i/M_{i-1}) - \Theta(M)$. Moreover, (4-3) implies $\|\lambda\| = (\Gamma, \Gamma)^{1/2}$.

Letting $u = (u_1, \ldots, u_r) \in \mathbb{R}^r$, we define the function

$$g_u : \mathbb{R}^r \backslash \{\mathbf{0}\} \to \mathbb{R}$$

where

$$g_u(x) = \frac{(x, u)}{\sqrt{(x, x)}}.$$

This function is called the Kempf function in [Zamora 2014]. We see that maximizing f among all one parameter subgroups inducing the Kempf filtration is equivalent to maximizing the Kempf function g_u on the open cone $C = \{x \in \mathbb{R}^r \mid x_1 > x_2 > \cdots > x_r\}$. By the Cauchy–Schwarz inequality, g_u attains global maximum at u. Since $\mu(M_1) > \mu(M_2/M_1) > \cdots > \mu(M/M_r)$, u is already in the cone C. Therefore, the one parameter subgroup constructed in Theorem 4.6 maximizes f.

5. A short example

In this example we construct an unstable representation of a bipartite quiver and compute the following:

(1) Several witnesses to the discrepancy.

(2) The Harder–Narasimhan filtration.

(3) A maximally destabilizing one parameter subgroup.

We compare the results of (1), (2) with Theorem 3.3, and demonstrate that not all optimal subrepresentations occur in the Harder–Narasimhan filtration. For (3), we present two approaches. The first approach applies Theorem 4.6 to the Harder–Narasimhan filtration obtained in (2), and the second is a sketch of the brute force procedure that can be applied to any representation of any reductive group. By doing so we hope to give readers a sense of the hardness of finding Kempf's one parameter subgroups in general. The set up of this example is motivated by [Chindris and Kline 2021].

Set up. Let us consider the following representation M of the following bipartite quiver over the field $\mathbb{C}$ of complex numbers:
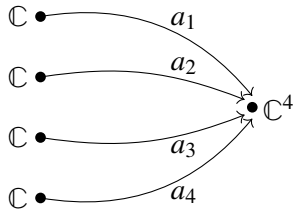
Label the four vertices on the left by x_1, x_2, x_3, x_4 from top to bottom. Label the single vertex on the right as y so that $M_{x_i} = \mathbb{C}$ for each $i = 1, \ldots, 4$, and $M_y = \mathbb{C}^4$.

We set $M(a_1)(1) = e_2$, $M(a_2)(1) = e_1$, $M(a_3)(1) = 2e_1$, and $M(a_4)(1) = e_3$ in $\mathbb{C}^4$, where $e_1, \ldots, e_4$ form the standard basis of $\mathbb{C}^4$. We let Θ and κ be two weights where $\Theta(x_i) = 4$, $\Theta(y) = -4$, and $\kappa(x_i) = \kappa(y) = 1$ for each i. It then follows that $\Theta(M) = 0$. We also let $\mu(-) = \Theta(-)/\kappa(-)$ be the corresponding slope.

We let $\mathrm{Sp}(e_1)$ be the largest subrepresentation M' of M such that $M'_y = \mathrm{Span}(e_1)$. Specifically,

$$\mathrm{Sp}(e_1)_{x_i} = \begin{cases} \mathbb{C} & \text{if } i = 2, 3, \\ 0 & \text{if } i = 1, 4, \end{cases}$$

and

$$\mathrm{Sp}(e_1)_y = \mathrm{Span}(e_1).$$

We define $\mathrm{Sp}(e_1, e_2)$, $\mathrm{Sp}(e_1, e_3)$, and $\mathrm{Sp}(e_1, e_2, e_3)$ similarly.

Witnesses to the discrepancy. We verify that the four subrepresentations $\mathrm{Sp}(e_1)$, $\mathrm{Sp}(e_1, e_2)$, $\mathrm{Sp}(e_1, e_3)$, $\mathrm{Sp}(e_1, e_2, e_3)$ are Θ-optimal in M with Θ value 4. For this, note that for any subrepresentation M' of M,

$$\Theta(M') \leq 4 \cdot (\#\{i \mid M'_{x_i} = \mathbb{C}\} - \dim(\mathrm{Span}(\{M(a_i)(1) \mid M'_{x_i} = \mathbb{C}\}))).$$

Therefore, for $\Theta(M')$ to be positive, M'_{x_2} and M'_{x_3} must be equal to $\mathbb{C}$. The only proper subrepresentations of M that are nonzero at x_2, x_3 with positive Θ values are exactly the four subrepresentations listed above. Since they all have the same Θ value 4, they are also Θ-optimal and $\mathrm{disc}(M, \Theta) = 4$.

The Harder–Narasimhan filtration. We verify that

$$0 \subsetneq \mathrm{Sp}(e_1) \subsetneq \mathrm{Sp}(e_1, e_2, e_3) \subsetneq M$$

is the Harder–Narasimhan filtration of M. For this, it is clear that $\mathrm{Sp}(e_1)$ is the only subrepresentation with the highest slope $4/3$. We may infer that $\mathrm{Sp}(e_1)$ is the scss subrepresentation of M. Moving on to $M/\mathrm{Sp}(e_1)$, since $\mathrm{Sp}(e_1)$ is Θ-optimal, we can infer $\mu(M'/\mathrm{Sp}(e_1)) \leq 0$ for any subrepresentation M' containing $\mathrm{Sp}(e_1)$. Since

$$\mu(\mathrm{Sp}(e_1, e_2)/\mathrm{Sp}(e_1)) = \mu(\mathrm{Sp}(e_1, e_3)/\mathrm{Sp}(e_1)) = \mu(\mathrm{Sp}(e_1, e_2, e_3)/\mathrm{Sp}(e_1)) = 0,$$

we infer that the quotient $\mathrm{Sp}(e_1, e_2, e_3)/\mathrm{Sp}(e_1)$ is the scss subrepresentation of $M/\mathrm{Sp}(e_1)$. The quotient $M/\mathrm{Sp}(e_1, e_2, e_3)$ is one dimensional on y and zero elsewhere. We conclude that $M/\mathrm{Sp}(e_1, e_2, e_3)$ is μ-semistable and that the filtration written above is the Harder–Narasimhan filtration of M.

Here we note that the two Θ-optimal subrepresentations $\mathrm{Sp}(e_1, e_2)$ and $\mathrm{Sp}(e_1, e_3)$ are not included in the filtration. We see in this example that not all optimal subrepresentations occur in the Harder–Narasimhan filtration. Also note that $\mathrm{Sp}(e_1)$ has the highest slope among all Θ-optimal subrepresentations. This agrees with the implications of Theorem 3.3 because $\mu(\mathrm{Sp}(e_1)) = 4/3 > 0 = \mu(\mathrm{Sp}(e_1, e_2, e_3)/\mathrm{Sp}(e_1))$. Namely, $\mathrm{Sp}(e_1)$ serves as M_l in the description of Theorem 3.3.

Maximal destabilizing one parameter subgroups. We now demonstrate a maximally destabilizing one parameter subgroup for M with respect to the weight Θ and the norm weighted by κ. The first thing to do is to reproduce the GIT set up introduced in Section 4B. To begin with, we have the vector space $\prod_{a \in Q_1} \mathrm{Hom}_{\mathbb{C}}(\mathbb{C}, \mathbb{C}^4) \simeq \mathbb{C}^{4\times4}$. Each column of a matrix in $\mathbb{C}^{4\times4}$ defines a map from a left vertex to the right vertex. We let the first column define a map from x_1 to y, and so on. In this way, the group $G = (\mathbb{C}^\times)^4 \times \mathrm{GL}(4)$ acts on $\mathbb{C}^{4\times4}$ via the rule

$$(t_1, t_2, t_3, t_4, A) \cdot U = A \cdot U \cdot \begin{pmatrix} t_1^{-1} & & & \\ & t_2^{-1} & & \\ & & t_3^{-1} & \\ & & & t_4^{-1} \end{pmatrix}$$

for all $(t_1, t_2, t_3, t_4, A) \in G$ and for all $U \in \mathbb{C}^{4\times4}$. There is the diagonal torus $D \simeq (\mathbb{C}^\times)^4$ in $\mathrm{GL}(4)$. Let $T = (\mathbb{C}^\times)^4 \times D \subset G$ be the maximal torus in G. We then have $\Gamma(T)_{\mathbb{Q}} \simeq \mathbb{Q}^8$. With the notations introduced in Theorem 4.6, we have

$$u_1 = \tfrac{4}{3}, \quad u_2 = 0, \quad u_3 = -4$$

by Remark 4.7. The maximally destabilizing one parameter subgroups in T is then on the positive ray of the point

$$\left(0, \tfrac{4}{3}, \tfrac{4}{3}, 0, \underbrace{\tfrac{4}{3}, 0, 0, -4}\right).$$
$$\underbrace{\phantom{0, \tfrac{4}{3}, \tfrac{4}{3}, 0,}}_{x_1, x_2, x_3, x_4} \underbrace{\phantom{\tfrac{4}{3}, 0, 0, -4}}_{y}$$

We now sketch a procedure using linear programming. Here we do not assume we know the Harder–Narasimhan filtration in the first place. Recall that Θ defines the character χ_Θ of G where

$$\chi_\Theta(t_1, t_2, t_3, t_4, A) = (t_1 t_2 t_3 t_4)^4 \cdot \det(A)^{-4}$$

for all $(t_1, t_2, t_3, t_4, A) \in G$. The representation M corresponds to the matrix:

$$\begin{pmatrix} 0 & 1 & 2 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \end{pmatrix}$$

Let us continue to use M to represent the matrix.

If $\lambda(t) = (t^{a_1}, \ldots, t^{a_8})$ is a one parameter subgroup in T, then

$$\lambda(t) \cdot M = \begin{pmatrix} 0 & t^{-a_2+a_5} & 2 \cdot t^{-a_3+a_5} & 0 \\ t^{-a_1+a_6} & 0 & 0 & 0 \\ 0 & 0 & 0 & t^{-a_4+a_7} \\ 0 & 0 & 0 & 0 \end{pmatrix}.$$

We see that $\lim_{t\to\infty} \lambda(t) \cdot M$ exists if and only if

$$\begin{aligned} -a_1 + a_6 &\geq 0, \\ -a_2 + a_5 &\geq 0, \\ -a_3 + a_5 &\geq 0, \\ -a_4 + a_7 &\geq 0. \end{aligned} \tag{5-1}$$

One also easily calculates that

$$\langle \chi_\Theta, \lambda \rangle = 4(a_1 + a_2 + a_3 + a_4) - 4(a_5 + a_6 + a_7 + a_8).$$

The norm on the set of one parameter subgroups of T is weighted by κ. In this case it is simply the standard norm. Therefore, to find one parameter subgroups in T that maximally destabilize M, we are maximizing the function

$$f = \frac{4(a_1 + a_2 + a_3 + a_4) - 4(a_5 + a_6 + a_7 + a_8)}{(a_1^2 + a_2^2 + a_3^2 + a_4^2 + a_5^2 + a_6^2 + a_7^2 + a_8^2)^{1/2}},$$

with the side constraints given by (5-1). With the help of SageMath, we calculated that f attains maximum on the positive ray of the point $\left(0, \frac{4}{3}, \frac{4}{3}, 0, \frac{4}{3}, 0, 0, -4\right)$.

Notice that the constraints depends only on the nonzero entries of M. That is, replacing the nonzero entries of M by any other nonzero scalars does not change the linear programming problem. The set of the positions of the nonzero entries of M is called the support of M.

We must be aware that up to this point, we are only calculating one parameter subgroups in T. In general, maximally destabilizing one parameter subgroups of an unstable point do not have to be in a particular maximal torus. To work with a particular maximal torus, one needs to enumerate all supports coming from the orbit of the unstable point, conduct linear programming for each support, and choose the maximum result. Often times the challenge is to determine the supports of points in the orbit. Beyond that the number of supports of points in the orbit can be quite large, leading to vast amount of linear programming required. As we can see, enumerating supports of $g \cdot M$ for all $g \in G$ in this relatively small example already gets tedious. We shall not continue this brute force procedure here, as the main point to illustrate the hardness of finding Kempf's one parameter subgroups in general, and to demonstrate the value of Theorem 4.6 is made.

Appendix: Algebraic complexity towards the stability of representations of quivers

In this section, we will demonstrate how algebraic complexity tools can be applied to computing the discrepancy, together with a witnesses, for a representation of an acyclic quiver.

Preliminaries. For starters, we let $\mathbb{F}$ be an infinite field of arbitrary characteristic. The space of $n \times n$ matrices over $\mathbb{F}$ is denoted $M(n, \mathbb{F})$. Let $\mathcal{B}$ be a subspace of $M(n, \mathbb{F})$. If U is a subspace of $\mathbb{F}^n$, we set

$$\mathcal{B}(U) = \sum_{B \in \mathcal{B}} BU,$$

where $\sum$ denote the sum of subspaces. We say U is a *c-shrunk subspace of* $\mathcal{B}$ if

$$\dim U - \dim \mathcal{B}(U) \geq c.$$

We say U is a *shrunk subspace of* $\mathcal{B}$ if U is a c-shrunk subspace of $\mathcal{B}$ for some $c \in \mathbb{Z}^+$. We define the *discrepancy of* $\mathcal{B}$ as the integer

$$\mathrm{disc}(\mathcal{B}) = \max_{c \in \mathbb{N}} \{\exists \text{ a } c\text{-shrunk subspace of } \mathcal{B}\}.$$

Remark A.1. For any subspace $\mathcal{B} \subset M(n, \mathbb{F})$, if $c = \mathrm{disc}(\mathcal{B})$, it is shown in [Huszar 2021, Lemma 2.1], that the intersection of c-shrunk subspaces is again a c-shrunk subspace. Therefore, there is an unique minimal c-shrunk subspace of $\mathcal{B}$.

The following is a summary of various results from [Ivanyos et al. 2018] and [Ivanyos et al. 2015] .

Theorem A.2 [Ivanyos et al. 2018, Theorem 1.5; 2015, Proposition 7, Lemma 9]. *Let $\mathcal{B}$ be a matrix space in $M(n, \mathbb{F})$ with $\mathrm{disc}(\mathcal{B}) = c$ (a priori unknown). There exists a deterministic algorithm using $n^{O(1)}$ arithmetic operations over $\mathbb{F}$ that returns a c-shrunk subspace of $\mathcal{B}$. Moreover, the c-shrunk subspace produced by this algorithm is the minimal c-shrunk subspaces of $\mathcal{B}$.*

Remark A.3. The original Theorem 1.5 from [Ivanyos et al. 2018] deals with how field elements are represented in a computer, which is not the main concern of this paper. Personal communications with the authors of [loc. cit.] confirmed that in terms of the number of operations needed, there is nothing special about the field of rational numbers stated in the theorem in [loc. cit.]. Any field with large enough set of elements to work with will suffice.

We now follow the recipe of [Huszar 2021] to build a matrix space from a given representation of an acyclic quiver, and establish the connection between the discrepancy of the matrix space and that of the representation. Recall we wrote a quiver as $Q = (Q_0, Q_1, t, h)$. Multiple arrows between any two vertices are allowed. We fix a weight $\theta : Q_0 \to \mathbb{Z}$ and a representation W of Q of dimension $\boldsymbol{d}$ with $\theta(\boldsymbol{d}) = 0$. If $p = a_j a_{j-1} \cdots a_1$ is a path in Q, we let $W(p)$ denote the composition $W(a_j) W(a_{j-1}) \cdots W(a_1)$.

Let $x_1, \ldots, x_n$ (resp. $y_1, \ldots, y_m$) be the vertices of Q where θ takes positive values (resp. negative values). We then define

$$\theta_+(x_i) = \theta(x_i) \quad (\text{resp. } \theta_-(y_j) = -\theta(y_j)),$$

$$N := \sum_i \theta_+(x_i) \boldsymbol{d}(x_i) = \sum_j \theta_-(y_j) \boldsymbol{d}(y_j),$$

and

$$M = \sum_j \theta_-(y_j) \quad (\text{resp. } M' = \sum_i \theta_+(x_i)).$$

For each $i \in [n]$ and each $j \in [m]$, we let

$$I_i^+ = \left\{ r \in \mathbb{N} \,\Big|\, \sum_{k=1}^{i-1} \theta_+(x_i) < r \le \sum_{k=1}^{i} \theta_+(x_i) \right\},$$

and

$$I_j^- = \left\{ q \in \mathbb{N} \,\Big|\, \sum_{k=1}^{j-1} \theta_-(y_j) < q \le \sum_{k=1}^{j} \theta_-(y_j) \right\}.$$

For each $q \in I_j^-$, $r \in I_i^+$, and each path p between x_i and x_j, we define the $M \times M'$ block matrix $A_{q,r}^{i,j,p}$ whose (q, r)-block is the $\boldsymbol{d}(y_j) \times \boldsymbol{d}(x_i)$ matrix representing $W(p) : \mathbb{R}^{d(x_i)} \to \mathbb{R}^{d(y_j)}$, and whose other blocks are zero block matrices of appropriate sizes. In this way, the same matrix representing $W(p)$ is placed in the various (q', r')-blocks of $M \times M'$ block matrices for $q' \in I_j^-$, and $r' \in I_i^+$. We set

$$\mathcal{A}_{W,\theta} = \text{Span}(\{A_{q,r}^{i,j,p} \mid i \in [n], j \in [m], q \in I_j^-, r \in I_i^+, p \in \mathcal{P}_{i,j}\}) \subset M(N, \mathbb{F}).$$

where $\mathcal{P}_{i,j}$ is the collection of paths from x_i to y_j.

We have an immediate relation described by the

Lemma A.4. *For any representation W, we have $\text{disc}(W, \theta) \le \text{disc}(\mathcal{A}_{W,\theta})$.*

Proof. Let $W' \subset W$ be a subrepresentation. Define $U = \bigoplus_i (W'_{x_i})^{\theta_+(x_i)}$. It then follows that

$$\theta(W') = \sum_i \theta_+(x_i) W'_{x_i} - \sum_j \theta_-(y_j) W'_{y_j}$$

$$= \dim U - \dim\left(\bigoplus_j W_{y_j}^{\theta_-(y_j)} \right)$$

$$\le \dim U - \dim\left(\bigoplus_{j,q} \sum_i \sum_{p \in \mathcal{P}_{ij}} W(p)(W'_{x_i}) \right)$$

$$= \dim U - \dim \mathcal{A}_{W,\theta}(U).$$

This finishes the proof of the lemma. $\square$

Relating the two discrepancies. In this section we relate better the two discrepancies $\mathrm{disc}(W, \theta)$, and $\mathrm{disc}(\mathcal{A}_{W,\theta})$. We will show that they are equal, and demonstrate how to translate a witness to one discrepancy into a witness to the other. For this we need to introduce a right action on $\bigoplus_i W_{x_i}^{\theta_+(x_i)}$ and on $\bigoplus_j W_{y_j}^{\theta_-(y_j)}$ by the paths of Q. For the following, homomorphisms and endomorphisms are taken in the category of representations of Q. In addition, we let $c = \mathrm{disc}(\mathcal{A}_{W,\theta})$ so that by a c-shrunk subspace of $\mathcal{A}_{W,\theta}$ we mean a subspace that witnesses $\mathrm{disc}(\mathcal{A}_{W\theta})$.

For any vertex $x \in Q_0$, we let P_x be the representation of Q where for each $y \in Q_0$, $(P_x)_y$ is the vector space with basis the paths from x to y. Here $(P_x)_x$ is one dimensional with basis the trivial path e_x. For any $a \in Q_1$, we define

$$P_x(a) : (P_x)_{ta} \to (P_x)_{ha}$$

by post composition.

For any representation V of Q, there is a natural identification $\mathrm{Hom}(P_x, V) \simeq V_x$. Letting

$$P_1 = \bigoplus_i P_{x_i}^{\theta_+(x_i)} \quad \left(\text{resp. } P_0 = \bigoplus_j P_{y_j}^{\theta_-(y_j)} \right),$$

we have

$$\mathrm{Hom}(P_1, W) \simeq \bigoplus_i W_{x_i}^{\theta_+(x_i)} \quad \left(\text{resp. } \mathrm{Hom}(P_0, W) \simeq \bigoplus_j W_{y_j}^{\theta_-(y_j)} \right).$$

Under the above identifications, a morphism $\varphi \in \mathrm{Hom}(P_0, P_1)$ induces a morphism $A(\varphi) : \bigoplus_i W_{x_i}^{\theta_+(x_i)} \to \bigoplus_j W_{y_j}^{\theta_-(y_j)}$ by precomposition. Moreover, if p is a path from x_i to y_j placed in the (q, r)-th component of $\mathrm{Hom}(P_0, P_1) \simeq \bigoplus_{i,j} P_{x_i}(y_j)^{\theta_-(y_j)\theta_+(x_i)}$, then the map it induces is precisely $A_{q,r}^{i,j,p}$. Therefore, for any subspace $U \subset \bigoplus_i W_{x_i}^{\theta_+(x_i)}$, $\mathcal{A}_{W,\theta}(U)$ is the same as the sum $\sum_\varphi A(\varphi)(U)$, taken over all $\varphi \in \mathrm{Hom}(P_0, P_1)$.

Similarly, $\mathrm{End}(P_1)$ (resp. $\mathrm{End}(P_0)$) induces a right action on $\bigoplus_i W_{x_i}^{\theta_+(x_i)}$ (resp. $\bigoplus_j W_{y_j}^{\theta_-(y_j)}$) by precomposition. We recall Remark A.1 that there is a unique minimal subspace that witnesses $\mathrm{disc}(\mathcal{A}_{W,\theta})$. Now we quote the

Lemma A.5 [Huszar 2021, Lemma 3.2]. *Let $U \subset \bigoplus_i W_{x_i}^{\theta_+(x_i)}$ be the minimal c-shrunk subspace of $\mathcal{A}_{W\theta}$. Then U (resp. $\sum_\varphi A(\varphi)(U)$) is a right $\mathrm{End}(P_1)$ (resp. $\mathrm{End}(P_0)$)-module. Namely, U (resp. $\sum_\varphi A(\varphi)(U)$) is closed under the right action of $\mathrm{End}(P_1)$ (resp. $\mathrm{End}(P_0)$) defined by precomposition.*

We now apply the lemma to establish the

Proposition A.6. *Let U be the minimal c-shrunk subspace of $\mathcal{A}_{W,\theta}$, and let $V = \sum_\varphi A(\varphi)(U)$. For each i, r (resp. j, q), let $\pi_{i,r} : \bigoplus_i W_{x_i}^{\theta_+(x_i)} \to W_{x_i}$ (resp. $\pi_{j,q} : \bigoplus_j W_{y_j}^{\theta_-(y_j)} \to W_{y_j}$) be the projection onto the (i, r)-th (resp. (j, q)-th) component. The following properties then hold:*

(1) $U = \bigoplus_i (W'_{x_i})^{\theta_+(x_i)}$, *where* $W'_{x_i} = \pi_{i,r}(U) \subseteq W_{x_i}$ *for any* $r \in I_i^+$.

(2) $V = \bigoplus_j (W'_{y_j})^{\theta_-(y_j)}$, *where* $W'_{y_j} = \pi_{j,q}(V) = \sum_{i,p \in \mathcal{P}_{i,j}} W(p)(W'_{x_i}) \subseteq W_{y_j}$ *for any* $q \in I_j^-$.

Proof. Let $u \in U$ and let p be a path from $x_{i'}$ to x_i in the (r', r)-th component of $\mathrm{End}(P_1) \simeq \bigoplus_{i,i'} (P_{x_{i'}})_{x_i}^{\theta_+(x_{i'})\theta_+(x_i)}$. The precomposition $u \circ p$ is then $W(p)(\pi_{i',r'}(u))$ placed in the r-th component

of $W(x_i)^{\theta_+(x_i)}$. Since U is a right $\text{End}(P_1)$-module, and the same path p can be placed in any (r', r)-th component of $\text{End}(P_1)$, the image $W(p)(\pi_{i'r'}(u))$ when placed in the r-th component of $W(x_i)^{\sigma_+(x_i)}$, is still in U for any $r' \in I_{i'}^+$ and any $r \in I_i^+$. In particular, taking $i' = i$, we see that for any $r_1, r_2 \in I_i^+$, $\pi_{i,r_1}(u)$ when placed at the r_2-th component of $W_{x_i}^{\theta_+(x_i)}$, is still in U. Basically, we obtain

$$\pi_{i,r_1}(U) = \pi_{i,r_2}(U) \text{ for any } r_1, r_2 \in I_i^+ \quad \text{and} \quad U = \bigoplus_i (\pi_{i,r}(U))^{\theta_+(x_i)} \text{ for any } r \in I_i^+.$$

This establishes property (1).

Using the fact that V is a right $\text{End}(P_0)$-module, we also get

$$V = \bigoplus_j (\pi_{j,q}(V))^{\theta_-(y_j)} \text{ for any } q \in I_j^-.$$

On the other hand, it is clear that

$$\sum_\varphi A(\varphi)(U) = \bigoplus_{j,q} \left(\sum_{i,p} \sum_r A_{q,r}^{i,j,p} \cdot \pi_{i,r}(U) \right)$$

$$= \bigoplus_{j,q} \left(\sum_{i,p} \sum_r W(p)\pi_{i,r}(U) \right)$$

$$= \bigoplus_j \left(\sum_{i,p \in \mathcal{P}_{i,j}} W(p)(W'_{x_i}) \right)^{\theta_-(y_j)}$$

This establishes property (2), finishing the proof of the proposition. $\qquad \square$

Theorem A.7. *Let $U \subset M(N, \mathbb{F})$ be the minimal c-shrunk subspace of $\mathcal{A}_{W,\theta}$. Let $\pi_{i,r} : \bigoplus_i W_{x_i}^{\sigma_+(x_i)} \to W_{x_i}$ be the projection onto the (i, r)-th component. For each $i \in [n]$, set*

$$W'_{x_i} = \pi_{i,r}(U).$$

For each $j \in [m]$, set

$$W'_{y_j} = \sum_{i,p \in \mathcal{P}_{i,j}} W(p)(W_{x_i}).$$

For each $y \in Q_0$ with $\theta(y) = 0$, set

$$W'_y = \sum_{\theta(x) \neq 0} \left(\sum_{p:x \to y} W(p)(W_x) \right).$$

Then $\{W'_x\}_{x \in Q_0}$ defined this way yields a subrepresentation W' with $\theta(W') = c$. In this case, we actually have

$$\text{disc}(W, \theta) = c.$$

Proof. We first verify that $\{W'_x\}_{x \in Q_0}$ forms a subrepresentation W' of W. We need to show that $W(a)(W'_{ta}) \subset W'_{ha}$ for all $a \in Q_1$. Due to the construction of W', it is more convenient to work with paths. We split the paths in Q into three major categories.

<u>Case 1</u>: The path starts at some $x_{i'}$. We consider three possibilities:

  <u>Case (1a)</u> The path ends at some x_i.

  <u>Case (1b)</u> The path ends at some y_j.

  <u>Case (1c)</u> The path ends at some y.

For case (1a), let $p : x_{i'} \to x_i$ be a path. We have seen in the proof of Proposition A.6 that $W(p)(W_{x_{i'}})$, when placed in any r-th component of $W_{x_i}^{\theta_+(x_i)}$, is still in U. By Proposition A.6(1), $W(p)(W'_{x_{i'}}) \subset W'_{x_i}$. This finishes case (1a). Case (1b) follows from the description of W'_{y_j} in Proposition A.6. Case (1c) follows from the construction of W'_y.

<u>Case 2</u>: The path starts at some $y_{j'}$. Again we consider three possibilities:

  <u>Case (2a)</u>: The path ends at some x_i.

  <u>Case (2b)</u>: The path ends at some y_j.

  <u>Case (2c)</u>: The path ends at some y.

For case (2a), if there are no paths from any $x_{i'}$ into y_j, then W'_{y_j} is the zero subspace, which is a trivial scenario. If there are paths from some $x_{i'}$ into y_j, the case is covered by case (1a). Case (2b) can be established using the fact that V is an $\text{End}(P_0)$-module like what we did for case (1a). Finally, case (2c) follows from the construction of W'_y as well.

<u>Case 3</u>: The path starts at some y'. We consider:

  <u>Case (3a)</u>: The path ends at some x_i.

  <u>Case (3b)</u>: The path ends at some y_j.

  <u>Case (3c)</u>: The path ends at some y.

If there are no paths from either $x_{i'}$ or $y_{j'}$ into y, then W'_y must be the zero subspace, which is a trivial scenario. If there is a path from some $x_{i'}$ or some $y_{j'}$ into y, case (3a) is covered by case (1a) and case (2a). Similarly case (3b) is covered by case (1b) and case (2b). Finally, case (3c) is covered by the construction of W'_y.

We have shown that $\{W'_x\}_{x \in Q_0}$ forms a subrepresentation W' of W. Next, note that

$$\theta(W') = \sum_i (\dim W'_{x_i}) \cdot \theta_+(x_i) - \sum_j (\dim W'_{y_j}) \cdot \theta_-(y_j) = \dim U - \dim \mathcal{A}_{W,\theta}(U)$$

by Proposition A.6. By Lemma A.4, it must be the case that $\theta(W') = \text{disc}(W, \theta) = \text{disc}(\mathcal{A}_{W,\theta})$. $\square$

We are now able to establish the algorithm of Theorem 3.5, which returns the discrepancy, together with a witness for any representation of an acyclic quiver over an infinite field.

Corollary A.8. *Let Q be an acylcic quiver and let W be a representation as above. There is a deterministic algorithm that finds the discrepancy* $\text{disc}(W, \theta)$, *together with a subrepresentation W' so that $\theta(W') =*

disc(W, θ). *If we set* $\Omega = \sum_{x \in Q_0} |\theta(x)|$, $K = \sum_{x \in Q_0} \dim W_x$, *and* P *as the number of paths in* Q, *then the algorithm has time complexity that is polynomial in* Ω, K, P.

Proof. We propose the following three steps to find the discrepancy of W, together with a witnessing subrepresentation. As step 1, we compute a basis for the matrix space

$$\mathcal{A}_{W,\theta} = \bigoplus_{i,j} \bigoplus_{q,r} \mathrm{Span}(A_{q,r}^{i,j,p} \mid p \in \mathcal{P}_{ij}),$$

where each $A_{q,r}^{i,j,p}$ is $W(p)$ placed in the (q, r)-block. If p has length l, then $W(p)$ involves $l - 1$ multiplications of matrices of sizes taken from a subset of $\{\dim W_x\}_{x \in Q_0}$. Iterating through paths in increasing lengths, computing $W(p)$ for all $p \in \mathcal{P}_{i,j}$ has time complexity polynomial in K, and linear in P. We then sort out linearly independent size $\dim W(y_j) \times \dim W(x_i)$ matrices among $\{W(p) \mid p \in \mathcal{P}_{i,j}\}$, which has time complexity polynomial in K, P. Finally, we put these linearly independent matrices into appropriate (q, r)-blocks for all $q \in I_j^+$, $r \in I_i^+$, and for all i, j. There are $M \times M'$ blocks so the time complexity to compute a basis for $\mathcal{A}_{W,\theta}$ is polynomial in Ω, K, P.

As step 2, apply the algorithm of Theorem A.2 to obtain a basis of the minimal c-shrunk subspace U of $\mathcal{A}_{W,\theta}$. The time complexity is polynomial in N, which in turn is bounded by a polynomial in $\Omega \cdot K$.

As step 3, we apply Theorem A.7 to derive a basis for W_x' for each $x \in Q_0$. A spanning set for each W_{x_i}' is obtained by truncating suitable parts of the basis of U we obtained earlier. We can then trim this spanning set down to a basis within time complexity polynomial in $\dim W_{x_i}$ and N, which is bounded by a polynomial in Ω, K. Since there are at most Ω many $x_i's$, the time complexity to obtain bases for all W_{x_i}' is bounded by a polynomial in Ω, K. For W_{y_j}', we can first compute $\sum_i \sum_{p \in P_{i,j}} W(p)(W_{x_i}')$. In this case, since each $W(p)$ is computed earlier, a spanning set of $W(p)(W_{x_i}')$ is obtained by multiplying $W(p)$ with the basis matrix of W_{x_i}'. Hence a spanning set for $W'(y_j)$ can be obtained within time complexity polynomial in K, P. Since there are at most $P \cdot \sum_i \dim W_{x_i}'$ vectors in the spanning set for each $W'(y_j)$, deriving a basis for W_{y_j}' has time complexity polynomial in K, P. Since the number of y_j is bounded by Ω, we may infer that computing bases for all W_{y_j}' has time complexity polynomial in Ω, K, P. Similarly, computing bases for all W_y' with $\theta(y) = 0$ has time complexity polynomial in Ω, K, P.

Since the three major steps are all within time complexity polynomial in Ω, K, P. The corollary is proved. $\square$

Acknowledgements

References

[Chindris and Kline 2021] C. Chindris and D. Kline, "Simultaneous robust subspace recovery and semi-stability of quiver representations", *J. Algebra* **577**:1 (2021), 210–236. MR Zbl

[Hille and de la Peña 2002] L. Hille and J. A. de la Peña, "Stable representations of quivers", *J. Pure Appl. Algebra* **172**:2-3 (2002), 205–224. MR Zbl

[Hoskins 2014] V. Hoskins, "Stratifications associated to reductive group actions on affine spaces", *Q. J. Math.* **65**:3 (2014), 1011–1047. MR Zbl

[Huszar 2021] A. Huszar, "Non-commutative rank and semi-stability of quiver representations", preprint, 2021. arXiv 2111.00039

[Ivanyos et al. 2015] G. Ivanyos, M. Karpinski, Y. Qiao, and M. Santha, "Generalized Wong sequences and their applications to Edmonds' problems", *J. Comput. System Sci.* **81**:7 (2015), 1373–1386. MR Zbl

[Ivanyos et al. 2018] G. Ivanyos, Y. Qiao, and K. V. Subrahmanyam, "Constructive non-commutative rank computation is in deterministic polynomial time", *Comput. Complexity* **27**:4 (2018), 561–593. MR Zbl

[Kempf 1978] G. R. Kempf, "Instability in invariant theory", *Ann. of Math.* (2) **108**:2 (1978), 299–316. MR Zbl

[King 1994] A. D. King, "Moduli of representations of finite-dimensional algebras", *Quart. J. Math. Oxford Ser.* (2) **45**:4 (= 180) (1994), 515–530. MR Zbl

[Mumford et al. 1994] D. Mumford, J. Fogarty, and F. Kirwan, *Geometric invariant theory*, 3rd ed., Ergebnisse der Math. (2) **34**, Springer, 1994. MR Zbl

[Zamora 2014] A. Zamora, "On the Harder–Narasimhan filtration for finite dimensional representations of quivers", *Geom. Dedicata* **170** (2014), 185–194. MR Zbl

ccheng@citadel.edu   *University of Missouri, Columbia, MO, United States*

Algebra & Number Theory

msp.org/ant

See inside back cover or msp.org/ant for submission instructions.

Algebra & Number Theory

Volume 18   No. 2   2024