ALGEBRAIC
COMBINATORICS

Eric Ramos, David Speyer & Graham White

FI–sets with relations

Volume 3, issue 5 (2020), p. 1079-1098.

<http://alco.centre-mersenne.org/item/ALCO_2020__3_5_1079_0>

Algebraic Combinatorics
Volume 3, issue 5 (2020), p. 1079–1098
https://doi.org/10.5802/alco.128

FI–sets with relations

Eric Ramos, David Speyer & Graham White

Abstract Let FI denote the category whose objects are the sets $[n] = \{1,\ldots,n\}$, and whose morphisms are injections. We study functors from the category FI into the category of finite sets. We write $\mathfrak{S}_n$ for the symmetric group on $[n]$. Our first main result is that, if the functor $[n] \mapsto X_n$ is "finitely generated" there is a finite sequence of integers m_i and a finite sequence of subgroups H_i of $\mathfrak{S}_{m_i}$ such that, for n sufficiently large, $X_n \cong \bigsqcup_i \mathfrak{S}_n/(H_i \times \mathfrak{S}_{n-m_i})$ as a set with $\mathfrak{S}_n$ action. Our second main result is that, if $[n] \mapsto X_n$ and $[n] \mapsto Y_n$ are two such finitely generated functors and $R_n \subset X_n \times Y_n$ is an FI–invariant family of relations, then the $(0,1)$ matrices encoding the relation R_n, when written in an appropriate basis, vary polynomially with n. In particular, if R_n is an FI–invariant family of relations from X_n to itself, then the eigenvalues of this matrix are algebraic functions of n. As an application of this theorem we provide a proof of a result about eigenvalues of adjacency matrices claimed by the first and last author. This result recovers, for instance, that the adjacency matrices of the Kneser graphs have eigenvalues which are algebraic functions of n, while also expanding this result to a larger family of graphs.

1. Introduction

We begin with a specific example of the sort of phenomenon we seek to explain. The Kneser graph $KG(n,r)$ has as vertices the r–element subsets of n and has an edge between two vertices if and only if the corresponding subsets are disjoint. Its adjacency matrix is computed in [8, Section 9.4] to have eigenvalues

$$\lambda_i := (-1)^i \binom{n-r-i}{r-i} \text{ for } 0 \leqslant i \leqslant r.$$

Moreover, each eigenvalue λ_i appears with multiplicity

$$\binom{n}{i} - \binom{n}{i-1}.$$

Therefore for each fixed r, we observe the following phenomena

- The total number of distinct eigenvalues is eventually independent of n. Specifically, there are eventually exactly $r+1$ such eigenvalues.
- The eigenvalues each agree with a function which is algebraic over the field $\mathbb{Q}(n)$. Specifically, these are the functions $(-1)^i\binom{n-r-i}{r-i}$ for $0 \leqslant i \leqslant r$.

Manuscript received 23rd May 2018, revised 10th February 2020 and 10th February 2020, accepted 6th May 2020.

Keywords. FI-modules, Representation Stability, Kneser graphs.

Acknowledgements. The first author was supported by NSF grant DMS-1704811. The second author was supported by NSF grant DMS-1600223.

ISSN: 2589-5486 http://algebraic-combinatorics.org/

- The multiplicity of each eigenvalue agrees with a polynomial in n. Specifically, these are the polynomials $\binom{n}{i} - \binom{n}{i-1}$.

Similar phenomena can be observed in the spectra of the adjacency matrices of Johnson graphs [1], as well as a variety of other examples (see Section 3.3). The main goal of this paper is to provide a uniform framework under which one can deduce the existence of these behaviors. We achieve this using the techniques of representation stability theory and related fields, as appearing in the works of Church, Ellenberg, Farb, Nagpal, Putman, Sam, Snowden, and many others [2, 3, 4, 10, 15].

Let FI denote the category whose objects are the sets $[n] = \{1, \ldots, n\}$, including the empty set $[0] = \varnothing$, and whose morphisms are injections. For any commutative ring k, an FI*–module* is a functor from FI to the category of k–modules. *In this paper, k will be a field of characteristic zero, which is henceforth fixed.* FI–modules were introduced by Church, Ellenberg, and Farb as a single unifying framework for a large collection of seemingly unrelated phenomena from topology, representation theory, and a variety of other subjects [2].

There has been a recent push in the literature to apply the theory of FI–modules to more traditionally combinatorial fields. In [7], Gadish introduces a theory of FI–posets, functors from FI to the category of posets. He then applied this framework to prove non–trivial facts about linear subspace arrangements. In [14], the first and third authors consider FI–graphs, functors from FI to the category of graphs. It is proven in that work that such families of graphs display a variety of asymptotic regularities in their enumerative, topological, and algebraic properties. This was expanded in the follow-up works [12, 13]

What is notable about the works mentioned in the previous paragraph is that they follow a common theme: both FI–posets and FI–graphs can be thought of as a pair of an FI–set with a relation. An FI*–set* is a functor $X_\bullet$ from FI to the category of finite sets. For an FI–set $X_\bullet$ we write X_n to denote its evaluation at n, while we use *transition map* to mean one of the maps $X_m \to X_n$ induced by the FI–structure. The *product* of two FI–sets $X_\bullet$ and $Y_\bullet$ is the FI–set $(X \times Y)_\bullet$ with $(X \times Y)_n = X_n \times Y_n$ and the obvious transition maps. A *relation* between FI–sets $X_\bullet$ and $Y_\bullet$ is any FI–set $R_\bullet$ with a natural inclusion into $(X \times Y)_\bullet$. In the case of graphs this relation is the edge relation, while in the case of posets it is the partial ordering. The purpose of this work is to use the language of FI–sets and relations to unify these topics, while also expanding our understanding of both.

We say that an FI–set $X_\bullet$ is *finitely generated in degree* $\leqslant d$ if, for all $n \geqslant d$, the elements of X_{n+1} are all in the image of some transition map from X_d. We will say that an FI–set is *finitely generated* if it is finitely generated in some degree. For instance, the assignment $X_n := [n]$ is finitely generated, while the assignment $X_n := 2^{[n]}$ is not.

Our first goal will be to prove the following structure theorem for FI–sets. For the remainder of this paper we write $\mathfrak{S}_n$ for the symmetric group on n letters.

Theorem 1.1. *Let $X_\bullet$ denote an* FI*–set finitely generated in degree $\leqslant d$. Then there exists a finite collection of integers $m_i \leqslant d$, and subgroups $H_i \subseteq \mathfrak{S}_{m_i}$, such that, for n sufficiently large, we have an isomorphism*

$$X_n \cong \bigsqcup_i \mathfrak{S}_n/(H_i \times \mathfrak{S}_{n-m_i})$$

as sets with an action of $\mathfrak{S}_n$.

Remark 1.2. The isomorphisms of Theorem 1.1 are natural, in the following sense. If $n \gg 0$ is big enough so that the result of Theorem 1.1 holds for the FI–set $X_\bullet$, and $f : [n] \to [m]$ is an injection of sets, then the transition map induced by f on $X_\bullet$ is given as follows. If π is an element of $\mathfrak{S}_n$, then the image of the coset $\pi \cdot (H_i \times \mathfrak{S}_{n-m_i})$ under the transition map induced by f is given by the coset $\pi' \cdot (H_i \times \mathfrak{S}_{m-m_i})$ defined by

$$\pi'(x) := \begin{cases} f \circ \pi \circ f^{-1}(x) & \text{if } x \text{ is in the image of } f \\ x & \text{otherwise.} \end{cases}$$

This fact will follow from the proof of Theorem 1.1, but will not be used in the paper.

Example 1.3. Let X_n be the set of ordered m–tuples of distinct elements of $[n]$, with transition maps defined coordinate–wise. Then $X_n \cong \mathfrak{S}_n/(\{e\} \times \mathfrak{S}_{n-m})$ as a set with $\mathfrak{S}_n$ action, so there is one term in the disjoint union, with $m_i = m$ and $H_i = \{e\}$.

Example 1.4. Let X_n be the set of unordered m–tuples of distinct elements of $[n]$, with the obvious induced maps. Then $X_n \cong \mathfrak{S}_n/(\mathfrak{S}_m \times \mathfrak{S}_{n-m})$ as a set with $\mathfrak{S}_n$ action, so there is one term in the disjoint union, with $m_i = m$ and $H_i = \mathfrak{S}_m$.

One can interpolate between Examples 1.3 and 1.4 by choosing intermediate subgroups of $\mathfrak{S}_m$, and can also take disjoint unions of this construction for different choices of m. Theorem 1.1 states that, once n is sufficiently large, all finitely generated FI–sets are built from these operations.

As previously stated, FI–sets become richer and more interesting when paired with a relation. Having proven the aformentioned structure theorem for finitely generated FI–sets, we turn our attention to properties of relations associated to these FI–sets.

The *linearization* $kX_\bullet$ is the FI–module where $(kX)_n$ is the free k–module on X_n and the transition maps are defined in the obvious way. For $x \in X_n$, we'll write e_x for the corresponding basis element of kX_n. Let $R_\bullet$ be a relation between $X_\bullet$ and $Y_\bullet$, meaning an FI–subset of $X_\bullet \times Y_\bullet$. Then R induces a sequence of maps

$$r_n : kX_n \to kY_n$$

by

$$r_n(e_x) = \sum_{(x,y) \in R_n} e_y.$$

The maps r_n commute with the $\mathfrak{S}_n$ action, but they do *not* form a map of FI–modules.

Our Main Theorem 1.5, roughly stated, says that for any finitely generated FI–sets $X_\bullet$ and $Y_\bullet$ and any relation $R_\bullet$ between them, the corresponding linear maps $r_n : kX_n \to kY_n$ are given by a matrix whose entries depend polynomially on n. e.g. Let $Z_\bullet$ be the FI–set where $Z_n = [n]$, with the obvious transition maps. We can linearize this to give an FI–module $kZ_\bullet$ where $(kZ_\bullet)_n = k^n$. We write $\{e_i\}_{i \in [n]}$ for the usual basis of $(kZ_\bullet)_n$. As an S_n–module, $k^n = \mathrm{Sp}(n) \oplus \mathrm{Sp}(n-1,1)$, where $\mathrm{Sp}(\lambda)$ is the Specht module. The regularity of this isotypic decomposition as n grows is an example of what is known as representation stability.

We have a relation, $R_\bullet$, on $Z_\bullet$, defined by $R_n := \{(i,j) \in [n]^2 \mid i \neq j\}$. We can linearize these relations to give maps $e_i \mapsto \sum_{j \in [n],\ j \neq i} e_j$ from $kZ_n \to kZ_n$. These maps do *not* give a map of FI–modules so the existing theory of FI–modules does not let us study them. However, for each n, this map is a map of symmetric group representations so, by Schur's lemma, it acts by scalars on $\mathrm{Sp}(n)$ and $\mathrm{Sp}(n-1,1)$. Explicitly, these scalars are $n-1$ on $\mathrm{Sp}(n)$ and -1 on $\mathrm{Sp}(n-1,1)$. We want to prove that this sort of simple algebraic dependence on n is what happens in general.

Our first goal, therefore, is to explain in what sense a family of maps $kX_n \to kX_n$, between different vector spaces of different sizes, can be given by a fixed matrix. We describe the relevant definitions briskly here; see Section 2.2 for the full details.

Given a positive integer n, a *partition* of n is a tuple of positive integers $\lambda = (\lambda_1, \ldots, \lambda_r)$ such that $\lambda_j \geqslant \lambda_{j+1}$ and $\sum_j \lambda_j = n$. The irreducible representations of

$\mathfrak{S}_n$ are in bijection with partitions of n in a standard manner, and we write $\mathrm{Sp}(\lambda)$ for the irreducible representation (over k) corresponding to the partition λ. We recall that $\mathrm{Hom}_{\mathfrak{S}_n}(\mathrm{Sp}(\lambda), \mathrm{Sp}(\lambda)) \cong k$. If W is a representation of $\mathfrak{S}_n$, then we write $W_\lambda := \mathrm{Hom}_{\mathfrak{S}_n}(\mathrm{Sp}(\lambda), W)$. So $W \mapsto W_\lambda$ is a functor from $\mathfrak{S}_n$–representations to vector spaces and we have a canonical isomorphism $W \cong \bigoplus_{|\lambda|=n} W_\lambda \otimes \mathrm{Sp}(\lambda)$. The summand $W_\lambda \otimes \mathrm{Sp}(\lambda)$ is called the *λ–isotypic component* of W. We will write α_λ for the injection $W_\lambda \otimes \mathrm{Sp}(\lambda) \to W$ and β_λ for the surjection $W \to W_\lambda \otimes \mathrm{Sp}(\lambda)$.

Our next task is to define an analogue of the λ–isotypic component for an FI–module $M_\bullet$. If $\lambda = (\lambda_1, \ldots, \lambda_r)$ is any partition of m, and $n \geqslant m + \lambda_1$, then we set $\lambda[n] = (n - m, \lambda_1, \ldots, \lambda_r)$. Let $m + \lambda_1 \leqslant p \leqslant q$. The inclusion of $[p] \hookrightarrow [q]$ induces an $\mathfrak{S}_p$–equivariant homomorphism from $M_p \to M_q$, and thus an $\mathfrak{S}_q$–equivariant homomorphism $\mathrm{Ind}_{\mathfrak{S}_p}^{\mathfrak{S}_q} M_p \to M_q$. We can take the $\lambda[q]$ isotypic component of this map, giving a map $\left(\mathrm{Ind}_{\mathfrak{S}_p}^{\mathfrak{S}_q} M_p\right)_{\lambda[q]} \to (M_q)_{\lambda[q]}$, and there is a natural inclusion $(M_p)_{\lambda[p]} \to \left(\mathrm{Ind}_{\mathfrak{S}_p}^{\mathfrak{S}_q} M_p\right)_{\lambda[q]}$. So for sufficiently large $p \leqslant q$, we obtain maps $(M_p)_{\lambda[p]} \to (M_q)_{\lambda[q]}$. We show (Theorem 2.8) that this map is an isomorphism for p large enough. So we can define $(M_\bullet)_\lambda$ to be the inductive limit $\lim_{p\to\infty} (M_p)_{\lambda[p]}$. There are many details to be checked here; we do so in Section 2.2.

Let $X_\bullet$ and $Y_\bullet$ denote finitely generated FI–sets and k a characteristic zero field. Let $R_\bullet$ denote a relation between them with associated maps $r_n : kX_n \to kY_n$. Since r_n is a map of $\mathfrak{S}_n$–modules, by Schur's lemma, it induces maps $r_{n,\lambda} : (kX_n)_{\lambda[n]} \to (kY_n)_{\lambda[n]}$. For n sufficiently large, we have isomorphisms $(kX_n)_{\lambda[n]} \cong (kX_\bullet)_\lambda$ and $(kY_n)_{\lambda[n]} \cong (kY_\bullet)_\lambda$, so we obtain a family of maps $r_{n,\lambda} : (kX_\bullet)_\lambda \longrightarrow (kY_\bullet)_\lambda$. In other words, the $r_{n,\lambda}$ are a family of linear transformations, varying in n, between two fixed vector spaces, unvarying in n.

Theorem 1.5. *In any bases for the vector spaces $(kX_\bullet)_\lambda$ and $(kY_\bullet)_\lambda$, the entries of $r_{n,\lambda}$ depend polynomially on n.*

This theorem has a number of more concrete consequences in the case where $X_\bullet = Y_\bullet$.

Corollary 1.6. *Let $X_\bullet$ denote a finitely generated* FI*–set, and let $R_\bullet$ denote a self-relation with associated maps $r_n : kX_n \to kX_n$. Then for $n \gg 0$:*

1. *the number N of distinct eigenvalues of r_n is unchanging in n;*
2. *there exists a finite list of functions f_1, ..., f_N, each algebraic over the field $\mathbb{Q}(n)$, such that the complete list of distinct eigenvalues of r_n is given by $f_1(n)$, ..., $f_N(n)$;*
3. *for each i, the function*

$$n \mapsto \textit{the algebraic multiplicity of } f_i(n) \textit{ as an eigenvalue of } r_n$$

agrees with a polynomial in n.

Remark 1.7. In [14, Theorem H] a version of the above theorem is claimed in the case of FI–graphs and the edge relation. In that work it is said that the theorem would be proven in this paper. In the final section of this work we will explain why [14, Theorem H] follows from Corollary 1.6.

A natural followup question related to the conclusions of Corollary 1.6 is how one can leverage these statements about eigenvalues to say something about the statistics of random walks being performed on FI–sets which have been paired with a transition relation. This is the topic of [13].

2. FI–modules and representation stability

2.1. FI–modules. The present work is largely concerned with structures we refer to as FI–sets. One of the primary tools we will use to study these objects, as well as one of the main motivations for considering FI–sets in the first place, are FI–modules.

FI–modules were introduced by Church, Ellenberg, and Farb in their seminal work [2]. It was later discovered that this concept arose in a variety of different, sometimes older, contexts such as the twisted commutative algebras of Sam and Snowden [15] and the study of polynomial functors (see [5, 6] for modern treatments). In this work we will follow the exposition of Church, Ellenberg, and Farb.

Definition 2.1. *Let* FI *denote the category whose objects are the sets* $[n] := \{1, \dots, n\}$ *and whose maps are injections. An* FI–module over k *is a (covariant) functor* $V_\bullet$ *from* FI *to the category of* k*–modules. In this paper,* k *is always a field of characteristic zero. We will often write* $V_n := V_\bullet([n])$ *for the* degree n piece of $V_\bullet$*, and* $f_* := V_\bullet(f)$*.*

Since $\mathrm{Hom}_{\mathrm{FI}}([n], [n])$ is the symmetric group $\mathfrak{S}_n$, each vector space V_n is an $\mathfrak{S}_n$–representation.

Our next objective will be to specialize to those objects which are finitely generated in the appropriate sense. One should note that the category of FI–modules and natural transformations is abelian, with abelian operations defined point–wise.

Definition 2.2. *A* submodule *of an* FI*–module* $V_\bullet$ *is an* FI*–module* $W_\bullet$*, such that* W_n *is a submodule of* V_n *for all* n*, and for all injections* $f : [n] \hookrightarrow [m]$ *one has* $W_\bullet(f) = V_\bullet(f)|_{W_n}$*. An* FI*–module* $V_\bullet$ *is said to be* finitely generated in degree $\leqslant d$ *if there is a finite subset of the disjoint union* $\bigsqcup_{n=0}^{d} V_n$ *which is not contained in any proper submodule of* $V_\bullet$*.*

We will now define the analogue of free modules for the FI–setting.

Definition 2.3. *Let* W *be a (left)* $k[\mathfrak{S}_m]$*–module. Then the* free FI–module on W*, or the* induced FI–module on W *is the* FI*–module* $M(W)_\bullet$ *defined by the assignments*

$$M(W)_n := k[\mathrm{Hom}_{\mathrm{FI}}([m], [n])] \otimes_{\mathfrak{S}_m} W,$$

where $k[\mathrm{Hom}_{\mathrm{FI}}([m], [n])]$ *is the free* k*–module with basis indexed by* $\mathrm{Hom}_{\mathrm{FI}}([m], [n])$*, viewed as a right* $k[\mathfrak{S}_m]$*–module in the obvious way. For any injection of sets* f*, the induced map* $M(W)(f)$ *is defined on pure tensors by post composition on the left tensor factor.*

We make the abbreviations $M(m)_\bullet = M(k\mathfrak{S}_m)_\bullet$ (here $k\mathfrak{S}_m$ is the regular representation of $\mathfrak{S}_m$) and $M(\lambda)_\bullet = M(\mathrm{Sp}(\lambda))_\bullet$.

Remark 2.4. It is an easily verifiable fact that there are isomorphisms

$$\mathrm{Hom}_{\mathrm{FI-mod}}(M(W)_\bullet, V_\bullet) \cong \mathrm{Hom}_{\mathfrak{S}_m}(W, V_m).$$

In particular, a map from $M(i)_\bullet$ into $V_\bullet$ is equivalent to choosing an element of V_i, and

$$\mathrm{Hom}_{\mathrm{FI-mod}}(M(\lambda)_\bullet, V_\bullet) \cong (V_{|\lambda|})_\lambda.$$

This uniquely defines $M(W)_\bullet$ via Yoneda's lemma.

Put another way, the functor $W \mapsto M(W)$ from $\mathfrak{S}_m$–representations to FI–modules is a left adjoint to the forgetful functor $V_\bullet \mapsto V_m$, hence the terminology of free.

Remark 2.5. Because k is a field of characteristic zero, the FI–module $M(W)_\bullet$ will always be projective. If we were to consider a general commutative ring k, then $M(W)_\bullet$ is projective if and only if W is a projective $k[\mathfrak{S}_n]$–module (See [11], for instance). Group algebras of finite groups over characteristic zero fields are semi–simple, so this is automatic in our setting.

An FI–module $V_\bullet$ is *finitely generated* if there exists a collection of non–integers, $(d_i)_{i\geqslant 0}$, all but finitely many equaling zero, and a surjection

$$\bigoplus_{i\geqslant 0} M(i)_\bullet^{d_i} \to V_\bullet \to 0$$

The following was first proven by Snowden in [16] when k was a field of characteristic 0, and later expanded to more general k by Church, Ellenberg, Farb and Nagpal in [2, 3].

Theorem 2.6 (Snowden, [16]; Church, Ellenberg, Farb, and Nagpal [2, 3]). *Let $V_\bullet$ be a finitely generated* FI*–module over a Noetherian ring k. Then every submodule of $V_\bullet$ is also finitely generated.*

The above Noetherian property is arguably the most powerful tool that one has access to when studying finitely generated FI–modules. We will see that it grants us surprisingly brief, although admittedly non–constructive, proofs of certain combinatorial facts. This philosophy can be seen in the context of FI–graphs and FI–posets in [14] and [7], respectively.

We take the opportunity to quote:

Proposition 2.7 (Church, Ellenberg, and Farb, [2]). *If $V_\bullet$ is finitely generated in degree $\leqslant d$ and $W_\bullet$ is finitely generated in degree $\leqslant e$, then*

(1) *$V_\bullet \oplus W_\bullet$ is generated in degree $\leqslant \max\{d, e\}$, where the directed sum is defined point–wise;*
(2) *$V_\bullet \otimes W_\bullet$ is generated in degree $\leqslant d + e$, where the tensor product is defined point–wise.*

2.2. Representation stability. We recall the notation from the introduction: $\mathrm{Sp}(\lambda)$ is the Specht module. For an $\mathfrak{S}_n$–representation W and a partition λ of n, we put $W_\lambda = \mathrm{Hom}_{\mathfrak{S}_n}(\mathrm{Sp}(\lambda), W)$. For a partition $\lambda = (\lambda_1, \dots, \lambda_r)$ and an integer n with $n - |\lambda| \geqslant \lambda_1$, we put $\lambda[n] = (n - |\lambda|, \lambda_1, \lambda_2, \dots, \lambda_r)$. When $n - |\lambda| < \lambda_1$, we put $\mathrm{Sp}(\lambda[n]) = 0$.

We also recall that $M(\lambda)$ is shorthand for $M(\mathrm{Sp}(\lambda))$.

Let $V_\bullet$ be a finitely generated FI–module. Church, Ellenberg, and Farb proved that $V_\bullet$ exhibits Representation Stability [2, Theorem 1.13]. This result has three parts, the one of which that is usually cited is Stability of Multiplicities, which states the following: There is a positive integer N and a sequence of nonnegative integers m_λ indexed by partitions, such that all but finitely many m_λ are 0, and

$$\dim(V_n)_{\lambda[n]} = m_\lambda \text{ for } n \geqslant N.$$

We will need a more precise statement which, as we will explain, is also part of Church, Ellenberg, and Farb's result.

Let λ be a partition, and let n be large enough that $\lambda[n]$ is defined. So by Remark 2.4, we have an isomorphism

$$\mathrm{Hom}_{\mathrm{FI-mod}}(M(\lambda[n])_\bullet, V_\bullet) \cong \mathrm{Hom}_{\mathfrak{S}_m}(\mathrm{Sp}(\lambda[n]), V_n) \cong (V_n)_{\lambda[n]}.$$

By the Pieri rule, there is a unique copy of $\mathrm{Sp}(\lambda[n+1])$ inside[(1)] $M(\lambda[n])_{n+1} = \mathrm{Ind}_n^{n+1} \mathrm{Sp}(\lambda[n])$ so by Remark 2.4 there is a unique up to scalar multiple nonzero map $(f_\bullet) : M(\lambda[n+1])_\bullet \to M(\lambda[n])_\bullet$. This induces a map

$$\mathrm{Hom}_{\mathrm{FI-mod}}(M(\lambda[n])_\bullet, V_\bullet) \to \mathrm{Hom}_{\mathrm{FI-mod}}(M(\lambda[n+1])_\bullet, V_\bullet).$$

[(1)] We will often abbreviate induction from $\mathfrak{S}_m$ to $\mathfrak{S}_n$ by Ind_m^n, and similarly write Res_m^n for restriction.

Church, Ellenberg, and Farb's result states that, for n sufficiently large, the vector spaces on the two sides of this map are of the same dimension. We require

Theorem 2.8. *For n sufficiently large, the map* $\mathrm{Hom}_{\mathrm{FI-mod}}(M(\lambda[n])_\bullet, V_\bullet) \to \mathrm{Hom}_{\mathrm{FI-mod}}(M(\lambda[n+1])_\bullet, V_\bullet)$ *is an isomorphism.*

Remark 2.9. This result, and our method of proof, is very similar to ideas from Sam and Snowden [15], particularly Section 2.2. We base our argument on Church, Ellenberg, and Farb [2] rather than Sam and Snowden in order to follow our general pattern of using the former's terminology, and because their paper is slightly earlier.

Proof. Let λ and μ be partitions with $|\lambda| > |\mu|$ and let $m \leqslant n$ with m and n large enough that $\lambda[m]$ and $\mu[n]$ are defined. So $(V_m)_{\lambda[m]} \otimes \mathrm{Sp}(\lambda[m])$ is the $\lambda[m]$–isotypic component of V_m, and likewise for $\mu[n]$. We claim that, for any transition map $V_m \to V_n$, the composition

$$(V_m)_{\lambda[m]} \otimes \mathrm{Sp}(\lambda[m]) \xrightarrow{\alpha_{\lambda[m]}} V_m \to V_n \xrightarrow{\beta_{\mu[n]}} (V_n)_{\mu[n]} \otimes \mathrm{Sp}(\mu[n])$$

is 0. To see this, note that this map must be $\mathfrak{S}_m$–equivariant, where we restrict the right hand side to a suitable $\mathfrak{S}_m \subset \mathfrak{S}_n$. But, by the Pieri rule, $\mathrm{Sp}(\lambda[m])$ does not occur in $(\mathrm{Sp}(\mu[n]))|_{\mathfrak{S}_m}$. The same argument also shows that, if $|\lambda| = |\mu|$ and $\lambda \neq \mu$, then the composite map $(V_m)_{\lambda[m]} \otimes \mathrm{Sp}(\lambda[m]) \longrightarrow (V_n)_{\mu[n]} \otimes \mathrm{Sp}(\mu[n])$ is 0.

Choose any total ordering of the set of partitions such that $|\lambda| < |\mu|$ implies $\lambda < \mu$. Define

$$V_n^{\geqslant\lambda} = \bigoplus_{\mu \geqslant \lambda} (V_n)_{\mu[n]} \otimes \mathrm{Sp}(\mu[n]) \subseteq V_n.$$

Then the above argument shows that the transition maps carry $V_m^{\geqslant\lambda}$ to $V_n^{\geqslant\lambda}$ for any $m \leqslant n$. So the $V_m^{\geqslant\lambda}$ form a submodule of $V_\bullet$, which we denote $V_\bullet^{\geqslant\lambda}$.

By Remark 2.4, we have isomorphisms

$$\mathrm{Hom}_{\mathrm{FI-mod}}(M(\lambda[n])_\bullet, V_\bullet) \cong (V_n)_{\lambda[n]} \cong \mathrm{Hom}_{\mathrm{FI-mod}}(M(\lambda[n])_\bullet, V_\bullet^{\geqslant\lambda}).$$

Also, since $V_\bullet^{\geqslant\lambda}$ is a submodule of the finitely generated FI–module $V_\bullet$, it is finitely generated itself. So we may, and do, replace $V_\bullet$ by $V_\bullet^{\geqslant\lambda}$. As a result, we may and do assume that $(V_n)_{\mu[n]} = 0$ for $\mu < \lambda$.

Now, suppose that $\mathrm{Hom}_{\mathrm{FI-mod}}(M(\lambda[n])_\bullet, V_\bullet) \to \mathrm{Hom}_{\mathrm{FI-mod}}(M(\lambda[n+1])_\bullet, V_\bullet)$ is not an isomorphism. Since both of these are vector spaces of dimension m_λ, this means that the map is not surjective. So there is some $U \subsetneq (V_{n+1})_{\lambda[n+1]}$ in which the image of $\mathrm{Hom}_{\mathrm{FI-mod}}(M(\lambda[n])_\bullet, V_\bullet)$ lies. Tracing through our isomorphisms, this means that all of our transition maps $(V_n)_{\lambda[n]} \to V_{n+1}$, followed by projection onto the $\lambda[n+1]$ isotypic component, land in $U \otimes \mathrm{Sp}(\lambda[n+1])$. Also, for every $\mu \neq \lambda$, the transition maps $(V_n)_{\mu[n]} \to V_{n+1}$ project to 0 in the $\lambda[n+1]$ isotypic component by our observations in the first paragraph of the proof. But [2, Proposition 3.3.3] states that the images of the transition maps from V_n span V_{n+1}. This contradiction establishes that $\mathrm{Hom}_{\mathrm{FI-mod}}(M(\lambda[n])_\bullet, V_\bullet) \to \mathrm{Hom}_{\mathrm{FI-mod}}(M(\lambda[n+1])_\bullet, V_\bullet)$ is an isomorphism after all. ☐

We define $(V_\bullet)_\lambda$ to be the inductive limit $\lim_{n\to\infty} \mathrm{Hom}_{\mathrm{FI-mod}}(M(\lambda[n])_\bullet, V_\bullet)$. So, by Theorem 2.8, $(V_\bullet)_\lambda \cong (V_n)_{\lambda[n]}$ for all sufficiently large n.

Fix a partition λ. Let $V_\bullet$ and $W_\bullet$ be two FI modules and let $r_n : V_n \to W_n$ be a sequence of maps of $\mathfrak{S}_n$–representations. For n large enough that $\lambda[n]$ is defined, r_n induces a map of vector spaces $(V_n)_{\lambda[n]} \to (W_n)_{\lambda[n]}$. For n sufficiently large, we have canonical isomorphisms $(V_n)_{\lambda[n]} \cong (V_\bullet)_\lambda$ and $(W_n)_{\lambda[n]} \cong (W_\bullet)_\lambda$, we obtain a composite map $(V_\bullet)_\lambda \to (W_\bullet)_\lambda$. This is the map we denote $r_{n,\lambda}$, which appears in Theorem 1.5.

The inclusion $M(\lambda[n+1]) \to M(\lambda[n])$ is only unique up to multiplication by a scalar. We fix choices of these scalars once and for all for the purpose of defining $(V_\bullet)_\lambda$. If we changed to other scalars, there would be a canonical isomorphism between the old and the new $(V_\bullet)_\lambda$, and the maps $r_{n,\lambda}$ above would be unchanged.

3. FI–sets

3.1. Elementary definitions and properties. In this section we define the fundamental object of study for this paper: FI–sets.

Definition 3.1. *An* FI*–set is a (covariant) functor from* FI *to the category of sets. We will usually denote* FI*–sets by $X_\bullet$ or $Y_\bullet$, where evaluation at $[n]$ is written X_n. If $X_\bullet$ is an* FI*–set, and $f : [n] \hookrightarrow [m]$ is an injection, we will generally write f_* to denote the induced map, or $X(f)$ if X is not clear from context. For any non–negative integers $n < m$, we will write $\iota_{nm} : [n] \hookrightarrow [m]$ for the standard inclusion. The* category of FI–sets *is that whose objects are* FI*–sets, and whose morphisms are natural transformations.*

We say that $Y_\bullet$ is a subset *of $X_\bullet$ if $Y_n \subseteq X_n$ for all n, and for any $f : [n] \hookrightarrow [m]$ one has $X(f)|_{Y_n} = Y(f)$. We say that $X_\bullet$ is* torsion–free *if $X(f)$ is injective for all choices of f.*

To the knowledge of the authors, this is the first paper which has formally considered FI–sets. That being said, related structures have been studied in the past. For instance, the first and third authors considered FI–graphs in [14], while Gadish studied FI–posets in [7].

Just as with FI–modules, we will begin by defining finite generation for FI–sets.

Definition 3.2. *We say that an* FI*–set $X_\bullet$ is* finitely generated in degree $\geqslant d$ *if there is a finite subset of $\bigsqcup_{i=0}^{d} X_i$ which is not contained in any proper* FI*–subset of $X_\bullet$. We say that $X_\bullet$ is finitely generated if it is finitely generated in some degree.*

Note, if $X_\bullet$ is finitely generated, then all the X_n are finite.

Definition 3.3. *For any* FI*–set $X_\bullet$, the* linearization $kX_\bullet$ *is the* FI*–module where $(kX_\bullet)_n$ is the free k–vector space on the set X_n, with the obvious maps.*

Linearization is a functor from FI–sets to FI–modules. We see that $X_\bullet$ is finitely generated if and only if $kX_\bullet$ is (and in the same degree). This yields some immediate consequences:

Proposition 3.4. *Let $X_\bullet$ be a finitely generated* FI*–set. For n sufficiently large, and $f : [n] \hookrightarrow [n+1]$ any injection, the map $f_* : X_n \to X_{n+1}$ is injective.*

Proof. The map $f_* : X_n \to X_{n+1}$ is injective if and only if the linearization $f_* : kX_n \to kX_{n+1}$ is, and the latter is true for $n \gg 0$ by the representation stability theorem of [2, Theorem 1.13]. ☐

Proposition 3.5. *Let $X_\bullet$ be a finitely generated* FI*–set. For n sufficiently large, the number of orbits of $\mathfrak{S}_n$ on X_n is a constant independent of n.*

Proof. The number of orbits of $\mathfrak{S}_n$ on X_n is the multiplicity of $\mathrm{Sp}(n)$ in kX_n. Since $\mathrm{Sp}(n) = \mathrm{Sp}(\varnothing[n])$, this is eventually constant by [2, Theorem 1.13]. ☐

For any fixed $n \geqslant 0$, every inclusion $f : [n] \to [n+1]$ induces a map on the orbit sets $X_n/\mathfrak{S}_n \to X_{n+1}/\mathfrak{S}_{n+1}$, and all of these maps are the same.

Proposition 3.6. *Let $X_\bullet$ be a finitely generated* FI*–set. For n sufficiently large, the map $X_n/\mathfrak{S}_n \to X_{n+1}/\mathfrak{S}_{n+1}$ described above is bijective.*

Proof. By Proposition 3.5, for n large enough, $|X_n/\mathfrak{S}_n| = |X_{n+1}/\mathfrak{S}_{n+1}|$. So, it is enough to show that the map is surjective for n sufficiently large. Suppose that $u \in X_{n+1}$ is such that the orbit $\mathfrak{S}_{n+1}u$ is not in the image of this map. Finite generation implies that $\cup_{\sigma\in\mathfrak{S}_{n+1}}\sigma \cdot X_n = X_{n+1}$ for n sufficiently large, a clear contradiction. □

Definition 3.7. *The previous proposition implies that we may define the inductive limit with respect to the maps defined above,*

$$X_\bullet/\mathfrak{S} := \lim_{n\to\infty} X_n/\mathfrak{S}_n.$$

Elements of $X_\bullet/\mathfrak{S}$ will be referred to as the stable orbits*, or just the* orbits *of $X_\bullet$.*

3.2. Induced FI–sets. In this section we discuss the properties of what we call induced FI–sets.

Definition 3.8. *Let m be a fixed non-negative integer, and let X be an $\mathfrak{S}_m$–set. For any $n \geqslant m$, we define $M(X)_n$ to be the set of ordered pairs (f, x) where $x \in X$ and $f \in \mathrm{Hom}_{\mathrm{FI}}([m], [n])$ is strictly increasing. For $n < m$, we set $M(X)_n = \varnothing$. Given an injection $g : [n] \to [p]$, we define $g_* : M(X)_n \to M(X)_p$ as follows: We can uniquely factor $g \circ f$ as $h \circ \sigma$ where $\sigma \in \mathfrak{S}_m$ and h is strictly increasing. We put*

$$g_*(f, x) = (h, \sigma x).$$

FI*–sets of the form $M(X)_\bullet$ will be called the* induced FI–sets*.*

Remark 3.9. There is an alternative definition of $M(X)$, for an $\mathfrak{S}_m$–set X, which we take the time to point out now as it will be used in the proof of Theorem 1.5. Any strictly increasing function $f : [m] \hookrightarrow [n]$ may be uniquely identified by its image in $[n]$. Therefore we may equivalently think of $M(X)_m$ as consisting of ordered pairs (K, x) where $K \subseteq [n]$ is of size m, and $x \in X$. With respect to this definition, the induced map of an injection $f : [n] \hookrightarrow [n']$ will map the pair (K, x) to the pair $(f(K), \sigma_{f,K} \cdot x)$, where $\sigma_{f,K} : [m] \to K$ is the bijection sending i to the preimage under f of the i–th smallest element of $f(K)$. Note that we have identified $\sigma_{f,K}$ with an element of $\mathfrak{S}_m$ by using the bijection between $[m]$ and K sending i to the i–th smallest element of K.

We have taken the time to outline both of these descriptions of induced FI–sets, as our original definition most closely replicates analogous constructions in the theory of FI–modules, while the second definition will prove to be a bit more convenient for computations later.

Example 3.10. Recall the FI–set $Z_\bullet$ defined in Example 1. We claim that $Z_\bullet$ is the induced FI–set $M(\{1\})$. Indeed, let $f : [1] \to [n]$ be an injection. Then we may assign to a pair $(f, 1)$ the image of $f(1) \in [n]$. This assignment defines the necessary isomorphism.

The following lemma is clear from the definition of the induced FI–set, and of induced representations.

Lemma 3.11. *Let X be an $\mathfrak{S}_m$–set. Then*

$$kM(X)_\bullet \cong M(kX)_\bullet.$$

We observe that Theorem 1.1 is straight-forward for induced FI–sets.

Lemma 3.12. *Let X be an $\mathfrak{S}_m$–set, and write*

$$X = \bigsqcup_{i\in I} \mathfrak{S}_m/H_i$$

for some indexing set I. Then for any $n \geqslant 0$,

$$M(X)_n = \bigsqcup_{i \in I} \mathfrak{S}_n/(H_i \times \mathfrak{S}_{n-m}),$$

where implicitly $\mathfrak{S}_n/(H_i \times \mathfrak{S}_{n-m})$ is empty for $n < m$.

Proof. We must show that the stabilizers of $M(X)_n$ are of the form $H_i \times \mathfrak{S}_{n-m}$. Indeed, given an element (f, x), where x is stabilized by H_i, for any $g \in \mathfrak{S}_n$, we have $g_*(f, x) = (f, x)$ if and only if $g \circ f = f \circ \sigma$ for $\sigma \in H_i$. This will occur if and only if $g|_{[m]} = \sigma$. So g stabilizes (f, x) if and only if $g \in H_i \times \mathfrak{S}_{n-m}$. □

One way to interpret Theorem 1.1 is that every finitely generated FI–set eventually "looks like" an induced FI–set. More precisely, if $X_\bullet$ is a finitely generated FI–set which is generated in degree $\leqslant d$, then there exists a collection $\{Y_i\}_{i=0}^{d}$, with Y_i a $\mathfrak{S}_i$–set, and an isomorphism of $\mathfrak{S}_n$–sets

$$X_n \cong \sqcup_i M(Y_i)_n,$$

for all $n \gg 0$. An analogous theorem has been known to be true about FI–modules since at least the work of Nagpal [9].

3.3. Motivating examples. In this section, we take time to write down a collection of motivating examples for the study of FI–sets. Our focus will be on constructing illustrative examples of FI–graphs. An FI–graph is a functor from FI to the category of graphs. In other words, an FI–graph is an FI–set of vertices, paired with a symmetric relation which dictates how these vertices are connected through edges (see [14]). Here, a symmetric relation of FI–sets is a relation of FI–sets $R_\bullet$ for which R_n is a symmetric relation for each n.

Our study of FI–graphs begins with Kneser graphs.

Example 3.13. For any fixed $n, r \geqslant 0$, the Kneser graph $KG_{n,r}$ has vertices indexed by the r–element subsets of $[n]$, with edges between two vertices if those subsets are disjoint.

The FI–graph $KG_{\bullet,r}$ has $G_n = KG_{n,r}$. For each injection f from $[m]$ to $[n]$, the corresponding transition map takes the vertex $\{a_1, \ldots, a_r\}$ of G_m to the vertex $\{f(a_1), \ldots, f(a_r)\}$ of G_n.

There are several minor ways in which this construction can be generalized, as in the following examples.

The vertices could be indexed by ordered r–tuples rather than by (unordered) subsets.

Example 3.14. For any fixed $r \geqslant 0$, define each graph G_n to have vertices indexed by the r–tuples of elements of $[n]$, with edges between two vertices if no element of $[n]$ appears in both r–tuples. For each injection f from $[n]$ to $[m]$, the corresponding transition map takes the vertex $(a_1, \ldots, a_r)$ of G_n to the vertex $(f(a_1), \ldots, f(a_r))$ of G_m.

As with Example 3.13, these graphs and transition maps form an FI–graph.

Future examples with the same vertex sets as Example 3.13 or 3.14 will use the same transition maps, without this being explicitly stated each time.

Rather than using ordered or unordered r–tuples, as in Examples 3.13 and 3.14, it is possible to care about only some of the order data.

Example 3.15. For any fixed $r \geqslant 0$ and subgroup H of the symmetric group $\mathfrak{S}_r$, define each graph G_n to have vertices indexed by orbits of r–tuples of elements of $[n]$ under the action of H. As in Examples 3.13 and 3.14, edges are placed between each pair of vertices labelled by disjoint r–tuples.

Other FI–graphs may be defined with the same vertex sets as in Examples 3.13, 3.14, or 3.15, but with different sets of edges. There may be multiple orbits of edges, and they may only exist from a certain degree onwards.

Example 3.16. Let the vertex set of G_n be indexed by r–element subsets of $[n]$, and let a_0 to a_r be positive integers or infinity. In G_n, there is an edge between two vertices which share exactly l elements if and only if $n \geqslant a_l$.

This example has $r+1$ orbits of pairs of vertices, determined by the size i of the intersection of the two labelling sets. For each of these orbits, there is an edge joining those two vertices from degree a_i onwards.

If the vertices are described by (ordered) r–tuples as in Example 3.14, then there are many more orbits of pairs of vertices — rather than these orbits being defined just by the size of the intersection, they also take into account which positions any equal entries occupy. As in Example 3.16, though, all that is required is to choose when edges appear in each vertex orbit. In this example, the orbits of pairs of vertices are a little more complicated.

Example 3.17. Let the vertex set of G_n be indexed by r–tuples of elements of $[n]$. For each integer l between 0 and r and each injection s from any l–element subset of $[r]$ to $[r]$, fix a_{ls} to be either a nonnegative integer or infinity. Because we are working with undirected graphs, we require that $a_{ls^{-1}} = a_{ls}$.

In G_n, there is an edge between two vertices whose labelling r–tuples have l entries in common in positions given by s exactly if $n \geqslant a_{ls}$.

The number of parameters a_i or a_{ls} required by Examples 3.16 and 3.17 is the number of orbits of pairs of vertices, in the sense of the minimal number of pairs of vertices required for any pair of vertices in any degree to be in their image under some transition map. Effectively, for each orbit i of pairs of vertices, we need to decide in which degree a_i pairs of vertices in this orbit are first connected by an edge. Once this happens, all other pairs of vertices in the same orbit in the graph G_{a_i} must also be connected by an edge, and likewise pairs of vertices in the image of these pairs in later graphs G_r, for $r > a_i$.

Disjoint sums of FI–sets are FI–sets, so the preceding examples may be combined to give larger ones. Such a construction will have additional orbits of pairs of vertices, allowing additional edges as in the following example.

Example 3.18. Choose nonnegative integers r, l, a_{11}, a_{12}, and a_{22}. Each graph G_n has a vertex for each subset of $[n]$ of size r and another vertex for each subset of size l. Color these vertices red and blue, respectively. There is an edge between

- two red vertices if their subsets have intersection of size a_{11}
- a red vertex and a blue vertex if their subsets have intersection of size a_{12}
- two blue vertices if their subsets have intersection of size a_{22}

Example 3.18 is the disjoint sum of two instances of Example 3.16, with additional edges added between red vertices and blue vertices. A more general example could be constructed with more parameters, both those used in Example 3.16, and new parameters for each orbit of pairs of vertices with one red and one blue.

Because the conditions on an FI–graph only involve maps from G_n to G_m with $m \geqslant n$, an FI–graph may be edited by removing all vertices and edges before a certain point.

Example 3.19. Let $G_\bullet$ be an FI–graph. Modify it by replacing each G_i by the empty graph, for $i = 0$ to $k-1$. Transition maps from these graphs are trivial. The resulting object is an FI–graph.

It is not possible to remove all vertices and edges from any graph after a nonempty graph G_n, because transition maps to the empty graph cannot be defined. The closest we can come is to crush the entire graph to a point, perhaps with a self–edge.

Example 3.20. Let $G_\bullet$ be an FI–graph. Modify it by replacing each G_i by a single vertex, for $i \geqslant k$. If there are any edges in any prior graph, then this vertex must have a self–edge. Transition maps to these single–vertex graphs map every vertex to the only vertex. The resulting object is an FI–graph.

Our desire to in general allow non–injective behavior of the sort described in Example 3.20 is why we allow graphs to have self–edges. If self–edges are forbidden, then this example is only allowed when there are no edges earlier in the FI–graph and in general vertices would only be able to map to the same vertex if they were not connected by an edge.

3.4. Proof of Theorem 1.1. We now prove Theorem 1.1. Let $X_\bullet$ be a finitely generated FI–set. For each $\mathcal{O}$ in $X_\bullet/\mathfrak{S}$, we get an FI–subset $X(\mathcal{O})$ of $X_\bullet$ corresponding to the elements which map to $\mathcal{O}$ under the maps $X_n \longrightarrow X_n/\mathfrak{S}_n \longrightarrow X_\bullet/\mathfrak{S}$, and we have $X_\bullet = \bigsqcup_{\mathcal{O}} X(\mathcal{O})$. So it is enough to prove the theorem for each $X(\mathcal{O})$. In other words, we may, and do, reduce to the case that $X_\bullet/\mathfrak{S}$ is a singleton.

By Proposition 3.6, for n large enough, the maps $X_n/\mathfrak{S}_n \to X_{n+1}/\mathfrak{S}_{n+1}$ are bijective. So, for n large enough, the action of $\mathfrak{S}_n$ on X_n is transitive.

Choose some k large enough for the action of $\mathfrak{S}_n$ on X_n to be transitive for all $n \geqslant k$. Choose some particular element $x \in X_k$. Let G_k be the stabilizer of x in the $\mathfrak{S}_k$ action on X_k. For all $n \geqslant k$, let ι_{kn} be the obvious inclusion of $[k]$ into $[n]$ and let G_n be the stabilizer of $\iota_{kn}(x)$ in $\mathfrak{S}_n$. We want to show that there is a nonnegative integer m and a subgroup $H \subseteq \mathfrak{S}_m$ such that, for n sufficiently large, the subgroup G_n of $\mathfrak{S}_n$ is conjugate to $H \times \mathfrak{S}_{n-m}$. Note that for $\sigma \in H$ and $\tau \in \mathfrak{S}_{n-m}$ we write $\sigma \times \tau$ for the permutation of $\mathfrak{S}_n$ which applies σ to $[m]$ and τ to $[n] - [m]$.

Lemma 3.21. *For all $n \geqslant \ell \geqslant k$, we have $G_\ell \times \mathfrak{S}_{n-\ell} \subseteq G_n$.*

Proof. For any $\sigma \in \mathfrak{S}_\ell$ and $\tau \in \mathfrak{S}_{n-\ell}$, we have $(\sigma \times \tau) \circ \iota_{\ell n} = \iota_{\ell n} \circ \sigma$. Now, suppose $\sigma \in G_l$ so $\sigma(\iota_{k\ell}(x)) = \iota_{k\ell}(x)$. Then

$$(\sigma \times \tau)\left(\iota_{kn}(x)\right) = (\sigma \times \tau) \circ \iota_{\ell n} \circ \iota_{k\ell}(x) = \iota_{\ell n} \circ \sigma\left(\iota_{k\ell}(x)\right) = \iota_{\ell n} \circ \iota_{k\ell}(x) = \iota_{kn}(x).$$

So $\sigma \times \tau$ stabilizes $\iota_{kn}(x)$, and thus lies in G_n. $\square$

Define $A_n \subseteq [n]$ to be the orbit of n under G_n.

Lemma 3.22. *For $n > k$, we have $A_{n+1} \supseteq A_n \cup \{n+1\}$.*

We remark that the statement is meaningful for $k = n$ but need not be true in that case.

Proof. By definition, $n+1 \in A_{n+1}$. So the task is to show that $A_n \subset A_{n+1}$.

By Lemma 3.21, G_{n+1} contains $G_{n-2} \times \mathfrak{S}_2$ and, in particular, contains the transposition $(n\ n+1)$. So $n+1$ and n are in the same $\mathfrak{S}_{n+1}$ orbit and $n \in A_{n+1}$. But also by Lemma 3.21, G_{n+1} contains $G_n \times \{e\}$. So the G_{n+1} orbit of n contains the G_n orbit of n. In other words, $A_n \subset A_{n+1}$ as required. $\square$

Let $B_n = [n] \backslash A_n$. Then Lemma 3.22 shows that $[k] \supseteq B_{k+1} \supseteq B_{k+2} \supseteq B_{k+3} \supseteq \cdots$. For n large enough, therefore, the subset B_n stabilizes at some subset B of $[k]$. Let $m = |B|$.

For a subset P of $[n]$, let $\mathfrak{S}_P$ be the subgroup of $\mathfrak{S}_n$ which fixes all elements of $[n] \backslash P$. Here is our final, key, lemma:

Lemma 3.23. *Let n be large enough that $2(n-k) > (n-m)$ and $|A_n| = n-m$. Then $\mathfrak{S}_{A_n} \subseteq G_n$.*

Proof. Let $a \in A_n$. It is enough to show that G_n contains the transposition $(a\ n)$.

Let $[k+1, n] = \{k+1, k+2, \dots, n\}$. By Lemma 3.21, we have $\mathfrak{S}_{[k+1,n]} \subseteq G_n$. Since $a \in A_n$, there is some element $\rho \in G_n$ mapping n to a. Then $\rho \mathfrak{S}_{[k+1,n]} \rho^{-1} = \mathfrak{S}_{\rho([k+1,n])}$ is in G_n as well. We have $[k+1,n] \subseteq A_n$, so $\rho([k+1,n]) \subseteq A_n$ and, since $2(n-k) > n-m = |A_n|$, the sets $[k+1,n]$ and $\rho([k+1,n])$ must overlap. So there is some $b \in [k+1,n] \cap \rho([k+1,n])$. Then the transpositions $(b\ n)$ and $(a\ b)$ are in G_n, so the transposition $(a\ b)(b\ n)(a\ b) = (a\ n)$ is as well. □

We are now ready to finish the proof. For n large enough that Lemma 3.23 holds, we know that $\mathfrak{S}_{A_n} \subseteq G_n \subseteq \mathfrak{S}_{A_n} \times \mathfrak{S}_B$. This means that G_n must be of the form $\mathfrak{S}_{A_n} \times H_n$ for some subgroup H_n of $\mathfrak{S}_B$. Moreover, by Lemma 3.21, we have $H_n \subseteq H_{n+1} \subseteq H_{n+2} \subseteq \cdots \subseteq \mathfrak{S}_B \cong \mathfrak{S}_m$ so, for n large enough, the subgroup H_n stabilizes. We take H to be this stable limit. □

Remark 3.24. The results of this section actually constrain the behavior of the stabilizer groups G_n quite severely, even when n is not yet 'large enough'. As we move from X_n to X_{n+1}, the groups G_n and G_{n+1} are related in one of the following ways.

- It may be that $\mathfrak{S}_{A_{n+1}} \subseteq G_{n+1}$, in which case the subgroup $H^{(n+1)}$ of $\mathfrak{S}_{B_{n+1}}$ contains the intersection $H^{(n)} \cap \mathfrak{S}_{B_{n+1}}$, bearing in mind that B_{n+1} may be smaller than B_n.
- Alternatively, it is possible that G_{n+1} does not contain $\mathfrak{S}_{A_{n+1}}$. This can only happen when the hypotheses of Lemma 3.23 are not yet satisfied. If this happens, then $\mathfrak{S}_{A_{n+1} \cup \{n+2\}}$ is contained in G_{n+2}. Example 3.25 gives an example of this behavior.

The second case of Remark 3.24 is why Lemma 3.23 requires that $2(n-k) > (n-m)$. The following example illustrates what may happen when n is not yet this large.

Example 3.25. Let X_n be empty for $n < 5$. Take the groups G_n for $n \geqslant 5$ to be

- Generated by the symmetric groups $\mathfrak{S}_2$ acting on $[2]$, $\mathfrak{S}_3$ acting on $\{3,4,5\}$ and S_{n-5} acting on $[n]\backslash[5]$ for $n \in \{5,6,7\}$
- Generated by $\mathfrak{S}_2$ acting on $[2]$, $\mathfrak{S}_3$ acting on $\{3,4,5\}$, $\mathfrak{S}_3$ acting on $\{6,7,8\}$, and the permutation $(3\,6)(4\,7)(5\,8)$ for $n = 8$.
- Generated by $\mathfrak{S}_2$ acting on $[2]$ and $\mathfrak{S}_{n-2}$ acting on $[n]\backslash[2]$ for $n > 8$

Observe the failure of Lemma 3.23 for G_8. The orbit A_8 is $\{3,4,5,6,7,8\}$, but not all of $\mathfrak{S}_{A_8}$ is contained in G_8. We do not give a complete construction of an FI–set with these stabilizer groups — the vertices may be taken to be appropriate cosets of the groups G_n.

The gist of Remark 3.24 and Example 3.25 is that to go from G_n to G_{n+1}, one may remove elements from B_n or increase the subgroup $H^{(n)}$. When n is small, it is also possible to have a wreath product factor appear in G_{n+1}. This factor is temporary, in that it will always further increase to a large symmetric group in G_{n+2}.

4. Relations of FI–sets

4.1. Elementary definitions and properties. In this section we turn our attention to relations defined by FI–sets. We recall the definition from the introduction:

Definition 4.1. *Let $X_\bullet$ and $Y_\bullet$ denote two* FI*–sets. Then a* relation between $X_\bullet$ and $Y_\bullet$ *is an* FI*–subset of the product $(X \times Y)_\bullet$.*

In Section 3.3 we examined a large collection of examples of FI–sets and relations.

Proposition 4.2. *Let $X_\bullet$, $Y_\bullet$ be finitely generated* FI*–sets, and let $R_\bullet$ be a relation between $X_\bullet$ and $Y_\bullet$. Then $R_\bullet$ is finitely generated.*

Proof. It suffices to prove that the linearization $kR_\bullet$ is finitely generated (for any choice of k). It is easily seen that

$$k(X \times Y)_\bullet \cong kX_\bullet \otimes kY_\bullet.$$

Proposition 2.7 implies that $kX_\bullet \otimes kY_\bullet$ is finitely generated, whence the same is true of $k(X \times Y)_\bullet$. The proposition now follows from Theorem 2.6. ☐

Example 4.3. Let $G_\bullet$ denote an FI–graph (see [14]). That is, $G_\bullet$ is a functor from FI to the category of simple graphs and graph homomorphisms. Then the FI–set encoding the edges of $G_\bullet$, $E(G_\bullet)$, can be viewed as a symmetric relation between the vertex FI–set and itself. In fact, understanding properties of $E(G_\bullet)$ is one of the main motivations of the present work. Theorem 1.5 can be seen as a vast generalization of Theorem H of [14] (see Section 4.4). Section 3.3 focused on giving a large collection of examples of specific FI–graphs.

Example 4.4. Let $P_\bullet$ denote an FI–poset with partial orderings $\leqslant_\bullet$ (see [7]). Then one has an FI–relation defined by

$$R_n = \{(x, y) \mid x \leqslant_n y\}.$$

FI–posets were used by Gadish in [7], where he showed that they have a variety of applications in studying representation stability phenomena arising from linear arrangements.

Given a relation $R_\bullet$ between two FI–sets $X_\bullet$ and $Y_\bullet$ one may associated a collection of maps $r_n : kX_n \to kY_n$. Namely, for any $x \in X_n$,

$$r_n(e_x) = \sum_{(x,y)\in R_n} e_y.$$

Critically, the collection $\{r_n\}_n$ does not necessarily extend to a morphism of FI–modules $k[X_\bullet] \to k[Y_\bullet]$. This can be seen, for instance, by having $X_n = Y_n = [n]$, and $R_n = \{(x, y) \mid x \neq y\}$. Despite this fact, we want to prove the maps r_n display a regularity as n varies.

Example 4.5. Once again let $G_\bullet$ be an FI–graph, and assume that the vertex set $V(G)_\bullet$ is finitely generated. If we chose our relation to be the edge relation, then the associated maps $r_n : \mathbb{Q}V(G_n) \to \mathbb{Q}V(G_n)$ are given by multiplication by the adjacency matrices of the associated graphs. These maps are studied in [14], where it is pointed out that they usually do not form a morphism of FI–modules.

Example 4.6. If we assume that $P_\bullet$ is an FI–poset, then the associated maps r_n are sometimes called the incidence matrix of the poset P_n. These are the matrices with rows and columns indexed by elements of P_n which have a 1 in position (x, y) whenever $x \leqslant y$, and a 0 otherwise. Note that unlike in the previous case, this matrix is not symmetric (unless the poset is trivial). The inverse of r_n is the Möbius function of the poset r_n.

4.2. A key diagram. In this section, we begin to detail the main construction used in the proof of Theorem 1.5. This construction does not make use of the FI–set structure in its early stages, and can be accomplished at the level of FI–modules. In the next section, we will specialize to the FI–set case, and complete the proof of Theorem 1.5.

Let $V_\bullet$ be a finitely generated FI–module and λ a partition. As explained in Section 2.2, we define $(V_\bullet)_\lambda$ to be $\lim_{n\to\infty} \operatorname{Hom}_{\mathrm{FI-mod}}(M(\lambda[n])_\bullet, V_\bullet)$. As we showed there, for n sufficiently large, the maps in this inductive limit are isomorphisms, so $(V_\bullet)_\lambda$ is canonically isomorphic to $(V_n)_{\lambda[n]}$ for any sufficiently large n. Any n which is sufficiently large for this purpose will be said to be *in the stable range*. We recall that the definition of $(V_\bullet)_\lambda$ required fixing once and for all embeddings $\mathrm{Sp}(\lambda[n]) \hookrightarrow \mathrm{Ind}_m^n \mathrm{Sp}(\lambda[m])$; we will use those same embeddings throughout this section.

By Frobenius reciprocity, the inclusion $\mathrm{Sp}(\lambda[n]) \hookrightarrow \mathrm{Ind}_m^n \mathrm{Sp}(\lambda[m])$ corresponds to an inclusion $\mathrm{Sp}(\lambda[m]) \hookrightarrow \mathrm{Res}_m^n \mathrm{Sp}(\lambda[n])$. We will denote this inclusion $\eta_{m,n}$.

Remark 4.7. Because it will be important later, we note that the η maps can be chosen to satisfy the following composition property:

$$\eta_{j,n} \circ \eta_{i,j} = \eta_{i,n}.$$

For instance, one may always start by only choosing the maps $\eta_{i,i+1}$, and defining the general $\eta_{i,j}$ by composition. We assume that this choice has been made in what follows.

Another subtle, but important point is that the image of $\eta_{m,n}$ in $\mathrm{Sp}(\lambda[n])$ will necessarily be invariant under the action of $\mathfrak{S}_{n-m}$, thought of as the symmetric group on the set $\{m+1,\dots,n\}$. Indeed, the usual branching rule implies that there is a unique copy of $\mathrm{Sp}(\lambda[m])$ inside $\mathrm{Res}_m^n \mathrm{Sp}(\lambda[n])$. This unique copy of $\mathrm{Sp}(\lambda[m])$ must be the image of $\eta_{m,n}$. On the other hand, Pieri's rule implies that the restriction of $\mathrm{Sp}(\lambda[n])$ to the product group $\mathfrak{S}_m \times \mathfrak{S}_{n-m}$ contains a unique copy of the tensor product of $\mathrm{Sp}(\lambda[m])$ with the trivial representation. These two facts in tandem imply that the image of $\eta_{m,n}$ must be fixed by the action of $\mathfrak{S}_{n-m}$, as desired.

As in Section 2.2, we define $V_n^{\geqslant\lambda}$ to be the subrepresentation of V_n spanned by the $\mu[n]$–isotypic pieces, where $|\mu| \geqslant |\lambda|$. We define $V_n^{>\lambda}$ to be the subrepresentation of $V_n^{\geqslant\lambda}$ spanned by the $\mu[n]$–isotypic components with $\mu \neq \lambda$. As observed in Section 2.2, the vector spaces $V_n^{\geqslant\lambda}$ span a sub–FI–module of $V_\bullet$, and the $V_n^{>\lambda}$ form a sub–FI–module of those.

Let $\iota_{mn} : [m] \to [n]$ be the standard inclusion $r \mapsto r$. Then we have transition maps $(\iota_{nm})_* : V_m^{\geqslant\lambda} \to V_n^{\geqslant\lambda}$ and $(\iota_{nm})_* : V_m^{>\lambda} \to V_n^{>\lambda}$ and hence we have a map on the subquotients $(\iota_{nm})_* : V_m^{\geqslant\lambda}/V_m^{>\lambda} \to V_n^{\geqslant\lambda}/V_n^{>\lambda}$.

Let $V_\bullet$ and $W_\bullet$ be two finitely generated FI–modules and suppose that, for all n, we have a map $r_n : V_n \to W_n$ of $\mathfrak{S}_n$ representations. Then, by Schur's lemma, the r_n induce linear maps $r_{n,\lambda} : (V_n)_{\lambda[n]} \to (W_n)_{\lambda[n]}$ and hence, for n sufficiently large, induce maps $(V_\bullet)_\lambda \to (W_\bullet)_\lambda$. Our subject in this section is how to compute those maps. We abbreviate $A = (V_\bullet)_\lambda$ and $B = (W_\bullet)_\lambda$. Finally, we recall the notation $\alpha_{\lambda[n]}$ for the inclusion $A \otimes \mathrm{Sp}(\lambda[n]) \to V_n$ and $\beta_{\lambda[n]}$ for the surjection $W_n \to B \otimes \mathrm{Sp}(\lambda[n])$.

The key technical lemma of this section is the following.

Lemma 4.8. *With $V_\bullet, W_\bullet, \lambda$ as above, and with $m \leqslant n \leqslant q$ in the stable range, all four maps obtained through composition from $A \otimes \mathrm{Sp}(\lambda[m])$ to $B \otimes \mathrm{Sp}(\lambda[q])$ in the following diagram are equal.*

$$\begin{tikzcd}
 & & W_q \arrow[r, "\beta_{\lambda[q]}"'] & B \otimes \mathrm{Sp}(\lambda[q]) \\
A \otimes \mathrm{Sp}(\lambda[n]) \arrow[r, "\alpha_{\lambda[n]}"] & V_n \arrow[r, "r_n"'] & W_n \arrow[u, "(\iota_{nq})_*"'] \arrow[r, "\beta_{\lambda[n]}"'] & B \otimes \mathrm{Sp}(\lambda[n]) \arrow[u, "\mathrm{Id} \otimes \eta_{n,q}"'] \\
A \otimes \mathrm{Sp}(\lambda[m]) \arrow[u, "\mathrm{Id} \otimes \eta_{m,n}"'] \arrow[r, "\alpha_{\lambda[m]}"] & V_m \arrow[u, "(\iota_{mn})_*"']
\end{tikzcd} \tag{1}$$

Remark 4.9. The two squares in the diagram (1) need *not* be commutative. Moreover, one should once again be aware that while A and B technically have n dependence, there are canonical isomorphisms between subsequent vector spaces so long as n is within the stable range. We therefore abuse notation in this and all following arguments.

Proof. The two maps arising from the bottom left square are equal modulo $V_n^{>\lambda}$. On the other hand, r_n maps $V_n^{>\lambda}$ to $W_n^{>\lambda}$, which is annihilated by $\beta_{\lambda[n]}$, and is mapped to $W_q^{>\lambda}$ by $(\iota_{nq})_*$. All of these facts imply that the choice of map from $A \otimes \mathrm{Sp}(\lambda[m])$ does not effect the overall composition.

Similarly, the two maps arising from the upper right square agree when restricted to $W_n^{\geqslant\lambda}$. Starting from $A \otimes \mathrm{Sp}(\lambda[m])$, all choices of maps land in this subspace. $\square$

Definition 4.10. *Let $V_\bullet, W_\bullet, \lambda$ be as above, and let $m \leqslant n \leqslant q$ be in the stable range. Then we write $\delta_{m,n,q} : A \otimes \mathrm{Sp}(\lambda[m]) \to B \otimes \mathrm{Sp}(\lambda[q])$ to denote the equal maps of Lemma* 4.8*.*

Our next goal will be to relate the map $\delta_{m,n,q}$ to $r_{n,\lambda}$. Composition in (4.8) along the path up–right–right–right–up yields the equality

$$\delta_{m,n,q} = r_{n,\lambda} \otimes \eta_{m,q}.$$

On the other hand, composition along the path right–up–right–up–right gives:

$$\delta_{m,n,q} = \beta_{\lambda[q]} \circ (\iota_{nq})_* \circ r_n \circ (\iota_{mn})_* \circ \alpha_{\lambda[m]}.$$

So we have:

$$r_{n,\lambda} \otimes \eta_{m,q} = \beta_{\lambda[q]} \circ (\iota_{nq})_* \circ r_n \circ (\iota_{mn})_* \circ \alpha_{\lambda[m]}. \tag{2}$$

We will find that in the case of FI–sets, the right hand side of the above equality is straight forward to compute. This will allow us to give an explicit description of $r_{n,\lambda}$.

Once and for all, fix some m in the stable range along with a vector $x \in \mathrm{Sp}(\lambda[m])$. For any $n \geqslant m$, choose some integer q such that $m \leqslant n \leqslant q$, and pick a linear functional $\psi : \mathrm{Sp}(\lambda[q]) \to \mathbb{Q}$ such that $\psi(\eta_{m,q}(x)) = 1$. We further impose the requirement that ψ is invariant with respect to the action of $\mathfrak{S}_{q-m}$, thought of as the automorphism group of the set $\{m+1, \ldots, q\}$. Note that this can be done via an averaging trick — where one replaces ψ with the function measuring ψ's average value along all $\mathfrak{S}_q$–conjugates — because $\mathfrak{S}_{q-m}$ acts trivially on the image of $\eta_{m,q}$ (see Remark 4.7).

Let $\{b_j\}$ be any fixed basis of B and $\{a_i\}$ a fixed basis for A. If we write $\{b_j^\vee\}$ to denote the dual basis of $\{b_i\}$, then we find that the (i,j)–th entry of $r_{n,\lambda}$ with respect to these bases is

$$\langle b_j^\vee \otimes \psi,\ \delta(a_i \otimes x)\rangle = \langle b_j^\vee \otimes \psi,\ \beta_{\lambda[q]} \circ (\iota_{nq})_* \circ r_n \circ (\iota_{mn})_* \circ \alpha_{\lambda[m]}(a_i \otimes x)\rangle. \tag{3}$$

Our goal in the next section will be to specialize this setup to FI–modules which arise from linearizations of FI–sets. We will see that in this setting the right hand side of (2) is computable enough for us to conclude Theorem 1.5.

4.3. The proof of Theorem 1.5.

In this section we specialize the setup in the previous section to FI–modules arising from linearizations of FI–sets. Let $X_\bullet$ and $Y_\bullet$ be finitely generated FI–sets, and let $R_\bullet$ be a relation between these sets. We will write $V_\bullet = kX_\bullet$, $W_\bullet = kY_\bullet$ and we write r_n for the map $V_n \to W_n$ induced by the relation R_n. For $x \in X_n$ and $y \in Y_n$ we will write $x \sim y$ to indicate $(x, y) \in R_n$.

Our first reduction will be to limit the total number of stable orbits of our FI–sets. The stable orbits of the relation $R_\bullet$ are subsets of products of stable orbits, one from $X_\bullet$ and one from $Y_\bullet$. It follows that the map r_n will split along such products. Thus, it suffices to understand the map r_n restricted to a chosen pair of orbits. We therefore may and do assume that $X_\bullet$ and $Y_\bullet$ have a unique stable orbit.

With this assumption in mind, Theorem 1.1 implies that for $n \gg 0$,

$$X_n = M(C)_n, \qquad Y_n = M(D)_n,$$

where C is an $\mathfrak{S}_r$–set for some r and similarly D is an $\mathfrak{S}_t$–set for some t. In particular, by Remark 3.9 we may think of V_n as having a basis of pairs (K, c), where K is an r–element subset of $[n]$, and $c \in C$ (see Definition 3.8 for how the action is defined on this basis). A similar description exists for W_n, which will have a basis of pairs (T, d), where $d \in D$ and T is a t–element subset of $[n]$.

For the remainder of this section, it will go without saying that K denotes a set of size r and T denotes a subset of size t.

We may write $\alpha_{\lambda[m]}(a_i \otimes x) = \sum_{(K,c)} \gamma_{K,c}(K, c)$, for some scalars $\gamma_{K,c}$, where $K \subseteq [m]$. Our job will be to compute

$$\langle b_j^\vee \otimes \psi,\ \beta_{\lambda[q]} \circ (\iota_{nq})_* \circ r_n \circ (\iota_{mn})_*(K, c)\rangle.$$

By definition we have $(\iota_{mn})_*(K, c) = (K, c)$, where on the right hand side K is thought of as a subset of $[n]$, and

$$r_n(K, c) = \sum_{(K,c)\sim(T,d),\ T\subset[n]} (T, d).$$

Thus, we have reduced the problem to needing to compute

$$\left\langle b_j^\vee \otimes \psi,\ \beta_{\lambda[q]}\left(\sum_{\substack{(K,c)\sim(T,d)\\ T\subset[n]\subseteq[q]}} (T, d)\right)\right\rangle = \sum_{\substack{(K,c)\sim(T,d)\\ T\subset[n]\subseteq[q]}} \left\langle b_j^\vee \otimes \psi,\ \beta_{\lambda[q]}\Big((T, d)\Big)\right\rangle. \tag{4}$$

We observe that ψ was constructed to be $\mathfrak{S}_{q-m}$–equivariant, and that $\beta_{\lambda[q]}$ is $\mathfrak{S}_q$–equivariant. This implies that the summand on the right hand side of (4) is unchanged by the action of $\mathfrak{S}_{q-m}$ on pairs (T, d). In particular, we may gather together those terms in the sum according to $S = T \cap [m]$ and $d \in D$. This yields the expression

$$\sum_{\substack{S\subseteq[m]\\ d\in D}} \phi_{S,d} \left|\{T \subseteq [n] \mid (K, c) \sim (T, d) \text{ and } T \cap [m] = S\}\right|,$$

where $\phi_{S,d}$ is some constant. We conclude our proof with the following lemma.

Lemma 4.11. *Using the notation of this section, the quantity $|\{T \subseteq [n] \mid (K, c) \sim (T, d)$ and $T \cap [m] = S\}|$ is either equal to 0 for all $n \gg m$ or to $\binom{n-m}{t-|S|}$ for all $n \gg m$.*

Proof. The key observation which will allow us to prove the lemma is the following. If there is some (T, d) with $(K, c) \sim (T, d)$ and $T \cap [m] = S$, then *every* choice of $T' \subseteq [n]$ with $T' \cap [m] = S$ and $|T'| = t$ has $(K, c) \sim (T', d)$. Indeed, one may find a permutation $\sigma \in \mathfrak{S}_{n-m}$ which maps T to T', and has $\sigma(T, d) = (T', d)$ (choose σ to map T to T' and to be strictly increasing on T). Moreover, because σ fixed $[m]$ we must have $\sigma(K, c) = (K, c)$. Thus, because our relation is equivariant under the action of the symmetric group, we conclude that $(K, c) \sim (T', d)$. This implies that the set in question either has size zero or $\binom{n-m}{t-|S|}$.

We have not quite proven our desired statement yet, however, as it is apriori possible that whether the relevant quantity is equal to zero or a polynomial at a given n, might itself be changing in n. To conclude the proof, we must show that for $n \gg 0$ there exists some set $T' \subseteq [n + 1]$ with $|T'| = t$, $(K, c) \sim (T', d)$, and $T' \cap [m] = S$ if and only if there exists some $T \subseteq [n]$ with the same properties. Put another way, once the relevant quantity agrees with a polynomial coefficient, it is then stuck doing so for all larger n. Indeed, this follows from the fact that the relation $R_\bullet$ is finitely generated and we have taken n to be in the stable range. □

Remark 4.12. In Section 4.2, it was noted that much of the groundwork for the proof of Theorem 1.5 was not special to the context of a relation between FI–sets. Indeed, if the linear maps r_n happened to actually form a morphism of FI–modules, then it can be shown that the entries of $r_{n,\lambda}$ will be constant in n. Therefore, the family of linear maps arising from a relation between FI–sets can be thought of as the next level of generality after morphisms of FI–modules. We therefore present the following question: Given two finitely generated FI–modules $V_\bullet, W_\bullet$ over k, and a collection of $\mathfrak{S}_n$–linear maps for each n $\{r_n : V_n \to W_n\}$, is there some natural condition that one can place on the maps r_n such that for $n \gg 0$, the associated $r_{n,\lambda}$ satisfy the conclusion of Theorem 1.5?

4.4. Applications of Theorem 1.5.

In this section, we consider applications of Theorem 1.5. In particular, we prove Corollary 1.6, and apply it to various cases.

The proof of Corollary 1.6. Let $X_\bullet$ be a finitely generated FI–set, and let $R_\bullet$ be a self relation. For any partition λ we write $r_{\lambda[n]}$ for the restriction of r_n to the $\lambda[n]$–isotypic piece of $\mathbb{Q}X_n$. Then, by Theorem 1.5, there exists a choice of bases such that for all $n \gg 0$ the maps $r_{\lambda[n]}$ take the form,

$$\begin{pmatrix} A_\lambda(n) & 0 & 0\dots 0 & 0 \\ 0 & A_\lambda(n) & 0\dots 0 & 0 \\ \vdots & & & \\ 0 & 0 & 0\dots 0 & A_\lambda(n) \end{pmatrix},$$

where $A_\lambda(n)$ is a square matrix of fixed (non–varying in n) dimension with entries in $\mathbb{Q}[n]$, and the total number of blocks is precisely $\dim_{\mathbb{Q}} \mathrm{Sp}(\lambda[n])$. Moreover, representation stability theory [2] implies that the total number of non–zero $r_{\lambda[n]}$ is a constant independent of n. Therefore, to understand the eigenvalues of r_n it suffices to understand the eigenvalues of $A_\lambda(n)$.

We may think of $A_\lambda(n)$ as being a matrix over the field $\mathbb{Q}(n)$. With this perspective, it becomes clear that we may factorize the characteristic polynomial of $A_\lambda(n)$, over some algebraic extension of $\mathbb{Q}(n)$, as

$$\chi_\lambda(n,x) = \prod_i (x - f_i(n))^{e_i},$$

where $e_i \geqslant 1$ are some integers, each $f_i(n)$ is some function which is algebraic over $\mathbb{Q}(n)$, and all of the $f_i(n)$ are distinct. This allows us to deduce that the eigenvalues of r_n are algebraic functions over $\mathbb{Q}(n)$ (in fact, they are integral over $\mathbb{Q}[n]$), as desired.

We next must argue that the functions $f_i(n)$, with i varying, only agree for finitely many values of n. In other words, if $f(n)$ and $g(n)$ are distinct algebraic functions, we must argue that $f(n) = g(n)$ for only finitely many n. Indeed, let $P(z,n)$ be the polynomial of minimal z degree with $P(f(n),n) = P(g(n),n) = 0$. By the minimality of the degree of P, the polynomial P is squarefree as a polynomial in z, as $\mathbb{Q}(n)$ is a separable field, so the discriminant $\Delta(n)$ of P with respect to z is a nonzero polynomial in n. For any n which is not a root of $\Delta(n)$, the roots of $P(z,n) = 0$ are distinct, so $f(n) \neq g(n)$ for such an n. Of course any non–zero polynomial in a single variable can only have finitely many roots, and so $f(n) \neq g(n)$ for all $n \gg 0$.

To conclude, we must show that the multiplicities of these distinct eigenvalues are equal to polynomials in n. This follows from the fact that the eigenvalues of each

$A_\lambda(n)$ have constant multiplicity, while the total number of $A_\lambda(n)$ which appear in the above matrix is precisely $\dim_{\mathbb{Q}} S(\lambda)_n$, a polynomial in n. □

Remark 4.13. We note for future use that, in the previous proof, we have proved that, even if we allow n to take real, non–integer values, there are only finitely many n for which any two of the eigenvalues $f_i(n), f_j(n)$ agree.

Once again calling upon the example of FI–graphs with the edge relation, we see that Corollary 1.6 implies Theorem H of [14]. Recall that for a graph G the Laplacian of G is the matrix $D - A$, where D is the diagonal matrix of degrees of vertices of G, and A is the adjacency matrix.

Theorem 4.14 (Theorem H of [14]). *Let $G_\bullet$ be a vertex–stable* FI*–graph, and let r_n denote either the adjacency matrix or Laplacian of G_n. We may write the distinct eigenvalues of r_n as,*

$$\lambda_1(n) < \lambda_2(n) < \cdots < \lambda_{l(n)}(n),$$

for some function $l(n)$. Then for all $n \gg 0$

1. *$l(n) = l$ is constant. In particular, the number of distinct eigenvalues of r_n is eventually independent of n;*
2. *for any i the function*
$$n \mapsto \lambda_i(n)$$
agrees with a function which is algebraic over $\mathbb{Q}(n)$.
3. *for any i the function*
$$n \mapsto \text{ the multiplicity of } \lambda_i(n)$$
agrees with a polynomial.

Proof. Note that a vertex stable FI–graph is one whose vertex FI–set is finitely generated. To make sense of the theorem statement, note that in this case r_n is symmetric, whence our matrix must have real eigenvalues. In particular, the functions $f_i(n)$ from the proof of Corollary 1.6 must be real–valued. Because we know that the $f_i(x)$ evaluate to distinct real numbers for all sufficiently large real numbers, x (Remark 4.13), it follows that we may order them using the usual order on $\mathbb{R}$. All of this put together imply Theorem H of [14] for the adjacency matrix.

To prove this statement for the Laplacian matrix, we note that it is shown in [14] that vertex–stable FI–graphs have vertex degrees which agree with polynomials in n for $n \gg 0$. Therefore, the Laplacian can be expressed as a $\mathbb{Q}[n]$–linear combination of relations. It follows that the Laplacian will satisfy the conclusion of Theorem 1.5, and therefore will also satisfy the conclusions of Corollary 1.6. □

Acknowledgements. The first author was supported by NSF grant DMS-1704811 and the second author was supported by NSF grant DMS-1600223. We would like to thank Tom Church for his helpful suggestions related to the proof of Theorem 1.5, and would like to thank our referees for their very careful readings.

References

[1] Andries E. Brouwer, Arjeh M. Cohen, and Arnold Neumaier, *Distance-regular graphs*, Springer-Verlag, 1989.

[2] Thomas Church, Jordan S. Ellenberg, and Benson Farb, *FI-modules and stability for representations of symmetric groups*, Duke Math. J. **164** (2015), no. 9, 1833–1910.

[3] Thomas Church, Jordan S. Ellenberg, Benson Farb, and Rohit Nagpal, *FI-modules over Noetherian rings*, Geom. Topol. **18** (2014), no. 5, 2951–2984.

[4] Thomas Church and Benson Farb, *Representation theory and homological stability*, Adv. Math. **245** (2013), 250–314.

[5] Aurélien Djament, *Des propriétés de finitude des foncteurs polynomiaux*, Fund. Math. **233** (2016), no. 3, 197–256.

[6] Aurélien Djament and Christine Vespa, *Foncteurs faiblement polynomiaux*, Int. Math. Res. Not. IMRN (2019), no. 2, 321–391.

[7] Nir Gadish, *Representation stability for families of linear subspace arrangements*, Adv. Math. **322** (2017), 341–377.

[8] Chris Godsil and Gordon F. Royle, *Algebraic graph theory*, Graduate Texts in Mathematics, vol. 207, Springer Science & Business Media, 2001.

[9] Rohit Nagpal, *FI-modules and the cohomology of modular representations of symmetric groups*, `https://arxiv.org/abs/1505.04294`, 2015.

[10] Andrew Putman, *Stability in the homology of congruence subgroups*, Invent. Math. **202** (2015), no. 3, 987–1027.

[11] Eric Ramos, *Homological invariants of FI-modules and FI_G-modules*, J. Algebra **502** (2018), 163–195.

[12] ———, *An application of the theory of FI-algebras to graph configuration spaces*, Math. Z. **294** (2020), no. 1-2, 1–15.

[13] Eric Ramos and Graham White, *Families of Markov chains with compatible symmetric-group actions*, `https://arxiv.org/abs/1810.08475`, 2018.

[14] ———, *Families of nested graphs with compatible symmetric-group actions*, Selecta Math. (N.S.) **25** (2019), no. 5, Paper No. 70 (42 pages).

[15] Steven V. Sam and Andrew Snowden, *GL-equivariant modules over polynomial rings in infinitely many variables*, Trans. Amer. Math. Soc. **368** (2016), no. 2, 1097–1158.

[16] Andrew Snowden, *Syzygies of Segre embeddings and Δ-modules*, Duke Math. J. **162** (2013), no. 2, 225–277.

Eric Ramos, Department of Mathematics, University of Oregon, Fenton Hall, Eugene, OR 97401, USA
E-mail : `eramos@uoregon.edu`

David Speyer, Department of Mathematics, University of Michigan, 530 Church St., Ann Arbor, MI 48109, USA
E-mail : `speyer@umich.edu`

Graham White, Department of Mathematics, Indiana University - Bloomington, Rawles Hall, Bloomington, IN 47405, USA
E-mail : `grrwhite@iu.edu`