A Dimension Theorem for Permutation Modules

Søren Riis
Queen Mary, University of London
Technical report (June 2005)

Abstract

It is shown that linear subspaces that are closed under the action of the symmetric group under quite general assumptions only have a limited set of possible vector space dimensions, and that these can be given by certain polynomial expressions.

More specifically, we show that for $k, r \in N$ there exists a finite set $\Gamma \subseteq Q[x]$ of polynomials over Q such that: for each $n \geq k+1$, for each field F (of any characteristic), for each FS_n-module V that can be generated by r vectors that each are invariant under the action of $S_{\{k+1,k+2,\ldots,n\}}$, and for each FS_n-submodule $U \subseteq V$ - there exists $p \in \Gamma$ such that $p(n) = dim(U)$. This is proven by two different methods.

1 Introduction

1.1 Background and motivation

In this short section I will attempt to give a very brief outline of the origin and motivation underlying the main theorem of this paper. It should, however, be emphasised that the results in this paper stand alone and can be viewed as results in pure algebra.

In Automated Theorem Proving, as well as in verification and many other areas of Computer Science, the specific problem in focus can often be reduced to a so called satisfiability problem. This satisfiability problem can, with only minor changes in notation, be written as a system of polynomial equations. For most practical problems, the polynomials (in the polynomial equations) can be chosen of low degree (e.g degree ≤ 3), while the number of variables is typically much larger.

In algebraic proof systems like, for example, the so-called Nullstellensatz Proof Systems or polynomial calculus we are given such a set Λ of low degree polynomials. They represent the assumptions (or axioms). The task is to derive other polynomials in the ideal $I(\Lambda)$ generated by these polynomials. Each such derived polynomial can be viewed as a "theorem". The ideal $I(\Lambda)$ consists of all "theorems" that logically follow from the axioms Λ. In this setting the 1 polynomial represents "false", and if $1 \in I(\Lambda)$, "false" is a "theorem", and the original set of assumptions has been shown to be inconsistent.

Let $T(\Lambda) \subseteq I(\Lambda)$ denote a set of "theorems" that can be derived from Λ using a certain limited amount of resources. Taking linear combinations of vectors usually requires very few additional resources. Thus the vectorspace $\widetilde{T}(\Lambda) := span\{T(\Lambda)\} \subseteq I(\Lambda)$ of vectors spanned by $T(\Lambda)$ essentially forms the set of "theorems" that can be derived from Λ using those limited resources. Since the resources are limited, and since we might even have decided only to focus on "theorems" of low complexity (i.e. degree $\leq d$ for some "small" number $d \in N$) in general $\widetilde{T}(\Lambda)$ is a proper subset of the ideal $I(\Lambda)$ of all "theorems".

In many applications the set Λ has certain symmetries. One such symmetry is that the set Λ might be invariant (as a set) under the action of the symmetric group S_n. In most of the applications, this symmetry appears because the "axioms" express a property (e.g. a graph theoretic property) that does not depend on the concrete representation of the problem (e.g. which labelling is chosen for vertex set of the graph). For many proof systems this symmetry is then inherited by the set $\widetilde{T}(\Lambda)$ of "theorems".

The research underlying this paper grew out of the "discovery" (that can be deduced from work by Ajtai [1]), that the structure of the set $\widetilde{T}(\Lambda)$ of "theorems" is of a very special and discrete nature. If the underlying field F (over which the polynomials in Λ are defined) has characteristic zero, the structure of $\widetilde{T}(\Lambda)$ is easy to describe. Since $\widetilde{T}(\Lambda)$ is a linear subspace closed under the action of the symmetric group, it can be shown - using standard results from the representation theory of the symmetric group - that $\widetilde{T}(\Lambda)$ is isomorphic to a direct sum of Specht modules. The vector space dimension of $\widetilde{T}(\Lambda)$ can be calculated using Hook's formula and Young's rule (see [7] for other results in this direction).

Our main concern is to what happens over fields of finite characteristic. What can be said about the vectorspace dimension of $\widetilde{T}(\Lambda)$? Which vectorspace dimensions can appear?

In this paper it is shown that only certain vectorspace dimensions can appear and that these can be expressed by polynomials (with rational coefficients). This result has a number of interesting consequences and applications in algebraic proof complexity.

It should, however, once more be emphasised that the main results in the paper are purely algebraic. In the rest of this paper I do not assume that the reader has any interest in, or knowledge of, algebraic proof complexity. I simply consider $\widetilde{T}(\Lambda)$ the set of "theorems" as a FS_n-submodule U of a suitable FS_n-module V.

1.2 Outline in general terms

Let V be a vectorspace over a field F and let S_n denote the symmetric group permuting n elements. We consider the cases where S_n acts on V (i.e. for $\pi \in S_n$, and for each $v \in V$ we define a vector $\pi(v) \in V$ subject to the rules: $(\pi_1\pi_2)(v) = \pi_1(\pi_2(v))$ and $1(v) = v$). We can view V as a FS_n-module. We say that a FS_n-submodule $U \subseteq V$ is generated by the vectors $v_1, v_2, \ldots v_r \in$

U if U is contained in each FS_n-submodule that contains $v_1, v_2, \ldots, v_r$. For $A \subseteq \{1, 2, \ldots n\}$, we let $S_A \subseteq S_n$ denote the subgroup of permutations $\pi \in S_n$ with $\pi(i) = i$ for each $i \in \{1, 2, \ldots n\} \setminus A$. A vector $v \in V$ is S_A-invariant if each $\pi \in S_A \subseteq S_n$ leaves v invariant (i.e. $\pi(v) = v$). For $k, r \in N$ we let $M_{F,k,r}$ denote the class of all FS_n-modules that can be generated by r (or fewer) generators, that all are invariant under the action of $S_{\{k+1,k+2,\ldots n\}}$. We will show:

Theorem (general version): For each $k, r \in N$ there exists a finite set $\Gamma_{k,r} \subset Q[x]$ of polynomials such that: for each field F, for each $n \geq k$, for each FS_n-module $V \in M_{F,k,r}$ each FS_n-submodule $U \subseteq V$ has a vector space dimension given by $dim(U) = p(n)$ for some $p \in \Gamma_{k,r}$.

I will prove the theorem using two different methods. The first proof relies on a number of central theorems from the representation theory of the symmetric groups. The second proof is more direct, and is based on a notion akin to that of a Grobner basis, combined with basic ideas from matroid theory, as well as a general combinatorial principle related to logic and finite model theory. Strictly speaking it is possible to present the second proof - in a more machinelike fashion - without making any explecit reference to Groebner basis, matroids, logic or finite models. However, I decided to include these links in order to make the presentation less dry.

1.3 Outline in basic terms

Let me explain the ideas underlying the statement of the theorem in less general terms.

The simplest illustration of the theorem appears if we consider $V := V_{F,n}$ the vector space spanned by the vectors $e_1, e_2, \ldots e_n$ over some field F. The symmetric group S_n acts naturally on $V_{F,n}$ via the definition $\pi(e_j) := e_{\pi(j)}$ which is then extended linearily so $\pi(\Sigma_j \lambda_j e_j) := \Sigma_j \lambda_j \pi(e_j)$. We say that a linear subspace $U \subseteq V_{F,n}$ is closed under the action of S_n if $\pi(U) \subseteq U$ for each $\pi \in S_n$. It is not hard to see that for each n there are two non-trivial linear subspaces $U \subseteq V_{F,n}$ which are closed under the action of S_n: the one-dimensional subspace generated by the vector $e_1 + e_2 + \ldots + e_n$, and the $(n-1)$-dimensional subspace consisting of all vectors $v = \Sigma_j \lambda_j e_j, \quad \lambda_j \in F, \quad j = 1, 2, \ldots, n$ where $\Sigma_j \lambda_j = 0$. If we include the trivial subspaces, we see that any S_n closed subspace of $V_{F,n}$ has vectorspace dimension belonging to the set: $\{0, 1, n-1, n\}$.

Let $V_{F,n}^k$ denote the k-fold tensor product of $V_{F,n}$. Thus $V_{F,n}^2$ can be viewed as the vector space spanned by the vectors $e_{i,j}$ where $i, j \in 1, 2, \ldots, n$. We notice that the vectorspace $V_{F,n}^k$ has dimension n^k, and that the symmetric group S_n acts naturally on $V_{F,n}^k$ via the action $\pi(e_{i_1,i_2,\ldots i_k}) := e_{\pi(i_1),\pi(i_2),\ldots \pi(i_k)}$. We say that a linear subspace $U \subseteq V_{F,n}^k$ is closed under the action of S_n if $\pi(U) \subseteq U$ for each $\pi \in S_n$. It can be shown using standard methods from the representation theory of the symmetric group (as explained in more details in [6]) that for the

field Q (or for the fields R, C or any other field of char$= 0$), any S_n closed subspace of $V_{Q,n}^2$ has a vectorspace dimension that belongs to the set:

$\Gamma_{Q,2} := \{0, 1, 2, n-1, n, n+1, 2n-2, 2n-1, 2n, 3n-3,$
$3n-2, 3n-1, \frac{n^2-3n}{2}, \frac{n^2-3n+2}{2}, \frac{n^2-3n+4}{2}, \frac{n^2-3n+6}{2},$
$\frac{n^2-n-2}{2}, \frac{n^2-n}{2}, \frac{n^2-n+2}{2}, \frac{n^2-n+4}{2}, \frac{n^2+n-4}{2}, \frac{n^2+n-2}{2},$
$\frac{n^2+n}{2}, \frac{n^2+n+2}{2}, \frac{n^2+3n-6}{2}, \frac{n^2+3n-4}{2}, \frac{n^2+3n-2}{2}, \frac{n^2+3n}{2},$
$n^2-3n+1, n^2-3n+2, n^2-3n+3, n^2-2n, n^2-2n+1, n^2-2n+2,$
$n^2-n-1, n^2-n, n^2-n+1, n^2-2, n^2-1, n^2\}$.

For each $n \geq 9$ this set $\Gamma_{Q,2}$ consists of exactly 40 distinct values. The set $\Gamma_{Q,2} = \Gamma_{R,2} = \Gamma_{C,2} = \Gamma_{F_0,2}$ is independent of the underlying field F_0 as long as it has characteristic 0. A priori it is not clear why the number of possible vectorspace dimensions can be bounded by a number that is independent of n. Also, there is no a priori reason why each element in the set of possible vectorspace dimensions is given by a polynomial in n.

The fact that there is a finite set $\Gamma_{Q,k}$ of polynomials (not just for $k = 2$, but for any k) in the case of fields of characteristic 0 was noticed in [7].

The fact that for each fixed Field F (even fields of finite characteristic) there exists a finite set $\Gamma_{F,k}$ of functions (not just for $k = 2$, but for any $k \in N$) follows relatively directly from the work by [1], as well from the later improvements by [5]. We show that the functions can actually be chosen as polynomials. Futhermore, we show that the finite set Γ of polynomials can be chosen *before* the field is given:

Theorem (Version 2): For any $k \in N$ there exists a finite set Γ_k of functions f, such that for any $n \in N$, for any field F (*any characteristic*), and for any linear subspace $U \subseteq V_{F,n}^k$ that is S_n-closed, the vectorspace dimension $dim(U)$ of U, is given by $f(n)$ for some $f \in \Gamma_k$. Furthermore, the functions f in Γ_k can be chosen to be polynomials with rational coeficients.

This version of the theorem is a special case of the general version of the theorem. To see this, notice that $V_{F,n}^k$ is a FS_n-module generated by the finite set of vectors $e_{i_1,i_2,\ldots i_k}$ for which $i_1, i_2, \ldots, i_k \in \{1, 2, \ldots k\}$. Moreover, each of those generators is invariant under $S_{\{k+1,k+2,\ldots,n\}}$, and thus $V_{F,n}^k$ is generated by a finite set of $S_{\{k+1,k+2,\ldots n\}}$-invariants.

Even though the general version of the theorem is in some sense much more general than the second version of the theorem, we will show that the general version of the theorem can be derived from version 2 of the theorem.

2 First proof of the Theorem

We are now ready to prove the theorem. The first proof is of a rather technical nature. In this section we prefer to rephrase the theorem using terminology from the representation theory of the symmetric group. All concepts and definitions can be found in [4].

Theorem (version 3): For any $k \in N$ there exists a finite set Γ of functions f, such that for each $n \in N$, for each partitioning α of k, for each field F and for each FS_n-submodule $U \subseteq M^{(n-k,\alpha)}$ the vector space dimension $dim(U)$ of U is given by $f(n)$ for some $f \in \Gamma$. Furthermore, the functions f in Γ can be chosen to be polynomials with rational coefficients.

Proof: First we show that the theorem is valid if the fields F are chosen to have characteristic zero. Fix any field F of characteristic zero. The module $M^{(n-k,\alpha)}$ is a direct sum of the irreducible Specht modules $S^{(n-|\beta|,\beta)}$ where β is a partition of $|\beta| \leq k$. The vector space dimension of each $S^{(n-|\beta|,\beta)}$ can be calculated using the Hook-formula, and it is straightforward to check that $p_\beta(n) := dim(S^{(n-|\beta|,\beta)})$ is a polynomial in n (for any fixed β). The multiplicity of each irreducible submodule is given by Young's rule, which states that the multiplicity $c_{\alpha,\beta} := [M^{(n-k,\alpha)} : S^{(n-|\beta|,\beta)}]$ of $S^{(n-|\beta|,\beta)}$ in $M^{(n-k,\alpha)}$ is given by the number of semi-standard $(n-|\beta|,\beta)$-tableaux of type $(n-k,\alpha)$. This number is clearly independent of n when $n \geq 2k$. According to standard results in the theory of modules (e.g. Jordan-Holder's Theorem), any submodule U is isomorphic to a direct sum of irreducible modules, i.e. is a direct sum of modules isomorphic to Specht modules. Thus for any $n \geq 2k$, and for any $U \subseteq M^{(n-k,\alpha)}$ the vector space dimension $dim(U)$ can be expressed as $\Sigma_\beta [U : S^{(n-|\beta|,\beta)}] p_\beta(n)$. Thus the theorem is valid if we let Γ consist of the collection of polynomials of the form $p(n) := \Sigma_\beta v_\beta p_\beta(n)$ where for each of the finitely many partitionings β of k, we have $v_\alpha \in 0, 1, \ldots c_{\alpha\,\beta}$. Notice that the number of polynomials in Γ is bounded from above by the number $\Pi_\beta(1 + c_{\alpha\,\beta}) < \infty$.

Now we consider the more interesting case where the fields are allowed to have finite characteristic. We show the first part of the theorem (that there exists a finite set Γ of functions) for $M^{(n-|\alpha|,\alpha)}$. We show this using induction on the partitioning $(n-|\alpha|,\alpha)$ with regards to the usual dominance ordering between partitionings. Notice that a partitioning $(n-|\beta_1|,\beta_1)$ dominates $(n-|\beta_2|,\beta_2)$ for one value of $n \geq 2max(|\beta_1|,|\beta_2|)$ if and only if this happens for all values of $n \geq 2max(|\beta_1|,|\beta_2|)$.

When $|\gamma| = 0$, the theorem is clearly satisfied by letting Γ_0 consist of the polynomials 0 and 1. Assume that we are given γ, and we know that the theorem is valid for all partitionings γ' with $(n-|\gamma'|,\gamma')$ dominating $(n-|\gamma|,\gamma)$. Let $U \subseteq M^{(n-|\gamma|,\gamma)}$ be an arbitrary submodule. According to James submodule theorem either $S^{(n-|\gamma|,\gamma)} \subseteq U$ or $U \subseteq S^{(n-|\gamma|,\gamma)^\perp}$. We consider each of these two cases separately.

Case 1: Assume $S^{(n-|\gamma|,\gamma)} \subseteq U$. A central theorem in James' theory states that each Specht module can be characterised as the kernel of a F_pS_n-homomorphism $\Psi : M^{(n-|\gamma|,\gamma)} \rightarrow Y_n$ where the F_pS_n module Y_n is of the form $\Sigma_{\gamma'} M^{(n-|\gamma'|,\gamma')}$ where $(n-|\gamma'|,\gamma)$ dominates $(n-|\gamma|,\gamma)$. Notice that the number of summands is independent of $n \geq 2|\gamma|$. Now, since Y_n is a finite direct sum of modules which each satisfies the first part of the theorem, Y_n itself must satisfy the first part of the theorem. This follows by straightforward application of standard results in the theory of modules (e.g. Jordan-Holder's Theorem). Thus there exists a finite set Γ' of functions such that for any submodule $U' \subseteq Y_n$ the

vector space dimension of U' is given by $f(n)$ for some function $f \in \Gamma'$. Now $U/ker(\Psi) \equiv \Psi(U) =: U' \subseteq Y_n$, and thus the vector space dimension $dim(U)$ is given by $dim(\Psi(U)) - dim(ker(\Psi)) = dim(U') - dim(S^{(n-|\gamma|,\gamma)})$. Let Γ_1 consist of all functions $h(n) = f(n) - dim(S^{(n-|\gamma|,\gamma)})$ where $f \in \Gamma'$, and where γ is one of the finitely many partitions with $(n-|\gamma|,\gamma)$ dominating $(n-|\alpha|,\alpha)$.

Case 2: Assume $U \subseteq S^{(n-|\gamma|,\gamma)^\perp}$.

Method 1: In this case we can use the fact that the assumption implies that $U^\perp \supseteq S^{(n-|\gamma|,\gamma)}$ and that we can then apply Case 1 and show that $dim(U^\perp)$ is given by $h(n)$ for some $h \in \Gamma_1$. But then $dim(U) = dim(M^{(n-|\gamma|,\gamma)}) - h(n)$ and the induction works if we let Γ_2 consist of Γ_1, together with all functions $r(n) := dim(M^{(n-|\gamma|,\gamma)}) - h(n)$ with $h \in \Gamma_1$.

Method 2: We can also handle case 2 without reference to case 1. To do this we use the work of James according to which there exists a module Y_n of the form $\Sigma_{\gamma'} M^{(n-|\gamma'|,\gamma')}$ where $(n-|\gamma'|,\gamma)$ dominates $(n-|\gamma|,\gamma)$ and a F_pS_n -homomorphism $\Psi : Y_n \to M^{(n-|\gamma|,\gamma)}$ such that $\Psi(Y_n) = S^{(n-|\gamma|,\gamma)^\perp}$. Let $U' := \Psi^{-1}(U) \subseteq Y_n$. As in 'case 1', there is a finite set Γ' of functions such that for any $n \geq 2|\gamma|$ and for any submodule $U' \subseteq Y_n$, the vector space dimension of U' is given by $f(n)$ for some function $f \in \Gamma'$. Now, since $U = \Psi(U') \equiv U'/ker(\Psi)$, we have $dim(U) = dim(U') - dim(ker(\Psi))$. Let Γ_2 consist of all functions $h(n) := f(n) - g(n)$ where $f, g \in \Gamma'$.

Whether we used method 1 or 2 we can see that there is a finite set Γ (namely $\Gamma_1 \cup \Gamma_2$) of functions such that the first part of the theorem is valid.

We want to show that the functions in Γ can be chosen to be polynomials with rational coefficients. First, notice that a corollary to the first part of the theorem (which we have already proved) is that the decomposition numbers $d_{\alpha\,\beta\,n} := [S^{(n-|\alpha|,\alpha)} : D^{(n-|\beta|,\beta)}]$ are bounded by a constant d (for example $d := |\Gamma|$) which is independent of n. Now, for each α for which $(n-|\alpha|,\alpha)$ is p-regular and we have the identity $S^{(n-|\alpha|,\alpha)} = D^{(n-|\alpha|,\alpha)} + \Sigma_\beta [S^{(n-|\alpha|,\alpha)} : D^{(n-|\beta|,\beta)}] D^{(n-|\beta|,\beta)}$, where the $\beta' s$ in the sum runs over all partitionings for which $(n-|\beta|\,\beta)$ is p-regular, and which dominates $(n-|\alpha|,\alpha)$. Notice that the property of being p-regular is independent of n for $n \geq 2|\alpha|$. From this we get

$Dim(D^{(n-|\alpha|,\alpha)}) = Dim(S^{(n-|\alpha|,\alpha)}) - \Sigma_\beta d_{\alpha\,\beta\,n} Dim(D^{(n-|\beta|,\beta)})$.

Further, $Dim(S^{(n-|\alpha|,\alpha)})$ is independent of the ground field, and as already noticed (using the Hook formula) it is given by a polynomial with rational coefficients. We show the second part of the theorem using the same induction as for the first part. Thus, for each $(n-|\beta|,\beta)$ dominating $(n-|\alpha|,\alpha)$, we can assume that there exists a finite set Γ_β of polynomials (over Q) such that for each n, we have $Dim(D^{(n-|\beta|,\beta)}) = p(n)$, for some $p \in \Gamma_\beta$. Thus for each fixed $n \geq 2|\alpha|$, there exist polynomials $p_\beta \in \Gamma_\beta$ and constants $c_\beta \in \{0, 1, \ldots d\}$ such that $Dim(D^{(n-|\alpha|,\alpha)}) = Dim(S^{(n-|\alpha|,\alpha)}) - \Sigma_\beta c_\beta p_\beta$. The second part of the theorem now follows easily from the fact that the dimension of a submodule equals the sum of the dimensions of its decomposition factors (taken with multiplicity), that all irreducible modules are isomorphic to some $D^{(n-|\beta|,\beta)}$, and that the height of $M^{(n-|\alpha|,\alpha)}$ according to the first part of the theorem is bounded by a constant which is independent of n. ♣

3 Second proof of the Theorem

3.1 General outline of the proof

For this proof we first introduce the notion of a *generating basis*. This is a notion that is somewhat similar to the well known notion of a Groebner basis. As in the definition of a Groebner basis, we need to fix (i.e. define) a term ordering. Recall that in the context of a Groebner basis, the notion of divisibility (of leading terms) plays a crucial role. In our setting we consider FS_n-modules, and these have no obvious multiplicative structure. Thus it is clear, that the notion of a Groebner basis does not apply directly to our situation. The idea behind our concept of a generating basis is to define a "divisibility relation" on the set of leading terms. This "divisiblity relation" is defined in terms of the group action S_n that acts on V. Roughly stated, a term t_1 "divides" another term t_2 if t_1 preceedes t_2 in the term order and there is a permutation $\pi \in S_n$ such that $\pi(t_1) = t_2$. For this notion to succeed, the term ordering is defined in such a manner that the for $u, v \in V$ with leading terms t_1 and t_2, and with t_1 "dividing" t_2, there exists a permutation $\pi \in S_n$ such that $\pi(u)$ has the leading term t_2.

This property allows us to show that any FS_n-submodule $U \subseteq V$ has a generating basis $\Im$ (to guide the explanation at this stage we introduce a few basic concept from matroid theory). Recall that our ultimate goal is to calculate the vector space dimension of FS_n-submodules $U \subseteq V$. The next step in our analysis is to show that this dimension is uniquely determined by the set of leading terms of a generating basis $\Im$ for U. Furthermore, and this is the main point of the construction, the set of leading terms of vectors in $\Im$ is very special, and has a description that, in some sense, does not depend on n. The number of non-equivalent generating basis can then be bounded by a constant. This number is independent of n, as well as the underlying Field F. Essentially, this provides another proof that the set of functions in Γ_k in version 1 of the theorem can be chosen to be finite.

The final step of the proof is to use basic ideas from logic and finite model theory to show that the dimension can actually be given by polynomial expressions. We prove a theorem that states that the number of elements of $L(<)$-definable substructures of finite models is given by a polynomial expression. In general, this polynomial expression does not always produce the correct answer. However, the set of values where it fails, is only finite. This is sufficient for our application. Alternatively, one can notice that for our application only quantifier-free $L(<)$-definable substructures need to be considered. And the polynomial expression is always correct for this class of structures.

From this analysis it follows that the set Γ_k of possible vector space dimensions can be chosen to consist solely of polynomials with rational coeficients. This completes the general outline of the proof.

3.2 Generating basis - the detailed construction

Let $\mathrm{V}_{F,n}^{o,k}$ denote the vectorspace spanned by the set of basis vectors of the form $e_{i_1,i_2,\dots i_k}$ with $i_1, i_2, \dots i_k \in \{1, 2, \dots n\}$ distinct. Thus $V_{F,n}^{o,k}$ is isomorphic to $M^{(n-k,1^k)}$ with elements from the field F. We define a term ordering $\preceq$ on the set of basis vectors $e_{i_1,i_2,\dots i_k} \in V_{F,n}^{o,k}$ by essentially considering the lexographic ordering of the set $\{i_1, i_2, \dots i_k\}$ as first priority and the lexographic ordering of the tuple $(i_1, i_2, \dots i_k)$ as second priority. In other words given two vectors $e_{i_1,i_2,\dots i_k} \in V_{F,n}^{o,k}$ and $e_{j_1,j_2,\dots j_k} \in V_{F,n}^{o,k}$ we decide which is the largest with respect to the ordering $\preceq$ by first considering the sets $\{i_1, i_2, \dots, i_k\} \subset \{1, 2, \dots, n\}$ and $\{j_1, j_2, \dots, j_k\} \subset \{1, 2, \dots, n\}$. The two sets are compared using the lexographic ordering (the maximal element of $\{i_1, i_2, \dots, i_k\}$ and $\{j_1, j_2, \dots, j_k\}$ decides which set is the largest in the ordering. If the sets contain the same largest element, the set with the second largest element is largest in the term ordering. If the second largest elements are identical, the third largest elements are compared etc.). If the sets are identical i.e. if $\{i_1, i_2, \dots, i_k\} = \{j_1, j_2, \dots, j_k\}$ we decide which of the tuples $(i_1, i_2, \dots i_k)$ and $(j_1, j_2, \dots j_k)$ are largest in the usual lexographic ordering.

Example: We have $e_{96,97,98} \preceq e_{98,97,96} \preceq e_{14,57,99} \preceq e_{99,14,57} \preceq e_{96,100,17}$.

A vector $v \in V_{F,n}^{o,k}$ can be written uniquely as $v = \Sigma_{i_1,i_2,\dots i_k} \lambda_{i_1,i_2,\dots,i_k} e_{i_1,i_2,\dots i_k}$ where the index $i_1, i_2, \dots, i_k \in \{1, 2, \dots n\}$ are distinct. The leading term $Lt(v)$ of the vector v is defined to be $\lambda_{i_1,i_2,\dots i_k} e_{i_1,i_2,\dots,i_k}$ where $e_{i_1,i_2,\dots i_k}$ is the largest basis vector (in the ordering) that have $\lambda_{i_1,i_2,\dots i_k} \neq 0$. The normalized leading term $NLt(v)$ of the vector v defined to be $e_{i_1,i_2,\dots,i_k}$ where $e_{i_1,i_2,\dots i_k}$ is the largest basis vector ordering that have $\lambda_{i_1,i_2,\dots i_k} \neq 0$.

Definition: Two k-tuples $(i_1, i_2, \dots, i_k)$ and $(j_1, j_2, \dots, j_k)$ have the same *order type* if for all $s, t \in \{1, 2, \dots k\}$ $i_s > i_t$ holds exactly when $j_s > j_t$ holds. We say a k-tuple $(i_1, i_2, \dots i_k)$ *dominates* a k-tuple $(j_1, j_2, \dots j_k)$ if we have $i_1 \geq j_1, \dots, i_k \geq j_k$ and the order type of $(i_1, i_2, \dots, i_k)$ is the same as the order type of $(j_1, j_2, \dots, j_k)$. We say a term $\lambda e_{i_1,i_2,\dots,i_k}$ *is divisible by* $\mu e_{j_1,j_2,\dots,j_k}$ if $(i_1, i_2, \dots i_k)$ dominates $(j_1, j_2, \dots, j_k)$ (and $\lambda, \mu \in F \setminus \{0\}$).

The divisibility relation just defined is refelxive and transitive. We refer to the relation as a *divisibility* relation because we want to make the link to Groebner basis more obvious.

Lemma(1): If v has $Lt(v) = \lambda e_{j_1,j_2,\dots,j_k}$ then for each $e_{i_1,i_2,\dots,i_k}$ that is divisible by $\lambda e_{j_1,j_2,\dots,j_k}$ there exists $\pi \in S_n$ and $\mu \in F \setminus \{0\}$ such that $\mu Lt(\pi(v)) = \lambda e_{i_1,i_2,\dots i_k}$.

Proof: First assume $i_1 > i_2 > \dots > i_k$. Since $(i_1, i_2, \dots i_k)$ dominates $(j_1, j_2, \dots, j_k)$ we have $j_1 > j_2 > \dots > j_k$ and $i_s \geq j_s$ for $s = 1, 2, \dots k$. Let $\pi = (i_k, j_k)(i_{k-1}, j_{k-1}) \dots (i_2, j_2)(i_1, j_1)$ where $(i, i) = 1$. I claim that for any vector v with leading term $\lambda e_{j_1,j_2,\dots,j_k}$, the vector $(i_k, j_k) \dots (i_1, j_1)v$ has leading

term $\lambda e_{i_1,i_2,\dots i_k}$. In other words I claim that any v with $Lt(v) = \lambda e_{j_1,j_2,\dots j_k}$ have $Lt(\pi(v)) = \lambda e_{i_1,i_2,\dots i_k}$. The claim follows using induction. First notice that $(i_1, j_1)v$ has leading term $\lambda e_{i_1,j_2,j_3,\dots j_k}$. By the same observation $(i_2, j_2)(i_1, j_1)v$ has leading term $\lambda e_{i_1,i_2,j_3,j_4,\dots j_k}$ and in general $(i_k, j_k)\dots(i_1, j_1)v$ has leading term $\lambda e_{i_1,i_2,\dots i_k}$.

The general case is proved in exactly the same fashion. Assume that $i_{\eta(1)} > i_{\eta(2)} > \dots > i_{\eta(k)}$ for some permutation of $\{1, 2, \dots k\}$ and proceede as before with $\pi := (i_{\eta(k)}, j_{\eta(k)})\dots(i_{\eta(1)}, j_{\eta(1)})$.♣

Definition: A set $I \subseteq U$ is called *independent* if for no distinct elements $v_1, v_2 \in I$ is $Lt(v_1)$ divisible by $Lt(v_2)$.

Proposition: Let $\Im$ denote the collection of independent sets $I \subseteq U$. Then the pair $(U, \Im)$ is a matroid.

Proof: Recall that a matroid is a pair $(U, \Im)$ where U is a set, and $\Im$ a non-empty set of subsets of U called independent sets, satisfying the two properties:

- If $I \in \Im$ and $J \subseteq I$, then $J \in \Im$.
- The *exchange axiom*: if $I_1, I_2 \in \Im$ with $|I_1| < |I_2|$, then there exists $v \in I_2 \setminus I_1$ with the property that $I_1 \cup \{v\} \in \Im$.

Definition: A set $B \subseteq U$ is a *generating basis* for U iff

1. For each $u \in U$ there exists $b \in B$ such that $Lt(u)$ is divisible by $Lt(b)$
2. For no distinct $b_1, b_2 \in B$ is $Lt(b_1)$ divisible by $Lt(b_2)$.

When condition 1 holds we say that B spans U. When condition 2 is satisfied we say B consists of independent vectors. Combining lemma(1) with this definition we get:

Lemma(2) : Let B be a generating basis for U. Then for any vector $u \in U$ there exists $b \in B$, $\lambda \in F \setminus \{0\}$ and a permutation $\pi \in S_n$ such that $Lt(v) = Lt(\pi(b))$.

Recall that an independent set $B \in \Im$ of a matroid $(U, \Im)$ is a called a basis if it is maximal (i.e. not properly contained in any independent set of $\Im$).

Proposition: A set $B \subseteq U$ is a generating basis if and only if $B \in \Im$ is a basis in the matroid $(U, \Im)$.

Proof: "if": Assume $B \in \Im$ is a maximal independent set. Let $u \in U$ and assume there is no $b \in B$ such that $Lt(u)$ is divisible by $Lt(b)$. But then the set $B \cup \{u\} \subseteq U$ is independent and thus $B \cup \{u\} \in \Im$. This violates the assumption that B was maximal.

"only if": Assume $B \subseteq U$ is a generating basis. Assume B is not maximal. Let $B \cup b \in \Im$ be a proper extension. Now since $b \in U$ and since B is a generating

basis, there exists $b' \in B$ such that $Lt(b)$ is divisible by $Lt(b')$. But this violates the assumption that $B \cup \{b\}$ is an independent set. ♣

We just noticed that a generating basis is a basis in the monoid $(U, \Im)$. Thus - according to a basic result in the theory of monoid - the number of elements in each generating basis is constant. This can also be shown directly: Assume $B_1, B_2 \subseteq U$ each is a generating basis. For each $b \in B_2 \subseteq U$ there exists $b' \in B_1$ such that $Lt(b)$ is divisible by $Lt(b')$. Now since B_2 is a generating basis and $b' \in U$ there exists $b'' \in B_2$ such that $Lt(b')$ is divisible by $Lt(b'')$. Now since $b, b'' \in B_2$ and b is divisible by b'' we must have $b = b''$ and thus that $Lt(b) = Lt(b')$. This argument shows that:

Proposition: Any two generating basis $B_1, B_2 \subseteq U$ have vectors with the same (normalised) leading terms. Especially B_1 and B_2 contain the same number of elements.

3.3 Irreducible tuples

In this section we consider the leading terms of vectors in any generating basis. We show that only terms $e_{i_1,i_2,\ldots,i_k}$ with index set $\{i_1, i_2, \ldots i_k\} \subseteq \{1, 2, \ldots a\} \cup \{n-b, n-b+1, \ldots n\}$ for some constants $a, b \in N$ can appear as index in leading term. It should be emphasised that we do not attept to find the optial values for a and b. For our purpose it suffices to prove that a and b can be chosen independently of n.

Definition: Let $a_1 > a_2 > \ldots > a_k$ be k elements from $\{1, 2, \ldots, n\}$. The tuple $(a_1, a_2, \ldots a_k)$ is *reducible* if there exists $1 \leq s < t \leq k$ such that $a_s > a_{s+1} + (t-s) > a_t > a_{t+1} + (t-s+1)$. The tuple $(a_1, a_2, \ldots a_k)$ is *irreducible* if it is not reducible. The set $\{i_1, i_2, \ldots, i_k\} \subseteq \{1, 2, \ldots n\}$ with k distinct elements is *reducible* if the ordered tuple $i_{r_1} > i_{r_2} > \ldots > i_{r_k}$ is reducible. The set is *irreducible* if it is not reducible.

The following lemma links this termonology to the set of possible leading terms in vectors that belong to a generating basis.

Lemma(3): Let $v \in V_{F,n}^{o,k}$ be a vector in a generating basis for some FS_n-submodule $U \subseteq V_{F,n}^{o,k}$. Let $e_{i_1,i_2,\ldots i_k} := NLt(v)$ denote the (normalized) leading term of v. Then the set $\{i_1, i_2, \ldots i_k\} \subseteq \{1, 2, \ldots n\}$ is irreducible.

Proof: Assume $v \in U$ is a vector with leading term $e_{i_1,i_2,\ldots,i_k}$ with $\{i_1, i_2, \ldots i_k\}$ reducible. Let $\tau \in S_k$ be a permutation such that $i_{\tau(1)} > i_{\tau(2)} > \ldots > i_{\tau(k)}$ and let $a_j := i_{\tau(j)}$ for $j = 1, 2, \ldots k$. Let $f_{a_1,a_2,\ldots a_k} := e_{i_1,i_2,\ldots i_k}$. Notice that $a_1 > a_2 > \ldots > a_k$ and that the tuple $(a_1, a_2, \ldots a_k)$ is reducible. Thus there exists $1 \leq s < t \leq k$ such that $a_s > a_{s+1} + (t-s) > a_t > a_{t+1} + (t-s+1)$. Consider the element $\eta \in FS_n$ given by

$$\eta := \Pi_{j=0}^{t-s}(a_t + j, a_{t+1} + j + 1)\Pi_{j=1}^{t-s}(1 - (a_{s+j}, a_t + j))$$

Now a careful calculation shows that $\eta(v) \in U$ has leading term that divides (and is strictly smaller than) the term $e_{i_1,i_2,\ldots i_k}$. Thus v cannot belong to a generating basis (for U or any other submodule). ♣

Lemma(4): Assume $a_1 > a_2 > \ldots > a_k$, assume $n \geq 6k$ and assume that $\{a_1, a_2, \ldots a_k\} \subseteq \{1, 2, \ldots n\}$ is irreducible. Then there exists $0 \leq s \leq k$ such that $\{a_1, a_2, \ldots a_k\} \subseteq \{1, 2, \ldots 3k - 3s\} \cup \{n, n-1, \ldots n-3s+1\}$. In general $\{a_1, a_2, \ldots a_k\} \subseteq \{1, 2, \ldots 3k\} \cup \{n, n-1, \ldots n-3k+1\}$. Any vector in a generating basis has normalised leading term $e_{i_1,i_2,\ldots i_k}$ with all index $i_1, i_2, \ldots i_k \in \{1, 2, \ldots 3k\} \cup \{n, n-1, \ldots n-3k+1\}$.

Proof: The first part is shown using induction on k. The second part follows by combining the first part with lemma(3) ♣

3.4 Background in logic

In this section we assume that the reader is familiar with basic concepts and ideas from basic logic and model theory. Good standard references are [2, 3] . It should, however, be noticed that any reference to logic can be easily avoided. Also, it should be noticed that we only need quite an easy special case of the main theorem in this section.

Lemma(5): Let $\alpha_0, \alpha_1, \alpha_2, \ldots, \alpha_{l-1}, \alpha_l \in \{0, 1, 2, 3 \ldots\}$ be $l+1$ given numbers. Consider the $l+1$ equations given by $\alpha_0 < x_1$, $x_1 + \alpha_1 < x_2$, $x_2 + \alpha_2 < x_3$, $\ldots, x_j + \alpha_j < x_{j+1}, \ldots, \ldots x_{l-1} + \alpha_{l-1} < x_l$ and $x_l + \alpha_l < n+1$. Let $s := \Sigma_{j=0}^{l} \alpha_j$. Over the structure $M_n = \{1, 2, \ldots n\}$ with $n \geq s + l - 1$ the number of solutions to the equations is given by the binomial expression $\begin{pmatrix} n-s \\ l \end{pmatrix}$. For a fixed list $\alpha_0, \alpha_1, \alpha_2, \ldots, \alpha_{l-1}, \alpha_l$ of numbers, the number of solutions is given by a polynomial in n.

For fixed $\alpha \in \{0, 1, 2, \ldots\}$ we let $<_\alpha$ denote the relation defined such that $x <_\alpha y$ if and only if $x + \alpha < y$. First we will show that the theory of a discrete linear order with end points (expressed in the first order language $L(<_0, <_1, <_2, \ldots)$) has quantifier elimination.

Let $L := L(\leq)$ be the first order language with one binary relation symbol $\leq$ (besides the relation symbol $=$ for equality). Let $(M, <)$ be a set and $<$ a total ordering on M. Thus $(M, <)$ is isomorphic to $\{1, 2, \ldots, |M|\}$ with the usual ordering between integers. Any $L(\leq)$-formula $\psi(x_1, x_2, \ldots, x_k)$ with free variables $x_1, x_2, \ldots, x_k$ determines naturally a set $A_{\psi,M} := \{(a_1, a_2, \ldots, a_k) : (M, <) \models \psi(a_1, a_2, \ldots, a_k)\}$.

Proposition: For each $\psi(x_1, x_2, \ldots, x_k)$ first order formula in the language $L(<_0, <_1, <_2, \ldots)$ there exists a *quantifier free* formula $\theta(x_1, x_2, \ldots, x_k)$ in the langage $L(<_0, <_1, <_2, \ldots)$ such that for all models M of a discrete linear ordering with endpoints the set $A_\psi = \{(a_1, a_2, \ldots, a_k) : M \models \psi(a_1, a_2, \ldots, a_k)$ equals the set $A_\theta := \{(a_1, a_2, \ldots, a_k) : M \models \theta(a_1, a_2, \ldots, a_k)$.

Combining this proposition with lemma(5) (and using the inclusion-exclusion principle) we get:

Theorem(FO): Let $\psi(x_1, x_2, \dots x_k)$ be first order formula expressed in the language $L(\leq)$, with free variables $x_1, x_2, \dots, x_k$. Then there exists a polynomial $p \in Q[x]$, such that the number of tuples $(a_1, a_2, \dots a_k)$ with $a_1, a_2, \dots a_k \in \{1, 2, \dots n\}$ and $M_n \models \psi(a_1, a_2, \dots . a_k)$ is given by $p(n)$ for all $n \in N \setminus E$ where E is a finite set of exceptions. The set E is empty if ψ is a quantifier free formula.

Example: Consider the formula $\psi(a_1, a_2) :\equiv (a_1 > a_2) \wedge (\forall x_1, x_2 \exists y(y \neq x_1 \wedge y \neq x_2 \wedge y \neq a_1 \wedge y \neq a_2)$. The number $A(n)$ of $(a_1, a_2) \in \{1, 2, \dots n\} \times \{1, 2, \dots n\}$ that have $\psi(a_1, a_2)$ is given by $A(n) = \frac{n(n-1)}{2}$ for $n = 5, 6, 7 \dots$ while $A(1) = A(2) = A(3) = A(4) = 0$. The function $A(n)$ with $n = 5, 6, 7 \dots$ is given by the polynomial $p(n) = \frac{n(n-1)}{2}$. However, the polynomial p does not count the number of solutions correct on the finite set $E = \{1, 2, 3, 4\}$. So, as "predicted" by the theorem, $A(n)$ is given by a polynomial together with a finite set E of exceptional values of $A(n)$.

Fix a natural number $k \in N$ and add: constants $c_1, c_2, \dots, c_{3k}$ and $d_0, d_1, \dots, d_{3k-1}$ to the language $L(\leq)$. We denote this new language $L(\leq, c_1, \dots c_{3k}, d_0, \dots, d_{3k-1})$ by L'.

Let Θ denote a set of axioms that insures that $c_1, c_2, \dots, c_{3k}$ are the $3k$ smallest elements, and ensures that $d_0, d_1, \dots, d_{3k-1}$ are the $3k$ largest elements. Include axioms that ensure that $c_1 < c_2 < \dots < c_{3k} < d_{3k-1} < \dots < d_1 < d_0$. Thus if $M = \{1, 2, \dots, n\}$ the axioms force $c_j = j$ and force $d_j = n - j$.

Lemma(5): Let $\psi(x_1, \dots x_k)$ be a quantifier free L'-formula. There exists a polynomial $p(x) \in Q[x]$ such that the set $A_{\psi,n} := \{(a_1, a_2, \dots, a_k) : (\{1, 2, \dots n\}, <) \models \Theta \wedge \psi(a_1, a_2, \dots, a_k)\}$ has cardinality $p(n)$ for all $n > 6k$.

We are now ready to combine these results and provide another proof of theorem(1). To this end we need to combine lemma(4) with:

Lemma(6): Let $U \subseteq V_{F,n}^{o,k}$ be any FS_n-submodule. Then the set $NLt(U)$ of normalised leading terms $e_{i_1,i_2,\dots,i_k}$is definable by a quantifier free L'-formula $\psi(x_1, \dots x_k)$. More specifically, there exists a quantifier free L'-formula $\psi(x_1, \dots x_k)$ such that $e_{i_1,i_2,\dots,i_k} \in NLt(U)$ if and only if $\psi(i_1, i_2, \dots, i_k)$.

Combining all of these we get:

Theorem (version 4): Let $k \in N$. There exists a finite set $\Gamma_k \subseteq Q[x]$ of polynomials such that: for each field F and for each $n \in N$ (with $n \geq k$) and each S_n-invariant linear subspace $U \subseteq V_{F,n}^{o,k}$ there exists $p \in \Gamma_k$ such that $dim(U) = p(n)$.

Combining this notice that the set $\Gamma_k \subseteq Q[x]$ consisting of the class of polynomials that appear as $p(n) := |A_{\psi,n}|$ for one of the finitely many inequivalent L'-formula $\psi(x_1, x_2, \ldots x_k)$ suffice as set in the theorem (version 4). In the theorem (version 4, or any other version), we did not make any assumptions about n (except that $n \geq k$) while in lemma 6 we assumed that $n \geq 6k$. We can mend this by, for example, adding the finite set $\{p_c\}_{0 \leq c \leq dim(V_{F,n}^{o,k})}$of constant polynomials with $p_c(x) := c$ for $c = 0, 1, \ldots 6k(6k-1)\ldots(5k+1)$.

4 Equivalence of the different versions of the Theorem

So far we have presented four versions of the theorem. In this section we will show that they are all equivalent in the sense that they can be (easily) derived from each other. The content of the four versions are however distinct and the general version of the theorem has a much broader scope than the other versions. From version 4 we easily get the following version:

Theorem (version 5): Let $k, r \in N$. There exists a finite set $\Gamma_{k,r} \subseteq Q[x]$ of polynomials such that: for each field F, and for each $n \in N$ (with $n \geq k$) and each S_n-invariant linear subspace $U \subseteq V_{F,n}^{o,k} \oplus \ldots \oplus V_{F,n}^{o,k}$ (r copies) there exists $p \in \Gamma_{k,r}$ such that $dim(U) = p(n)$.

Version 4 implies version 5: We prove this using induction on $r \in N$. For $r = 1$ version 4 or 5 are clearly equivalent. For $r > 1$ let $\Gamma_{k,r} := \Gamma_k \oplus \ldots \oplus \Gamma_k$ (r copies). Let $U \subseteq V_{F,n}^{o,k} \oplus \ldots \oplus V_{F,n}^{o,k}$ be any FS_n-submodule. Let $\psi : V_{F,n}^{o,k} \oplus \ldots \oplus V_{F,n}^{o,k} \to V_{F,n}^{o,k}$ be the projection on the first summand. Now the induction assumption applies to $\psi(U) \subseteq V_{F,n}^{o,k}$ as well as $Ker(\psi) \cap U \subset 0 \oplus V_{F,n}^{o,k} \oplus \ldots \oplus V_{F,n}^{o,k}$ and thus there exists $p_1 \in \Gamma_k$ and $p_2 \in \Gamma_k \oplus \Gamma_k \oplus \ldots \oplus \Gamma_k \equiv \Gamma_{k,r-1}$ such that $p_1(n) = dim(\psi(U))$ and $p_2(n) = dim(Ker(\psi) \cap U)$. Now $dim(U) = dim(\psi) + dim((Ker(\psi) \cap U)) = p_1(n) + p_2(n)$ and so since $p := p_1 + p_2 \in \Gamma_{k,r}$ we have $dim(U) = p(n)$.

Version 5 implies the general version: Let $\Gamma_{k,r} \subset Q[x]$ in the general version of the theorem consist of $\Gamma_{k,r} \ominus \Gamma_{k,r}$. In the general version of the theorem we define for $k, r \in N$, the class $M_{F,k,r}$ of all FS_n-modules that can be generated by r (or fewer) generators, that are all invariant under the action of $S_{\{k+1,k+2,\ldots n\}}$. Let $V \in M_{F,k,r}$ be any FS_n-module and assume that it is generated by generators $g^{(1)}, g^{(2)}, \ldots, g^{(r)} \in V$ that each is $S_{\{k+1,k+2,\ldots,n\}}$ invariant. Next, consider $M := V_{F,n}^{o,k} \oplus \ldots \oplus V_{F,n}^{o,k}$ (r copies). Let $\{e_{i_1,i_2,\ldots,i_k}^{(v)} : v = 1, 2, \ldots r$ and $i_1, i_2, \ldots, i_k \in \{1, 2, \ldots n\}$ with $|\{i_1, i_2, \ldots, i_k\}| = k\}$ be the natural basis for M, and assume that S_n acts on M in the usual way. We define a FS_n-homomorphism $\psi : M \to V$ by letting $\psi(e_{1,2,3,\ldots,k}^{(v)}) = g^{(v)}$ for $v = 1, 2, \ldots, r$. This map is a well defined FS_n-homomorphism because each

of $g^{(1)}, g^{(2)}, \ldots, g^{(r)} \in V$ is $S_{\{k+1,k+2,\ldots,n\}}$-invariant. Notice that ψ is surjective (since $g^{(1)}, g^{(2)}, \ldots, g^{(r)}$ generates V).

Let $U \subseteq V$ be any FS_n-submodule. Then the FS_n-submodule $\psi^{-1}(U) \subseteq M$ has dimension $dim(\psi^{-1}(U)) = p(n)$ for some $p \in \Gamma_{k,r}$ and $dim(Ker(\psi)) = q(n)$ for some $q \in \Gamma_{k,r}$. But then $dim(U) = dim(\psi^{-1}(U)) - dim(Ker(\psi)) = p(n) - q(n)$. All that is left is to notice that the polynomial $p - q \in \Gamma_{k,r} \ominus \Gamma_{k,r}$.

Version 3 implies version 4: Clear since $M^{(n-k,1^k)}$ over the field F is isomorphic to $V^{o,k}_{F,n}$. **Version 2 implies version 4:** Clear since $V^{o,k}_{F,n}$ is a FS_n-submodule of $V^k_{F,n}$.

General version implies version 2,3,4 and 5: Each of the FS_n-modules $V^k_{F,n}$, $M^{(n-k,\alpha)}$ with $k = |\alpha|$, $V^{o,k}_{F,n}$ and $V^{o,k}_{F,n} \oplus \ldots \oplus V^{o,k}_{F,n}$ is generated by a finite number $c \in N$ of $S_{\{k+1,k+2,\ldots,n\}}$-invariant generators (that can be chosen independent of n). This completes the proof that the five versions of the theorems are "equivalent".

5 Acknowledgements

I would like to thank Peter Cameron for inviting me to present parts of this work in his combinatorics seminar at Queen Mary, University of London.

References

[1] M Ajtai. Symmetric systems of linear equations modulo p. Technical Report TR95-015, Electronic Colloquium, 1994. http://www.eccc.uni-trier.de/eccc.

[2] J. Barwise. An introduction to first-order logic. In Barwise, editor, *Handbook of Mathematical Logic*, pages 5–46. North-Holland, Amsterdam.

[3] W Hodges. *Model Theory*. Cambridge University Press, 1993.

[4] G. James. *The representation theory of the symmetric groups*, volume 682 of *Lecture notes in Mathematics*. Spinger-Verlag, 1978.

[5] J. Krajicek. Uniform families of polynomial equations over a finite field and structures admitting an euler characteristic of definable sets. *Proc. London Mathematical Society*, 81(3):257–284, 2000.

[6] S. Riis. Classifications of fs_n-submodules of $v_f, n \otimes v_f, n$. Technical report, Queen Mary, University of London, 2005.

[7] S. Riis and M. Sitharam. Uniformly generated submodules of permutation modules. *Journal and of Pure and Applied algebra*, 160:285–318, 2001.