BRICS

Basic Research in Computer Science

Uniformly Generated Submodules of Permutation Modules

Søren Riis
Meera Sitharam

BRICS Report Series RS-98-20

ISSN 0909-0878 September 1998

See back inner page for a list of recent BRICS Report Series publications.
Copies may be obtained by contacting:

BRICS
Department of Computer Science
University of Aarhus
Ny Munkegade, building 540
DK–8000 Aarhus C
Denmark
Telephone: +45 8942 3360
Telefax: +45 8942 3255
Internet: BRICS@brics.dk

BRICS publications are in general accessible through the World Wide Web and anonymous FTP through these URLs:

`http://www.brics.dk`
`ftp://ftp.brics.dk`
This document in subdirectory `RS/98/20/`

Uniformly Generated Submodules of Permutation Modules

Søren Riis[*][†]
Meera Sitharam [‡][§]

September 1998

Abstract

This paper is motivated by a link between algebraic proof complexity and the representation theory of the finite symmetric groups. Our perspective leads to a series of non-traditional problems in the representation theory of S_n.

Most of our technical results concern the structure of "uniformly" generated submodules of permutation modules. We consider (for example) sequences W_n of submodules of the permutation modules $M^{(n-k,1^k)}$ and prove that if the modules W_n are given in a uniform way - which we make precise - the dimension $p(n)$ of W_n (as a vector space) is a single polynomial with rational coefficients, for all but finitely many "singular" values of n. Furthermore, we show that $\dim(W_n) < p(n)$ for each singular value of $n \geq 4k$. The results have a non-traditional flavor arising from the study of the irreducible structure of the submodules W_n *beyond* isomorphism types.

We sketch the link between our structure theorems and proof complexity questions, which can be viewed as special cases of the famous NP vs. co-NP problem in complexity theory. In particular, we focus on the efficiency of proof systems for showing

[*] The International PhD Research School at BRICS, Aarhus, Denmark; Email: smriis@daimi.aau.dk

[†] Part of this work was done while visiting the Fields Institute, Toronto, Canada

[‡] CISE Department, University of Florida, Gainesville, FL 32611-6120; Email: sitharam@cise.ufl.edu

[§] Supported in part by NSF Grant CCR 94-09809.

membership in polynomial ideals, for example, based on Hilbert's Nullstellensatz.

Keywords: Finite Groups, Representation Theory of the Symmetric Group, Polynomial Ideals, Algebraic Proof Complexity Lower Bounds, Complexity Theory.

Subject Classification: 20C30, 05E10, 68Q15, 13Cxx

I Introduction and Motivation

Consider the question whether there exists a proof of the Riemann conjecture which uses less than k printed pages? Or consider the same question for the Poincare conjecture? This kind of question is not only well-defined (if the "proof" is within some fixed axiomatization of ZFC), but may seem trivial in the sense that it only involves checking finitely many possibilities. I.e, it is a so-called finite decision problem, and in that sense, is no different in character than asking: is there a group of order n with a specific algebraic property? However, we can now ask whether this *search* - for a proof of length n in ZFC for varying input conjectures, and varying values of n, or for a group of order n with a well-defined algebraic property - can be carried out feasibly by a computer. This can be seen as a version of the famous P vs. NP question. This and other questions about the complexity of finite decision problems play a substantial role in the foundations of contemporary computer science. Moreover, they are generally considered among the deepest mathematical problems for the next century (see, for example, [16]).

I.1 Hilbert's Nullstellensatz and Algebraic Proofs

All finite decision problems in NP (not just the earlier example about ZFC proofs) require decisions about the existence of short "proofs," in an elementary proof system. These proofs are not to be confused with the ZFC proofs in the example, and are alternatively also called "easily checkable witnesses, or certificates". As a result, the study of lengths and complexity of proofs in elementary proof systems is draw considerable motivation from another famous problem: the NP vs. co-NP problem. In terms of the examples given above, one version of this problem is to ask whether there is a short proof - in an appropriate proof system - of the *non*-existence of a group of order n with some algebraic property, or of the fact that a ZFC proof of size n does *not* exist for an input conjecture.

One class of proof systems that are studied in this context are the so-called algebraic proof systems. Such systems have been studied intensively within recent years. The systems we will consider was first introduced in [4]. These systems arise from the following observation. All *NP* decision problems can be phrased as deciding the existence of 0/1 solutions to systems of (multilinear) polynomial equations. As in the examples given earlier, if the decision problems are parametrized by n, then the resulting polynomial systems are also parametrized by n. We can think of $\bar{Q}_n$ as, for example, the finite system of polynomial equations corresponding to the question about the existence of groups of size n with some algebraic property. If we include the polynomials $x^2 - x$ in $\bar{Q}_n$ (one for each variable x), we see (as also observed in [4]) that $1 \in (\bar{Q})_n$ if and only if there is no group of size n possessing a specific algebraic property.

This suggests (and this was indeed suggested in [4]) that we consider elementary, algebraic proof systems designed for proving ideal membership. As mentioned earlier, an elementary proof system should provide easily checkable certificates witnessing the fact being proved. One natural way of witnessing ideal membership of a polynomial R in the ideal generated by the polynomials $Q_1, Q_2, \ldots, Q_l$, denoted $(Q_1, Q_2, \ldots, Q_l)$, is to provide a list of multiplying polynomials $P_j, j \in \{1, 2, \ldots, l\}$ such that $\Sigma_{j=1}^{l} P_j Q_j = R$. Such a list of polynomials constitute what is now called a *Nullstellensatz Proof (NS-proof)* of $R \in (Q_1, Q_2, \ldots, Q_l)$. The complexity of the proof is reflected in the size/degree of the polynomials $P_j, j \in \{1, 2, \ldots, l\}$. See also [5] for bounds on this degree. The degree of the NS-proof is usually defined as the maximal degree of the polynomials $P_j, j \in \{1, 2, \ldots, l\}$. This proof system is too weak for results about NS-proof complexity to have any direct impact on the *NP* vs. co-*NP* problem. Other related algebraic proof systems (for example the so-called *Polynomial Calculus* proof system) are in general preferable, and can be shown to be stronger than NS-proofs. Although results of this paper are applicable to most algebraic proof systems, inorder to illustrate our main points it suffices to focus on NS-proofs.

It should be mentioned that another important reason for studying algebraic proof systems is that many automated theorem provers are based on some elementary proof system for proving ideal membership, and there seems little doubt that computer assisted proofs will play a considerable role in future mathematics.

I.2 Link to Symmetric Group Representations

The link to the Representation theory is heavily inspired (but technically independent of) the pioneering work by M. Ajtai [1], [2] and [3]. Our paper is also strongly motivated by an earlier result by the authors in [14], which considers a large class of finite decision problems which includes all of the examples given earlier. These problems have the form: “is there a model or finite structure of size n satisfying a given existential second order sentence ψ ?” Hence it is natural to study the algebraic proof complexity of showing nonexistence of models of size n satisfying this type of sentence ψ.

Furthermore, a translation method developed in [14] shows a 1-1 correspondence between the models of ψ of size n and 0/1 points in special algebraic varieties $V_{n,\psi}$, given by systems of polynomial equations $\bar{Q}_{n,\psi}$, which are closed under the action of the symmetric group S_n and, moreover, are uniformly given in n. While we shall not dwell on this 1-1 correspondence here, it should be emphasized that it is sufficiently direct that one can read off the models from the 0/1 points on the variety $V_{n,\psi}$.

To study the complexity of algebraic proofs showing nonexistence of models of size n for ψ, as discussed in the last subsection, one can study for example, the degree of Nullstellensatz multiplying polynomials that witness that the constant function 1 belongs to the ideal $(\bar{Q}_{n,\psi})$. Now, since the variety $V_{n,\psi}$ is closed under the action of S_n, so is the ideal $(Q_{n,\psi})$. This, not surprisingly, affects the degree of Nullstellensatz multiplying polynomials or indeed the complexity of any algebraic proof of $1 \in (Q_{n,\psi})$, and thereby closely links algebraic proof complexity questions to natural questions about symmetric group representations that are of independent interest. Most of this paper directly addresses these latter representation theory questions, although their bearing on algebraic proof complexity issues is briefly sketched in Section VII.

Note: Since the motivating application of our results concerns polynomial ideals (closed under the action of the finite symmetric groups), we find it natural to use the language of polynomial rings to phrase all of our results on S_n representations. Hence, for example, permutation modules and their submodules will be viewed as consisting of polynomials from certain polynomial rings ♣

I.3 Brief Summary of Results

In this section, we present a series of theorems that illustrate the flavor of the technical results in the paper. Readers unfamiliar with the terminology used in the representation theory of S_n may refer to Section II and [9].

Fix a field $\mathbb{F}$ of characteristic 0. For each $n \in \mathbb{N}$, consider the space $\Pi_{n,d}$ of polynomials of degree at most d in the ring $\mathbb{F}[x_{11}, x_{12}, \ldots, x_{1n}, x_{21}, \ldots, x_{nn}]$, i.e, $\mathbb{F}[x_{ij} : 1 \leq i, j \leq n]$. For convenience, usually, we first state and prove results for the larger vector space $\mathcal{V}_{n,d}$ of formal polynomials of degree $\leq d$. In a formal polynomial, monomials like $x_{ij}x_{kl}$ and $x_{kl}x_{ij}$ are considered distinct.

We let the symmetric group S_n act on $\mathcal{V}_{n,d}$ in the natural way. If, for example, $P := x_{12}x_{34} - 3x_{23} + 1$ and $\pi \in S_n$ we let $\pi(P) = x_{\pi(1)\ \pi(2)}x_{\pi(3)\ \pi(4)} - 3x_{\pi(2)\ \pi(3)} + 1$. In other words, we can consider $\mathcal{V}_{n,d}$ as an $\mathbb{F}S_n$-module.

Recall that a $\mathbb{F}S_n$-submodule of $\mathcal{V}_{n,d}$ is a linear subspace $W \subseteq \mathcal{V}_{n,d}$ which is closed under S_n. In this paper, we will mainly be concerned with such $\mathbb{F}S_n$-submodules. Notice that $\Pi_{n,d}$ is a quotient $\mathbb{F}S_n$-module of $\mathcal{V}_{n,d}$, obtained by identifying formal monomials (like $x_{ij}x_{kl}$ and $x_{kl}x_{ij}$) which defines the same monomial. First we show (using standard results from the representation theory of the symmetric group):

Theorem 1A: *For any $d \in \mathbb{N}$, there exists a finite collection A_d of functions $f : \mathbb{N} \to \mathbb{N}$ such that for any n and any $\mathbb{F}S_n$-submodule $W \subseteq \mathcal{V}_{n,d}$, (or $\subseteq \Pi_{n,d}$), there is $f \in A_d$ such that the dimension of W (as a linear vector space) is given by $f(n)$.*

Furthermore for any $d \in \mathbb{N}$, all the functions f in A_d are actually polynomial functions with rational coefficients.

Corollary: *Let $W_n \subseteq \mathcal{V}_{n,d}$ (or $\subseteq \Pi_{n,d}$) be an arbitrary sequence of submodules. Then there exists an infinite set $B \subseteq \mathbb{N}$ and a single polynomial function $p \in \mathbb{Q}[z]$ such that* $\dim(W_n) = p(n)$ *for all $n \in B$.*

Theorem 1A expresses two remarkable facts: (1) there exists a constant C_d such that for any n, the linear subspaces $W \subseteq \mathcal{V}_{n,d}$ (or $\subseteq \Pi_{n,d}$) which are closed under the action of S_n have at most C_d different vector space dimensions, (2) these C_d different dimensions can be given as polynomials in n. We note that C_d grows super-exponentially in d. For example, C_1 is 64, and a rough estimate shows (see below) that C_2 is somewhere between $10,000,000$ and $20,000,000,000$.

In general there are infinitely many different linear subspaces which have W_n closed under the action of S_n. There are for example infinitely many different linear subspaces W_n of polynomials of degree ≤ 2 (in variables $x_{11}, x_{12}, \ldots, x_{1n}, x_{21}, \ldots, x_{nn}$) which have W_n closed under the action of S_n (for more details see the example in section IV, which shows this indeed is the case for $n \geq 8$). Theorem 1A says that there are only finitely many (as it turns out at most $20,000,000,000$) different choices of vectorspace dimensions for W_n. The linear spaces W_n can thus typically be "rotated" in infinitely many different ways.

Next we consider *formal expressions* obtained by formal sums over $\mathcal{V}_{n_0,d}$, for some fixed n_0, for example: $P_{exp} := 1 + \sum_{j=1} x_{1j} + 3 \sum_{i=1} \sum_{j=1} x_{2i}x_{j5}$. In this example n_0 is at least 5 because a monomial like x_{15} must belong to $\mathcal{V}_{n_0,d}$. The expression allows us to define a sequence of polynomials given by the expression:

$$P_n := (P_{exp})_n := 1 + \sum_{j=1}^{n} x_{1j} + 3 \sum_{i=1}^{n} \sum_{j=1}^{n} x_{2i}x_{j5},$$

for any $n \geq 5$ (or $\geq n_0$ in general). We say the expression P_{exp} has *support* $\{1, 2, 5\}$, i.e $\{1, 2, 5\}$ are the describing indices in the expression.

The *support size* of P_{exp} is $3 = |\{1, 2, 5\}|$. We call a formal expression P_{exp} *ultrasmall* if it has support size at most $4d$. Later, we extend this definition of ultrasmall to other spaces than $\mathcal{V}_{n,d}$ (and $\Pi_{n,d}$). An element (here a polynomial) $E \in \mathcal{V}_{n,d}$ is called *ultrasmall* if there exists an ultrasmall formal expression P_{exp} such that $E = P_n$. Notice that for $n > 4d$, an ultrasmall element (polynomial) $E \in \mathcal{V}_{n,d}$ has a unique ultrasmall formal expression P_{exp} such that $E = P_n$. When it is clear from the context, sometimes we refer to the support size of P_{exp} also as the support size of E.

Theorem 2A: *Every submodule $W \subseteq \mathcal{V}_{n,d}$ (or $\subseteq \Pi_{n,d}$) is generated as an $\mathbb{F}S_n$-submodule by a collection of ultrasmall expressions.*

Furthermore the ultrasmall expressions can be chosen such that each of them generates an irreducible submodule.

The significance of Theorem 2A lies in the fact that it clarifies the structure and decomposition of $\mathbb{F}S_n$-modules *beyond* isomorphism types. It follows from existing decomposition theorems, Jordan-Hölder's Theorem, and the fact that the modules we consider in this paper all are semi-simple (when $\mathbb{F}$ has characteristic 0) that

1. every $\mathbb{F}S_n$-submodule can be uniquely (up to isomorphism) decomposed into a direct sum of irreducible modules (isomorphic to the so-called Specht modules);
2. each Specht module is (independent of any field characteristic) generated cyclically by a so-called polytabloid.

The polytabloids generating the Specht modules have ultrasmall support size (when defined in the obvious way). However, it should be noted that since an isomorphism may not, in general, preserve the property of being generated by ultrasmalls, it is not clear whether the *actual* irreducibles in the decomposition are themselves generated by ultrasmalls. All we know from the general theory is that each irreducible is *isomorphic* to an object which can be defined by very few (i.e. $\leq 4d$) parameters. Theorem 2A shows that each irreducible submodule is not only isomorphic to a submodule generated by ultrasmall generators (which follows from the general theory), but that each irreducible submodule itself is generated by ultrasmall objects. We clarify this point further using an Example in Section III.

Now consider the case where we are given a *uniform* sequence $W_n \subseteq \mathcal{V}_{n,d}$ of $\mathbb{F}S_n$-submodules. The word "uniform" is used here in an informal sense. Intuitively, this means that each W_n only depends on n in a straightforward manner. We could, for example, define the sequence W_n by letting W_n denote the smallest $\mathbb{F}S_n$-module which contains a given finite list of ultrasmall elements $(E_1)_n, \ldots, (E_v)_n$. For example, the sequence W_n of $\mathbb{F}S_n$-modules generated by $E_n := 1 + \sum_{j=1}^{n} x_{1j} + 3\sum_{i=1}^{n}\sum_{j=1}^{n} x_{2i}x_{j5}$ is given in a uniform way. Later in the paper, we give a precise definition of different methods of generating uniform sequences of modules.

From Theorem 1A, we know that there exists a finite collection of polynomials A_d such that for each $n \in \mathbb{N}$ there exists $p \in A_d$ such that $\dim(W_n) = p(n)$. If the family W_n is given in a uniform way (which we later will define), it is tempting to conjecture that there is a single polynomial $p \in A_d$ which expresses the dimension of W_n for *all* $n \geq 8d$. Later, we give examples showing that this is not true in general. However, we show:

Theorem 4A: *Let $W_n \subseteq \mathcal{V}_{n,d}$ (or $\subseteq \Pi_{n,d}$) be a uniformly generated sequence of $\mathbb{F}S_n$-submodules. Then there exists a single polynomial $p \in \mathbb{Q}[z]$ and a finite set $B \subseteq \mathbb{N}$ such that*

(1) $\dim(W_n) = p(n)$ *for all $n \in \mathbb{N} \setminus B$.*

(2) $\dim(W_n) < p(n)$ *for all $n \in B$ for which $n \geq 8d$.*

In the process of proving this result, we show various uniform versions of Theorem 2A. In particular, we employ the notion of a *generalized formal expression* over $\mathcal{V}_{n_0,d}$, for a fixed n_0. Such expressions are formal expressions which have coefficients in the field $\mathbb{F}(x)$ of rational functions over $\mathbb{F}$, instead (as formal expressions) of have coefficients in the field $\mathbb{F}$. For example, the expressions $T_{gen} := (z^2-3z+4)\sum_i\sum_j x_{ij}x_{j3} - (z^3+7z^2-3z+2)\sum_j x_{j5} + 3zx_{14}$ and $E_{gen} := 17\sum_i x_i + z\sum_j y_j$ are both generalized formal expressions. The support size of T_{gen} is $4 = |\{1,3,4,5\}|$ (which is smaller than $4d = 8$) and the support size of E_{gen} is 0, hence they are both *generalized ultrasmall expressions.*

Theorem 3A: *Let $W_n \subseteq \mathcal{V}_{n,d}$ (or $\subseteq \Pi_{n,d}$) be a uniformly generated family of $\mathbb{F}S_n$-submodules. Then there exists a fixed set Γ_{gen} (independent of n) of generalized ultrasmall expressions such that the corresponding generalized ultrasmall elements in Γ_n generate W_n, for all $n \geq 8d$. Furthermore, each generalized ultrasmall in Γ_{gen} for each value of $n \geq 8d$ is either zero or generates an irreducible module.*

Moreover, for each generalized ultrasmall element $E \in \Gamma_{gen}$ there exists a fixed partition β such that each E_n (for $n \geq 8d$) either is zero, or generates an irreducible module which is isomorphic to the Specht module $S^{(n-|\beta|,\beta)}$.

The height of the module W_n (i.e. the number of irreducible factors) is a fixed constant C for n sufficiently large. The height of W_n is bounded by C from above for all values of $n \geq 8d$. For certain singular values of

n the height of W_n might drop (i.e. take a value strictly less than C) however there are only finitely many such singular values.

Essentially combining Theorem 3A and Theorem 4A we obtain corollaries that are useful for proving algebraic proof complexity gaps and bounds. For example:

Corollary: *If a uniformly generated module sequence W_n is irreducible for some sufficiently large n, then W_n is irreducible for all $n \geq 8d$. Moreover, there exists a fixed partition β with $|\beta| \leq 2d$ such that for each $n \geq 8d$ W_n is either zero or is isomorphic to the Specht module $S^{(n-|\beta|,\beta)}$.*

Corollary: *If a uniformly generated module sequence W_n is strictly contained in the entire module $\mathcal{V}_{n,d}$ for sufficiently large n, then it is not equal to $\mathcal{V}_{n,d}$ for any $n \geq 8d$.*

In a later section, we sketch the link between these results and algebraic proof complexity. To strengthen this link, we consider more general methods of defining uniform sequences, with similar results. Other methods give dual results. For example, the sequence V_n defined by $V_n := W_n^{\perp}$, where W_n is a uniformly generated sequence (in the sense we just considered), is *not* a uniformly generated sequence in general. However the sequence V_n satisfies the obvious dual versions of Theorem 3A and Theorem 4A where the height (as well as the vector space dimension) might increase (rather than drop) at singular values of n. In [15], we use these results to obtain a new class of theorems that provide gaps and lower bounds on algebraic proof complexity of propositional formulae.

II Background on Finite Symmetric Group Representations

Let $M^{(n-k,1^k)}$ be the permutation module from the representation theory of the symmetric group [9]. Recall that this $\mathbb{F}S_n$-module is the vector space over $\mathbb{F}$ spanned by tabloids for the partition: $(n-k,1,1,\ldots,1)$, with k one's, written as $(n-k,1^k)$. In general, there is a permutation module M^λ associated with each partition $\lambda = (\lambda_1, \lambda_2, \ldots)$ which satisfies $\sum_i \lambda_i = n$ and $\lambda_1 \geq \lambda_2 \geq \ldots$; and the diagram $[\lambda]$ is $\{\lambda_{ij} : i,j \in \mathbb{Z}, 1 \leq i, 1 \leq j \leq \lambda_i\}$; a row (or column) of the diagram corresponds to fixing i (or j). A λ-tableau t is one of the $n!$ lists $L_1, L_2, \ldots$ of ordered subsets of $\{1,\ldots,n\}$, with $|L_i| = \lambda_i$; and a λ-tabloid $\{t\}$ is an equivalence class of λ-tableaux obtained by viewing the L_i as unordered subsets. There are $n(n-1)(n-2)\ldots(n-k+1)$ tabloids for the partition $(n-k,1^k)$, with $(n-k)!$ tableaux associated with each tabloid, and S_n acts on $M^{(n-k,1^k)}$ in the natural way (see [9]). There is a useful dominance (partial) ordering $\trianglerighteq$ on partitions: $\lambda \trianglerighteq \mu$ provided, for all m, $\sum\limits_{l=1}^{m} \lambda_l \geq \sum\limits_{l=1}^{m} \mu_l$.

The permutation module $M^{(n-k,1^k)}$ can be viewed as the vector space spanned by the vectors $\{e_{i_1,i_2,\ldots,i_k} : i_1, i_2, \ldots, i_k \in \{1,2,\ldots,n\}$ distinct$\}$. The action of a permutation $\pi \in S_n$ is given by: $\pi(e_{i_1,i_2,\ldots,i_k}) := e_{\pi(i_1),\pi(i_2),\ldots,\pi(i_k)}$.

For any partition λ (except $\lambda = (n)$), and for any field $\mathbb{F}$ of any characteristic, the permutation module M^λ is reducible and can be written as a Specht series whose factors are isomorphic to the Specht modules S^β, each of which is also associated with a partition β and is cyclically generated by a so-called polytabloid associated with a β-tableau. The multiplicity of isomorphic copies of a given Specht Module S^β in the Specht series of a given permutation module can be calculated by The Littlewood-Richardson rule or the Young rule [9]. In this paper, we only consider the case where the field $\mathbb{F}$ has characteristic 0, and in this case the Specht modules are irreducible [9], and hence the Specht series is in fact a composition series. Moreover, for characteristic 0, all modules we consider are semi-simple, and the Jordan-Hölder decomposition [8] is not just a composition series, but in fact a direct sum of irreducibles which is unique up to isomorphism. The total number of irreducibles in this direct sum is called the *height* of W. Next, we state three lemmas that will be used in the following sections. Lemma 1 is directly from [9], while Lemma 2 and Lemma 3 follow (by arguments given in the proof of Theorem 1B) from basic results in [9].

Lemma 1: *Let λ and μ be partitions of n. If $\lambda \not\trianglerighteq \mu$, then for any λ-tableau t, and any element f of S^μ, $\kappa_t f = 0$, where the signed column sum κ_t is the element of the group ring or group algebra $\mathbb{F}S_n$, obtained by summing over permutations that fix the columns of t, attaching the signature sign to each permutation. Furthermore, for $\lambda = \mu$, $\kappa_t f = +/- \kappa_t t$ is a polytabloid that generates S^λ. See [9] for the required definitions.*

It follows from the standard theory that the multiplicity of $S^{(n-k',m_1',m_2',\ldots)}$ in $M^{(n-k,m_1,m_2,\ldots)}$ is independent of n for $n \geq 2k$ (for more details see the proof of Theorem 1B). More specifically we have

Lemma 2: *Let α_n denote the partition $(n-k, n_2, \ldots, n_s)$ where $\sum\limits_{j=2}^{s} n_j = k$, and β_n denote the partition $(n-k', m_2, \ldots, m_s)$ where $\sum\limits_{j=2}^{s} m_j = k'$. Then the multiplicity* $\mathrm{Mult}(S^{\beta_n}, M^{\alpha_n})$ *of S^{β_n} in the decomposition of M^{α_n} is given by Young's rule as the number of semi-standard β_n-tableaux of type α_n (see [9]) and is independent of n for $n \geq 2k$.*

The dimension of each Specht Module S^{β_n}, for $\mathbb{F}$ of any characteristic, can be calculated by use of the hook formula: $\frac{n!}{\text{product of the hook lengths for}\beta_n}$ [9]. From this we get (see the proof of Theorem 1B for details):

Lemma 3: *Let β_n be defined as in Lemma 2. There exists a polynomial $p \in \mathbb{Q}[z]$ such that* $\dim(S^{\beta_n}) := p(n)$ *for all $n \geq 2k$.*

We will illustrate the latter two lemmas by an example which will additionally allow us to calculate the exact number of polynomials needed in A_1 and A_2 of Theorem 1A, as well as give the idea behind the proofs of Theorems 1A, 1B and 1C.

Example: Following the notation in [9], and employing the Littlewood-Richardson rule (or Young's rule), we use the equation $[n-2][1][1] = [n] + 2[n-1,1] + [n-2,1^2] + [n-2,2]$ to express the fact that $M^{(n-2,1^2)}$ decomposes into a direct sum of one isomorphic copy of $S^{(n)}$, two isomorphic copies of $S^{(n-1,1)}$, $S^{(n-2,1^2)}$ and one copy of $S^{(n-2,2)}$. Thus we obtain the following.

$[n-1][1] = [n] + [n-1,1]$

$[n-2][1][1] = [n] + 2[n-1,1] + [n-2,1^2] + [n-2,2]$

$[n-3][1][1][1] = [n] + 3[n-1,1] + 3[n-2,2] + 3[n-2,1^2] + 2[n-3,2,1] + [n-3,3] + [n-3,1^3]$

$[n-4][1][1][1][1] = [n] + 4[n-1,1] + 6[n-2,2] + 6[n-2,1^2] + 4[n-3,3] + 8[n-3,2,1] + 4[n-3,1^3] + [n-4,4] + 3[n-4,3,1] + 2[n-4,2^2] + 3[n-4,2,1^2] + [n-4,1^4]$

Using the hook formula we obtain:

$\dim(S^{(n)}) = 1$

$\dim(S^{(n-1,1)}) = n-1$

$\dim(S^{(n-2,2)}) = n(n-3)/2$

$\dim(S^{(n-2,1^2)}) = (n-1)(n-2)/2$

$\dim(S^{(n-3,3)}) = n(n-1)(n-5)/6$

$\dim(S^{(n-3,2,1)}) = n(n-2)(n-4)/3$

$\dim(S^{(n-3,1^3)}) = (n-1)(n-2)(n-3)/6$

$\dim(S^{(n-4,4)}) = n(n-1)(n-2)(n-7)/24$

$\dim(S^{(n-4,3,1)}) = n(n-1)(n-3)(n-6)/8$

$\dim(S^{(n-4,2^2)}) = n(n-1)(n-4)(n-5)/12$

$\dim(S^{(n-4,2,1^2)}) = n(n-2)(n-3)(n-5)/8$ and finally,

$\dim(S^{(n-4,1^4)}) = (n-1)(n-2)(n-3)(n-4)/24$

Now let us calculate A_1 from Theorem 1A. First, notice that we can write $\mathcal{V}_{1,n}$ as a direct sum of $M^{(n)}$, $M^{(n-1,1)}$ and $M^{(n-2,1^2)}$. These three sums arise from the constants, the elements of $\mathcal{V}_{1,n}$ spanned by x_{ii}, and the elements spanned by x_{ij} where $i \neq j$. This gives us a decomposition of

$\mathcal{V}_{1,n}$ into three isomorphic copies of $S^{(n)}$, three copies of $S^{(n-1,1)}$, and one copy each of $S^{(n-2,1^2)}$ and $S^{(n-2,2)}$. We take A_1 to consist of polynomials of the form:

$$p(n) = b_0 + b_1(n-1) + b_2(n-1)(n-2)/2 + b_3 n(n-3)/2$$

where $b_0, b_1 \in \{0,1,2,3\}$ and where $b_2, b_3 \in \{0,1\}$.

It follows using Jordan-Hölder's Theorem [8] that there is a unique decomposition of W as a direct sum of irreducible modules, and all the submodules of W are embedded (up to isomorphism) as the various partial sums of these irreducibles. Hence the polynomials in A_1 suffice to capture all submodule dimensions. We get an upper bound of $64 (= 4^2 \cdot 2^2)$ on the number of polynomials in A_1. An explicit check shows that all these 64 polynomials are distinct.

Now consider $\mathcal{V}_{2,n}$. This space can be written as a direct sum of $M^{(n)}$ (constant polynomials) two copies of $M^{(n-1,1)}$ (from the polynomials x_{ii} and $x_{jj}x_{jj}$), of 7 copies of $M^{(n-2,1^2)}$ (from x_{ij}, $x_{ii}x_{ij}, x_{ii}x_{ji}, x_{ij}x_{ii}, x_{ii}x_{jj}$, $x_{ij}x_{ij}$, and $x_{ij}x_{ji}$ where $i \neq j$), of 6 copies of $M^{(n-3,1^3)}$ (from $x_{ii}x_{jk}$, $x_{ij}x_{ik}, x_{ij}x_{ki}, x_{ji}x_{ik}, x_{ji}x_{ki}$, and $x_{jk}x_{ii}$ for i, j, k distinct) and finally one copy of $M^{(n-4,1^4)}$ (from $x_{ij}x_{kl}$ where i, j, k, l are distinct).

Thus we have a decomposition of $\mathcal{V}_{2,n}$ into
$[n] + 2[n-1][1] + 7[n-2][1][1] + 6[n-3][1][1][1] + [n-4][1][1][1][1]$
$= [n] + 2([n] + [n-1,1]) + 7([n] + 2[n-1,1] + [n-2,1^2] + [n-2,2]) +$
$6([n] + 3[n-1,1] + 3[n-2,2] + 3[n-2,1^2] + 2[n-3,2,1] + [n-3,3] +$
$[n-3,1^3]) + ([n] + 4[n-1,1] + 6[n-2,2] + 6[n-2,1^2] + 4[n-3,3]$
$+8[n-3,2,1] + 4[n-3,1^3] + [n-4,4] + 3[n-4,3,1] + 2[n-4,2^2] + 3[n-$
$4,2,1^2] + [n-4,1^4])$
$= 17[n] + 36[n-1,1] + 31[n-2,1^2] + 31[n-2,2] + 20[n-3,2,1] + 10[n-3,3]$
$+10[n-3,1^3] + [n-4,4] + 3[n-4,3,1] + 2[n-4,2^2] + 3[n-4,2,1^2] + [n-4,1^4]$.

This decomposition gives an upper bound of $332,720,898,048 = (18 \cdot 37 \cdot 32 \cdot 32 \cdot 21 \cdot 11 \cdot 11 \cdot 2 \cdot 4 \cdot 3 \cdot 4 \cdot 2)$ on the number of polynomials in A_2. To calculate the exact number, it is necessary to determine the number of distinct polynomials in this collection. A rough estimate shows that this number lies somewhere between $10,000,000$ and $20,000,000,000$.

Again, using the same arguments as in the case of $\mathcal{V}_{n,1}$, it follows that the polynomials in A_2 actually suffice for $\mathcal{V}_{n,2}$. $\clubsuit$

III Dimension theorems (non-uniform case)

The ideas illustrated by the above Example allow us to prove a more general version of Theorem 1A.

Theorem 1B: *For any $k, t \in \mathbb{N}$ there exists a finite collection $A_{k,t}$ of polynomials $p \in \mathbb{Q}[z]$ such that for any n and any FS_n-submodule $W \subseteq \oplus_{j=1}^{t} M^{(n-m_j,1^{m_j})}$ with $m_j \leq k$, there is $p \in A_{k,t}$ such that the dimension of W (as a linear vector space) is given by $p(n)$.*

Proof: As explained in the previous section, for characteristic 0, the permutation module $M^{(n-m,1^m)}$ can be written uniquely as a direct sum of irreducible modules. More specifically, we have $M^{(n-m,1^m)} = \oplus_{j=1}^{\mu} S_j$ where the S_j's are isomorphic to Specht Modules. For each $\beta = (n - |\beta'|, \beta') \trianglerighteq (n-m, 1^m)$ the module $S^{(n-|\beta'|,\beta')}$ appears with multiplicity $\mathrm{Mult}(S^\beta, M^\alpha)$ given by Young's rule. We claim (as stated in Lemma 2) that this is independent of n (as long as $n \geq 2m$). The multiplicity $\mathrm{Mult}(S^\beta, M^\alpha)$, for $\alpha = (n-m, 1^m)$ is the number of semi-standard tableaux which have shape β and which have $n-m$ 1's, one 2, one 3, ..., and one m. It follows, therefore, $\mathrm{Mult}(S^\beta, M^{(n-m,1^m)})$ for $\beta = (n - |\beta'|, \beta')$ is independent of n for $n \geq 2m$. The module $\oplus_{j=1}^{t} M^{(n-m_j,1^{m_j})}$ with $m_j \leq k$ can also be written uniquely (up to isomorphism) as a direct sum of irreducible Specht modules, and $\mathrm{Mult}(S^\beta, \oplus_{j=1}^{t} M^{(n-m_j,1^{m_j})})$ with $m_j \leq k$ is just $\sum_{j=1}^{t} \mathrm{Mult}(S^\beta, M^{(n-m_j,1^{m_j})})$. This number, which we denote $c_{\beta'}$ is independent of n for $n \geq 2k$.

The dimension of the Specht Module $S^\beta = S^{(n-|\beta'|,\beta')}$ is given by the hook formula:
$\frac{n!}{\text{product of the hook lengths for}\beta}$. The hook lengths for $\beta = (n - |\beta'|, \beta')$ can be split into two disjoint groups: the hook lengths for the first row of the diagram β, and the rest. The product of the hook lengths in the first row is of the form: $(n - 2|\beta'|)! \prod_{j \in B} (n-j)$ where $B \subseteq \{0, 1, \ldots, 2k' - 1\}$ have size $|B| = |\beta'|$. The product of the remaining hook lengths is a constant $C_{\beta'}$ which depends only on β'.

Thus, as claimed in Lemma 3, the dimension of $S^{(n-|\beta'|,\beta')}$ is given by

$$p_{\beta'}(n) := \frac{n!}{C_{\beta'}(n - 2|\beta'|')! \prod_{j \in B} (n-j)}$$

which is a polynomial in n. Now take $A_{k,t}$ to be the finite set of polynomials (in $\mathbb{Q}[z]$) of the form:

$$\sum_{\{\beta' : (n-|\beta'|,\beta') \geq (n-k,1^k)\}} b_{\beta'} p_{\beta'}(n)$$

where $0 \leq b_{\beta'} \leq c_{\beta'}$.

As in the example of the previous section, the partial sums, of the unique direct sum of irreducibles gives all of its submodules up to isomorphism. This ensures that the polynomials in $A_{k,t}$ exactly capture the dimensions of all submodules of $\oplus_{j=1}^{t} M^{(n-m_j,1^{m_j})}$ with $m_j \leq k$. ■

This theorem allows us to generalize Theorem 1A to a larger class of vector spaces than $\mathcal{V}_{n,d}$ which have many different variable types. Let $\Pi_{n,d}(r_1,\ldots,r_u)$ denote the space of polynomials of degree $\leq d$ built from u different variable types $x^{(1)}_{i_1,i_2,\ldots,i_{r_1}},\ldots,$ $x^{(u)}_{i_1,i_2,\ldots,i_{r_u}}$, where $i_1, i_2, \in \{1,2,\ldots,n\}$. These are polynomials of degree at most d in the ring $\mathbb{F}[x_{j,e_j} : 1 \leq j \leq u, e_j \in \{1,\ldots,n\}^{r_j}]$, where $\mathbb{F}$ is any field of characteristic 0. Clearly, the corresponding larger vector space $\mathcal{V}_{n,d}(r_1,\ldots,r_u)$ – obtained by treating, for example, the monomials $x^{(j)}_{e_j}x^{(i)}_{e_i}$ $x^{(i)}_{e_i}x^{(j)}_{e_j}$ as distinct – is an $\mathbb{F}S_n$-module under the natural action of S_n. The space $\mathcal{V}_{n,d}$ defined in the Introduction is thus the same as $\mathcal{V}_{n,d}(2)$. The space $\mathcal{V}_{n,d}(2,2)$ consists of polynomials in two types of variables: variables $x^{(1)}_{ij}$ and $x^{(2)}_{ij}$, $i,j \in \{1,2,\ldots,n\}$ (or simply x_{ij} and y_{ij}, $i,j \in \{1,2,\ldots,n\}$).

Theorem 1C: *For any $d, r_1, r_2, \ldots, r_u \in \mathbb{N}$ there exists a finite collection $A_{d,r_1,r_2,\ldots,r_u}$ of polynomials $p \in \mathbb{Q}[z]$ such that for any n and any $\mathbb{F}S_n$-submodule*
$W \subseteq \mathcal{V}_{n,d}(r_1,r_2,\ldots,r_u)$ (or $\subseteq \Pi_{n,d}(r_1,r_2,\ldots,r_u)$), there is a polynomial $p \in A_{d,r_1,\ldots,r_u}$ such that the dimension of W (as a linear vector space) is given by $p(n)$.

Proofs of Theorem 1A and Theorem 1C: There is a straightforward embedding of $\mathcal{V}_{n,d}(r_1,\ldots,r_u)$ (and of the quotient module $\Pi_{n,d}(r_1,\ldots,r_u)$) into the direct sum: $\oplus_{j=1}^{t}$ $M^{(n-m_j,1^{m_j})}$ with $m_j \leq k$, where $k := d\max\{r_1,r_2,\ldots,r_u\}$, and where $t := t(d,r_1,r_2,\ldots,r_u)$ is sufficiently large. More specifically, as in the previous Example, we choose t large enough to account for all possible order-types of monomial indices. Thus Theorem 1C follows from Theorem 1B. Theorem 1A is a special case of Theorem 1C. ■

Corollary: *Let $d, r_1, r_2, \ldots, r_u \in \mathbb{N}$. For any sequence $W_n \subseteq \mathcal{V}_{n,d}(r_1,r_2,\ldots,r_u)$ of $\mathbb{F}S_n$-submodules, there exists a polynomial $p \in A_{d,r_1,r_2,\ldots,r_u} \subseteq \mathbb{Q}[z]$ and an infinite set B such that* $\dim(W_n) = p(n)$, *for all $n \in B$.*

IV Decomposition Theorems (non-uniform case)

In this section, we give decomposition theorems which have a somewhat different emphasis than standard results in the representation theory of the symmetric group. We give an *explicit* characterization of all submodules $W \subseteq M^{(n-k,1^k)}$. Not just in terms of structure up to isomorphism, but also including a precise description of the generators of all the submodules. We use an example to illustrate the difference from the traditional analysis.

Example: Consider $M^{(n-2,1^2)}$. It can be uniquely decomposed into a direct sum of: one isomorphic copy of $S^{(n)}$, two isomorphic copies of $S^{(n-1,1)}$, one copy of $S^{(n-2,1^2)}$ and one copy of $S^{(n-2,2)}$. One concrete realization of this decomposition (viewing $M^{(n-2,1^2)} := \text{span}(\{e_{ij} : i, j \in \{1, 2, \ldots, n\}, i \neq j\})$) consists of the subspaces:

$S^{(n)} := \{\sum\limits_{ij} \ \lambda e_{ij} : \lambda \in \mathbb{F}\}$

$S'^{(n-1,1)} := \{\sum\limits_{ij} \ \lambda_i e_{ij} : \lambda_i \in \mathbb{F} \wedge \sum\limits_{i} \lambda_i = 0\}$

$S''^{(n-1,1)} := \{\sum\limits_{ij} \ \lambda_j e_{ij} : \lambda_j \in \mathbb{F} \wedge \sum\limits_{j} \lambda_j = 0\}$

$S^{(n-2,2)} := \{\sum\limits_{ij} \ \lambda_{ij} e_{ij} : \lambda_{ij} = \lambda_{ji} \wedge \sum\limits_{i} \lambda_{ij} = 0 \text{ for } j = 1, 2, \ldots, n\}$

$S^{(n-2,1^2)} := \{\sum\limits_{ij} \ \lambda_{ij} e_{ij} : \lambda_{ij} = -\lambda_{ji} \wedge \sum\limits_{i} \lambda_{ij} = 0 \text{ for } j = 1, 2, \ldots, n\}$

This decomposition is unique except that the two copies of $S^{(n-1,1)}$ can be "rotated" arbitrarily. More specifically, for every $a, b, c, d \in \mathbb{F}$ with $ad - bc \neq 0$, $S'_{a,b} := \{\bar{v} : \ a\bar{v}_1 + b\bar{v}_2, \ \bar{v}_1 \in S'^{(n-1,1)} \wedge \bar{v}_2 \in S''^{(n-1,1)}\}$ and $S''_{c,d} := \{\bar{v} : \ c\bar{v}_1 + d\bar{v}_2, \ \bar{v}_1 \in S'^{(n-1,1)} \wedge \bar{v}_2 \in S''^{(n-1,1)}\}$ we obtain the decomposition:
$M^{(n-2,1^2)} = S^{(n)} \oplus S'_{a,b} \oplus S''_{c,d} \oplus S^{(n-2,2)} \oplus S^{(n-2,1^2)}$.

This shows that although the submodules of $M^{(n-2,1^2)}$ have only finitely many dimensions and isomorphism types, $M^{(n-2,1^2)}$ contains infinitely many different $\mathbb{F}S_n$-submodules. However, it is straightforward (if one uses the fact that each S^α is irreducible) to show that any decomposition of $M^{(n-2,1^2)}$ into irreducibles is of this form.

Now consider the decomposition $M^{(n-2,1^2)} = S^{(n)} \oplus S'^{(n-1,1)} \oplus S''^{(n-1,1)} \oplus S^{(n-2,2)} \oplus S^{(n-2,1^2)}$. Consider the following formal expressions using formal sums over $M^{(n_0-2,1^2)}$ for some fixed $n_0 \geq 4$:

$E_{1,exp} := \sum\limits_{ij} e_{ij}$

$E_{2,exp} := \sum\limits_{i} e_{i1} - \sum\limits_{i} e_{i2}$

$E_{3,exp} := \sum\limits_{j} e_{1j} - \sum\limits_{j} e_{2j}$

$E_{4,exp} := e_{13} - e_{14} + e_{24} - e_{23} + e_{31} - e_{41} + e_{42} - e_{32}$, and

$E_{5,exp} := e_{13} - e_{14} + e_{24} - e_{23} - e_{31} + e_{41} - e_{42} + e_{32}$.

The corresponding elements $E_{i,n} \in M^{(n-2,1^2)}$ - obtained by restricting the scope of the formal sums in $E_{i,exp}$ to $\{1, 2, \ldots, n\}$ - generate, respectively, $S^{(n)}, S'^{(n-1,1)}, S''^{(n-1,1)}, S^{(n-2,2)}$, and $S^{(n-2,1^2)}$. Notice that the elements $E_{i,n}$ are ultrasmall because they have support size $\leq 4 = (2k)$. ♣

Remark: The above example indicates that the decomposition of $M^{(n-2,1^2)}$ into irreducible submodules (not just up to isomorphism) has the property that the irreducibles are each generated by an ultrasmall element.

This is significant because although it is known that the Specht modules are generated by the so-called polytabloids which are ultrasmall, it is not immediately clear that the property of being generated by ultrasmalls is preserved under arbitrary isomorphisms. ♣

Our next theorem states that in fact, this is always the case, and any irreducible module is generated by an ultrasmall element.

Note: We extend the definitions of (generalized) formal expressions and (generalized) ultrasmall formal expressions, in the natural way, to expressions constructed using formal sums over $\mathcal{V}_{n_0,d}(r_1,\ldots,r_u)$, for a fixed n_0. The corresponding (generalized) elements are in $\mathcal{V}_{n,d}(r_1,\ldots,r_u)$) for any n. Ultrasmall elements, in this context, have support size at most $2d\max\{r_1,r_2,\ldots,r_u\}$. Furthermore, as described in the above example, taking $M^{(n-l,1^l)} := \text{span}(\{e_{i_1,\ldots,i_l} : i_j \in \{1,2,\ldots,n\}, i_j \neq i_m \text{ for } j \neq m\})$, we define generalized formal expressions constructed using formal sums over $\oplus_{j=1}^{t} M^{(n_0-m_j,1^{m_j})}$ with $m_j \leq k$, where typically, $k := d\max\{r_1,r_2,\ldots,r_u\}$, and where $t := t(d,r_1,r_2,\ldots,r_u)$ is sufficiently large, with the resulting generalized elements being in $\oplus_{j=1}^{t} M^{(n-m_j,1^{m_j})}$, for any n. Ultrasmall elements, in this context, have support size at most $2k$. ♣

Theorem 2B: *For every $t,k \in \mathbb{N}$, every $\mathbb{F}S_n$-submodule W of $\oplus_{j=1}^{t} M^{(n-m_j,1^{m_j})}$ with $m_j \leq k$, is generated by ultrasmalls, each of which generates an irreducible submodule.*

Theorem 2C: *For any $d,r_1,r_2,\ldots,r_u \in \mathbb{N}$, every $\mathbb{F}S_n$-submodule $W \subseteq \mathcal{V}_{n,d}(r_1,r_2,\ldots,r_u)$ (or $\Pi_{n,d}(r_1,r_2,\ldots,r_u)$) is generated by ultrasmall elements (polynomials). The ultrasmall elements (polynomials) can be chosen such that they each generates an irreducible submodule.*

First, we refine the notion of support for a formal expression E_{exp} (and the corresponding sequences of elements E_n). We say E_{exp} has (a,b)-*support* if there exists a set A of size $\leq a$ such that any formal sum in E_{exp} has at most b parameters that are *not* in A. Notice that any E_{exp} has $(0,k)$-support. An expression E_{exp} is ultrasmall if and only if it has $(2k,0)$-support. Notice that (a,b)-support implies (a',b')-support provided $a' \geq a$ and $b' \geq b$.

Proof: We show Theorem 2B. The proofs of Theorem 2C (and in particular Theorem 2A) follow directly. Without loss of generality, we can assume W is irreducible (otherwise write $W := W_1 \oplus W_2 \oplus \ldots \oplus W_r$ where each W_j, $j = 1,2,\ldots,r$ is irreducible, and find ultrasmall generators for each W_j). Let E_n be a generator for W. Assume E_{exp} is the corresponding formal expression containing formal sums. To show that W is generated by an ultrasmall (i.e. an element of $(2k,0)$-support), we first show a property that even reducible modules possess. We refer to the process behind the following lemma as compression. The compression consists of replacing each generator by generators of smaller support.

Lemma 2D: *If any $\mathbb{F}S_n$-module W is generated by a set of generators that have (a,b)-support ($a \leq n-2, b \geq 1$), then in fact, W is generated by elements that have $(a+2,b-1)$-support.*

Proof of Lemma 2D: Assume E is a generator of (a,b)-support ($a \leq n-2, b \geq 1$). It suffices to show that there exists a collection of generators $F_1, \dots, F_u$ which have $(a+2,b-1)$-support and which together generate the same submodule as E. Without loss of generality we can assume that $A := \{1,2,\dots,a\}$ has the property that any term H (i.e. every abstract sum) in E_{exp}, the formal expression corresponding to E, contains at most b parameters not in A.

For every $i,j \in \{a+1,a+2,\dots,n\}$ consider $E_{ij} := (1-(ij))E$, where, as usual, (ij) denotes a 2-cycle in S_n, and $(1-(ij))$ is an element of the group ring or group algebra of S_n over $\mathbb{F}$ of characteristic 0. Also let $E_* := \sum_{\delta \in S_{\{a+1,a+2,\dots,n\}}} \delta E$, where $S_{\{a+1,a+2,\dots,n\}}$ is the subgroup of S_n that fixes $\{1,\dots,a\}$. Notice that each E_{ij} has $(a+2,b-1)$-support ($A \cup \{i,j\}$ is the witnessing set for this support), and it is not hard to see that E_* has $(a,0)$-support.

To complete the proof of the lemma, it suffices to show that $\{E_{ij} : i,j \in \{a+1,a+2,\dots,n\}\} \cup \{E_*\}$ generates exactly the same submodule as E, and in particular, it suffices to show that E can be derived from or generated by $\{E_{ij} : i,j \in \{a+1,a+2,\dots,n\}\} \cup \{E_*\}$.

First, notice that

$$(n-a)!E = E_* + \sum_{\delta \in S_{\{a+1,a+2,\dots,n\}}} (1-\delta)E \qquad I$$

Second, notice that $(1-\delta)$ where $\delta \in S_{\{a+1,a+2,\dots,n\}}$ can be written as a linear combination of $\delta'(1-(ij))$ where $i,j \in \{a+1,a+2,\dots,n\}$ and $\delta' \in S_{\{a+1,a+2,\dots,n\}}$. To see this, write

$$\delta = (i_1,j_1)(i_2,j_2)\dots(i_u,j_u)$$

and

$$(1-\delta) = (1-(i_1j_1)) + (i_1,j_1)(1-(i_2,j_2)) + \dots + (i_1,j_1)\dots(i_{u-1},j_{u-1})(1-(i_u,j_u))$$

Substituting in (I), and dividing by $(n-a)!$ ($\mathbb{F}$ has characteristic 0) we get the required derivation of E from $\{E_{ij} : i,j \in \{a+1,a+2,\dots,n\}\} \cup \{E_*\}$.

To complete the proof of the theorem, notice that an irreducible W is generated by a generator of $(0,k)$-support. Iterating Lemma 2D k times, it follows that W is generated by a generator of $(2k,0)$-support. ■

Remark. To appreciate the significance of the theorem, notice that not only are ultrasmalls a natural class of generators, they are uniquely suited to the task of general decomposition presented here. These theorems are sensitive to this definition of ultrasmalls, and the property of being generated by ultrasmalls is not preserved under arbitrary isomorphisms. For example, Theorem 2A, 2B and 2C would all fail if we did not allow, say, expressions with sums over repeated indices such as $\sum\limits_i x_{ii}$ in the the definition of ultrasmall. ♣

V Decomposition Theorems (uniform case)

We have shown that there exists a finite set $p_1, p_2, ..., p_v \in \mathbb{Q}[z]$ of polynomials such that for each sequence W_n of submodules (of some fixed $\mathbb{F}S_n$-module), there is a sequence of indices $j(n) \in \{1, 2, \ldots, v\}$ such that $\dim(W_n) = p_{j(n)}(n)$, for all n.

Take a finite collection Γ_{exp} of formal expressions over $\oplus_{j=1}^{t} M^{(n_0-m_j,1^{m_j})}$ with $m_j \leq k$, for some k, t, (or over $\mathcal{V}_{n_0,d}(r_1, \ldots, r_u)$, for some $r_1, \ldots, r_u$) for some fixed n_0; for any n, let Γ_n be the corresponding collection of elements of $\oplus_{j=1}^{t} M^{(n-m_j,1^{m_j})}$ with $m_j \leq k$, obtained from Γ_{exp}.

The module sequence $W_n \subseteq \oplus_{j=1}^{t} M^{(n-m_j,1^{m_j})}$ with $m_j \leq k$, (resp. $\mathcal{V}_{m,d}(r_1, \ldots, r_u)$) generated by Γ_n is said to be *uniformly generated from* Γ_{exp}, or from Γ_n, if it is clear from the context that Γ_n is obtained from a fixed collection of formal expressions, Γ_{exp}, for all n. In this case, we refer to both Γ_{exp} and Γ_n as the collection of generators.

Analogously, we also define module sequences that are uniformly generated by a set of generalized formal expressions Γ_{gen}.

If the sequence W_n is given thus in a uniform way, it is natural to expect that this uniformity is reflected in the sequence $j(n)$. In particular, if the uniformity condition on W_n is strong, it seems reasonable to expect that $j(n)$ is independent of n (i.e. $j(n)$ is a constant).
The next example shows that this is not generally the case:

Example: Consider $\mathcal{V}_{n,1}(1)$, i.e. the linear vector space of polynomials in the variables $x_1, x_2, \ldots, x_n$ of degree ≤ 1. Let W_n be the submodule generated by:
$E := 17x_1 - \sum\limits_{j=1}^{n} x_j$.
Let $E_1 := \frac{1}{17}(1 - (12))E = x_1 - x_2$ and let $E_2 := \frac{1}{(n-1)!} \sum\limits_{\delta \in S_n} \delta(E) = (17 - n) \sum\limits_{j=1}^{n} x_j$.

From this it is not difficult to see that $\dim(W_n) = n$ for $n \neq 17$, while $\dim(W_n) = n - 1$ for $n = 17$. Notice that W_n is reducible for all $n \neq 17$. More specifically, each W_n, $n \neq 17$ is isomorphic to a direct orthogonal sum of two irreducible modules which are isomorphic to $S^{(n)}$ and $S^{(n-1,1)}$ respectively. For the singular value $n = 17$, the decomposition factor $S^{(n)}$ vanishes and W_{17} becomes irreducible and isomorphic to $S^{(16,1)}$. ♣

Next we give a more involved example:

Example: Consider $\mathcal{V}_{n,1}(1,1)$. This module consists of all polynomials of degree ≤ 1 in the variables x_i and y_j, $1 \leq i, j \leq n$.

Let W_n be the submodule generated by:

$E := 17x_1 - \sum_{j=1}^{n} x_j + \sum_{j=1}^{n} y_j - 13y_2$.

and

$E' := 19x_1 - \sum_{j=1}^{n} x_j + \sum_{j=1}^{n} y_j - 23y_2$

The module W_n contains $x_1 - x_2$, and $y_1 - y_2$ each of which generate orthogonal submodules, isomorphic to $S^{(n-1,1)}$. The remaining part of W_n is spanned, as a vector space, by $E_1 := (17-n) \sum_{j=1}^{n} x_j + (n-13) \sum_{j=1}^{n} y_j$ and $E_2 := (19-n) \sum_{j=1}^{n} x_j + (n-23) \sum_{j=1}^{n} y_j$.

These two vectors are linearly independent except when $n = 18$. Thus $\dim(W_n) = 2n$ for all $n \neq 18$, while the dimension $\dim(W_n)$ "drops" to $2n-1$ for $n = 18$. To illustrate what happens, notice that W_n is, in fact, generated by the pairwise orthogonal generators $G_1 := x_1 - x_2, G_2 := y_1 - y_2, G_3 := 5 \sum_{j=1}^{n} x_j + \sum_{j=1}^{n} y_j$ and $G_4 := (n-18) \sum_{j=1}^{n} x_j - 5(n-18) \sum_{j=1}^{n} y_j$. For $n \neq 18$ each of those generators generates irreducible submodules isomorphic to $S^{(n-1,1)}, S^{(n-1,1)}, S^{(n)}$ and $S^{(n)}$ respectively. When $n = 18$, the generator G_4 becomes zero and the "height" of W_n drops from 4 to 3. ♣

In each of the examples, there exists a single polynomial $p(n)$ ($= n$, resp. $= 2n$) which gives the correct value of the dimension W_n for all but finitely many "singular" values of n. In each example there was only one singular value. It turns out that the structure of the singularities is closely related to the phenomenon of complexity gaps in algebraic complexity theory [15]. In fact, it turns out that singular values of n (which arise from the translations of logical propositions as we defined it in [14]) corresponds to values of n for which there exists an "sporadic" Nullstellensatz proof of the proposition. Intuitively, the proof is "sporadic" in the sense that it does not fall into the general class of proofs which essentially are all based on "proof ideas" which are independent of n (see [15] for more details).

Each of the examples illustrates our main technical result which is a uniform version of the decomposition in Theorem 2B: for any module sequence W_n generated uniformly from a set of formal expressions, there exists a set of generalized ultrasmall formal expressions which for each value of $n \geq 4k$, gives $\mathbb{F}S_n$-module elements that generate pairwise orthogonal, irreducible $\mathbb{F}S_n$-modules. For all but its singular values, the set generates W_n. At the singular values, it generates a submodule of W_n. Moreover, each generalized generator generates submodules which are isomorphic to $S^{(n-|\beta|,\beta)}$ for some fixed k-partition β (which is independent of n). At each singular value, one or more of the generators in the set generates the zero module. Whenever this happens, the height as well as the dimension of W_n "drops" and becomes strictly smaller than $p(n)$.

In this section, we set up the machinery needed to explain these phenomena. First, we prove a uniform version of the compression Lemma 2D.

Lemma 3D: *Take a finite collection of generalized formal expressions of support size $\leq l$ that uniformly generate $W_n \subseteq \oplus_{j=1}^t M^{(n-m_j,1^{m_j})}$ with $m_j \leq k$, (resp. $\mathcal{V}_{n,d}(r_1, r_2, \ldots, r_u)$) for $n \geq l$. There exists a fixed set of generalized ultrasmall expressions that uniformly generate W_n for each $n \geq \max\{2k, l+1\}$ (resp. $n \geq \max\{l+1, 2d \max\{r_1, \ldots, r_u\}\}$).*

Remark. It turns out that even if the original collection were to consist of ordinary formal expressions, the final collection in Lemma 3D may have to contain generalized ultrasmall expressions. ♣

Proof. We prove the lemma for $W_n \subseteq \oplus_{j=1}^t M^{(n-m_j,1^{m_j})}$ with $m_j \leq k$; the proof for $W_n \subseteq \mathcal{V}_{n,d}(r_1, r_2, \ldots, r_u)$ is virtually identical.

Without loss, we assume that the sequence W_n is generated by a single generalized expression, say

$$E := \ldots + p(n) \sum_{m_1}^{n} \sum_{m_2}^{n} x_{m_1,m_2,4,6} + \ldots,$$

where p is a rational function in $\mathbb{F}(z)$. To avoid unnecessary complications we always assume that all rational functions $p(n)$ are defined (i.e. have non-zero denominators) for $n \geq 2k$.

At start, we assume nothing about the support of E: all we know is that it has $(0, k)$-support, and has support size l; without loss, the support is restricted to $\{1, \ldots, l\}$. First we show (essentially by the same argument as in the proof of Lemma 2D) that we actually can generate the sequence W_n, $n \geq l+1$ by generalized expressions which are ultrasmall i.e. have support size $2k$.

For each $i, j \in \{1, 2, \ldots, n\}$ consider the generalized element $E_{ij} := (1-(ij))E$. Notice that $E_{ij} = 0$ for $i, j \geq l+1$ and that $E_{ij} = (j, j')E_{ij'}$ for $j, j' \geq l$. Thus we actually only need to consider E_{ij} for $i, j \in \{1, 2, \ldots, l+1\}$ (which is independent of n). We also consider $E_\sigma := \sum_{\delta \in S_n} \delta E$. Notice that

$$E_\sigma := \ldots + n!p(n)/n(n-1)(n-2)\sum_{m_1}^{n}\sum_{m_2}^{n}\sum_{m_3}^{n}\sum_{m_4}^{n} x_{m_1,m_2,m_3,m_4} + \ldots.$$

Let $E_* := 1/(n-k)!E_\sigma$. This process is uniform in n and we notice that that E_* also has a corresponding generalized formal expression:

$$E_{\sigma,exp} := \frac{1}{(z-k)!}(\ldots+z!p(z)/z(z-1)(z-2)\sum_{m_1}\sum_{m_2}\sum_{m_3}\sum_{m_4} x_{m_1,m_2,m_3,m_4}+\ldots).$$

As in the proof of Lemma 2D, and additionally using the fact that $(1-(ij))E$ for $j \notin \{1,\ldots,l+1\}$ can be obtained as $(l+1,j)(1-(i,l+1))E$, for any $n \geq l+1$, we can replace E by the set of expressions $\{E_{ij},\ i,j \in \{1,2,\ldots,l+1\}\} \cup E_*$. I.e, E and this collection both generate exactly the same submodule (for each fixed value of $n \geq l+1$). All the elements of this collection have $(2,k-1)$-support, and support size at most l.

As in the proof of Lemma 2D, we repeat this procedure. After iterating the procedure k times, we get generalized generators which have $(2k,0)$-support, and without loss, their support is restricted to $\{1,\ldots,2k\}$. At this point, notice that there are fixed, finitely many generalized ultrasmall expressions in this collection, independent of n, and the collection generates the same module as E for $n \geq \max\{l+1,2k\}$. ■

To get a complete analogy of Theorem 2B, we need to show that the generalized ultrasmall expressions obtained from Lemma 3D can in fact be modified so that each generates an irreducible $\mathbb{F}S_n$-module for all n. One cannot, as in the proof of Theorem 2B, a priori decompose W_n into irreducibles and proceed, since it is not clear that the same irreducible decomposition extends uniformly to the next n, and whether each irreducible in the decomposition is a member of a sequence generated uniformly in n. Instead, we rely on a crucial observation: the collection, call it Φ_{gen}, of generalized ultrasmall expressions given by Lemma 3D - when closed under the natural set of operations:

$$\gamma \in S_{4k} \text{ and } [\sum_{\delta \in S_n^u}]\delta \qquad *$$

for all subgroups S_n^u fixing $u \subseteq \{1,\ldots,2k\}$ - generates the sequence of modules W_n in a highly uniform manner. In particular, the next two lemmas show a remarkable fact: for any n, all ultrasmall elements in W_n with support in $\{1,\ldots,2k\}$ is in the vector space spanned by Φ^*_{gen} (the closure of Φ_{gen} under the operations $*$), i.e, arbitrary permutations from S_n are not necessary.

Although Lemma 3E and Lemma 3F are not directly used, they provide the intuition and motivation for the machinery that is used for proving the main result of the section.

Lemma 3E: *Consider an ultrasmall element F (with support in $\{1,\ldots,2k\}$) which is generated by a collection Φ_{gen} of ultrasmall generalized expressions, for some n, say n'. Then F is in fact in the linear span of $\Phi^*_{n'}$.*

Proof: Notice that if

$$\sum_{G\in\Phi_{gen}}\sum_{\delta\in S_n} c_\delta^{(G)}\delta G_{n'} = F$$

with each $c_\delta^G \in \mathbb{F}$, then if we apply $\sum\limits_{\delta\in S_n^u}$ to both sides, where $u \subseteq \{1,\ldots,2k\}$ is the support of F, then the right hand side remains a scalar multiple of F. The left hand side, however, is an $\mathbb{F}$-linear combination of elements in $\Phi^*_{n'}$. ■

Consider the space $\mathcal{G}$ of generalized generators with support in $\{1,2,\ldots,4k\}$. We view $\mathcal{G}$ as a $\mathbb{F}(z)S_{4k}$-module. More specifically, we view $\mathcal{G}$ as a linear vector space with each primitive element and formal sum being treated as an independent basis element, and with coefficients in the fraction field $\mathbb{F}(z)$ of rational functions over the field $\mathbb{F}$. Since $\mathbb{F}$ has characteristic zero, so does $\mathbb{F}(z)$. Notice that $\mathcal{G}$ is isomorphic to a direct sum of $\mathbb{F}(z)S_{4k}$-permutation modules: $\oplus_{j=1}^t M^{(4k-m_j,1^{m_j})}$ with $m_j \leq k$, for some t (resp. isomorphic to $\mathcal{V}_{4k,d}(r_1,\ldots,r_u)$ for some $r_1,\ldots,r_u$, where, as usual, $k = d\max\{r_1,\ldots,r_u\}$).

Consider two generalized expressions, say $E := \Sigma_{ijl}x_{ijl}$ and $F := (n-17)\Sigma_{ijl}x_{ijl}$. The generators E, F are *proportional* in $\mathcal{G}$ and thus they actually generate the same $\mathbb{F}(z)S_{4k}$ submodule (namely the submodule consisting of all expressions $r(z)\Sigma_{ijl}x_{ijl}$ where $r(z)$ is a rational function). The expressions E and F generate the same $\mathbb{F}S_n$-submodule sequence $W_n \subseteq M^{(n-k,1^k)}$ except for $n = 17$, where $F_n = 0$. In other words, the generators E_n and F_n generate the same $\mathbb{F}S_n$-submodule W_n (i.e. for all "non-singular" values of $n \geq 2k$, where neither E_n nor F_n is 0). The forward direction of the next lemma follows from this observation, and the reverse direction follows directly from Lemma 3E.

Lemma 3F: *Let Φ_{gen} and Γ_{gen} be a collection of generalized ultrasmall elements of $\mathcal{G}$ that are closed under the operations in (*). Then if Φ_{gen} and Γ_{gen} generate the same $\mathbb{F}(z)S_{4k}$-module, they also generate the same $\mathbb{F}S_n$-module for all values of n except finitely many singular values. Conversely, if Γ_{gen} and Φ_{gen} generate the same $\mathbb{F}S_n$ module for infinitely many values of $n \geq 4k$, then in fact, they generate the same $\mathbb{F}(z)S_{4k}$-module.*

Next, we define a formal *inner product* on $\mathcal{G}$. The inner product takes values in the fraction field $\mathbb{F}(z)$. The inner product (E,F) of two formal expression $E, F \in \mathcal{G}$ is defined to be the rational function obtained from the natural inner product in $\oplus_{j=1}^t M^{(n-m_j,1^{m_j})}$ with $m_j \leq k$, (resp.

$\mathcal{V}_{n,d}(r_1,\ldots,r_u)$ with $k = d\max\{r_1,\ldots,r_u\}$) of E_n and F_n, for $n \geq 4k$. For example, the inner product of $E := \Sigma_{jk}x_{1jk}$ and $F := \Sigma_{ijl,i\neq j}\ x_{ijl}$ is $n(n-1) \in \mathbb{F}(n)$. By linear extension, this defines a unique inner product in $\mathcal{G}$. Notice that the inner product is S_{4k}-invariant i.e. $(E,F) = (\delta E, \delta F)$ for each $E,F \in \mathcal{G}$ and for each $\delta \in S_{4k}$.

We say $E,F \in \mathcal{G}$ generate orthogonal submodules if for each $\delta \in S_{4k}$ we have $(E,\delta(F)) = 0$. The next lemma shows that orthogonal $\mathbb{F}(z)S_{4k}$-modules generated by ultrasmall generalized expressions remain orthogonal for all n, when viewed as $\mathbb{F}S_n$-modules. The proof follows immediately from the definition of the inner product on $\mathcal{G}$, and from the fact that E and F are ultrasmall.

Lemma 3G: *Let E and F be generalized ultrasmall expressions that generate orthogonal $\mathbb{F}(z)S_{4k}$ submodules of $\mathcal{G}$. Then E_n and F_n generate orthogonal $\mathbb{F}S_n$-modules for all $n \geq 4k$, where E_n and F_n are well-defined $\mathbb{F}S_n$-module elements (i.e, where none of the coefficients has a zero denominator).*

Next, we formalize the notion of "singular" values and how they can be "removed" meaningfully. We consider two types of singular values, zeroes and poles. We say that E is a generalized expression with *a zero* at $n = n_0$ when the $\mathbb{F}S_{n_0}$-module element E_{n_0} is 0. (A collection Φ_{gen} of generalized expressions is said to have a singular value whenever one of its elements has a singular value). In this case, there exists $r \in \mathbb{N}$ such that $E' := \frac{1}{(n-n_0)^r}E$ is a generalized generator (with coefficients being rational functions) with no singularity at n_0. Clearly we can iterate this idea and remove the (at most finitely many) zeroes of any generalized generator E. Equally, by multiplying by $(n-n_0)^r$, for suitable r, we could potentially also remove *poles* or singular values n_0, where E becomes undefined – i.e, one of its coefficients has a denominator that becomes zero at n_0. Note that we generally avoid poles altogether by assuming that our generalized expressions give well-defined $\mathbb{F}S_n$-module elements for all $n \geq 2k$. To see this assumption is reasonable, notice that the reduction in the proof of Lemma 3D only creates poles for $n < 2k$. Notice, however, that the reduction in the proof Lemma 3D can very well create generalized generators which vanish at various (at most finitely many) values of n. In general there is no way of to avoid the creation of zeroes (for $n \geq 4k$) during the compression process described in the proof of Lemma 3D.

Observe that when the singular values (zeroes or poles) of E are removed to give E', no new zeroes or poles are created, and the two generalized expressions are proportional (when considered as $\mathbb{F}(z)S_{4k}$-elements in $\mathcal{G}$), so they generate the same submodule of $\mathcal{G}$. Thus, using Lemma 3F and 3G we get the following.

Lemma 3H: *Let E' be a generalized generator obtained from E after removing singularities and exceptional values. Then E and E' generate* *sequences W_n and W'_n which are identical except for finitely many values of n. Similarly, if E and F are generalized ultrasmall expressions that generate orthogonal $\mathbb{F}(z)S_{4k}$- submodules of $\mathcal{G}$, then after removing singularities and exceptional values, the resulting E' and F' continue to generate orthogonal submodules of $\mathcal{G}$, and E'_n and F'_n generate orthogonal $\mathbb{F}S_n$-modules for all $n \geq 4k$.*

Finally, we are ready to prove the two main lemmas which are used to manipulate the set Φ_{gen} of generalized ultrasmall expressions obtained as a result of Lemma 3D. These manipulations are then used to prove a the uniform version of Theorem 2B (and Theorem 2C).

Lemma 3I: *Let Φ_{gen} be a collection of ultrasmall generalized formal expressions that generate a $\mathbb{F}(z)S_{4k}$-submodule $\tilde{W}$ of $\mathcal{G}$, and assume that the $\mathbb{F}S_n$-module elements corresponding to Φ_{gen} are all well-defined for all values of $n \geq 4k$. Then:*

1. *There exists a finite collection Γ_{gen} of ultrasmall generalized formal expressions that generate modules that form an orthogonal irreducible decomposition of $\tilde{W}$. For all but finitely many singular values of Γ_{gen}, the $\mathbb{F}S_n$ module U_n generated by Γ_n is well-defined and is identical to the $\mathbb{F}S_n$ module W_n generated by Φ_n. At the singular values, $U_n \subseteq W_n$.*

2. *There is a collection Δ_{gen} of ultrasmall generalized formal expressions that form an orthogonal irreducible decomposition of $\tilde{W}^{\perp}$ in $\mathcal{G}$. I.e, the collection $\Gamma_{gen} \cup \Delta_{gen}$ generates an orthogonal irreducible decomposition of $\mathcal{G}$ which is isomorphic to the direct sum of permutation modules $\oplus_{j=1}^{t} M^{(4k-m_j,1^{m_j})}$ with $m_j \leq k$, (resp. $\mathcal{V}_{4k,d}(r_1,\ldots,r_u)$, where $k := d\max\{r_1,\ldots,r_u\}$). Moreover, the collection Δ_{gen} has no singular values; Δ_n generates an $\mathbb{F}S_n$-module that is contained in $W_n^{\perp}$ for each $n \geq 4k$; and for n that are non-singular for Γ_{gen}, Δ_n in fact generates exactly $W_n^{\perp}$.*

3. *For all $n \geq 4k$, if all singular values has been removed from Γ_{gen}, to give Γ'_{gen}, the corresponding module U'_n generated by Γ'_n contains W_n we have $U'_n \supseteq W_n$; moreover the collection $\Gamma'_n \cup \Delta_n$ generates an irreducible decomposition of $\oplus_{j=1}^{t} M^{(n-m_j,1^{m_j})}$ with $m_j \leq k$.*

4. *There is a collection Ψ_{gen} of ultrasmall generators (not necessarily pairwise orthogonal) such that for each $n \geq 4k$, each element of Ψ_n either generates an irreducible submodule or is identically zero. Furthermore, for each $n \geq 4k$ (also for singular values of Ψ_{gen}), Ψ_n generates exactly W_n.*

Proof: Since we are working over characteristic 0, we can obtain an orthogonal irreducible decomposition of $\tilde{W}$ using the standard process akin to Gram-Schmidt orthogonalization. Using the non-uniform compression of Lemma 2D (putting $n = 4k$), we can compress the generator of each irreducible since it has $(0, k)$-support and Lemma 2D not only applies to $\mathbb{F}S_n$ module elements (for any fixed n), but also to $\mathbb{F}(z)S_{4k}$-module elements, since $\mathbb{F}(z)$ is a field of characteristic 0. We take the resulting collection of ultrasmalls - that generate an irreducible decomposition of $\tilde{W}$ - to be Γ_{gen}. By Lemma 3G, the $\mathbb{F}S_n$-modules generated by elements of Γ_n continue to remain orthogonal to each other for all values of $n \geq 4k$ where they are defined.

Moreover, the orthogonalization and the compression processes ensure that each $F_i \in \Gamma_{gen}$ has no poles (for $n \geq 2k$) and gives a well-defined $\mathbb{F}S_n$-module element $F_{i,n}$ and can be expressed as a well-defined $\mathbb{F}$-linear combination of the elements of Φ_n, for all values of $n \geq 4k$. The zeroes of Φ_{gen} is contained in the set of zeroes of Γ_{gen}, and while the zeroes of Γ_{gen} need not coincide with zeroes of Φ_{gen}, they do indicate a collapse in the irreducible decomposition structure of W_n, for that specific n. This collapse happens, for example, when some independent $\mathbb{F}(z)S_{4k}$-module elements in Φ_{gen} become dependent in Φ_n.

Viceversa, however, for certain singular values of Γ_{gen}, certain $E_{i,n} \in \Phi_n$ may not be expressible an $\mathbb{F}$-linear combination of the elements in Γ_n. So the most we can say is that the module U_n generated by Γ_n is a submodule of the module W_n generated by Φ_n for all $n \geq 4k$. However, proper containment occurs only at certain (finitely many) singular values of Γ_{gen}. I.e, the $\mathbb{F}S_n$-modules U_n and W_n generated by Γ_n and by Φ_n remain exactly the same for all but finitely many $n \geq 4k$.

This proves (1).

Similarly, to prove (2), we construct an orthogonal irreducible decomposition of $\tilde{W}^{\perp}$ by finding a maximal set of expressions that generate $\mathbb{F}(z)S_{4k}$-modules orthogonal to each other and to the elements in $\tilde{W}$, and perform the compression of Lemma 2D on them to make them ultrasmall. Next, we remove all singular values of these ultrasmall expressions and call the resulting collection Δ_{gen}. The maximality of the set forces each ultrasmall expression to generate an irreducible module, and forces the collection Δ_{gen} to generate all of $\tilde{W}^{\perp}$. Since Γ_{gen} gives an orthogonal irreducible decomposition of $\tilde{W}$ and Δ_{gen} of $\tilde{W}^{\perp}$, the entire collection $\Gamma_{gen} \cup \Delta_{gen}$ gives an orthogonal irreducible decomposition of the complete module $\mathcal{G}$, which is isomorphic to $\oplus_{j=1}^{t} M^{(4k-m_j,1^{m_j})}$ with $m_j \leq k$. By Lemma 3G and Lemma 3H, and since Δ_{gen} consists of ultrasmall expressions, orthogonality is preserved for all values of $n \geq 4k$, and thus Δ_n generates an $\mathbb{F}S_n$-module that is orthogonal to W_n and hence contained in $W_n^{\perp}$.

To prove (3), first notice that since the elements of $\Gamma_n \cup \Delta_n$ are ultrasmall, by Lemma 3G and Lemma 3H, they continue to generate orthogonal $\mathbb{F}S_n$-modules for all $n \geq 4k$. We first show that in addition, they generate an irreducible decomposition of $\oplus_{j=1}^{t} M^{(n-m_j,1^{m_j})}$ with $m_j \leq k$, for all but finitely many singular values of Γ_{gen} (Δ_{gen} is constructed without singular values). This follows from the facts:

(a) $\Gamma_{gen} \cup \Delta_{gen}$ generates a complete irreducible decomposition of $\oplus_{j=1}^{t} M^{(4k-m_j,1^{m_j})}$,

(b) (for $n \geq 4k$), the heights of $\oplus_{j=1}^{t} M^{(4k-m_j,1^{m_j})}$, and $\oplus_{j=1}^{t} M^{(n-m_j,1^{m_j})}$ are exactly the same when $m_j \leq k$,

(c) (at nonsingular values n of Γ_{gen}), none of the elements in $\Gamma_n \cup \Delta_n$ is identically 0, and finally,

(d) (for $n \geq 4k$), the elements of $\Gamma_n \cup \Delta_n$ are orthogonal and hence distinct.

Now, $\Gamma'_{gen} \cup \Delta_{gen}$ also generates a complete orthogonal irreducible decomposition of
$\oplus_{j=1}^{t} M^{(4k-m_j,1^{m_j})}$ since it consists of $\mathbb{F}(z)S_{4k}$-module elements that are proportional to those in $\Gamma_{gen} \cup \Delta_{gen}$. Moreover, since $\Gamma'_{gen} \cup \Delta_{gen}$ has no singular values, the same arguments used above for $\Gamma_{gen} \cup \Delta_{gen}$ now hold for *all* $n \geq 4k$. Finally, since Δ_n generates a module contained in $W_n^{\perp}$, it follows that the module U'_n generated by Γ'_n contains the module W_n for all $n \geq 4k$.

To prove (4), we construct Ψ_{gen} step by step, starting with Γ_{gen} and adding to it successively at the zeroes n_0 of Γ_{gen}. We consider 3 cases of zeroes.

When U_{n_0}, the module generated by Γ_{n_0} is equal to W_{n_0}, i.e, a collapse in W_{n_0} coincides with a singular value of Γ_{gen} at n_0, (in this case, U'_{n_0}, the module generated by Γ'_{n_0} properly contains W_{n_0}), no modification is made to Ψ_{gen}.

When U_{n_0}, the module generated by Γ_{n_0} is properly contained in W_{n_0}, and U'_{n_0}, the module generated by Γ'_{n_0} is equal to W_{n_0}, the zero at n_0 alone is removed from Ψ_{gen}, i.e, those $F_i \in \Gamma_{gen}$ that have a zero at n_0 are multiplied by $1/(n-n_0)^{r_i}$ for an appropriate value of r_i.

When both U_{n_0} is properly contained in W_{n_0} and W_{n_0} is properly contained in in U'_{n_0}, then there must exist, for example, $F_{i_1}, F_{i_2}, \ldots, F_{i_r}$ in Γ_{gen} which generate $\mathbb{F}S_{n_0}$-modules isomorphic to the same Specht module S^{β}, such that $F_{i_1,n_0}, F_{i_2,n_0}, \ldots, F_{i_r,n_0} \notin W_{n_0}$, but for linear combination $a_1F_{i_1,n_0} + a_2F_{i_2,n_0} + \ldots + a_rF_{i_r,n_0} \in W_{n_0}$,and generate an irreducible element which is isomorphic to S^{β}.

Next, remove the zero at n_0 alone from each of the elements $F_{i_1}, F_{i_2}, \ldots, F_{i_r} \in \Gamma_{gen}$ and denote the resulting element $F_{i_1}^{n_0}, F_{i_2}^{n_0}, \ldots, F_{i_r}^{n_0}$. Now the generalized ultrasmall expression $a_1F_{i_1}^{n_0} + a_2F_{i_2}^{n_0} + \ldots + a_rF_{i_r}^{n_0}$ is added to Ψ_{gen}.

Notice that the last addition destroys the orthogonality of elements in Ψ_{gen}, for example, at a value of n that is nonsingular for $F_{i_1}, F_{i_2}, \ldots, F_{i_r}$, the collection Ψ_n contains all the nonzero module elements $F_{i_1}, F_{i_2}, \ldots, F_{i_r}$ and $a_1 F_{i_1}^{n_0} + a_2 F_{i_2}^{n_0} + \ldots + a_r F_{i_r}^{n_0}$.

However, after going through all the zeroes of Γ_{gen} and adding generalized ultrasmall expressions as described above, we obtain Ψ_{gen} which generates exactly W_n for all $n \geq 4k$, and each of it members generates an irreducible for all values of $n \geq 4k$.

■

The next Lemma shows a crucial fact: not only does each ultrasmall in Γ_n and Δ_n always generate irreducible modules for all $n \geq 4k$, in fact, it generates a highly uniform sequence of irreducible modules that are isomorphic, in a sense, to the "same" Specht module $S^{(n-|\gamma|,\gamma)}$, for some *fixed* partition γ.

Lemma 3J: *Let $\tilde{W}$, Φ_{gen}, Γ_{gen} and Δ_{gen} be as in Lemma 3I. Then for each $F_i \in \Gamma_{gen} \cup \Delta_{gen}$ (F_i' after removing singularities), there is a unique partition $\beta_i := (4k - |\gamma_i|, \gamma_i)$, with $|\gamma_i| \leq k$, such that F_i and F_i' generate the same $\mathbb{F}S_{4k}$-module isomorphic to the Specht module S^{β_i}. For each n that is nonsingular for F_i, both F_i and F_i' generate the same $\mathbb{F}S_n$-module isomorphic to the Specht module $S^{\beta_{n,i}}$, where $\beta_{n,i} = (n - |\gamma_i|, \gamma_i)$. At F_i's singular values F_i is zero, while F_i' continues to generate an $\mathbb{F}S_n$-module isomorphic to the "same" Specht module $S^{\beta_{n,i}}$.*

Proof: Since $\mathcal{G}$ is isomorphic to $\oplus_{j=1}^{t} M^{(4k-m_j,1^{m_j})}$ with $m_j \leq k$, and $\mathbb{F}(z)$ has characteristic 0, each $F_i \in \Gamma_{gen} \cup \Delta_{gen}$ generates an irreducible module isomorphic to a Specht module S^{β_i}, with $\beta_i := (4k - |\gamma_i|, \gamma_i)$, where $|\gamma_i| \leq k$. By Lemma 3I, at F_i's nonsingular values, F_i generates an $\mathbb{F}S_n$ module isomorphic to some Specht module $S^{\beta_{n,i}}$, with $\beta_{n,i} := (n - |\gamma_{n,i}|, \gamma_{n,i})$, where $|\gamma_{n,i}| \leq k$.

The idea of the proof is based on the following. We know from Lemma 3I that $\Gamma'_{gen} \cup \Delta_{gen}$ generates a complete irreducible decomposition of $\oplus_{j=1}^{t} M^{(4k-m_j,1^{m_j})}$, and $\Gamma'_n \cup \Delta_n$ gives a complete irreducible decomposition of $\oplus_{j=1}^{t} M^{(n-m_j,1^{m_j})}$ for all n. These two decompositions have a bijective correspondence g. I.e, for each copy of some Specht module $S^{(n-|\gamma|,\gamma)}$ in the latter decomposition, there is a distinct corresponding copy of the Specht module $S^{(4k-|\gamma|,\gamma)}$ in the former decomposition, and vice versa. However, we need to show is that the Specht modules $S^{(n-|\gamma_{n,i}|,\gamma_{n,i})}$ (generated by the F_i's in Γ_{gen}) are all the same $S^{(n-|\gamma_i|,\gamma_i)}$ (or 0), independent of n. I.e, we need to show that the bijective correspondence g between the decompositions is very well-behaved, and in fact extends directly to the generating ultrasmalls in $\Gamma'_{gen} \cup \Delta_{gen}$ itself. In other words, the generating ultrasmalls do not generate wildly different irreducibles for different n's, or in other words, g does not allow irreducibles to jump around among the generating ultrasmalls. To show this, we use a simple property of Specht modules given by Lemma 1, and the structure of generalized ultrasmalls, embodied in the following claim. The claim then allows us to use a type of pigeon-hole principle based on the bijective correspondence g.

Claim: There are at most finitely many $n \geq 4k$ where $(n-|\gamma_i|, \gamma_i) \not\trianglerighteq \beta_{n,i}$. Moreover, for any m, there are at most finitely many $n \geq m$ where $(n-|\gamma_{m,i}|, \gamma_{m,i}) \not\trianglerighteq \beta_{n,i}$.

Proof of Claim: First notice that the signed column sums κ_t and $\kappa_{t'}$ for a β_i-tableau t and an $(n-|\gamma_i|, \gamma_i)$-tableau t' are exactly same, for any $n \geq 4k$. Thus, by Lemma 1, for any $n \geq 4k$ where $(n-|\gamma_i|, \gamma_i) \not\trianglerighteq \beta_{n,i}$, the sum $\kappa_t F_{n,i} = 0$, for any β_i-tableau t, since $S^{\beta_{n,i}}$ is isomorphic to the irreducible module generated by $F_{n,i}$. Since the coefficients in F_i are all rational functions in n, there can only be finitely many values of n where $\kappa_t F_{n,i} = 0$, unless $\kappa_t F_{n,i}$ is identically zero, which is not the case, since by Lemma 1, $\kappa_t F_i$ is isomorphic to a polytabloid that generates S^{β_i}. Therefore, there can only be finitely many values of $n \geq 4k$ where $(n-|\gamma_i|, \gamma_i) \not\trianglerighteq \beta_{n,i}$. For all other values of n, either $(n-|\gamma_i|, \gamma_i) \triangleright \beta_{n,i}$, or $(n-|\gamma_i|, \gamma_i) = \beta_{n,i}$. The proof of the second part of the claim goes through exactly the same way, replacing β_i by $\beta_{m,i} := (m-|\gamma_{m,i}|, \gamma_{m,i})$, and γ_i by $\gamma_{m,i}$ everywhere. This completes the proof of the Claim.

Let κ_t be the signed column sum of a β_i-tableau t. Let Q_i be the set of $n \geq 4k$ where $\kappa_t F_{n,i} = 0$. Clearly Q_i includes all singular values of F_i. We consider 2 cases for values of n.

Case 1: First we consider $n \notin \bigcup_{j:F_j \in \Gamma_{gen} \cup \Delta_{gen}} Q_j$. We show that for all such n, in fact the required property holds, i.e, $\beta_{n,i} = (n-|\gamma_i|, \gamma_i)$, or in other words, $\gamma_i = \gamma_{n,i}$. Assume, to the contrary, that this property does not hold for some such n_0. Using the definition of Q_i, and using the proof of the Claim, this would imply that $(n_0-|\gamma_i|, \gamma_i) \triangleright \beta_{n_0,i}$. Since n_0 is nonsingular for $\Delta_{gen} \cup \Gamma_{gen}$, using Lemma 3I, we know that $\Delta_{n_0} \cup \Gamma_{n_0}$ gives an irreducible decomposition of $\oplus_{j=1}^{t} M^{(n_0-m_j,1^{m_j})}$ with $m_j \leq k$, just as $\Delta_{gen} \cup \Gamma_{gen}$ gives an irreducible decomposition of $\oplus_{j=1}^{t} M^{(4k-m_j,1^{m_j})}$. As mentioned towards the beginning of the proof, these two decompositions have a bijective correspondence g. But we assumed that $F_i \in \Gamma_{gen} \cup \Delta_{gen}$ generates an $\mathbb{F}(z)S_{4k}$-module isomorphic to $S^{\beta_i=(4k-|\gamma|,\gamma)}$, whereas $F_{n_0,i}$ generates an $\mathbb{F}S_{n_0}$-module isomorphic to $S^{\beta_{n_0,i}}$, where $(n_0-|\gamma_i|, \gamma_i) \triangleright \beta_{n_0,i}$. Therefore, in order to preserve the bijective correspondence g, there must be another $F_l \in \Delta_{gen} \cup \Gamma_{gen}$ such that F_l generates an $\mathbb{F}(z)S_{4k}$-module isomorphic to a Specht module S^{α_1} while $F_{n_0,l}$ generates an $\mathbb{F}S_{n_0}$-module isomorphic to a Specht module S^{α_2} where $\alpha_1 \not\trianglerighteq \alpha_2$, which, using the Claim, contradicts the choice of n_0 to be outside the set $\bigcup_{j:F_j \in \Gamma_{gen} \cup \Delta_{gen}} Q_j$.

Case 2: Next we turn to $n \in \bigcup\limits_{j:F_j \in \Gamma_{gen} \cup \Delta_{gen}} Q_j$, and show that for all such n, the required property holds, i.e, we show that

$$\beta_{n,i} = (n - |\gamma_i|, \gamma_i), \tag{i}$$

if n is a nonsingular value of F_i, and if n is a singular value of F_i (so F_i generates the 0 module at n), we use Lemma 3I, take $S^{\beta'_{n,i}}$ to be the Specht module generated by F'_i after removing singularities, and show that

$$\beta'_{n,i} = (n - |\gamma_i|, \gamma_i). \tag{ii}$$

Assume the contrary (to (i) or (ii)) and let m be a counterexample value of n. Let Q be the set of i such that F_i has a singular value at m. First, we show that it must hold for $i \notin Q$ (resp. $i \in Q$) that:

$$\beta_{m,i} \trianglerighteq (m - |\gamma_i|, \gamma_i) \text{ (resp. } \beta'_{m,i} \trianglerighteq (m - |\gamma_i|, \gamma_i)). \tag{iii}$$

Say for some $i \notin Q$, it holds contrary to (iii) that $\beta_{m,i} \not\trianglerighteq (m - |\gamma_i|, \gamma_i)$. By the proof of Case 1, there are infinitely $n \geq m$ with $n \notin \bigcup\limits_{j:F_j \in \Gamma_{gen} \cup \Delta_{gen}} Q_j$, for which in fact $\beta_{n,i} = (n - |\gamma_i|, \gamma_i)$, it follows that there are infinitely many $n \geq m$ where $(n - |\gamma_{m,i}|, \gamma_{m,i}) \not\trianglerighteq \beta_{n,i}$, contradicting the second part of the Claim. This shows (iii) for $i \notin Q$. The same proof of (iii) goes through for $i \in Q$, due to the following reason. We know that F'_i and F_i generate the same $\mathbb{F}(z)S_{4k}$-module due to which $\beta'_i = \beta_i = (4k - |\gamma_i|, \gamma_i)$. Therefore the proof of Case 1 goes through also for $\beta'_{n,i}$. I.e, for $n \notin \bigcup\limits_{j:F_j \in \Gamma_{gen} \cup \Delta_{gen}} Q_j$, $\beta'_{n,i} = \beta_{n,i} = (n - |\gamma_i|, \gamma_i)$.

Now we continue with the proof Case 2 by contradiction, recalling that m is a counterexample value of $n \in \bigcup\limits_{j:F_j \in \Gamma_{gen} \cup \Delta_{gen}} Q_j$ and Q is the set of i such that F_i has a singular value at m.

From the proofs of Lemma 3I and 3J, it follows that the set $\{F'_i : i \in Q\} \cup \{F_i : i \notin Q\}$ (takes the place of $\Gamma'_{gen} \cup \Delta_{gen}$ and) gives an irreducible decomposition of $\oplus_{j=1}^{t} M^{(4k-m_j, 1^{m_j})}$ with $m_j \leq k$, just as $\{F'_{m,i} : i \in Q\} \cup \{F_{m,i} : i \notin Q\}$ gives an irreducible decomposition of $\oplus_{j=1}^{t} M^{(m-m_j, 1^{m_j})}$. Now, as in the proof of Case 1, we exploit the bijective correspondence g between the two irreducible decompositions. I.e, we conclude that if there is one $i \notin Q$ with $\beta_{m,i} \triangleright (n - |\gamma_i|, \gamma_i)$, or if there is an $i \in Q$ with $\beta'_{m,i} \triangleright (n - |\gamma_i|, \gamma_i))$, then in fact there must be another $l \notin Q$ (resp. $l \in Q$) with $\beta_{m,l} \not\trianglerighteq (m - |\gamma_l|, \gamma_l)$ (resp. $\beta'_{m,l} \not\trianglerighteq (m - |\gamma_l|, \gamma_l)$), which would cause a contradiction to (iii). ■

We are now ready to state the main result of the section, whose proof follows directly from Lemma 3D, Lemma 3I and Lemma 3J.

Theorem 3B (resp. 3C): *For any k, t, take a finite collection of generalized formal expressions of support size $\leq l$ that uniformly generate $W_n \subseteq \oplus_{j=1}^{t} M^{(n-m_j,1^{m_j})}$ with $m_j \leq k$, (resp. $\mathcal{V}_{n,d}(r_1, r_2, \ldots, r_u)$) for $n \geq l$. There exists a fixed set Γ_{gen} of generalized ultrasmall expressions such that the corresponding generalized ultrasmall elements Γ_n generate W_n for each $n \geq \max\{4k, l+1\}$ (resp. $n \geq \max\{l+1, 4d \max(\{r_j, j = 1, 2, \ldots, u\}\}$).*

Furthermore for each $n \geq \max\{4k, l+1\}$ (resp. $\geq \max\{l+1, 4d \max(\{r_j\})$) each generalized ultrasmall in Γ_n generates either zero or an irreducible module.

If we drop the condition of Γ_n having to generate W_n for singular values of n, we can choose Γ_{gen} such that the generators in Γ_{gen} generate pairwise orthogonal, irreducible $\mathbb{F}S_n$-submodules (for each $n \geq \max\{4k, l+1\}$ (resp. $\geq \max\{l+1, 4d \max(\{r_j, j = 1, 2, \ldots, u\}\}$).

In both cases, for each generator $F_i \in \Gamma_{gen}$, there exists a unique γ_i with $|\gamma_i| \leq k$ such that $F_{n,i}$ generates either 0 or an $\mathbb{F}S_n$-module that is isomorphic to the Specht module $S^{\beta_{n,i}}$, where $\beta_{n,i} = (n - |\gamma_i|), \gamma_i)$.

The following corollaries are straightforward.

Corollary 3K: *Let W_n be as in Theorem 3B. If W_n is irreducible for some sufficiently large n, then W_n is irreducible (or zero) for each $n \geq 4k$. Moreover, there exists a fixed partition γ with $|\gamma| \leq k$ such that each W_n is either zero or is isomorphic to the Specht module $S^{(n-|\gamma|,\gamma)}$.*

Corollary 3L: *Let W_n be as in Theorem 3B. If it is strictly contained in the entire module $\oplus_{j=1}^{t} M^{(n-m_j,1^{m_j})}$ with $m_j \leq k$, i.e, it does not take maximal dimension for sufficiently large n, then it is does not take maximal dimension for any $n \geq 4k$.*

VI Dimension Theorems (uniform case)

Now we are ready to prove our main Dimension theorem.

Theorem 4B (resp. 4C): *For any k, t, take a finite collection of generalized formal expressions of support size $\leq l$ that uniformly generate $W_n \subseteq \oplus_{j=1}^{t} M^{(n-m_j,1^{m_j})}$ with $m_j \leq k$, (resp. $\mathcal{V}_{n,d}(r_1, r_2, \ldots, r_u)$ or $\subseteq \Pi_{n,d}(r_1, r_2, \ldots, r_u)$) for $n \geq l$. There exists a single polynomial $p \in \mathbb{Q}[z]$, and a finite set $B \subseteq \mathbb{N}$ such that*

(1) $\dim(W_n) = p(n)$ *for all* $n \in \mathbb{N} \setminus B$.

(2) $\dim(W_n) < p(n)$ *for all* $n \in B$, *for which* $n \geq 4k$ *(resp.* $n \geq 2dr$*).*

Proof: By Theorem 3B, we know that there is a collection Γ_{gen} of generalized ultrasmall expressions F_i that generate a sequence of pairwise orthogonal (and hence distinct) irreducibles isomorphic to Specht modules $S^{(n-|\gamma_i|,\gamma_i)}$, where γ_i depends only on i, (not on n), for all but finitely many singular values of n. Furthermore, for these nonsingular values, Γ_n generates exactly W_n. Now (1) follows from a straightforward application of Lemma 3.

At the singular values of Γ_{gen} some of the F_i's generate the zero module. By Lemma 3I, after removing the singular values, the resulting expressions $F_i' \in \Gamma_{gen}'$ generate pairwise orthogonal irreducibles isomorphic to Specht modules $S^{(n-|\gamma_i|,\gamma_i)}$, for all $n \geq 4k$. Hence it is clear that the height of the module U_n' generated by Γ_n' is constant for all values of $n \geq 4k$, and by using Lemma 3 as in (1), we see that its dimension is the polynomial $p(n)$ for all $n \geq 4k$. Furthermore, U_n' is the same as W_n for nonsingular values n of Γ_{gen} and *contains* W_n for singular values. Hence the dimension and height of W_n always drop at the singular values of Γ_{gen} for $n \geq 4k$. ■

Remark: Theorem 4B shows that the dual problem where W_n is given as the solutions to uniformly generated homogeneous linear equations (closed under S_n) has the dimension increasing and the height increasing at singular values. An interesting corollary (keeping our previous examples in mind) is that for uniformly generated sequences W_n the sequence $W_n^{\perp}$ is in general NOT generated by generalized expressions. ♣

VII Relationship to Nullstellensatz Proofs

We now briefly describe another method of generating uniform families W_n of $\mathbb{F}S_n$-submodules of $\mathcal{V}_{n,d}$. It will follow that Theorem 3A, Theorem 3B, Theorem 4A and Theorem 4B remain valid for these notions of uniformity. We use this to give examples of NS-proof complexity results.

One method of generating a uniform family W_n is to start with a finite collection of generators $E_{1n} = (E_{1,exp})_n, \ldots, E_{vn} = (E_{v,exp})_n$ (ultrasmalls) and then define $W_n \subseteq \mathcal{V}_{n,d}$ ($W_n \subseteq \Pi_{n,d}$) to be the smallest submodule that contains $E_{1n}, \ldots, E_{vn}$ and is closed under other operations such as multiplication in $\mathcal{V}_{n,d}$ (or $\Pi_{d,n}$). In other words, if $E \in W_n$ and $F \in \mathcal{V}_{n,d}$ ($\in \Pi_{n,d}$) are such that $EF \in \mathcal{V}_{n,d}$ (or $\in \Pi_{n,d}$), then in fact, $EF \in W_n$.

This method allows us to define (in a uniform way) $W_{d_1,d_2,n} \subseteq \mathcal{V}_{n,d}$, $d_2 \leq d_1 \leq d$, the module consisting of the polynomial module elements

$$\{E \in \mathcal{V}_{n,d_2} : \ \exists F_{1n}, ..., F_{vn} \text{ of degree } \leq d_1 \text{ such that } \sum_{j=1}^{v} E_{jn}F_{jn} = E\}.$$

Informally, the polynomials in $W_{d_1,d_2,n}$ consist of the collection of elements of $\mathcal{V}_{n,d}$ that have degree $\leq d_2$ and that have Nullstellensatz multiplying polynomials of degree $\leq d_1$ that witness their membership in the ideal generated by $E_{1n}, \ldots, E_{vn}$. Theorem 3A and Theorem 4A are valid for this method of defining uniform families $W_{d_1,d_2,n}$ of $\mathbb{F}S_n$-submodules, by the following Lemma.

Lemma 4: *Fix two numbers d_1, d_2 with $d_1 \geq d_2$. Let $\bar{Q}$ be a collection of polynomials (of degree $\leq d_2$) given by formal expressions. For each n, let $\bar{Q}_n$ denote the closure of the expressions $\bar{Q}$ under S_n. Let $W_{d_1,d_2,n}$ denote the polynomials in $\Pi_{d_2,n}(r_1, \ldots, r_u)$ of degree $\leq d_2$ which can be proved by a NS-proof of degree $\leq d_1$ to belong to the ideal $(\bar{Q}_n)$. Let Ψ_{gen} consists of all linear combinations of polynomial expressions in $\bar{Q}$ but where we also close these under multiplication by monomials (whenever the result has degree $\leq d_1$). Then the space $W_{d_1,d_2,n}$ is generated by the generating polynomial expressions in Ψ_n.*

Corollary: *The sequence $W_{d_1,d_2,n}$ as defined in Lemma 4 is a uniform sequence of $\mathbb{F}S_n$-submodules.*

This shows that we can apply our structural results to the modules $W_{d_1,d_2,n}$. We get:

Theorem 5: *Let the sequence $W_{d_1,d_2,n}$ be as defined in Lemma 4. There exists a polynomial p with rational coefficients such that the vector space dimension of $W_{d_1,d_2,n}$ is given by $p(n)$ for all but finitely many values of n.*

Now let us return to the examples in the introduction.

Theorem 6: *Let ϕ be any sentence in the language of ZFC (ϕ could, for example, be the Riemann Conjecture or the Poincare Conjecture). Let $\bar{Q}_n \subseteq \Pi_{d_1,n}(r_1, \ldots, r_u)$ be an S_n-closed system of polynomial expressions which which has a solution if and only if there is a ZFC-proof of ϕ which uses at most n symbols. (Such a system of polynomial expressions can be shown to exist by combining standard methods of logic with the results in [14]). Then for no d_1, and $d_2 \geq 1$ does $W_{d_1,d_2,n}$ (as defined in Lemma 4) contain all polynomials of degree $\leq d_2$ (assuming $n \geq 2d_2\max(\{r_1, r_2, \ldots, r_u\})$).*

Proof (outline): We know from the contrapositive of Corollary 3L that if $W_{d_1,d_2,n}$ contains all polynomials of degree $\leq d_2$, i.e, if it takes maximal dimension for some n, then it in fact contains all such polynomials for for all sufficiently large values of n. Now ZFC can prove this fact, because the results in this paper are provable in naive set theory and thus are provable in ZFC. If there is $n \geq 2d_2\max(\{r_1, r_2, \ldots, r_u\})$ such that $W_{d_1,d_2,n}$ has maximal dimension, ZFC can verify this and hence ZFC can prove the fact: "$1 \in W_{d_1,d_2,n}$ for all sufficiently large values of n." But by the definition of $\bar{Q}_n$ and $W_{d_1,d_2,n}$ this means that ZFC can prove that "there is no ZFC proof of size n for ϕ for any value of n," or, in other words ZFC can prove that, "there is no ZFC proof of ϕ." This statement however can only be true (and this is provable in ZFC) if ZFC is consistent. Thus the assumption implies that ZFC can prove its own consistency. This is in contradiction with Gödel's second incompleteness theorem. In other words $W_{d_1,d_2,n}$ never takes maximal dimension. ■

In general, it is unclear which polynomial functions $n \rightarrow \dim(W_{d_1,d_2,n})$ can appear in this context. Theorem 6 (which was heavily based on Gödels second incompleteness theorem) shows that we can exclude the polynomial $n \rightarrow \dim(\Pi_{d_2,n}(r_1,\dots,r_u))$. Are there other polynomials which can be excluded? Even if we only consider there case where $d_2 = 2$ the number of potential polynomials is enormous (somewhere between 10^{14} and 10^{20}, if we work in $\mathcal{V}_{n,2}(2,2)$).

At the moment, we have very little understanding about which polynomial functions occur and whether this has any significance. And how robust are these questions? Is the answer very sensitive to the exact formalization of the provability predicate within ZFC? We believe it is quite tractable to compute (on modern computers) the concrete polynomial function which express the vector space dimension of spaces like $W_{d_1,d_2,n}$.

In the next section, we pose a series of concrete (but more abstract) questions we would like to answer.

VIII Open problems

The first question relates to Theorem 3B. We would like to show that for any uniformly generated family W_n, there exists a family Γ_{gen} of ultrasmall generalized generators generating pairwise orthogonal irreducible modules, which together generate exactly W_n for each $n \geq 4k$. At the moment, we have to either drop the property of orthogonality or have Γ_{gen} generate W_n only for sufficiently large n. More specifically we ask:

Question *Assume we are given a finite collection of generalized formal expressions that uniformly generate $W_n \subseteq \oplus_{j=1}^{t} M^{(n-m_j,1^{m_j})}$ with $m_j \leq k$, (resp. $\mathcal{V}_{n,d}(r_1,r_2,\dots,r_u)$ or $\subseteq \Pi_{n,d}(r_1,r_2,\dots,r_u)$). Is it always the case that there exists a family of ultrasmall generalized generators that generate orthogonal irreducible modules and together generate W_n for each $n \geq 4k$?*

This problem is important in getting a full understanding of the behavior of the submodules W_n. The missing key question is: to what extent can the modules W_n be built from irreducibles which do not "rotate" relative to the given generators.

Over fields of finite characteristic, there are still many unanswered questions. It is, for example, not clear if the analogous versions of Theorem

1A,B,C hold. However (based on the work by Ajtai [1]) we conjecture:

Conjecture 1A: *For each prime q and for each k there exists a finite set $A_{q,d}$ of functions $f : \mathbb{N} \to \mathbb{N}$ such that for any n and any $\mathbb{F}S_n$ submodule $W \subseteq M^{(n-k,1^k)}$ there exists $f \in A_{q,d}$ such that* $\dim(W) = f(n)$.

In fact, one can strengthen this conjecture.

Conjecture 1B: *For each prime q and for each k there exists $n_0, l \in \mathbb{N}$ and polynomial functions $p_0, p_1, \dots, p_{q^l-1} \in \mathbb{Q}[x]$ such that for each $n \geq n_0$ with $n \equiv r$* modulo *q^l, and each $\mathbb{F}S_n$ submodule $W \subseteq M^{(n-k,1^k)}$, it holds that* $\dim(W) = p_r(n)$.

In fact, we suggest that the conjecture is valid when $q^l \geq k+1$. In its strongest form we conjecture:

Conjecture 1C: *Conjecture 1B is valid when $q^l \geq k+1$ and when $n \geq c(q)k$ where $c(q)$ is some function which only depends on q (based on [15] we suggest that $c(q) = (7+q^2)$ will do).*

Theorem 2A, Theorem 2B, and Theorem 2C all fail over fields of finite characteristics. This follows from the fact that for $q = 2$ the $\mathbb{F}_2S_n$-submodule $W := \{E : E = \Sigma_{i<j} a_{ij}x_{ij} + b_{ij}x_{ij}$ where $\forall i,j\ a_{ij} = b_{ij}$ or $\forall i,j\ a_{ij} + b_{ij} = 1\}$ is only generated by elements of support size n (for example $E = \Sigma_{i<j}\ x_{ij}$). This suggests modifying and extending the definition of generalized ultrasmall expressions.

Moreover, Theorem 3A, Theorem 3B and Theorem 3C also fail over fields of finite characteristic. Based on [15] we believe however that the following modification is valid:

Conjecture 2A: *For any k and for any uniformly generated sequence $W_n \subseteq M^{(n-k,1^k)}$, there exists polynomial functions $p_0, p_1, \dots, p_{q^l-1} \in \mathbb{Q}[x]$ (where $q^l \geq k+1$) and there exists $n_0 \in \mathbb{N}$ such that for all $n \geq n_0$ with $n \equiv r$* modulo *q^l we have* $\dim(W_n) = p_r(n)$.

Conjecture 2B: *Conjecture 2A is valid for $n_0 \geq c(q)k$.*

More interesting questions remain for fields of characteristic 0. Is it possible to improve the upper bound on "n sufficiently large" in Theorem 3A, Theorem 3B and Theorem 3C? Given an upper bound on the smallest n that is nonsingular for Γ_{gen}, i.e, where W_n (in Theorem 3A, Theorem 3B and Theorem 3C) decomposes into irreducibles in the the same way as it decomposes for all sufficiently large n.

An upper bound of say $4k$ (or any constant times k) has profound consequences in showing linear complexity gaps for proofs of membership in ideals generated by *general* S_n-closed polynomial systems. The gaps would apply to algebraic proof systems like the Nullstellensatz proof system and Polynomial Calculus proof system.

Note: the upper bound of 2^k achieved in this paper implies a complexity jump from constant degree Nullstellensatz proofs to logarithmic degree Nullstellensatz proofs. Furthermore, Corollary 3K and 3L provide linear complexity gaps for algebraic proofs of ideal membership in certain classes of S_n-closed polynomial systems. ♣

IX Concluding Remarks

In [14], we show that most natural decision problems translate to the question of deciding membership in the ideals generated by uniform, S_n-closed polynomial systems. The main theorems of this paper remain valid under a larger class of notions of uniformity. In [15], we use these notions of uniformity to show gaps and lower bounds on the complexity of algebraic proofs of ideal membership [1], [7], [4], [6], for S_n-closed, uniformly generated polynomial systems.

Another interesting use of the results in this paper is based on the following observation. The singularities n at which some irreducible component of a uniformly generated module vanishes corresponds to "sporadic" algebraic proofs which use very specific properties of n and which cannot be generalized to general values of n. A similar phenomenon of a ono-to-one correspondence between singular (or exceptional) objects and efficient (but sporadic) propositional proofs was first discovered in [12] and [13]).

References

[1] Ajtai, M.: The independence of the modulo p counting principles. Proceedings of the 26th ACM Symposium on Theory of Computing, 402-411 (1994)

[2] Ajtai, M.: On the existence of modulo p cardinal functions, in: Feasible Mathematics II, eds. P. Clote and J. Remmel, Birkhauser. 1-14 (1994)

[3] Ajtai, M.: Symmetric systems of linear equations modulo p, TR94-015 of the Electronic Colloquium on Computational Complexity (1994)

[4] Beame, P., Impagliazzo, R., Krajicek, J., Pitassi, T., Pudlak, P.: Lower bounds on Hilbert's Nullstellensatz and propositional proofs. Proceedings of the London Mathematical Society **73(3)** 1-26 (1996)

[5] Brownawell, D.: Bounds for the degrees in the Nullstellensatz. Annals of Math., second series, **121(3)**, 577-592, (1987)

[6] Buss, S., Krajicek, J,. Pitassi, T., Razborov, A., Sergal, J.: Polynomial bound on Nullstellensatz for counting principles. To appear in Computational Complexity (1997)

[7] Clegg, M., Edmonds, J., Impagliazzo, R.: Using the Groebner basis algorithm to find proofs of unsatisfiability. Proceedings of the 28th ACM Symposium on Theory of Computing, 174-183 (1996)

[8] Isaacs, I.M: Algebra, A graduate course, Brooks Cole Publishers, (1994)

[9] James, G.: The Representation Theory of the Symmetric Groups. Lecture Notes in Mathematics., Vol. 682, Springer-Verlag, (1978)

[10] Krajicek, J.: On the degree of ideal membership proofs from uniform families of polynomials over a finite field (manuscript)

[11] Musili, C.: Representations of Finite Groups. Text and Readings in Mathematics, Hindustan Book Agency, India (1993)

[12] Riis, S.: Count(q) does not imply Count(p). Annals of Pure and Applied Logic, 90(1-3):1-56, (1997)

[13] Riis, S.: Count(q) versus the pigeon-hole principle. Archive for Mathematical Logic **36** 157-188 (1997)

[14] Riis, S., Sitharam, M.: Generating Hard Tautologies using Predicate Logic and the Symmetric Group, Appeared in the 5th Workshop on Logic, Language, Information and Computation (WoLLIC) (1998)

[15] Riis, S., Sitharam, M.: Manuscript in preparation. Incomplete version appear as technical report TR97-048 of the *Electronic Colloquium on Computational Complexity,* http://www.eccc.uni-trier.de/pub/eccc

[16] Smale, S.: Important mathematical problems for the next century. Mathematical Intelligencer, Spring 1998

Recent BRICS Report Series Publications

RS-98-20 **Søren Riis and Meera Sitharam.** ***Uniformly Generated Submodules of Permutation Modules*****. September 1998. 35 pp.**

RS-98-19 **Søren Riis and Meera Sitharam.** ***Generating Hard Tautologies Using Predicate Logic and the Symmetric Group*****. September 1998. 13 pp.**

RS-98-18 **Ulrich Kohlenbach.** ***Things that can and things that can't be done in PRA*****. September 1998. 24 pp.**

RS-98-17 **Roberto Bruni, José Meseguer, Ugo Montanari, and Vladimiro Sassone.** ***A Comparison of Petri Net Semantics under the Collective Token Philosophy*****. September 1998. 20 pp. To appear in** ***4th Asian Computing Science Conference,*** **ASIAN '98 Proceedings, LNCS, 1998.**

RS-98-16 **Stephen Alstrup, Thore Husfeldt, and Theis Rauhe.** ***Marked Ancestor Problems*****. September 1998.**

RS-98-15 **Jung-taek Kim, Kwangkeun Yi, and Olivier Danvy.** ***Assessing the Overhead of ML Exceptions by Selective CPS Transformation*****. September 1998. 31 pp. To appear in the proceedings of the** ***1998 ACM SIGPLAN Workshop on ML,*** **Baltimore, Maryland, September 26, 1998.**

RS-98-14 **Sandeep Sen.** ***The Hardness of Speeding-up Knapsack*****. August 1998. 6 pp.**

RS-98-13 **Olivier Danvy and Morten Rhiger.** ***Compiling Actions by Partial Evaluation, Revisited*****. June 1998. 25 pp.**

RS-98-12 **Olivier Danvy.** ***Functional Unparsing*****. May 1998. 7 pp. This report supersedes the earlier report BRICS RS-98-5. Extended version of an article to appear in** ***Journal of Functional Programming*****.**

RS-98-11 **Gudmund Skovbjerg Frandsen, Johan P. Hansen, and Peter Bro Miltersen.** ***Lower Bounds for Dynamic Algebraic Problems*****. May 1998. 30 pp.**

RS-98-10 **Jakob Pagter and Theis Rauhe.** ***Optimal Time-Space Trade-Offs for Sorting*****. May 1998. 12 pp.**

RS-98-9 **Zhe Yang.** ***Encoding Types in ML-like Languages (Preliminary Version)*****. April 1998. 32 pp.**