Permutation Group Algebras

Julian D. Gilbey

School of Mathematical Sciences

Queen Mary, University of London

Mile End Road

London E1 4NS, England

December 2001

Abstract

We consider the permutation group algebra defined by Cameron and show that if the permutation group has no finite orbits, then no homogeneous element of degree one is a zero-divisor of the algebra. We proceed to make a conjecture which would show that the algebra is an integral domain if, in addition, the group is oligomorphic. We go on to show that this conjecture is true in certain special cases, including those of the form $H \operatorname{Wr} S$ and $H \operatorname{Wr} A$, and show that in the oligormorphic case, the algebras corresponding to these special groups are polynomial algebras. In the $H \operatorname{Wr} A$ case, the algebra is related to the shuffle algebra of free Lie algebra theory. We finish by considering some integer sequences which arise from certain of these groups.

1 Introduction

Let G be a permutation group on an (infinite) set Ω. Cameron [2] defined a commutative, associative, graded algebra $A(G)$ which encodes information about the action of G on finite subsets of Ω. It is known that this algebra has zero divisors if G has any finite orbits. The question of what happens when G has no finite orbits is the subject of several conjectures due to Cameron [2], and we will be exploring two of them. The first is:

Conjecture 1.1. If G has no finite orbits, then ε is a prime element in $A(G)$.

Here ε is a certain element in the degree one component of the algebra, defined in section 2. The following weaker conjecture would follow from this, as we explain below.

Conjecture 1.2. If G has no finite orbits, then $A(G)$ is an integral domain.

The first conjecture would give us insight into the following question. If the number of orbits of G on unordered k-element subsets of Ω is n_k, then for which groups does $n_k = n_{k+1} < \infty$ hold? We will not study this question directly here; more information can be found in [2] and [3, sect. 3.5].

We first show that no homogeneous element of degree one in the algebra is a zero-divisor. Unfortunately, it is not obvious how to extend this argument

1

to higher degrees. We then go on to give a conjecture which would, if proven, yield a proof of the weaker conjecture 1.2, and show that it holds in two interesting classes of permutation groups. It also turns out in these two cases that the algebra $A(G)$ is a polynomial algebra, and we determine an explicit set of polynomial generators. It will follow that the stronger conjecture also holds in these cases. Although these results do not help to answer the question raised in the previous paragraph (as in these cases, $n_k < n_{k+1}$ for all k), they do provide further evidence to support the conjectures.

Finally, using the inverse Euler transform, Cameron [5] determined the number of polynomial generators of each degree which would be needed for certain of these algebras if they were actually polynomial algebras. Some of these sequences appear in The On-Line Encyclopedia of Integer Sequences [12] in the context of free Lie algebras. Our work gives an explanation for the sequences observed and the connection with free Lie algebras.

2   The graded algebra of a permutation group

We now give the definition of the algebra under consideration. Let G be a permutation group acting on Ω. Let K be a field of characteristic 0 (either $\mathbb{Q}$ or $\mathbb{C}$ will do). Define $V_n(G)$ to be the K-vector space of all functions from n-subsets of Ω to K which are invariant under the natural action of G on n-subsets of Ω. Define the graded algebra

$$A(G) = \bigoplus_{n=0}^{\infty} V_n(G)$$

with multiplication defined by the rule that for any $f \in V_m(G)$ and $g \in V_n(G)$, the product $fg \in V_{m+n}(G)$ is such that for any $(m+n)$-subset $X \subseteq \Omega$,

$$(fg)(X) = \sum_{\substack{Y \subseteq X \\ |Y|=m}} f(Y)g(X \setminus Y).$$

It is easy to check that, with this multiplication, $A(G)$ is a commutative, associative, graded algebra.

If G has any finite orbits, then this algebra contains zero-divisors. For let $X \subseteq \Omega$ be a finite orbit, $|X| = n$, and let $f \in V_n(G)$ be the characteristic function of this set (so $f(X) = 1$ and $f(Y) = 0$ for $Y \neq X$); then clearly $f^2 = 0$.

Considering Conjecture 1.2, it is clear that there are no zero-divisors in $V_0(G)$, as multiplying by an element of $V_0(G)$ is equivalent to multiplying by an element of K.

We also note that if there is a zero-divisor in $A(G)$, so we have $fg = 0$ with $0 \neq f, g \in A(G)$, then we can consider the non-zero homogeneous components of f and g with lowest degree; say these are f_m of degree m and g_n of degree n respectively. Then the term of degree $m+n$ in fg will be precisely $f_m g_n$, and as $fg = 0$, we must have $f_m g_n = 0$. So we may restrict our attention to considering homogeneous elements, and showing that for any positive integers m and n, we cannot find non-zero $f \in V_m(G)$ and $g \in V_n(G)$ with $fg = 0$.

Furthermore, we will show in the next section that $V_1(G)$ contains no zero-divisors as long as G has no finite orbits, so in particular, the element $\varepsilon \in V_1(G)$

defined by $\varepsilon(x) = 1$ for all $x \in \Omega$ is a non-zero-divisor. So if f is a homogeneous zero-divisor of degree m, with $fg = 0$, and g is homogeneous of degree $n > m$, we also have $(\varepsilon^{n-m} f)g = 0$, so $\varepsilon^{n-m} f \neq 0$ is a zero-divisor of degree n. Thus, if we wish, we can restrict our attention to showing that, for each positive integer n, we cannot find non-zero $f, g \in V_n(G)$ with $fg = 0$.

Turning now to the stronger Conjecture 1.1, we see that the second conjecture follows from this (as in [2]). For if $fg = 0$, with f and g homogeneous and non-zero, and $\deg f + \deg g$ is minimal subject to this, then $\varepsilon \mid fg$, so we can assume $\varepsilon \mid f$ by primality. Thus $f = \varepsilon f'$, and $\deg f' = \deg f - 1$. Thus $\varepsilon f' g = 0$, which implies $f'g = 0$ by the above, contrary to the minimality of $\deg f + \deg g$.

3 The degree one case

We intend to prove the following theorem.

Theorem 3.1. *If G has no finite orbits, then $V_1(G)$ contains no zero-divisors.*

In order to prove this theorem, we will make use of a technical proposition, which is based on a theorem of Kantor [8]. We first quote a version of Kantor's theorem, as we will have use for it later.

Proposition 3.2. *Let $0 \leqslant e < f \leqslant d - e$. Let X be a set with $|X| = d$. We define (E, F) for subsets $E, F \subset X$ with $|E| = e$ and $|F| = f$ by*

$$(E, F) = \begin{cases} 1 & \text{if } E \subset F \\ 0 & \text{otherwise,} \end{cases}$$

and the matrix $M = ((E, F))$, where the rows of M are indexed by the e-subsets of X and the columns by the f-subsets.
Then $\operatorname{rank} M = \binom{d}{e}$.

The extension of this result is as follows.

Proposition 3.3. *Let $0 \leqslant e < f \leqslant d - 2e$. Let X be a set with $|X| = d$, and let $E_0 \subset X$ with $|E_0| = e$ be a distinguished subset of X. Let w be a weight function on the $(f - e)$-subsets of X with values in the field K, satisfying the condition that $w(X') = 1$ whenever X' is an $(f - e)$-subset of X such that $X' \nsubseteq E_0$. We define (E, F) for subsets $E, F \subset X$ with $|E| = e$ and $|F| = f$ by*

$$(E, F) = \begin{cases} w(F \setminus E) & \text{if } E \subset F \\ 0 & \text{otherwise,} \end{cases}$$

and the matrix $M = ((E, F))$, where the rows of M are indexed by the e-subsets of X and the columns by the f-subsets.
Then $\operatorname{rank} M = \binom{d}{e}$.

Proof of theorem 3.1. Let $g \in V_1(G)$ with $g \neq 0$, and assume $h \in V_n(G)$ with $n \geqslant 1$ and $gh = 0$ (the $n = 0$ case has been dealt with in section 2). We must show that $h = 0$, so that for any $Y \subset \Omega$ with $|Y| = n$, we have $h(Y) = 0$. We assume that a set Y has been fixed for the remainder of this proof.

Since $g \neq 0$, there exists some (infinite) orbit $\Delta \subseteq \Omega$ on which g is non-zero; multiplying by a scalar if necessary, we may assume that $g(\delta) = 1$ for all $\delta \in \Delta$. Pick $X \subset \Omega$ with $|X| = 3n + 1$, $Y \subset X$ and $X \setminus Y \subset \Delta$.

Now for any $(n+1)$-subset $F \subset X$, we have $(hg)(F) = 0$ as $gh = hg = 0$, so that

$$(hg)(F) = \sum_{\substack{E \subset F \\ |E|=n}} h(E)g(F \setminus E) = 0.$$

This can be thought of as a system of linear equations in the unknowns $h(E)$ for $E \subset X$, $|E| = n$, with the matrix $M = (m_{EF})$ given by $m_{EF} = g(F \setminus E)$ if $E \subset F$, and $m_{EF} = 0$ otherwise.

This is precisely the situation of the proposition if we let $e = n$, $f = n + 1$ (so that $f - e = 1$), $d = 3n + 1$, $E_0 = Y$ and $w(\alpha) = g(\alpha)$; note that $w(\alpha) = 1$ whenever $\alpha \notin E_0$. (We write $g(\alpha)$ instead of the more correct $g(\{\alpha\})$; no confusion should arise because of this.) Thus $\operatorname{rank} M = \binom{d}{e}$ and the system of equations has a unique solution, which must be $h(E) = 0$ for all $E \subset X$ with $|E| = n$, as this is a possible solution. In particular, this means that $h(Y) = 0$, and since Y was chosen arbitrarily, it follows that $h = 0$.

Hence g is not a zero-divisor. $\square$

Proof of proposition 3.3. Let $R(E)$ be the row of M corresponding to E. M has $\binom{d}{e}$ rows, so we must show that the rows are linearly independent. We thus assume that there is a linear dependence among the rows of M, so

$$R(E^*) = \sum_{E \neq E^*} a(E)R(E) \tag{1}$$

for some e-set E^* and some $a(E) \in K$. We first note that $R(E^*)$ itself is non-zero: this follows as we can pick some $F \supset E^*$ with $F \setminus E^* \nsubseteq E_0$; for this F, we have $(E^*, F) = 1$.

Let Γ be the subgroup of $\operatorname{Sym}(X)$ which stabilises E_0 pointwise and E^* setwise. If $\sigma \in \Gamma$, then

$$(E^\sigma, F^\sigma) = \begin{cases} w((F \setminus E)^\sigma) = w(F \setminus E) & \text{if } E \subset F \\ 0 & \text{otherwise;} \end{cases}$$

either way, $(E^\sigma, F^\sigma) = (E, F)$. (For the result $w((F \setminus E)^\sigma) = w(F \setminus E)$, note that both sides are equal to 1 unless $F \setminus E \subseteq E_0$, in which case σ fixes this set pointwise.) Thus (1) implies that, for all F,

$$(E^*, F) = (E^*, F^\sigma) = \sum_{E \neq E^*} a(E^\sigma)(E^\sigma, F^\sigma)$$

$$= \sum_{E \neq E^*} a(E^\sigma)(E, F).$$

Thus

$$R(E^*) = \sum_{E \neq E^*} a(E^\sigma)R(E).$$

It follows that

$$|\Gamma|\, R(E^*) = \sum_{\sigma \in \Gamma} \sum_{E \neq E^*} a(E^\sigma) R(E)$$

$$= \sum_{E \neq E^*} R(E) \sum_{\sigma \in \Gamma} a(E^\sigma). \tag{2}$$

We now consider the orbits of Γ on the e-subsets of X, excluding E^*. The e-sets E_1 and E_2 will lie in the same orbit if and only if $E_1 \cap E_0 = E_2 \cap E_0$ and $|E_1 \cap E^*| = |E_2 \cap E^*|$. Thus every orbit is described by a subset $E' \subseteq E_0$ and an integer $0 \leqslant i \leqslant e - 1$. (We cannot have $i = e$, as we are excluding E^* from consideration.) Clearly not all possible pairs (E', i) will actually correspond to an orbit (it is not hard to see that necessary and sufficient conditions for this are $|E' \cap E^*| \leqslant i \leqslant \min\{e - 1, e - |E' \setminus E^*|\}$), so that whenever we consider or sum over such pairs below, we implicitly restrict attention to those which correspond to an orbit. In such cases, we write $\mathcal{E}(E', i)$ for the orbit. Also, for each such pair, pick some $E(E', i) \in \mathcal{E}(E', i)$. Then (2) implies

$$|\Gamma|\, R(E^*) = \sum_{(E',i)} \sum_{E \in \mathcal{E}(E',i)} R(E) \sum_{\sigma \in \Gamma} a(E^\sigma)$$

$$= \sum_{(E',i)} \sum_{E \in \mathcal{E}(E',i)} R(E) \sum_{\sigma \in \Gamma} a(E(E',i)^\sigma)$$

$$= \sum_{(E',i)} \sum_{\sigma \in \Gamma} a(E(E',i)^\sigma) \sum_{E \in \mathcal{E}(E',i)} R(E)$$

so that

$$R(E^*) = \sum_{(E',i)} b(E', i) \sum_{E \in \mathcal{E}(E',i)} R(E) \tag{3}$$

with $b(E', i) \in K$, and clearly not all of the $b(E', i)$ can be zero as $R(E^*)$ is not zero.

We define a total order on the pairs (E', i) as follows. Extend the partial order given by $\subseteq$ on the subsets of E_0 to a total order $\leqslant$, and then define $(E', i) \leqslant (E'', j)$ if $E' < E''$ or $E' = E''$ and $i \leqslant j$. We now proceed to derive a contradiction by showing that (3) leads to a system of linear equations for the $b(E', i)$ which is triangular under this total order, with non-zero diagonal entries, and deduce that all of the $b(E', i)$ must be zero.

Let $(\bar{E}, n)$ be a pair corresponding to an orbit. Since $2e + f \leqslant d$, there exists an f-set $F(\bar{E}, n)$ satisfying $F(\bar{E}, n) \cap E_0 = \bar{E}$ and $|F(\bar{E}, n) \cap E^*| = n$. (Simply take $E(\bar{E}, n)$ and adjoin $f - e$ points lying in $X \setminus (E_0 \cup E^*)$.) As $n \leqslant e - 1$, it follows that $F(\bar{E}, n) \not\supseteq E^*$, so $(E^*, F(\bar{E}, n)) = 0$. Hence by (3), we have

$$0 = \sum_{(E',i)} b(E', i) \sum_{E \in \mathcal{E}(E',i)} (E, F(\bar{E}, n)) \tag{4}$$

for all such pairs $(\bar{E}, n)$.

We note that $F(\bar{E}, n) \cap E_0 = \bar{E}$, and further that $E \in \mathcal{E}(E', i)$ implies that $E \cap E_0 = E'$; thus for the term $(E, F(\bar{E}, n))$ in equation (4) to be non-zero, where $E \in \mathcal{E}(E', i)$, we require $E' \subseteq \bar{E}$, hence also $E' \leqslant \bar{E}$. Furthermore, if

$(E, F(\bar{E}, n)) \neq 0$, we must have $i \leqslant n$ as $E \subset F(\bar{E}, n)$. Thus if $(\bar{E}, n) < (E', i)$, we have

$$\sum_{E \in \mathcal{E}(E', i)} (E, F(\bar{E}, n)) = 0. \tag{5}$$

Also, there is an e-set $E \subset F(\bar{E}, n)$ satisfying $E \cap E^* = F(\bar{E}, n) \cap E^*$ and $E \cap E_0 = F(\bar{E}, n) \cap E_0 = \bar{E}$; just take the union of $\bar{E}$ with $F(\bar{E}, n) \cap E^*$ and sufficiently many remaining points of $F(\bar{E}, n)$. For each such E, we have $F(\bar{E}, n) \setminus E \not\subseteq E_0$, so $(E, F(\bar{E}, n)) = 1$. Since K has characteristic zero, we deduce that

$$\sum_{E \in \mathcal{E}(\bar{E}, n)} (E, F(\bar{E}, n)) \neq 0, \tag{6}$$

as the sum is over all sets of precisely this form.

It then follows from (4) and (5) that for each pair $(\bar{E}, n)$:

$$0 = \sum_{(E', i) \leqslant (\bar{E}, n)} b(E', i) \sum_{E \in \mathcal{E}(E', i)} (E, F(\bar{E}, n)).$$

Now this is a system of linear equations in the unknowns $b(E', i)$ which is lower triangular. Also, by (6), the diagonal entries are non-zero. It follows that the unique solution to this system is that all of the $b(E', i)$ are zero, which provides the required contradiction to equation (3) above. $\qquad \square$

4 Oligomorphic-type cases: our conjecture

4.1 Ramsey orderings on orbits of n-sets

Cameron proved the following Ramsey-type result in [3, Prop. 1.10].

Lemma 4.1. *Suppose that the n-sets of an infinite set X are coloured with r colours, all of which are used. Then there is an ordering $c_1, \ldots, c_r$ of the colours and infinite subsets $X_1, \ldots, X_r$, such that X_i contains an n-set of colour c_i but no set of colour c_j for $j > i$.*

We use this as the inspiration for the following definition. If G is a permutation group on Ω, we say that the orbits of G on n-sets of Ω can be *Ramsey ordered* if, given any finite $N > n$, there is an ordering of the orbits c_α, $\alpha \in \mathcal{A}$, where $\mathcal{A}$ is a well-ordered set, and a corresponding sequence of (possibly infinite) subsets $X_\alpha \subseteq \Omega$ with $|X_\alpha| \geqslant N$, and such that X_α contains an n-set in the orbit c_α but no n-set in an orbit c_β for $\beta > \alpha$. (We can take $\mathcal{A}$ to be a set of ordinals with the $\in$-ordering if we wish; this is the reason for using Greek letters.) This pair of sequences forms a *Ramsey ordering*. While the particular Ramsey ordering may depend on N, we do not usually mention N unless we have to. The reader may think throughout of N having a very large finite value. It turns out that this makes certain constructions below simpler than if we required the X_α to be infinite sets.

Not every permutation group has such an ordering. For example, in the regular action of $\mathbb{Z}$ on $\mathbb{Z}$, there is no set with more than two elements, all of

whose 2-subsets are in the same orbit, so there cannot be a Ramsey ordering on 2-subsets. However, Cameron's result implies that if G is *oligomorphic* (that is, there are only finitely many orbits on n-sets for each n), then the orbits of G on n-sets can be Ramsey ordered for each n.

It turns out that Ramsey orderings on n-sets naturally yield Ramsey orderings on m-sets whenever $m < n$.

Proposition 4.2. *Let G be a permutation group acting on an infinite set Ω. Let $m < n$ be positive integers, and assume that the n-set orbits of G can be Ramsey ordered, say c_α and X_α with $\alpha \in \mathcal{A}$ are a Ramsey-ordering with $N \geqslant m + n$. Then this ordering induces a Ramsey ordering on the m-set orbits as follows. There is a subset $\mathcal{B} \subseteq \mathcal{A}$ and a labelling of the m-set orbits as d_β, $\beta \in \mathcal{B}$, such that for each $\beta \in \mathcal{B}$, an m-set in the orbit d_β appears in X_β, and that for each $\alpha \in \mathcal{A}$, X_α contains no m-sets in the orbit d_β for $\beta > \alpha$.*

We call the ordering of orbits d_β, $\beta \in \mathcal{B}$ together with the corresponding sets X_β given by this proposition the *induced Ramsey ordering*. Note that we use the same parameter N in both orderings.

The proof uses the following application of Kantor's theorem (Proposition 3.2 above), shown to me by Peter Cameron.

Lemma 4.3. *Let $m < n$ be positive integers, and let X be a finite set with $|X| \geqslant m + n$. Let the m-sets of X be coloured with colours from the set $\mathbb{N}$. Given an n-subset of X, we define its* colour-type *to be the multiset of colours of its $\binom{n}{m}$ m-subsets. Then the number of distinct m-set colours used in X is less than or equal to the number of distinct colour-types among the n-subsets of X.*

Proof. We note that only a finite number of colours appear among the m-subsets of X, as they are finite in number. Without loss of generality, we may assume that the colours used are precisely $1, 2, \ldots, s$.

As in Kantor's theorem (Proposition 3.2), we let M be the incidence matrix of the m-subsets versus n-subsets of X. By that theorem, as $m < n$ and $|X| \geqslant m + n$, this matrix has rank $\binom{|X|}{m}$, which equals the number of rows in the matrix. Thus, by the rank-nullity theorem, M represents an injective linear transformation.

Now for each $i = 1, \ldots, s$, let v_i be the row vector, with entries indexed by the m-subsets of X, whose j-th entry is 1 if the j-th m-subset has colour i, and 0 if it does not. Then $v_i M$ is a row vector, indexed by the n-subsets of X, whose k-th entry is the number of m-subsets of the k-th n-subset which have colour i.

Consider now the matrix M' whose rows are $v_1 M, \ldots, v_s M$. Note that the k-th column of this matrix gives the colour-type of the k-th n-subset of X. Its rank is given by

$$\operatorname{rank} M' = \dim \langle v_1 M, \ldots, v_s M \rangle = \dim \langle v_1, \ldots, v_s \rangle = s,$$

as M represents an injective linear transformation, and the s vectors $v_1, \ldots, v_s$ are clearly linearly independent. Now since the row rank and column rank of a matrix are equal, we have $s = \operatorname{rank} M' \leqslant$ number of distinct columns in M', which is the number of n-set colour-types in X. Thus the number of m-set colours appearing in X is less than or equal to the number of n-set colour-types in X, as we wanted. $\square$

Proof of Proposition 4.2. Let c_α be any n-set orbit, and let X be a representative of this orbit. We observe that the multiset of m-set orbits represented by the $\binom{m}{n}$ m-subsets of X is independent of the choice of X in this orbit. (For let $\bar{X}$ be another representative of the orbit c_α, with $\bar{X} = g(X)$, where $g \in G$. Then the set of m-subsets of X is mapped to the set of m-subsets of $\bar{X}$ by g, and so the multisets of m-set orbits represented by these two sets are identical.) In particular, we may say that an n-set orbit contains an m-set orbit, meaning that any representative of the n-set orbit contains a representative of the m-set orbit.

We first claim that every m-set orbit appears in some X_α: take a representative of an m-set orbit, say $Y \subset \Omega$. Adjoin a further $n - m$ elements to get an n-set $\bar{X}$. This n-set lies in some orbit, so there is a representative of this orbit in one of the X_α, say $X \subset X_\alpha$. Then this X_α contains a representative of our m-set orbit by the above argument, as we wished to show.

Now if $Y \subset \Omega$ is a representative of an m-set orbit, we set

$$\beta_Y = \min \{\, \alpha : g(Y) \subset X_\alpha \text{ for some } g \in G \,\}.$$

Note that this implies that the m-set orbit containing Y is contained in c_{β_Y} but not in c_α for any $\alpha < \beta_Y$. We set $\mathcal{B} = \{\, \beta_Y : Y \subset \Omega \text{ and } |Y| = m \,\}$, and if Y is an m-set, then we set d_{β_Y} to be the orbit of Y. We claim that $\mathcal{B}$ satisfies the conditions of the proposition with this orbit labelling. Certainly an m-set in the orbit d_{β_Y} appears in X_{β_Y} for each Y, by construction, and for each $\alpha \in \mathcal{A}$, X_α contains no m-sets in the orbit d_β for $\beta > \alpha$, again by construction. However, for d_{β_Y} to be well-defined, we require that $\beta_{Y_1} \neq \beta_{Y_2}$ if Y_1 and Y_2 lie in distinct orbits. We now show this to be the case by demonstrating that given any $\alpha_0 \in \mathcal{A}$, there can only be one m-set orbit appearing in c_{α_0} which has not appeared in any c_α with $\alpha < \alpha_0$.

So let $\alpha_0 \in \mathcal{A}$, and let $X \subseteq X_{\alpha_0}$ have size $m + n$ and contain an n-set in the orbit c_{α_0}. By the observation we made above, namely that the m-set orbits appearing in an n-set are independent of the choice of the n-set in its n-set orbit, it suffices to show that our set X contains at most one new m-set orbit. To use the lemma, we colour the m-subsets of X as follows. If Y is an m-set with $\beta_Y < \alpha_0$, then Y is given colour 1. Those $Y \subset X$ with $\beta_Y = \alpha_0$ are given the colours 2, 3, $\ldots$, with a distinct colour per m-set orbit. (Note that any $Y \subset X$ has $\beta_Y \leqslant \alpha_0$, as all n-subsets of X_{α_0} lie in orbits c_α with $\alpha \leqslant \alpha_0$.)

We now consider the possible colour-types of the n-sets of X. Note first that since the m-sets in a given m-set orbit all have the same colour, the colour-type of an n-set depends only upon the n-set orbit in which it lies. There is some n-subset of X in the orbit c_{α_0} by construction, and this has a certain colour-type. Any other n-subset $\tilde{X} \subset X$ is either in the same orbit c_{α_0}, and so has the same colour-type, or it is in some other orbit c_α with $\alpha < \alpha_0$. In the latter case, every m-subset $Y \subset \tilde{X}$ must have $\beta_Y \leqslant \alpha < \alpha_0$, and so it has colour 1. Thus the colour-type of such an n-set must be the multiset $[1, 1, \ldots, 1]$.

If every n-subset of X is in the orbit c_{α_0}, then there is only one colour-type, and so there can only be one m-set colour in X by the lemma, that is, only one m-set orbit with $\beta_Y = \alpha_0$. On the other hand, if X contains an n-set in an orbit c_α with $\alpha < \alpha_0$, then there are at most two colour-types in X: the all-1 colour-type and the colour-type of c_{α_0}. Thus, by the lemma, X contains at most two m-set colours. Colour 1 appears in c_α, and so there is at most one

other colour present, that is, there is at most one m-set orbit with $\beta_Y = \alpha_0$. Thus d_{β_Y} is well-defined on m-set orbits, and we are done. □

4.2 The Ramsey-ordering conjecture

Let G be a permutation group on Ω and let m and n be positive integers. Let d be an m-set orbit and e an n-set orbit. If c is an $(m+n)$-set orbit, then we say that c *contains a $d \cup e$ decomposition* if an $(m+n)$-set X in the orbit c can be written as $X = X_m \cup X_n$ with X_m in d and X_n in e. We can easily show using a theorem of P. M. Neumann that if G has no finite orbits, then for every pair (d, e), there exists an $(m+n)$-set orbit c containing a $d \cup e$ decomposition, as follows.

Neumann [9] proved the following: Let G be a permutation group on Ω with no finite orbits, and let Δ be a finite subset of Ω. Then there exists $g \in G$ with $g\Delta \cap \Delta = \varnothing$. It follows trivially that if Y and Z are finite subsets of Ω, then there exists $g \in G$ with $gY \cap Z = \varnothing$ (just take $\Delta = Y \cup Z$). In our case, let X_m and X_n be representatives of d and e respectively. Then there exists $g \in G$ with $gX_m \cap X_n = \varnothing$, and $gX_m \cup X_n$ is an $(m+n)$-set with the required decomposition, hence we can take c to be its orbit.

We will be considering groups G which have a Ramsey ordering on their $(m+n)$-set orbits. Let c_α, $\alpha \in \mathcal{A}$ be the ordering on $(m+n)$-sets, and let d_β, $\beta \in \mathcal{B}$ and e_γ, $\gamma \in \mathcal{C}$ be the induced Ramsey orderings on m- and n-sets respectively (where we assume N is sufficiently large). We then define

$$\beta \vee \gamma = \min \{\, \alpha : c_\alpha \text{ contains a } d_\beta \cup e_\gamma \text{ decomposition} \,\}.$$

Here is our main conjecture.

Conjecture 4.4. Let G be a permutation group on Ω with no finite orbits and for which the orbits on n-sets can be Ramsey ordered for every n. Then given positive integers m and n, there exists some Ramsey ordering of the orbits on $(m+n)$-sets with $N \geqslant 2(m+n)$, say c_α, $\alpha \in \mathcal{A}$ with corresponding sets $X_\alpha \subseteq \Omega$, which induces Ramsey orderings d_β, $\beta \in \mathcal{B}$ and e_γ, $\gamma \in \mathcal{C}$ on the m-set orbits and n-set orbits respectively, and which satisfies the following conditions for all $\beta, \beta' \in \mathcal{B}$ and $\gamma, \gamma' \in \mathcal{C}$:

$$\beta \vee \gamma < \beta' \vee \gamma \text{ if } \beta < \beta' \quad \text{and} \quad \beta \vee \gamma < \beta \vee \gamma' \text{ if } \gamma < \gamma'.$$

Note that the conditions of this conjecture also imply that if $\beta < \beta'$ and $\gamma < \gamma'$, then $\beta \vee \gamma < \beta \vee \gamma' < \beta' \vee \gamma'$, so that $\beta \vee \gamma \leqslant \beta' \vee \gamma'$ implies that either $\beta < \beta'$ or $\gamma < \gamma'$ or $(\beta, \gamma) = (\beta', \gamma')$.

Given this conjecture, it is easy to show that $A(G)$ is an integral domain for such groups. For if $fg = 0$ with $0 \neq f \in V_m(G)$ and $0 \neq g \in V_n(G)$, let β_0 be such that $f(d_\beta) = 0$ for $\beta < \beta_0$ but $f(d_{\beta_0}) \neq 0$, and let γ_0 be such that $g(e_\gamma) = 0$ for $\gamma < \gamma_0$ but $g(e_{\gamma_0}) \neq 0$. (We write $f(d_\beta)$ to mean the value of $f(Y)$ where Y is any representative of the orbit d_β, and so on.) Letting $\alpha_0 = \beta_0 \vee \gamma_0$, we can consider $fg(c_{\alpha_0})$. Now since $fg = 0$, this must be zero, but we can also determine this explicitly. Letting X be a representative of c_{α_0}, we have

$$fg(c_{\alpha_0}) = fg(X) = \sum_{\substack{Y \subset X \\ |Y|=m}} f(Y)g(X \setminus Y).$$

Every term in the sum is of the form $f(d_\beta)g(e_\gamma)$ where $d_\beta \cup e_\gamma$ is a decomposition of c_{α_0}, so that $\beta \vee \gamma \leqslant \alpha_0 = \beta_0 \vee \gamma_0$. But by the conjecture, this implies that except for terms of the form $f(d_{\beta_0})g(e_{\gamma_0}) \neq 0$, every term either has $\beta < \beta_0$ so that $f(d_\beta) = 0$, or $\gamma < \gamma_0$ so that $g(e_\gamma) = 0$, and hence every one of these terms is zero. Since there exist terms of the form $f(d_{\beta_0})g(e_{\gamma_0})$ by the choice of α_0, we must have $fg(c_{\alpha_0}) \neq 0$. But this contradicts $fg = 0$, and so $A(G)$ is an integral domain.

Recall from section 2 that we can assume $m = n$ when showing that $A(G)$ is an integral domain (that is, $fg = 0$ where $f, g \in V_n(G)$ implies $f = 0$ or $g = 0$); hence we can restrict ourselves to proving the conjecture in the case $m = n$ if this is easier.

5   Special cases (I): Wreath-S-like groups

5.1   Notational conventions

We gather here some notation that we will be using for the rest of this paper.

We will make use of the lexicographical order on finite sequences and multisets, which we define as follows. Let $(X, <)$ be a totally ordered set. If $x = (x_1, \ldots, x_r)$ and $y = (y_1, \ldots, y_s)$ are two ordered sequences of elements of X, then we say that x is lexicographically smaller than y, written $x <_{\text{lex}} y$, if there is some t with $x_i = y_i$ for all $i < t$, but either $x_t < y_t$ or $r + 1 = t \leqslant s$. If we now take a finite multiset of elements of X, say M, we write $\text{seq}(M)$ to mean the sequence obtained by writing the elements of M (as many times as they appear in M) in decreasing order. Then if M_1 and M_2 are finite multisets, we define $M_1 <_{\text{lex}} M_2$ to mean $\text{seq}(M_1) <_{\text{lex}} \text{seq}(M_2)$. Note that $<_{\text{lex}}$ is a total order on the set of finite multisets, for $\text{seq}(M_1) = \text{seq}(M_2)$ if and only if $M_1 = M_2$. If we need to explicitly list the elements of a multiset, we will write $[x_1, x_2, \ldots]$. We write $M_1 + M_2$ for the multiset sum of the multisets M_1 and M_2, so if $M_1 = [x_1, \ldots, x_r]$ and $M_2 = [y_1, \ldots, y_s]$, then $M_1 + M_2 = [x_1, \ldots, x_r, y_1, \ldots, y_s]$.

In the following sections, we will talk about a set of *connected blocks* for a permutation group, the idea being that every orbit will correspond to a multiset or sequence of connected blocks. The choice of terminology will be explained below, and is not related to blocks of imprimitivity. Also, the individual words "connected" and "block" have no intrinsic meaning in the context of the definitions in this paper. Every connected block has a positive integral weight (for which we write $\text{wt}(\Delta)$), and the weight of a sequence or multiset of connected blocks is just the sum of weights of the individual connected blocks. We well-order the connected blocks of each weight, and denote the connected blocks of weight i by $\Delta_i^{(j)}$, where j runs through some well-ordered indexing set. Without loss of generality, we assume that $\Delta_1^{(1)}$ is the least connected block of weight 1. We then define a well-ordering on all connected blocks by $\Delta_i^{(j)} < \Delta_{i'}^{(j')}$ if $i < i'$ or $i = i'$ and $j < j'$. Using this ordering, we can then talk about the lexicographic ordering on sequences or multisets of connected blocks.

5.2   Wreath-S-like groups

Our prototypical family of groups for this class of groups are those of the form $G = H \operatorname{Wr} S$, where H is a permutation group on Δ and $S = \operatorname{Sym}(\mathbb{Z})$, the

symmetric group acting on a countably infinite set (we take the integers for convenience). The action is the imprimitive one, so G acts on $\Omega = \Delta \times \mathbb{Z}$. We extract those features of this group which are necessary for the proof below to work.

Definition 5.1. We say that a permutation group G on Ω is *wreath-S-like* if there is a set of connected blocks $\{\Delta_i^{(j)}\}$ and a bijection ϕ from the set of orbits of G on finite subsets of Ω to the set of all finite multisets of connected blocks, with the bijection satisfying the following conditions (where we again blur the distinction between orbits and orbit representatives):

(i) If $Y \subset \Omega$ is finite, then $\mathrm{wt}(\phi(Y)) = |Y|$.

(ii) If $Y \subset \Omega$ is finite and $\phi(Y) = [\Delta_{i_1}^{(j_1)}, \ldots, \Delta_{i_k}^{(j_k)}]$, we can partition Y as $Y = Y_1 \cup \cdots \cup Y_k$ with $|Y_l| = i_l$ for each l. Furthermore, if $Z \subseteq Y$ and $Z = Z_1 \cup \cdots \cup Z_k$, where $Z_l \subseteq Y_l$ for each l, then we can write $\phi(Z)$ as a sum of multisets $\phi(Z) = M_1 + \cdots + M_k$, where $\mathrm{wt}(M_l) = |Z_l|$ for each l and $M_l = [\Delta_{i_l}^{(j_l)}]$ if $Z_l = Y_l$.

Note that condition (ii) implies that $\phi(Y_l) = [\Delta_{i_l}^{(j_l)}]$ for $j = 1, 2, \ldots, k$. Essentially, this condition means that subsets of Y correspond to "submultisets" of $\phi(Y)$ in a suitable sense.

In the case of $G = H \,\mathrm{Wr}\, S$ mentioned above, we take the connected blocks of weight n to be the orbits of the action of H on n-subsets of Δ. Then every orbit of G can be put into correspondence with a multiset of H-orbits as follows. If $Y \subset \Omega$ is an orbit representative, then $\phi(Y) = [\,\pi_i(Y) : \pi_i(Y) \neq \varnothing\,]$, where the π_i are projections: $\pi_i(Y) = \{\,\delta : (\delta, i) \in Y\,\}$, and we identify orbits of H with orbit representatives. Note that $\mathrm{wt}(\phi(Y)) = |Y|$ as required, and that condition (ii) is also satisfied; in fact, in the notation of the condition, we have $M_l = [\Delta_{i_l'}^{(j_l')}]$ for each l, for some appropriate i_l' and j_l'.

Another example is the automorphism group of the random graph. The random graph is the unique countable homogeneous structure whose age consists of all finite graphs. It is also known as the Fraïssé limit of the set of finite graphs; see Cameron [3] for more information on homogeneous structures and Fraïssé's theorem. We take the set of connected blocks to be the isomorphism classes of finite connected graphs, where the weight of a connected block is the number of vertices in it. Any orbit can be uniquely described by the multiset of connected graph components in an orbit representative. Condition (i) is immediate, as is condition (ii). Note, however, that there are examples in this scenario where M_l may not be a singleton. For example, if $Y = P_2$ is the path of length 2 (with three vertices), so that $\phi(Y) = [P_2]$, and $Z \subset Y$ consists of the two end vertices of the path, then $\phi(Z) = [K_1, K_1]$.

This prototypical example explains the choice of terminology: the basic units in this example are the connected graphs, so we have called our basic units connected blocks, both to suggest this example and that of strongly connected components in tournaments as considered in section 6 below.

Cameron [4, Sec. 2] has shown that $A(G)$ is a polynomial algebra if G is an oligomorphic wreath-S-like group, from which it follows that $A(G)$ is an integral domain in this case. It also follows that ε is a prime element, so both

Conjectures 1.1 and 1.2 hold in this case. The argument that $A(G)$ is a polynomial algebra in the oligomorphic case is similar to that presented below for wreath-A-like groups, only significantly simpler.

We now show, using a new argument based on Ramsey-orderings, that $A(G)$ is an integral domain in the wreath-S-like case, even without the assumption that G is oligomorphic. This will also provide a basis for the arguments presented in the next section for wreath-A-like groups.

Theorem 5.2. *If G is wreath-S-like, then $A(G)$ is an integral domain.*

Proof. We claim that in such a situation, the conditions of Conjecture 4.4 are satisfied, and hence $A(G)$ is an integral domain.

Following the requirements of the conjecture, let m and n be positive integers and pick any integer $N \geqslant 2(m+n)$. Denote the inverse of ϕ by ψ and let α run through all multisets of connected blocks of total weight $m+n$, then we set $c_\alpha = \psi(\alpha)$ and let X_α be an N-set in the orbit $\psi(\alpha + [\Delta_1^{(1)}, \ldots, \Delta_1^{(1)}])$, where the second multiset has $N - (m+n)$ copies of $\Delta_1^{(1)}$. We claim that this gives a Ramsey ordering of the orbits on $(m+n)$-sets, where the multisets are ordered lexicographically (which gives a well-ordering on the multisets). Firstly, every $(m+n)$-set orbit appears among the list by hypothesis, as ψ is a bijection. Secondly, by construction, there is an $(m+n)$-subset of X_α in the orbit $\psi(\alpha)$, namely partition X_α as in condition (ii) of the definition, and remove all of the elements corresponding to the copies of $\Delta_1^{(1)}$ added. This subset will then map to α under ϕ, by condition (ii). Finally, any $(m+n)$-subset of X_α can be seen to correspond to a multiset lexigraphically less than or equal to α, again using condition (ii) and the fact that $\Delta_1^{(1)}$ is the least connected block, so the subset will be in an orbit c_β with $\beta \leqslant_{\mathrm{lex}} \alpha$, as required.

We note that the induced Ramsey orderings on m-set orbits and n-set orbits are given by precisely the same construction. Specifically, let β be a multiset with $\mathrm{wt}(\beta) = n$. Then the orbit corresponding to the multiset β first appears in X_{α_0} where $\alpha_0 = \beta + [\Delta_1^{(1)}, \ldots, \Delta_1^{(1)}]$. For assume that an n-set Z in the orbit $\psi(\beta)$ appears in X_α. As we have $\phi(Z) = \beta$, β must be a "submultiset" of α in the sense of condition (ii), and it is clear that the lexicographically smallest such α is the one given by adjoining an appropriate number of copies of $\Delta_1^{(1)}$ to β. It is not difficult to show that $\beta \vee \gamma$ is precisely the multiset $\beta + \gamma$, and that $\beta <_{\mathrm{lex}} \beta'$ implies $\beta + \gamma <_{\mathrm{lex}} \beta' + \gamma$, and therefore $\beta \vee \gamma <_{\mathrm{lex}} \beta' \vee \gamma$; similarly, $\gamma <_{\mathrm{lex}} \gamma'$ implies $\beta \vee \gamma <_{\mathrm{lex}} \beta \vee \gamma'$. (The argument is similar to that of Theorem 6.2 below.) Thus the conditions of the conjecture are satisfied by this Ramsey ordering, and hence $A(G)$ is an integral domain. $\square$

6 Special cases (II): Wreath-A-like groups

We can now apply the same ideas used for the wreath-S-like case to the next class of groups, although the details are more intricate. The only essential difference between these two classes is that here we deal with ordered sequences of connected blocks instead of unordered multisets of connected blocks. We first define this class of groups and show that their algebras are integral domains. We then show that in the oligomorphic case, they have a structure similar to that of shuffle algebras, and deduce that they are polynomial rings. With this

information, we then look at some integer sequences which arise from this family of groups.

6.1 Wreath-A-like groups

If we have two finite sequences $S_1 = (x_1, \ldots, x_r)$ and $S_2 = (y_1, \ldots, y_s)$, then we write $S_1 \oplus S_2 = (x_1, \ldots, x_r, y_1, \ldots, y_s)$ for their concatenation.

Definition 6.1. We say that a permutation group G on Ω is *wreath-A-like* if there is a set of connected blocks $\{\Delta_i^{(j)}\}$ and a bijection ϕ from the set of orbits of G on finite subsets of Ω to the set of all finite sequences of connected blocks, with the bijection satisfying the following conditions:

(i) If $Y \subset \Omega$ is finite, then $\mathrm{wt}(\phi(Y)) = |Y|$.

(ii) If $Y \subset \Omega$ is finite and $\phi(Y) = (\Delta_{i_1}^{(j_1)}, \ldots, \Delta_{i_k}^{(j_k)})$, we can partition Y as an ordered union $Y = Y_1 \cup \cdots \cup Y_k$ with $|Y_l| = i_l$ for each l. Furthermore, if $Z \subseteq Y$ and $Z = Z_1 \cup \cdots \cup Z_k$, where $Z_l \subseteq Y_l$ for each l, then we can write $\phi(Z)$ as a concatenation of sequences $\phi(Z) = S_1 \oplus \cdots \oplus S_k$ where $\mathrm{wt}(S_l) = |Z_l|$ for each l, and $S_l = (\Delta_{i_l}^{(j_l)})$ if $Z_l = Y_l$.

As in the wreath-S-like case, condition (ii) implies that $\phi(Y_l) = (\Delta_{i_l}^{(j_l)})$ for $l = 1, 2, \ldots, k$.

Our prototypical family of groups for this class of groups are those of the form $G = H \operatorname{Wr} A$, where H is a permutation group on Δ, and A is the group of all order-preserving permutations of the rationals. Again, the wreath product action is the imprimitive one, so G acts on $\Omega = \Delta \times \mathbb{Q}$. As before, we take the connected blocks of weight n to be the orbits of the action of H on n-subsets of Δ. Then every orbit of G can be put into correspondence with a unique sequence of H-orbits as follows. If $Y \subset \Omega$ is an orbit representative, we can apply an element of the top group A to permute Y to a set of the form $(\Delta_1 \times \{1\}) \cup (\Delta_2 \times \{2\}) \cup \cdots \cup (\Delta_t \times \{t\})$, where each Δ_i is non-empty. Each of the Δ_i is a representative of some H-orbit, so we set $\phi(Y) = (\Delta_1, \Delta_2, \ldots, \Delta_t)$, again blurring the distinction between orbits and orbit representatives. It is again easy to see that conditions (i) and (ii) of the definition hold in this case.

Another example is the automorphism group of the random tournament. In this context, a tournament is a complete graph, every one of whose edges is directed, and the random tournament is the Fraïssé limit of the set of finite tournaments. A tournament is called strongly connected if there is a path between every ordered pair of vertices. It can be shown quite easily that every tournament can be decomposed uniquely as a sequence of strongly connected components, where the edges between components are all from earlier components to later ones. So here we take our set of connected blocks to be the isomorphism classes of finite strongly connected tournaments (and again, the weight of a connected block is the number of vertices in it), and if T is a finite subset of the random tournament, we set $\phi(T)$ to be the sequence of strongly connected components of T. Again, it is not difficult to see that conditions (i) and (ii) hold. Also, as in the case of the random graph, it may be that a sub-tournament has more components that the original tournament; for example, the cyclically-oriented 3-cycle is strongly connected, but any 2-element subset of it consists of two strongly connected 1-sets.

A third example is the automorphism group of the "generic pair of total orders". This is the Fraïssé limit of the class of finite sets, where each finite set carries two (unrelated) total orders, which can be taken as $a_1 < a_2 < \cdots < a_n$ and $a_{\pi(1)} \prec a_{\pi(2)} \prec \cdots \prec a_{\pi(n)}$ for some permutation $\pi \in S_n$. Thus orbits of the Fraïssé limit are described by permutations. We can take the connected blocks for this group to be the permutations $\pi \in S_n$ for which there exists no k with $0 < k < n$ such that π maps $\{1, \ldots, k\}$ to itself. The details of this example are not hard to check.

Theorem 6.2. *If G is wreath-A-like, then $A(G)$ is an integral domain.*

Proof. The proof runs along very similar lines to that of Theorem 5.2. If α is a sequence of connected blocks, we write $[\alpha]$ to denote the multiset whose elements are the terms of the sequence with their multiplicities. We define an ordering on sequences by $\alpha < \beta$ if $[\alpha] <_{\text{lex}} [\beta]$ or $[\alpha] = [\beta]$ and $\alpha >_{\text{lex}} \beta$.

Again, we show that the conditions of Conjecture 4.4 are satisfied in this case. Let m and n be positive integers and let N be a positive integer with $N \geqslant 2(m+n)$. Denoting the inverse of ϕ by ψ and letting α run through all sequences of connected blocks of total weight $m+n$, we set $c_\alpha = \psi(\alpha)$ and let X_α be a N-set in the orbit $\psi(\alpha \oplus (\Delta_1^{(1)}, \ldots, \Delta_1^{(1)}))$, where the second sequence has $N - (m+n)$ copies of $\Delta_1^{(1)}$. We claim that this gives a Ramsey ordering of the orbits on $(m+n)$-sets, where the sequences are ordered as described in the previous paragraph. Firstly, every $(m+n)$-set orbit appears in the list by hypothesis, as ψ is a bijection. Secondly, by construction, there is an $(m+n)$-subset of X_α in the orbit $\psi(\alpha)$, namely partition X_α as in condition (ii) of the definition, and remove all of the elements corresponding to the copies of $\Delta_1^{(1)}$ appended. This subset will then map to α under ϕ, by condition (ii).

To show the final condition of Ramsey orderings, we must show that any $(m+n)$-subset of X_α is in an orbit corresponding to a sequence less than or equal to α. Using the notation of condition (ii), we let $\alpha = (\Delta_{i_1}^{(j_1)}, \ldots, \Delta_{i_k}^{(j_k)})$ and $X_\alpha = X_1 \cup \cdots \cup X_k \cup X_{k+1} \cup \cdots \cup X_r$, where $X_{k+1}, \ldots, X_r$ correspond to the appended copies of $\Delta_1^{(1)}$. Consider a subset $Y = Y_1 \cup \cdots \cup Y_r \subset X_\alpha$ with $|Y| = m+n$. If $Y_l \neq X_l$ for some l with $X_l \neq \Delta_1^{(1)}$, then clearly $[\phi(Y)] <_{\text{lex}} [\alpha]$, as $\text{wt}(S_l) < i_l$, and the only new connected blocks which can be used are copies of $\Delta_1^{(1)}$, which is the least connected block. So the remaining case to consider is where some of the $\Delta_{i_l}^{(j_l)}$ are equal to $\Delta_1^{(1)}$, and for some or all of those, $Y_l = \varnothing$, whereas $Y_s = X_s$ for some $s > k$. But in such a case, while we have $[\phi(Y)] = [\alpha]$, it is clear that $\phi(Y) \geqslant_{\text{lex}} \alpha$. So in either case, we have $\phi(Y) \leqslant \alpha$, or equivalently $Y \leqslant c_\alpha$, as required.

We note that the induced Ramsey orderings on m-set orbits and n-set orbits are given by precisely the same construction; in particular, the orbit given by the sequence β first appears in X_α, where $\alpha = \beta \oplus (\Delta_1^{(1)}, \ldots, \Delta_1^{(1)})$.

Finally, we must show that the remaining conditions of the conjecture are satisfied by this Ramsey ordering. We will only show that $\beta < \beta'$ implies $\beta \vee \gamma < \beta' \vee \gamma$; the other condition follows identically. We first deduce an explicit description of $\beta \vee \gamma$.

A *shuffle* of two sequences, say $(x_1, \ldots, x_r)$ and $(y_1, \ldots, y_s)$, is a sequence $(z_1, \ldots, z_{r+s})$ for which there is a partition of $\{1, 2, \ldots, r+s\}$ into two disjoint sequences $1 \leqslant i_1 < i_2 < \cdots < i_r \leqslant r+s$ and $1 \leqslant j_1 < j_2 < \cdots < j_s \leqslant r+s$

with $z_{i_k} = x_k$ for $1 \leqslant k \leqslant r$ and $z_{j_k} = y_k$ for $1 \leqslant k \leqslant s$.

We first show that $\beta \vee \gamma$ is the lexicographically greatest shuffle of β with γ; this is not difficult although the argument is a little intricate. We let α_0 be this greatest shuffle and note that $[\alpha_0] = [\beta] + [\gamma]$. Now let α be any sequence of connected blocks for which c_α contains a $d_\beta \cup e_\gamma$ decomposition; we must show that $\alpha_0 \leqslant \alpha$. (Here d_β and e_γ are the orbits on m-sets and n-sets corresponding to β and γ respectively.)

We let $\alpha = (A_1, \ldots, A_k)$ be this sequence of connected blocks, and let Y be a representative of the orbit c_α. Write Y as an ordered union $Y = Y_1 \cup \cdots \cup Y_k$ as in condition (ii) of the definition of wreath-A-like groups. Then any decomposition of c_α into two subsets can be written as

$$c_\alpha = Z \cup Z' = (Z_1 \cup \cdots \cup Z_k) \cup (Z'_1 \cup \cdots \cup Z'_k),$$

where $Y_l = Z_l \cup Z'_l$ as a disjoint union for each l. Now if we require $\phi(Z) = \beta$ and $\phi(Z') = \gamma$, this means that the sequences $S_1 \oplus \cdots \oplus S_k$ and $S'_1 \oplus \cdots \oplus S'_k$ corresponding to Z and Z' respectively, as given by condition (ii), must equal β and γ respectively. If $\{Z_l, Z'_l\} = \{Y_l, \varnothing\}$, then $[S_l] + [S'_l] = [A_l]$ by condition (ii), but if not, then $[S_l] + [S'_l] <_{\mathrm{lex}} [A_l]$ by comparing weights. As $M_1 <_{\mathrm{lex}} M_2$ implies $M_1 + M <_{\mathrm{lex}} M_2 + M$ for any multisets M_1, M_2 and M, it follows that $[\beta] + [\gamma] \leqslant_{\mathrm{lex}} [\alpha]$ with equality if and only if $\{Z_l, Z'_l\} = \{Y'_l, \varnothing\}$ for each l, that is, $[\alpha_0] \leqslant_{\mathrm{lex}} [\alpha]$ with equality if and only if α is a shuffle of β and γ. And if α is such a shuffle, then $\alpha \leqslant_{\mathrm{lex}} \alpha_0$ by construction, so $\alpha_0 \leqslant \alpha$, as required.

Given this, we can now show that if $\beta < \beta'$, then $\beta \vee \gamma < \beta' \vee \gamma$. We first consider the case that $[\beta] <_{\mathrm{lex}} [\beta']$, from which it follows that $[\beta] + [\gamma] <_{\mathrm{lex}} [\beta'] + [\gamma]$. Since $[\beta \vee \gamma] = [\beta] + [\gamma]$ and $[\beta' \vee \gamma] = [\beta'] + [\gamma]$, we deduce that $[\beta \vee \gamma] <_{\mathrm{lex}} [\beta' \vee \gamma]$, so $\beta \vee \gamma < \beta' \vee \gamma$.

Now consider the other possible case, namely $[\beta] = [\beta']$ but $\beta >_{\mathrm{lex}} \beta'$. Note that $[\beta \vee \gamma] = [\beta' \vee \gamma]$ in this case, so we must show that $\beta \vee \gamma >_{\mathrm{lex}} \beta' \vee \gamma$. We let $\beta = (\Delta_1, \ldots, \Delta_r)$, $\beta' = (\Delta'_1, \ldots, \Delta'_r)$ and $\gamma = (E_1, \ldots, E_s)$ in the following. We also let $\alpha = \beta \vee \gamma = (A_1, \ldots, A_{r+s})$ and $\alpha' = \beta' \vee \gamma = (A'_1, \ldots, A'_{r+s})$. Recalling that $\beta \vee \gamma$ is the lexicographically greatest shuffle of β and γ, we can construct $\beta \vee \gamma$ by using the following merge-sort algorithm (written in pseudo-code).

function *MergeSort*(β, γ)
      { *We have* $\beta = (\Delta_1, \ldots, \Delta_r)$ *and* $\gamma = (E_1, \ldots, E_s)$ }
   $i \leftarrow 1$
   $j \leftarrow 1$
   while $i \leqslant r$ **or** $j \leqslant s$ **do**
     if $(i > r)$ **then** { $A_{i+j-1} \leftarrow E_j$; $j \leftarrow j+1$ }
     else if $(j > s)$ **then** { $A_{i+j-1} \leftarrow \Delta_i$; $i \leftarrow i+1$ }
     else if $(E_j \geqslant \Delta_i)$ **then** { $A_{i+j-1} \leftarrow E_j$; $j \leftarrow j+1$ }
     else { $A_{i+j-1} \leftarrow \Delta_i$; $i \leftarrow i+1$ }
   od
   return $\alpha = (A_1, \ldots, A_{r+s})$

Observe what happens if we run the algorithm on the pairs (β, γ) and (β', γ). Assume that $\Delta_i = \Delta'_i$ for $i < i_0$, but that $\Delta_{i_0} > \Delta'_{i_0}$. Then they will run identically as long as $i < i_0$. When $i = i_0$, they will both continue taking terms from γ until $E_j < \Delta_{i_0}$ or γ is exhausted. Once this happens, the (β, γ) algorithm will take Δ_{i_0} next, so $A_{i_0+j-1} = \Delta_{i_0}$, but the (β', γ) algorithm will

take $\max\{\Delta'_{i_0}, E_j\}$, so $A'_{i_0+j-1} = \max\{\Delta'_{i_0}, E_j\} < \Delta_{i_0} = A_{i_0+j-1}$. Thus we have $\beta \vee \gamma >_{\text{lex}} \beta' \vee \gamma$, so $\beta \vee \gamma < \beta' \vee \gamma$ as required.

It follows that $A(G)$ is an integral domain, as we wanted. $\qquad\square$

6.2 Shuffle algebras

In the oligomorphic case, we can do better: the algebra $A(G)$ is actually a polynomial algebra if G is an oligomorphic wreath-A-like group. We show this by noting strong similarities between our algebra and standard shuffle algebras, and using well-known properties of shuffle algebras, in particular that the Lyndon words form a polynomial basis for the shuffle algebra.

We start by briefly recalling the key facts we will need. We take these results from Reutenauer's book on free Lie algebras [11]. The references to definitions, theorems and so forth are to his book.

Let T be an alphabet. Although Reutenauer sometimes assumes the alphabet to be finite, it will be clear that all of the results we use below work equally well in the infinite case: since words are always of finite length and we only ever work with finitely many words at once, we can always restrict attention to the finite subset of T containing the letters in use.

We write T^* for the set of words in the alphabet T. We write $K\langle T\rangle$ for the K-vector space with basis T^*. If we use the concatenation product (where the product of two words is just their concatenation), then this is the ring of non-commuting polynomials over T. But there is another product that we can define on words, and by extension on $K\langle T\rangle$, called the *shuffle product*. This is explained in section 1.4 of Reutenauer, and we now essentially quote parts of it.

Let $w = a_1 \cdots a_n$ be a word of length n in T^*, and let $I \subseteq \{1, \ldots, n\}$. We denote by $w|I$ the word $a_{i_1} \cdots a_{i_k}$ if $I = \{i_1 < i_2 < \cdots < i_k\}$; in particular, $w|I$ is the empty word if $I = \varnothing$. (Such a word $w|I$ called a *subword* of w.) Note that when

$$\{1, \ldots, n\} = \bigcup_{j=1}^{p} I_j,$$

then w is determined by the p words $w|I_j$ and the p subsets I_j.

Given two words u_1 and u_2 of respective lengths n_1 and n_2, their *shuffle product*, denoted by $u_1 \shuffle u_2$, is the polynomial

$$u_1 \shuffle u_2 = \sum w(I_1, I_2),$$

where the sum is taken over all pairs (I_1, I_2) of disjoint subsets of $\{1, \ldots, n\}$ with $I_1 \cup I_2 = \{1, \ldots, n\}$ and $|I_j| = n_j$ for $j = 1$, 2, and where the word $w = w(I_1, I_2)$ is defined by $w|I_j = u_j$ for $j = 1$, 2. Note that $u_1 \shuffle u_2$ is a sum of words of length n, each with the same multiset of letters, and so is a homogeneous polynomial of degree n. Note also that the empty word, denoted by 1, is the identity for the shuffle product, that the shuffle product is commutative and associative, and that it is distributive with respect to addition. Thus $K\langle T\rangle$ with the shuffle product is a commutative, associative algebra, called the *shuffle algebra*.

Using the associative and distributive properties of the shuffle product, we can also give an expression for the shuffle product of the words $u_1, \ldots, u_p$, of

respective lengths $n_1, \ldots, n_p$; their shuffle product is the polynomial

$$u_1 \sqcup\!\sqcup \cdots \sqcup\!\sqcup u_p = \sum w(I_1, \ldots, I_p),$$

where now the sum is taken over all p-tuples $(I_1, \ldots, I_p)$ of pairwise disjoint subsets of $\{1, \ldots, n\}$ with $\bigcup_{i=1}^{p} I_j = \{1, \ldots, n\}$ and $|I_j| = n_j$ for each $j = 1, \ldots, p$, and where the word $w = w(I_1, \ldots, I_p)$ is defined by $w|I_j = u_j$ for each $j = 1, \ldots, p$.

A word appearing in the shuffle product $u_1 \sqcup\!\sqcup \cdots \sqcup\!\sqcup u_p$ is called a *shuffle* of $u_1, \ldots, u_p$. Note that this is consistent with the definition of shuffle we used in the proof of Theorem 6.2 above. As an example, if $a, b, c \in T$, then $ab \sqcup\!\sqcup ac = abac + 2aabc + 2aacb + acab$, and $aabc$ and $acab$ are both shuffles of ab and ac.

The next definition we need is that of a Lyndon word. Assume that our alphabet T is totally ordered. Then a *Lyndon word* in T^* is a non-empty word which is lexicographically smaller than all of its nontrivial proper right factors; in other words, w is a Lyndon word if $w \neq 1$ and if for each factorisation $w = uv$ (concatenation product) with $u, v \neq 1$, one has $w <_{\text{lex}} v$.

An alternative categorisation of Lyndon words is as follows (Corollary 7.7 in Reutenauer). Given a word $w = a_1 \cdots a_n$ of length n, we can define the rotation operator ρ by $\rho(w) = a_2 \cdots a_n a_1$. Then a word w of length $n \geqslant 1$ is Lyndon if and only if $w <_{\text{lex}} \rho^k(w)$ for $k = 1, \ldots, n-1$, which is to say that w is primitive (it does not have the form $w = u^r$ for some $r > 1$) and that it is lexicographically smaller than any rotation (cyclic permutation) of itself. It follows that Lyndon words are in bijective correspondence with primitive necklaces; see [11, Chap. 7] for more information.

A key property of Lyndon words is that every word $w \in T^*$ can be written *uniquely* as a decreasing product of Lyndon words, so $w = l_1^{r_1} \cdots l_k^{r_k}$, where $l_1 >_{\text{lex}} \cdots >_{\text{lex}} l_k$ and $r_1, \ldots, r_k \geqslant 1$. (This follows from Theorem 5.1 and Corollary 4.4, and can also easily be proved directly—see section 7.3.)

Finally, Theorem 6.1 states that the shuffle algebra $K\langle T \rangle$ is a polynomial algebra generated by the Lyndon words, and that for each word w, written as a decreasing product of Lyndon words $w = l_1^{r_1} \cdots l_k^{r_k}$ as in the previous paragraph, one has

$$S(w) \stackrel{\text{def}}{=} \frac{1}{r_1! \cdots r_k!} l_1^{\sqcup\!\sqcup r_1} \sqcup\!\sqcup \cdots \sqcup\!\sqcup l_k^{\sqcup\!\sqcup r_k} = w + \sum_{\substack{[u]=[w] \\ u <_{\text{lex}} w}} \alpha_u u, \tag{7}$$

for some non-negative integers α_u, where $l^{\sqcup\!\sqcup r}$ means $l \sqcup\!\sqcup \cdots \sqcup\!\sqcup l$ with r terms in the product, and, in this context, $[u]$ means the multiset of letters in the word u.

Note that it is equation (7) which proves that $K\langle T \rangle$ is a polynomial algebra: the set T^* is a K-vector space basis for $K\langle T \rangle$, and given any finite multiset M of elements of T, the matrix relating the basis elements $\{\, w : w \in T^* \text{ and } [w] = M \,\}$ to $\{\, S(w) : w \in T^* \text{ and } [w] = M \,\}$ is unitriangular when the words are listed in lexicographic order, so that $\{\, S(w) : w \in T^* \,\}$ also forms a basis for $K\langle T \rangle$. This argument is true whether T is finite or infinite.

We can now apply this to our case of oligomorphic wreath-A-like permutation groups. Let G acting on Ω be such a group, as in Definition 6.1 above. We obviously take our alphabet T to be the set of connected blocks of the action (as given by the definition of wreath-A-like groups), so that T^* corresponds

bijectively to the set of orbits of G on finite subsets of Ω. The alphabet T has the standard ordering defined on connected blocks, and the set T^* can then be ordered either by the lexicographic order (denoted $<_{\text{lex}}$) or by the order we defined at the start of Theorem 6.2 (denoted $<$).

Clearly $A(G)$ can be regarded as a K-vector space, with the set of characteristic functions of finite orbits as basis. We will identify the connected block sequence $w = (\Delta_{i_1}^{(j_1)}, \ldots, \Delta_{i_k}^{(j_k)})$ with the characteristic function of the corresponding orbit, writing w for both. Via this correspondence, we can identify $A(G)$ with $K\langle T \rangle$ as vector spaces. The grading on $A(G)$ induces a grading on $K\langle T \rangle$: the homogeneous component $V_n(G)$ is identified with the subspace of $K\langle T \rangle$ spanned by $\{ w \in T^* : \text{wt}(w) = n \}$. We then consider the product that the vector space $K\langle T \rangle$ inherits via this identification. Let $v \in T^*$ be another connected block sequence. We write $v \,\overline{\sqcup}\, w$ for the product in $A(G)$ and the induced product in $K\langle T \rangle$. The notation is designed to indicate that this product is related to the shuffle product, as we will see, and we call it the *complete shuffle product*. (It is also somewhat related to the infiltration product on $K\langle T \rangle$; see [11, sect. 6.3].) Recalling the definition of multiplication in $A(G)$, we see that for any finite subset $X \subset \Omega$ with $|X| = \text{wt}(v) + \text{wt}(w)$,

$$(v \,\overline{\sqcup}\, w)(X) = \sum_{\substack{Y \subseteq X \\ |Y| = \text{wt}(v)}} v(Y) w(X \setminus Y).$$

But $v(Y)$ is none other than the characteristic function which has value 1 if $\phi(Y) = v$ and 0 otherwise, and similarly for $w(Y \setminus X)$. So we have

$$(v \,\overline{\sqcup}\, w)(X) = |\{ Y \subseteq X : \phi(Y) = v, \ \phi(X \setminus Y) = w \}|.$$

Thus, setting $u = \phi(X)$ and writing $u \to v \cup w$ if there is a $Y \subseteq X$ with $\phi(Y) = v$ and $\phi(X \setminus Y) = w$, we have

$$v \,\overline{\sqcup}\, w = \sum_{u \in T^*} \beta_u u,$$

where $\beta_u > 0$ if $u \to v \cup w$ and $\beta_u = 0$ otherwise.

Now we can characterise those u for which $u \to w \cup v$ quite easily. Firstly, consider the case that $[u] = [w] + [v]$, that is, the set of connected blocks of u is the same as those of w and v combined. Then $u \to w \cup v$ if and only if u is a shuffle of w and v, by condition (ii) of Definition 6.1, as in the proof of Theorem 6.2. In fact, the terms in $w \,\overline{\sqcup}\, v$ with $[u] = [w] + [v]$ will be precisely $w \sqcup v$, which is easy to see. Now consider those terms with $[u] \neq [w] + [v]$. If $[u] <_{\text{lex}} [w] + [v]$, then it is easy to see that we cannot have $u \to w \cup v$, but it may be possible otherwise. We deduce that our product is given by:

$$w \,\overline{\sqcup}\, v = w \sqcup v + \sum_{\substack{\text{wt}(u) = \text{wt}(w) + \text{wt}(v) \\ [u] >_{\text{lex}} [w] + [v]}} \beta_u u \tag{8}$$

for some non-negative integers β_u.

Now given $w = l_1^{r_1} \cdots l_k^{r_k}$ written as a (concatenation) product of decreasing Lyndon words, we can consider the complete shuffle product as we did for the

normal shuffle product above:

$$
\begin{aligned}
\bar{S}(w) &\overset{\text{def}}{=} \frac{1}{r_1! \cdots r_k!}\, l_1^{\boxdot r_1} \boxdot \cdots \boxdot l_k^{\boxdot r_k} \\
&= \frac{1}{r_1! \cdots r_k!}\, l_1^{\sqcup\!\sqcup r_1} \sqcup\!\sqcup \cdots \sqcup\!\sqcup l_k^{\sqcup\!\sqcup r_k} + \sum_{\substack{\text{wt}(u)=\text{wt}(w) \\ [u] >_{\text{lex}} [w]}} \beta_u u \\
&= w + \sum_{\substack{[u]=[w] \\ u <_{\text{lex}} w}} \alpha_u u + \sum_{\substack{\text{wt}(u)=\text{wt}(w) \\ [u] >_{\text{lex}} [w]}} \beta_u u \\
&= w + \sum_{\substack{\text{wt}(u)=\text{wt}(w) \\ u > w}} \alpha_u u,
\end{aligned}
\tag{9}
$$

where the α_u and the β_u are non-negative integers. To get the second line, we have repeatedly used equation (8) to reduce the complete shuffle product to a normal shuffle product. Observe that $\text{wt}(l_1^{r_1} \cdots l_k^{r_k}) = \text{wt}(w)$, hence the sum is over words with $\text{wt}(u) = \text{wt}(w)$, and with $[u] >_{\text{lex}} [w]$, since $>_{\text{lex}}$ is transitive and $[u_1] >_{\text{lex}} [u_2]$ implies $[u_1] + [u] >_{\text{lex}} [u_2] + [u]$ for any word u. That the β_u are non-negative is easy to see, and it is not that much harder to see that they are integral, although we do not need this. In the third line, we have used equation (7), and in the last line, we have set $\alpha_u = \beta_u$ in the case that $[u] >_{\text{lex}} [w]$, and used the relation on words (sequences) defined in the previous section, namely $u > w$ if $[u] >_{\text{lex}} [w]$ or $[u] = [w]$ and $u <_{\text{lex}} w$.

It is also important to note that in our case, the set $\{\, u : \text{wt}(u) = \text{wt}(w) \,\}$ is finite, as there are only finitely many connected blocks of each weight, the same number as the number of orbits on sets of size $\text{wt}(w)$, so that the sums in equation (9) are all finite.

We now see, as above, that the matrix relating $\{\, w : w \in T^* \text{ and } \text{wt}(w) = n \,\}$ to $\{\, \bar{S}(w) : w \in T^* \text{ and } \text{wt}(w) = n \,\}$ is unitriangular when the words of weight n are listed in the order we have defined. It follows that the $\bar{S}(w)$ form a vector space basis for $A(G) = K\langle T \rangle$, and hence the set of Lyndon words is a set of polynomial generators for $A(G)$. We summarise these results as a theorem.

Theorem 6.3. *If G is an oligomorphic wreath-A-like permutation group, then $A(G)$ is a polynomial ring, and the generators are those characteristic functions on orbits corresponding to Lyndon words as described above.* $\square$

We can now deduce:

Corollary 6.4. *If G is an oligomorphic wreath-A-like permutation group, then the element $\varepsilon \in V_1(G)$ is prime in $A(G)$.*

Proof. We have $e = \Delta_1^{(1)} + \cdots + \Delta_1^{(r)}$, where the $\Delta_1^{(j)}$ are the orbits on 1-sets. As each of the $\Delta_1^{(j)}$ is a Lyndon word, $A(G) = K[\Delta_1^{(1)}, \ldots, \Delta_1^{(r)}, \Delta_2^{(1)}, \ldots]$. It follows that we can replace the polynomial generator $\Delta_1^{(1)}$ by ε (as they are linearly related), giving $A(G) = K[\varepsilon, \Delta_1^{(2)}, \ldots, \Delta_1^{(r)}, \Delta_2^{(1)}, \ldots]$. It is clear, since we then have $A(G)/(\varepsilon) \cong K[\Delta_1^{(2)}, \ldots, \Delta_1^{(r)}, \Delta_2^{(1)}, \ldots]$, that $A(G)/(\varepsilon)$ is an integral domain, so ε is prime in $A(G)$. $\square$

6.3   Integer sequences, necklaces and free Lie algebras

Theorem 6.3 leads us to revisit some counting questions. Cameron [5] considered the following question. If the algebra $A(G)$ corresponding to an "interesting" oligomorphic group G were polynomial, what would be the sequence counting the number of polynomial generators of each degree? From knowledge of the dimension of each homogeneous component of $A(G)$, the answer can be determined using the inverse Euler transform. Now that we have an explicit description of the polynomial generators in the wreath-A-like case, an examination of the sequences observed might yield some interesting new information about those sequences.

The two sequences we will consider are those arising from the groups $S_2 \operatorname{Wr} A$ and $A \operatorname{Wr} A$, both of which appear in the On-Line Encyclopedia of Integer Sequences [12]. There are some obvious generalisations to other groups, as we observe below. The n-th homogeneous component of the group $S_2 \operatorname{Wr} A$ has dimension F_{n+1} (a Fibonacci number, where $F_0 = 0$ and $F_1 = 1$), and so the sequence counting the number of generators of degree n is A006206, beginning $1, 1, 1, 1, 2, 4, 5, 8, 11, 18, \ldots$. By our result, the n-th term of this sequence gives the number of Lyndon words of weight n (starting with $n = 1$) in the alphabet $T = \{\Delta_1, \Delta_2\}$, where Δ_1 and Δ_2 have respective weights 1 and 2.

Similarly, for the group $A \operatorname{Wr} A$, the n-th homogeneous component has dimension 2^{n-1} for $n \geqslant 1$, and the sequence counting the number of generators of degree n is A059966, beginning $1, 1, 2, 3, 6, 9, 18, 30, \ldots$. (Note that the paper quoted above had sequence A001037 by mistake, this being the inverse Euler transform of the closely related sequence (2^n).) This sequence then counts the number of Lyndon words of weight n in the alphabet $T = \{\Delta_1, \Delta_2, \ldots\}$, where Δ_i has weight i.

The Encyclopedia entry gives a different explanation, however: this sequence lists the dimensions of the homogeneous components of the free Lie algebra with one generator of each degree 1, 2, 3, etc. The connection between these two descriptions of this sequence is easy to describe, using [11, Thm. 4.9]. Let T be an alphabet whose letters each have a positive integral degree/weight (we use these terms interchangeably in this section), and where there are only finitely many letters of each possible weight. There is a basis of the free Lie algebra on the alphabet T (viewed as a vector space) given by $\{ P_w : w \in T^* \text{ Lyndon} \}$, where $P_a = a$ if $a \in T$, and $P_w = [P_u, P_v]$ otherwise, where $w = uv$ with v being the lexicographically smallest nontrivial proper right factor of w (see [11, Thm. 5.1]). Note that it trivially follows by induction that the degree of the homogeneous polynomial P_w is $\operatorname{wt}(w)$. Thus the dimension of the homogeneous component of degree n of the free Lie algebra on the alphabet T is the number of Lyndon words in T^* of weight n. It follows that we can also describe the two sequences above as either the number of Lyndon words of weight n in the alphabets $\{\Delta_1, \Delta_2\}$ and $\{\Delta_1, \Delta_2, \ldots\}$ respectively, or as the number of primitive necklaces of weight n in these symbols, or as the dimension of the homogeneous component of degree n of the free Lie algebras on these sets. This obviously generalises to other wreath-A-like groups.

We may ask other counting questions based on these ideas. We start with an alphabet of weighted letters T (again with only finitely many letters of each weight). The primary questions arising are how to transform between the three

sequences:

$$a_n = \text{number of letters of weight } n \text{ in } T,$$
$$w_n = \text{number of words of weight } n \text{ in } T^*,$$
$$l_n = \text{number of Lyndon words of weight } n \text{ in } T^*.$$

(Of course, l_n can also be regarded as the number of primitive necklaces of weight n in this alphabet.) In our context, a_n is the number of connected blocks of weight n in our wreath-A-like group, w_n gives the dimension of the homogeneous component of weight n in $A(G)$ and l_n gives the number of polynomial generators of weight n in $A(G)$. We use the notation and some of the ideas presented in Bernstein and Sloane's paper on integer sequences [1].

The transformation between (a_n) and (w_n) can be effected by INVERT, as every word is an ordered sequence of letters:

$$1 + \sum_{n=1}^{\infty} w_n x^n = \frac{1}{1 - \sum_{n=1}^{\infty} a_n x^n}.$$

The transformation between (w_n) and (l_n) is performed using EULER, as every word is a product of a decreasing sequence of Lyndon words, so can be identified with a multiset of Lyndon words:

$$1 + \sum_{n=1}^{\infty} w_n x^n = \prod_{n=1}^{\infty} \frac{1}{(1 - x^n)^{l_n}}.$$

It follows that we can transform between (a_n) and (l_n) using a variant of WEIGH:

$$1 - \sum_{n=1}^{\infty} a_n x^n = \prod_{n=1}^{\infty} (1 - x^n)^{l_n}. \tag{10}$$

Most of the six possible conversions between (a_n), (w_n) and (l_n) are straightforward given these formulæ; the two which are harder are converting (w_n) and (a_n) to (l_n). Inverting the EULER transform is explained in [1]; we apply the same idea to convert from (a_n) to (l_n).

Given a sequence (a_n), we introduce the auxiliary sequence (c_n) defined by the equation $1 - \sum_{n=1}^{\infty} a_n x^n = \exp\left(-\sum_{n=1}^{\infty} c_n x^n/n\right)$. Using the generating functions $A(x) = \sum_{n=1}^{\infty} a_n x^n$ and $C(x) = \sum_{n=1}^{\infty} c_n x^n$, we can perform standard manipulations using the defining equation for (c_n) to deduce that $C(x) = xA'(x) + C(x)A(x)$. It follows that

$$c_n = na_n + \sum_{k=1}^{n-1} c_k a_{n-k}. \tag{11}$$

Now substituting $\exp\left(-\sum c_n x^n/n\right)$ for $1 - \sum a_n x^n$ in equation (10), taking logarithms and expanding as a power series gives the coefficient of x^n/n to be $c_n = \sum_{d|n} d\, l_d$. Finally, Möbius inversion gives

$$l_n = \frac{1}{n} \sum_{d|n} \mu(n/d) c_d. \tag{12}$$

Thus we have an effective way of calculating the number of Lyndon words of a given weight given the number of letters of each possible weight.

As an interesting example of this process, let us consider our favourite group, $G = S_2 \operatorname{Wr} A$. In this case, recall that we have $T = \{\Delta_1, \Delta_2\}$, so $a_1 = a_2 = 1$ and $a_n = 0$ for $n \geqslant 3$. Then the sequence (c_n) is calculated by equation (11): we have $c_1 = 1$ and $c_2 = 3$. For $n \geqslant 3$, we have $c_n = c_{n-1} + c_{n-2}$, so (c_n) is the standard Lucas sequence (L_n): 1, 3, 4, 7, 11, 18, ... We can now calculate the sequence (l_n): the first few terms are as we predicted: 1, 1, 1, 1, 2, 2, 4, 5, ..., and a general formula is $l_n = \frac{1}{n} \sum_{d|n} \mu(n/d) L_d$, as is given in the Encyclopedia entry for A006206. One interesting thing to observe is that if p is prime, then we have $l_p = (\mu(1)L_p + \mu(p)L_1)/p = (L_p - 1)/p$. It follows that the Lucas sequence satisfies $L_p \equiv 1 \pmod{p}$ for all primes p, a known result (see Hoggart and Bicknell [7]), but somewhat surprising in this context.

The description of our sequence A006206 in the Encylcopedia is "aperiodic binary necklaces [of length n] with no subsequence 00, excluding the sequence '0'." Our description is that it counts primitive necklaces of weight n in the alphabet $\{\Delta_1, \Delta_2\}$. These are easily seen to be equivalent: if we replace every Δ_1 by the symbol 1 and every Δ_2 by the symbols 10 (in clockwise order, say), then we will get a primitive (aperiodic) binary necklace with no subsequence 00 whose length equals the weight of the necklace we started with, and we can perform the inverse transformation equally simply (as we are excluding the necklace 0). We can do the same with the group $S_n \operatorname{Wr} A$, enabling us to count the number of primitive binary necklaces of length n with no subsequence $00 \cdots 0$ (with n zeros) and excluding the necklace 0.

Now let us apply these ideas to the case $G = A \operatorname{Wr} A$. Firstly, the auxilary sequence turns out to be $c_n = 2^n - 1$, and the sequence (l_n) is given by $l_n = \sum_{d|n} \mu(n/d)(2^d - 1)$. This can be simplified using the result $\sum_{d|n} \mu(n/d) = [n = 1]$, where we are using Iverson's convention that if P is a predicate, then $[P] = 1$ if P is true and 0 otherwise. So we have $l_n = \sum_{d|n} \mu(n/d)2^d - [n = 1]$. The sequence given by $\sum_{d|n} \mu(n/d)2^d$ is sequence A001037, and so our sequence differs from it by 1 in the $n = 1$ term only, yielding the observed sequence A059966. We can also give a necklace description of this sequence as above: it is the number of primitive binary necklaces of length n excluding the necklace 0—the sequence A001037 is essentially the same, but does not exclude the necklace 0, so it it also counts the number of binary Lyndon words of *length* n. (These are the descriptions of this sequence given in the Encyclopedia.) Finally, as above, if we consider the term l_p for p prime, we see that $l_p = ((2^p - 1) - 1)/p = 2(2^{p-1} - 1)/p$, so for $p > 2$, we deduce Fermat's little theorem for base 2, that is $2^{p-1} \equiv 1 \pmod{p}$.

An investigation of those sequences of non-negative integers (b_n) for which $\frac{1}{n} \sum_{d|n} \mu(n/d)b_d$ is a non-negative integer for all n has been undertaken by Puri and Ward [10], who call them exactly realizable. We can thus add to their work a class of exactly realizable sequences: those which are of the form (c_n), where (c_n) is given by equation (11) for some sequence of non-negative integers (a_n). A particular family of such sequences is given by $a_i = 1$ for $1 \leqslant i \leqslant n$ and $a_i = 0$ for $i > n$; these are sometimes known as "generalised Fibonacci sequences", and have been discussed by Du [6] (where this sequence is called ϕ_n). It would be interesting to know whether new congruence identities can be discovered by applying this technique to some of the sequences identified there or to sequences

produced by other wreath-A-like groups.

7   Non-oligomorphic groups

Throughout this paper, we have mostly focused on oligomorphic groups, proving results in general where there was no problem in doing so. In this final section, we consider briefly the issues arising in the non-oligomorphic case.

As has already been pointed out above, the group $\mathbb{Z}$ acting regularly on $\mathbb{Z}$ does not have a Ramsey ordering on 2-sets, so much of what we did above will not help us to understand the algebra $A(\mathbb{Z})$. It is easy to construct other similar examples.

A more difficult question is whether we have even got the "right" definition of the algebra $A(G)$ in the non-oligomorphic case. The definition we have been using was introduced specifically to study the behaviour of oligomorphic groups. There are two finiteness conditions which can be imposed on the algebra we consider.

Firstly, we have taken the direct sum $A(G) = \bigoplus_{n=0}^{\infty} V_n(G)$, which is the direct limit as $N \to \infty$ of the vector spaces $\bigoplus_{n=0}^{N} V_n(G)$ (with the obvious direct maps). We could have instead taken the cartesian sum $\sum_{n=0}^{\infty} V_n(G)$, being the inverse limit of the same family of vector spaces (with the obvious inverse maps).

Secondly, and independently of the first choice, we could either take $V_n(G)$ to be the vector space of all functions from n-subsets of Ω to K which are fixed by G, as we have until now, or we could take it to be the subspace of this consisting of those functions which assume only finitely many distinct values on n-sets. (The latter idea was suggested to me by Peter Cameron.) Note, though, that if there are infinitely many orbits on n-sets, this vector space will still have uncountable dimension. It is not hard to check that if we use the latter definition, the multiplication in the algebra is still well-defined. Also, this distinction does not exist in the oligomorphic case. (Another seemingly plausible choice, those functions in $V_n(G)$ which are non-zero on only finitely many orbits of G, can fail to produce a well-defined multiplication: consider, for example, the case of e^2 with our favourite non-oligomorphic group, $\mathbb{Z}$: it takes the value 2 on every 2-set.)

Thus we have four plausible algebras to choose from, and it is not clear which is the "correct" one to use. More work is still required in this area.

Acknowledgements

References

[1] M. Bernstein and N. J. A. Sloane, *Some canonical sequences of integers*, Linear Algebra and Applications **226/228** (1995), 57–72.

[2] Peter J. Cameron, *Orbits of permutation groups on unordered sets, II*, J. London Math. Soc. (2) **23** (1981), 249–264.

[3] ———, *Oligomorphic Permutation Groups*, Cambridge University Press, 1990.

[4] ———, *The algebra of an age*, Model Theory of Groups and Automorphism Groups (David M. Evans, ed.), Cambridge University Press, 1997, pp. 126–133.

[5] ———, *Sequences realized by oligomorphic permutation groups*, Journal of Integer Sequences **3** (2000), Article 00.1.5.

[6] B.-S. Du, *A simple method which generates infinitely many congruence identities*, Fibonacci Quart. **27** (1989), 116–124.

[7] V. E. Hoggart Jr. and M. Bicknell, *Some congruences of the fibonacci numbers modulo a prime p*, Math. Mag. **47** (1974), 210–214.

[8] William M. Kantor, *On incidence matrices of finite projective and affine spaces*, Mat Z. **124** (1972), 315–318.

[9] Peter M. Neumann, *The structure of finitary permutation groups*, Arch. Math. (Basel) **27** (1976), 3–17.

[10] Yash Puri and Thomas Ward, *Arithmetic and growth of periodic orbits*, Journal of Integer Sequences **4** (2001), Article 01.2.1.

[11] Christophe Reutenauer, *Free Lie Algebras*, Oxford University Press, 1993.

[12] N. J. A. Sloane, *The On-Line Encyclopedia of Integer Sequences*, available at `http://www.research.att.com/~njas/sequences/`, 2001.