Séminaire Lotharingien de Combinatoire **84B** (2020)
Article #79, 12 pp.

Proceedings of the 32nd Conference on Formal Power Series and Algebraic Combinatorics (Online)

Word Measures on Symmetric Groups

Liam Hanany*[0] and Doron Puder†[0]

[0]*School of Mathematical Sciences, Tel Aviv University, Tel Aviv, Israel*

Abstract. Fix a word w in a free group F_r on r generators. A w-random permutation in the symmetric group S_n is obtained by sampling r independent uniformly random permutations $\sigma_1, \ldots, \sigma_r \in S_n$ and evaluating $w(\sigma_1, \ldots, \sigma_r)$. In (Puder 2014, Puder–Parzanchevski 2015) it was shown that the average number of fixed points in a w-random permutation is $1 + \theta\left(n^{1-\pi(w)}\right)$, where $\pi(w)$ is the smallest rank of a subgroup $H \leq F_r$ containing w as a non-primitive element. We show that $\pi(w)$ plays a role in estimates of other natural families of characters. In particular, we show that for all $s \geq 2$, the average number of s-cycles is $\frac{1}{s} + O\left(n^{-\pi(w)}\right)$.

Keywords: symmetric groups, free groups, random permutations

1 Introduction

Let $\mathbf{F}_r$ denote the free group on r generators. A word $w \in \mathbf{F}_r$ induces a mapping on any finite group, $w : G^r \to G$, by substituting the letters of w with elements of G. This map defines a distribution on the group G: the push forward of the uniform distribution on G^r. Equivalently, this distribution is the normalized number of times each element in G is obtained by a substitution in w. We call this distribution the w-measure on G. For example, if $w = xyxy^{-2}$, a w-random element in G is $ghgh^{-2}$ where g, h are independent, uniformly random elements of G. In this paper, we restrict our attention to word measures on the symmetric groups S_n.

More concretely, we study these measures using non-abelian Fourier theory: given a character[1] χ of G, we compute $\mathbb{E}_w[\chi]$, the average value of this character under the w-measure on G. Word measures are constant on conjugacy classes of G, i.e. are class functions on the group G, and every class function is a linear combination of irreducible characters. Therefore, the expressions $\mathbb{E}_w[\chi]$, running over all characters χ, uniquely determine the resulting word measure.

Our focus is on the following class functions of S_n. Given $k_1, \ldots, k_\ell \in \mathbb{Z}_{\geq 1}$, denote

$$\chi_{k_1,\ldots,k_\ell}(\sigma) \stackrel{\text{def}}{=} \#\text{fix}\left(\sigma^{k_1}\right) \cdot \ldots \cdot \#\text{fix}\left(\sigma^{k_\ell}\right), \tag{1.1}$$

*liamhanany@mail.tau.ac.il. Supported by Israel Science Foundation (ISF) via grant 1071/16
†doronpuder@gmail.com. Supported by Israel Science Foundation (ISF) via grant 1071/16
[1]A character χ of the group G is given by $\text{tr} \circ \rho$, where $\rho : G \to \text{GL}_d(\mathbb{C})$ is some homomorphism.

where $\#\mathrm{fix}(\tau)$ is the number of fixed points of the permutation τ. When we write $\mathbb{E}_w\left[\chi_{k_1,\ldots,k_\ell}\right]$ there is also a suppressed parameter n. In fact, $\mathbb{E}_w\left[\chi_{k_1,\ldots,k_\ell}\right]$ is a map $\mathbb{N}\to\mathbb{Q}$, where n is mapped to the average value of this character under the w-measure on S_n.

For every word $w\in\mathbf{F}_r$ and $k_1,\ldots,k_\ell\in\mathbb{Z}_{\geq 1}$, the expectation $\mathbb{E}_w\left[\chi_{k_1,\ldots,k_\ell}\right]$ is a rational function of n, for large enough n: this is essentially a result of Nica [8], and see also Section 4 and especially Remark 31 in [4]. For example, $\mathbb{E}_{xyx^{-1}y^{-1}}\left[\chi_{1,2}\right]=3+\frac{4\left(n^4-9n^3+23n^2-13n-1\right)}{n(n-1)(n-2)(n-3)(n-5)}$ for all $n\geq 6$. In particular, for large enough n, $\mathbb{E}_w\left[\chi_{k_1,\ldots,k_\ell}\right]$ can be written as a Laurent series in n. Our main goal in this paper is to estimate the leading terms of this Laurent series expansion. The special case of $\chi_1=\#\mathrm{fix}(\sigma)$, the average number of fixed points, was studied in [9, 10]. These papers show a connection between $\mathbb{E}_w\left[\chi_1\right]$ and invariants of w as an element of the free group.

In order to explain these invariants, we need a few notions from combinatorial group theory and the study of free groups. A generating set of minimal size of a finitely generated free group is called a *basis*. An element $w\in\mathbf{F}$ is called *primitive* if it belongs to a basis of $\mathbf{F}$. The rank of the free group $\mathbf{F}$, denoted $\mathrm{rk}\mathbf{F}$, is the size of a basis of $\mathbf{F}$. The classical Nielsen-Schreier theorem states that subgroups of free groups are free. The primitivity rank of a word, which plays an important role in this paper, was first introduced in [9]:

Definition 1.1. The primitivity rank $\pi(w)$ of a word $w\in\mathbf{F}_r$ is the minimal rank of a subgroup $H\leq\mathbf{F}_r$ containing w as a non-primitive element. If there are no such subgroups, set $\pi(w)=\infty$. We also consider the set of critical subgroups of w defined as

$$\mathrm{Crit}(w)=\left\{H\leq\mathbf{F}\,\middle|\,\mathrm{rk}H=\pi(w),H\ni w\text{ and }w\text{ non}-\text{primitive in }H\right\}.$$

For example, $\pi(w)=0\iff w=1$ as the trivial word is contained in the trivial subgroup but not as a primitive element. Words with $\pi(w)=1$ are precisely proper powers and if $u\in\mathbf{F}_r$ is not a proper power and $m\geq 2$, then $\mathrm{Crit}(u^m)=\left\{\left\langle u^d\right\rangle\,\middle|\,d\mid m,\ d<m\right\}$. Finally, $\pi(w)=\infty$ if and only if w is primitive in $\mathbf{F}_r$, and in any other case $\pi(w)\leq r$ [9, Lemma 4.1]. The set $\mathrm{Crit}(w)$ is always finite [10, Section 4]. We can now state the aforementioned result from [10].

Theorem 1.2 ([10, Theorem 1.8])**.** *For every $w\in\mathbf{F}_r$*

$$\mathbb{E}_w\left[\#\mathrm{fix}(\sigma)\right]=1+\frac{|\mathrm{Crit}(w)|}{n^{\pi(w)-1}}+O\left(\frac{1}{n^{\pi(w)}}\right).$$

Since the expected number of fixed points in a uniformly random permutation is 1, the theorem can be restated as

$$\mathbb{E}_w\left[\chi_1\right]=\mathbb{E}_{\mathrm{unif}}\left[\chi_1\right]+\frac{|\mathrm{Crit}(w)|}{n^{\pi(w)-1}}+O\left(\frac{1}{n^{\pi(w)}}\right),$$

where $\mathbb{E}_{\text{unif}}[f]$ is the expectation of a function $f\colon S_n \to \mathbb{R}$ with respect to the uniform distribution on S_n. In this paper we prove the following generalization of Theorem 1.2:

Theorem 1.3. *For every non-power $w \in \mathbf{F}_r$ and every $k_1, \ldots, k_\ell \in \mathbb{Z}_{\geq 1}$, there exists a positive integer $C_{k_1,\ldots,k_\ell} \in \mathbb{N}$ such that*

$$\mathbb{E}_w\left[\chi_{k_1,\ldots,k_\ell}\right] = \mathbb{E}_{\text{unif}}\left[\chi_{k_1,\ldots,k_\ell}\right] + \frac{C_{k_1,\ldots,k_\ell}\cdot|\text{Crit}(w)|}{n^{\pi(w)-1}} + O\left(\frac{1}{n^{\pi(w)}}\right).$$

Moreover, the constant $C_{k_1,\ldots,k_\ell}$ is equal to the scalar product

$$\left\langle \chi_{k_1,\ldots,k_\ell}, \chi_1 - 1\right\rangle = \frac{1}{n!}\sum_{\sigma\in S_n}\chi_{k_1,\ldots,k_\ell}(\sigma)\cdot(\chi_1(\sigma)-1),$$

which is independent of n for all large enough n.

Note that the exclusion of powers in the statement of the theorem is necessary, but these expected values can still be understood. Indeed, $\chi_{k_1,\ldots,k_\ell}(\sigma^t) = \chi_{tk_1,\ldots,tk_\ell}(\sigma)$. Hence, we can still obtain an approximation for the expected value of a power using the theorem.

An interesting corollary of Theorem 1.3 deals with the expected number of cycles of a given size. Fix $s \in \mathbb{N}$. Let $\text{Cyc}_s(\sigma)$ denote the number of cycles of length s in the permutation σ. The expected number of such cycles in a uniformly random permutation is $\frac{1}{s}$. For every large enough n, Cyc_s is a fixed linear combination of the characters from (1.1). For example, $\text{Cyc}_2 = \frac{\chi_2-\chi_1}{2}$. If $s \geq 2$, then for every large enough n, $\langle \text{Cyc}_s, \chi_1 - 1\rangle = 0$. Therefore, Theorem 1.3 yields,

Corollary 1.4. *Let $s \geq 2$. For every non-power $w \in \mathbf{F}_r$,*

$$\mathbb{E}_w[\text{Cyc}_s] = \frac{1}{s} + O\left(\frac{1}{n^{\pi(w)}}\right).$$

It is a well-known fact, c.f. [6], that there are families of irreducible characters $\xi = \{\xi_n\}_{n\geq n_0}$ (ξ_n being an irreducible character of S_n), with dimension polynomial in n and which are finite linear combinations of the functions $\chi_{k_1,\ldots,k_\ell}$. Such a family of irreducible characters corresponds to a family of Young diagrams, given by a Young diagram $\mathcal{D}$ of size n_0, so that ξ_n corresponds to the Young diagram obtained from $\mathcal{D}$ by adding $n - n_0$ boxes to the first row. The first few examples of such characters are $1, \chi_1 - 1, \frac{\chi_2+\chi_{1,1}}{2} - 2\chi_1, \frac{\chi_{1,1}-\chi_2}{2} - \chi_1 + 1$, corresponding to the young diagrams $(n), (n-1,1), (n-2,2)$ and $(n-2,1,1)$, respectively.

In fact, by orthogonality of irreducible characters, Theorem 1.3 is equivalent to that for such a family of irreducible characters ξ, if $\xi \neq 1, \chi_1 - 1$ then

$$\mathbb{E}_w[\xi] = O\left(\frac{1}{n^{\pi(w)}}\right). \tag{1.2}$$

We conjecture the following much stronger bound:

Conjecture 1.5. *Let ξ be a family of irreducible characters of S_n as above. Then,*

$$\mathbb{E}_w[\xi] = O\left(\frac{1}{(\dim \xi)^{\pi(w)-1}}\right). \tag{1.3}$$

Whenever $\xi \neq 1, \chi_1 - 1$, the dimension $\dim \xi$ is a polynomial function of n, of degree greater than 1. Thus, the conjectural bound (1.3) is stronger (for non-powers) than (1.2). The conjecture holds for words of primitivity rank 1, namely, for proper powers (this follows from [8] and from [4, Section 4]). Another known special case of this conjecture is the commutator $[x,y] = xyx^{-1}y^{-1}$: indeed, $\pi([x,y]) = 2$ and already in 1896 Frobenius [2] showed that $\mathbb{E}_{[x,y]}[\chi] = \frac{1}{\dim \chi}$ for every finite group G and every irreducible character χ of G. Moreover, given two class functions $f_1, f_2 \colon G \to \mathbb{R}$ and an irreducible character χ of G, a simple application of Schur's Lemma gives $\langle f_1 * f_2, \chi \rangle = \frac{\langle f_1, \chi \rangle \langle f_2, \chi \rangle}{\dim \chi}$. If $w_1 \in \mathbf{F}(x_1, \ldots, x_k), w_2 \in \mathbf{F}(x_{k+1}, \ldots, x_r)$ are two words generated by disjoint sets of letters, then the $w_1 w_2$-measure on G is the convolution of the w_1- and the w_2-measures, and by the corollary of Schur's Lemma, $\mathbb{E}_{w_1 w_2}[\chi] = \frac{\mathbb{E}_{w_1}[\chi] \cdot \mathbb{E}_{w_2}[\chi]}{\dim \chi}$. On the other hand, $\pi(w_1 w_2) = \pi(w_1) + \pi(w_2)$ [9, Lemma 6.8]. Hence, knowing the conjecture for two such words implies the claim for their product. In particular, this implies the conjecture for every product of disjoint commutators and powers, that is, for every word of the form

$$w = [x_1, y_1] \cdot [x_2, y_2] \cdot \ldots \cdot [x_r, y_r] \cdot z_1^{k_1} \cdot \ldots \cdot z_m^{k_m} \in \mathbf{F}(x_1, \ldots, x_r, y_1, \ldots, y_r, z_1, \ldots, z_m),$$

with $r, m \in \mathbb{N}_{\geq 0}$ and $k_1, \ldots, k_m \in \mathbb{Z}$.

2 The main ideas in the proofs

In this section we give an overview of the main ideas of our work, which together lead to our main result: Theorem 1.3. Full proofs and further details are given in the full version of this paper.

2.1 Generalizations of the object of study

The quantities we wish to study are of the form

$$\mathbb{E}_w\left[\chi_{k_1,\ldots,k_\ell}\right] = \mathbb{E}_{\sigma_1,\ldots,\sigma_r \in S_n}\left[\#\mathrm{fix}\left(w^{k_1}(\sigma_1,\ldots,\sigma_r)\right) \cdot \ldots \cdot \#\mathrm{fix}\left(w^{k_\ell}(\sigma_1,\ldots,\sigma_r)\right)\right].$$

Assume that w is written in the ordered basis $B = \{b_1, \ldots, b_r\}$ of $\mathbf{F}_r$. Choosing a uniformly random r-tuple of permutations from S_n is the same as choosing at random an

homomorphism $\varphi\colon \mathbf{F}_r \to S_n$, as $\varphi(b_1),\ldots,\varphi(b_r)$ is a uniformly random r-tuple of permutations. Replacing the letters of w by the permutations $\varphi(b_1),\ldots,\varphi(b_r)$, we obtain the permutation $\varphi(w)$. Hence,

$$\mathbb{E}_w\left[\chi_{k_1,\ldots,k_\ell}\right] = \mathbb{E}_{\varphi\in\mathrm{Hom}(\mathbf{F}_r,S_n)}\left[\#\mathrm{fix}\left(\varphi\left(w^{k_1}\right)\right)\cdot\ldots\cdot\#\mathrm{fix}\left(\varphi\left(w^{k_\ell}\right)\right)\right]. \tag{2.1}$$

Following [10], the first step in our analysis is to generalize the function we study. This generalization is crucial for the next steps. The most straightforward generalization is to consider quantities of the form

$$\mathbb{E}_{\varphi\in\mathrm{Hom}(\mathbf{F}_r,S_n)}\left[\#\mathrm{fix}\left(\varphi\left(w_1\right)\right)\cdot\ldots\cdot\#\mathrm{fix}\left(\varphi\left(w_\ell\right)\right)\right], \tag{2.2}$$

for arbitrary words $w_1,\ldots,w_\ell \in \mathbf{F}_r$. Next, we generalize from fixed points of a word to *common* fixed points of several words, or, equivalently, to common fixed points of subgroups: note that given a finite set of words $w_1,\ldots,w_t \in \mathbf{F}_r$, an element $i \in [n]$ is a *common* fixed point of all the permutations $\varphi(w_1),\ldots,\varphi(w_t)$ if and only if it is a common fixed point of all the permutations in the subgroup $\varphi(H) \leq S_n$ where $H = \langle w_1,\ldots,w_t\rangle \leq \mathbf{F}_r$. For (a finitely generated) $H \leq \mathbf{F}_r$ we denote by $\#\mathrm{fix}(\varphi(H))$ the number of common fixed points of $\varphi(H)$. We extend the function we wish to study to quantities of the form

$$\mathbb{E}_{\varphi\in\mathrm{Hom}(\mathbf{F}_r,S_n)}\left[\#\mathrm{fix}\left(\varphi\left(H_1\right)\right)\cdot\ldots\cdot\#\mathrm{fix}\left(\varphi\left(H_\ell\right)\right)\right], \tag{2.3}$$

where $H_1,\ldots,H_\ell \leq \mathbf{F}_r$ are f.g. (finitely generated) subgroups of $\mathbf{F}_r$.

If $H, H' \leq \mathbf{F}_r$ are conjugate subgroups then $\#\mathrm{fix}(\varphi(H)) = \#\mathrm{fix}(\varphi(H'))$. Therefore, (2.3) depends, in fact, on a *multiset of conjugacy classes of f.g. subgroups* of $\mathbf{F}_r$. We shall work in the category of these objects, which we denote $\mathcal{MOCC}(\mathbf{F}_r)$.

Finally, assume that there are two multisets of f.g. subgroups $H_1,\ldots,H_\ell \leq \mathbf{F}_r$ and $J_1,\ldots,J_m \leq \mathbf{F}_r$, and that there is a map $f\colon [\ell] \to [m]$, such that $H_i \leq J_{f(i)}$ for all $1 \leq i \leq \ell$. Let $\left\{\varphi_j\colon J_j \to S_n\right\}_{j=1}^m$ be independent, uniformly random homomorphisms. Our final generalization of the object of study is to

$$\mathbb{E}_{\left\{\varphi_j\in\mathrm{Hom}\left(J_j,S_n\right)\right\}_{j=1}^m}\left[\#\mathrm{fix}\left(\varphi_{f(1)}\left(H_1\right)\right)\cdot\ldots\cdot\#\mathrm{fix}\left(\varphi_{f(\ell)}\left(H_\ell\right)\right)\right]. \tag{2.4}$$

As described in the sequel, much of the technique in our proofs relies on this generalization of (2.1) to (2.2), (2.3) and (2.4). Next, we give a geometric description of the category of $\mathcal{MOCC}(\mathbf{F}_r)$ which makes many of our definitions more straightforward and many of our lemmas more intuitive.

2.2 Multi core graphs

Let $B = \{b_1,\ldots,b_r\}$ be a basis of $\mathbf{F}_r$, and consider the bouquet X_B of r circles with distinct labels from B and arbitrary orientations and with wedge point o. Then $\pi_1(X_B,o)$

is naturally identified with $\mathbf{F}_r$. The notion of (B-labeled) core graphs, introduced in [11], refers to (usually) finite, connected graphs with no leaves, that come with a graph morphism to X_B which is an *immersion*, namely, locally injective. In other words, this is a finite connected graph with no leaves, with edges that are directed and labeled by the elements of B, such that for every vertex v and every $b \in B$, there is at most one incoming b-edge and at most one outgoing b-edge at v. We stress that multiple edges between two vertices and loops at vertices are allowed.

There is a natural one-to-one correspondence between finite B-labeled core graphs and conjugacy classes of f.g. subgroups of $\mathbf{F}_r$ – see [11, 3, 9, 10] for more details. Here, we consider core graphs which are not necessarily connected:

Definition 2.1. Let B be a basis of $\mathbf{F}_r$. A B-labeled *multi core graph* is a disjoint union of finitely many finite core graphs. In other words, this is a finite graph, not necessarily connected, with no leaves, and which comes with an immersion to X_B. We denote the set of B-labeled multi core graphs by $\mathcal{M}\mathrm{u}\mathcal{C}\mathcal{G}_B(\mathbf{F}_r)$.

Because a connected core graph corresponds to a conjugacy class of f.g. subgroups of $\mathbf{F}_r$, a multi core graph corresponds to a multiset of such objects. Therefore, every basis B of $\mathbf{F}_r$ gives a one-to-one correspondence

$$\begin{array}{ccc} \mathcal{M}\mathrm{u}\mathcal{C}\mathcal{G}_B(\mathbf{F}_r) = & & \mathcal{M}\mathcal{O}\mathcal{C}\mathcal{C}(\mathbf{F}_r) = \\ \left\{\begin{array}{c} B-\text{labeled} \\ \text{multi core graphs} \end{array}\right\} & \longleftrightarrow & \left\{\begin{array}{c} \text{finite multisets of conjugacy classes} \\ \text{of f.g. subgroups of } \mathbf{F}_r \end{array}\right\}. \end{array} \tag{2.5}$$

Definition 2.2. A morphism $\eta\colon \Gamma \to \Delta$ between B-labeled multi core graphs is a graph-morphism which commutes with the immersions p, q to X_B.

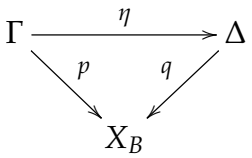

In particular, the morphism η is itself an immersion, and it preserves the orientations and labels of the edges. To get a description of η in terms of subgroups, assume that Γ consists of ℓ components $\Gamma_1, \ldots, \Gamma_\ell$ and that Δ consists of m components $\Delta_1, \ldots, \Delta_m$. Let $f\colon [\ell] \to [m]$ be the induced map on connected components, so $\eta(\Gamma_i) \subseteq \Delta_{f(i)}$. For every $i \in [\ell]$, pick an arbitrary vertex $v_i \in \Gamma_i$ and let $H_i = \pi_1(\Gamma_i, v_i)$. As η is an immersion, it induces *injective* maps in the level of the fundamental groups: indeed, any non-backtracking cycle in Γ is mapped to a non-backtracking cycle in Δ. Therefore, η can be thought of as the injection, for all $i \in [\ell]$

$$H_i \hookrightarrow \pi_1\left(\Delta_{f(i)}, \eta(v_i)\right). \tag{2.6}$$

Now we are very close to the situation that led to (2.4), but in order to define a (2.4)-like quantity that depends on η, we still need to conjugate the images in (2.6) so that they all sit in the same subgroups in the conjugacy class of subgroups of Δ_j. Formally, pick an arbitrary vertex $v_j \in \Delta_j$ for all $j \in [m]$ and let $J_j = \pi_1(\Delta_j, v_j)$. For every $i \in [\ell]$, let $x_i \in \mathbf{F}_r$ satisfy $x_i \left[\pi_1\left(\Delta_{f(i)}, \eta(v_i)\right)\right] x_i^{-1} = J_{f(i)}$. So now $x_i H_i x_i^{-1} \leq J_{f(i)}$, and we are now able to define another form of (2.4) which comes from the morphism η:

Definition 2.3. In the above notation, let $\left\{\varphi_j \colon J_j \to S_n\right\}_{j=1}^m$ be independent, uniformly random homomorphisms. Define

$$\Phi_\eta(n) \stackrel{\text{def}}{=} \mathbb{E}_{\left\{\varphi_j \in \mathrm{Hom}\left(J_j, S_n\right)\right\}_{j=1}^m} \left[\#\mathrm{fix}\left(\varphi_{f(1)}\left(x_1 H_1 x_1^{-1}\right)\right) \cdot \ldots \cdot \#\mathrm{fix}\left(\varphi_{f(\ell)}\left(x_\ell H_\ell x_\ell^{-1}\right)\right)\right].$$

Note that the values of Φ_η do not depend on our choice of basepoints in the components of Γ and Δ. As explained in [10, Section 6] for the simpler case analyzed there, $\Phi_\eta(n)$ can also be given the following topological interpretation. Let $\widehat{\Delta_n}$ be a random n-sheeted covering space of Δ. Then $\Phi_\eta(n)$ is equal to the average number of lifts of η to $\widehat{\Delta_n}$.

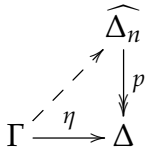

We end this subsection with three important invariants of multi core graphs.

Definition 2.4. Let $\Gamma \in \mathcal{M}u\mathcal{C}\mathcal{G}_B(\mathbf{F}_r)$ be a multi core graph and $\mathcal{H} = \left\{H_1^{\mathbf{F}_r}, \ldots, H_\ell^{\mathbf{F}_r}\right\}$ the corresponding multiset in $\mathcal{MOCC}(\mathbf{F}_r)$. We denote by $rk\mathcal{H} = \mathrm{rk}\Gamma$ the sum of ranks of $H_1, \ldots, H_\ell$, by $\chi(\Gamma) = \chi(\mathcal{H})$ the Euler characteristic of Γ, and by $c(\Gamma) = c(\mathcal{H})$ the number of connected components of Γ (which is ℓ in the current notation). These three quantities are related by $\mathrm{rk}\Gamma + \chi(\Gamma) = c(\Gamma)$.

2.3 Free and algebraic morphisms

A subgroup H of a free group $\mathbf{F}$ is called a *free factor* of $\mathbf{F}$, and $\mathbf{F}$ a *free extension* of H, denoted $H \stackrel{*}{\leq} \mathbf{F}$, if it is generated by some subset of a basis of $\mathbf{F}$. Equivalently, this means that there is another subgroup $K \leq \mathbf{F}_r$, such that $\mathbf{F} = H * K$. The useful notion of an algebraic extensions of free groups is defined as follows (see [7] for a survey):

Definition 2.5. Let H be a subgroup of the free group $\mathbf{F}$. Then $\mathbf{F}$ is an *algebraic extension* of H, denoted $H \leq_{\mathrm{alg}} \mathbf{F}$, if there is no intermediate proper free factor of $\mathbf{F}$. Namely, if whenever $H \leq J \stackrel{*}{\leq} F$, we have $J = \mathbf{F}$.

Given a morphism of *connected* core graphs, we may say it is free (algebraic) if the induced map in the level of fundamental groups gives a free (algebraic, respectively) extension of groups. A crucial ingredient of our argument is to find the right generalizations of these notions to morphisms of multi core graphs. We start with free extensions.

Definition 2.6. If $H_1, \ldots, H_\ell$ are subgroups of the free group J, we say that J is a *free extension* of the multiset $\{H_1, \ldots, H_\ell\}$, denoted $\{H_1, \ldots, H_\ell\} \overset{*}{\leq} J$, if J decomposes as a free product

$$J = \left(*_{i=1}^{\ell} j_i H_i j_i^{-1} \right) * K$$

for some conjugate subgroup $j_i H_i j_i^{-1}$ of H_i (so $j_i \in J$) and some subgroup $K \leq J$.

The definition of a free morphism $\eta : \Gamma \to \Delta$ of multi core graphs is very similar, except that one needs first to choose arbitrarily some subgroup from the conjugacy class of every component of Δ, and then find a suitable subgroup for every component of Γ mapping to it. This can be done similarly to the manner by which we defined Φ_η in Definition 2.3, and we give the precise definition in the full version of the paper. The following theorem states some properties of free morphisms.

Proposition 2.7.

1. *Every* injective *morphism of multi core graphs is free. In particular, the identity morphism if free.*
2. *The composition of two free morphisms is free.*
3. *If* $\bullet \xrightarrow[\varphi]{} \bullet \xrightarrow[\psi]{\eta} \bullet$ *is a composition of morphisms with* ψ *free, then* $\Phi_\varphi = \Phi_\eta$.

We move to defining our generalization of the notion of algebraic extensions.

Definition 2.8. Let $\eta \colon \Gamma \to \Delta$ be a morphism of multi core graphs $\Gamma, \Delta \in \mathcal{M}\mathrm{u}\mathcal{C}\mathcal{G}_B(\mathbf{F}_r)$. We say that η is *algebraic* if whenever $\Gamma \xrightarrow{\eta_1} \Sigma \xrightarrow{\eta_2} \Delta$ is a decomposition of η with η_2 free, we have that η_2 is an isomorphism.

This notion can also be described in algebraic terms, namely, when a free group $\mathbf{F}$ is considered to be an algebraic extension of a multiset of its subgroups $\{H_1, \ldots, H_\ell\}$ – we elaborate in the full paper. The following theorem lists some important properties of algebraic morphisms.

Theorem 2.9.

1. *Every algebraic morphism of multi core graphs is surjective.*
2. *The composition of two algebraic morphisms is algebraic.*

3. *The identity morphism is algebraic.*

4. *Let $\eta : \Gamma \to \Delta$ be a morphism of multi core graphs. Then there is a unique decomposition*

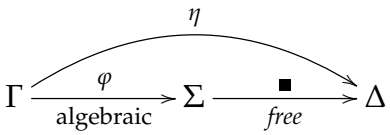

$\eta = \psi \circ \varphi$ such that φ is algebraic and ψ is a free[2].

2.4 The length of a morphism

In view of Proposition 2.7, for every morphism $\eta : \Gamma \to \Delta$ of multi core graphs, if $\overline{\eta}$ marks the induced morphism from Γ to the image of η in Δ, then $\Phi_\eta \equiv \Phi_{\overline{\eta}}$. Therefore, in order to analyze the values of the function Φ, it is enough to consider *surjective* morphisms.

Every surjective morphism $\eta : \Gamma \to \Delta$ can be thought of as a partition of the vertices of Γ given by the fibers of η: the vertices of Δ correspond to the blocks of this partition, and there is a b-edge from the block V to the block U if and only if there a b-edge from some vertex $v \in V$ to some vertex $u \in U$. Note that not every partition of the vertices of Γ corresponds to a moprhism, as the resulting graph might not be a core graph (and have, say, two distinct b-edges emanating from the same vertex). However, every partition of the vertices of Γ can "generate" a legitimate morphism using "Stallings foldings", as we now explain.

Given a B-labeled directed finite graph, which may not be a multi core graph, we can "fold" it until it becomes a multi core graph. Every folding step consists of identifying two edges e_1 and e_2 with the same label, either emanating from the same vertex or entering the same vertex, and then gluing them to a single edge and identifying their other ends to a single vertex. We continue folding until no such coincidences exist, in which case we have a multi core graph[3]. The resulting multi core graph does not depend on the order of folding steps (see [11]).

Definition 2.10. A merging-step of a multi core graph is a gluing together of two vertices of this graph followed by folding. Let $\eta : \Gamma \to \Delta$ be a surjective morphism of B-labeled multi core graphs. The *length of* η, denoted $\rho_B(\eta)$, is the smallest number of merging-steps which lead from Γ to Δ to create η.

There is also a natural "algebraic", basis-independent version of a distance between multisets of conjugacy classes of subgroups. This algebraic distance gives rise to an

[2]To be precise, this decomposition is unique up to an isomorphism of the intermediate multi core graph which preserves the two decompositions.

[3]In some sources, the folding process also includes leaf-pruning. In our situation we never introduce leaves in the process, so it does not matter.

algebraic distance of morphisms of $\mathcal{M}u\mathcal{C}\mathcal{G}_B(\mathbf{F}_r)$, denoted $\rho(\eta)$, we define in the full paper. The following theorem sums up the main properties of these two distances.

Theorem 2.11. *Let $\eta\colon \Gamma \to \Delta$ be a morphism of multi core graphs. Let $\Sigma = \text{Image}(\eta)$ denote the image of η in Δ, and let $\eta = \iota \circ \overline{\eta}$ be the decomposition of η to a surjective and an injective morphisms $\Gamma \xrightarrow{\overline{\eta}} \Sigma \xhookrightarrow{\iota} \Delta$. Then*

$$\rho(\eta) = \rho_B(\overline{\eta}) + [\text{rk}(\Delta) - \text{rk}(\Sigma)].$$

In particular,

$$\rho(\eta) = \text{rk}(\Delta) - \text{rk}(\Gamma) \iff \rho_B(\overline{\eta}) = \text{rk}(\Sigma) - \text{rk}(\Gamma) \iff \eta \text{ is free} \iff \overline{\eta} \text{ is free.}$$

2.5 Möbius inversions of Φ

Recall that our goal is to estimate $\Phi_\eta(n)$ for some morphism $\eta\colon \Gamma \to \Delta$ of multi core graphs, and that we may assume that η is surjective. The next component is to analyze several Möbius inversions of this function inside a finite poset which we now define.

Definition 2.12. Let $\eta\colon \Gamma \twoheadrightarrow \Delta$ be a surjective morphism of multi core graphs. Let $\mathcal{D}\text{ecomp}(\eta)$ denote the poset of decompositions of η into two surjective morphisms $\Gamma \xrightarrow{\eta_1} \Sigma \xrightarrow{\eta_2} \Delta$, where the latter decomposition is considered identical to the morphism $\Gamma \xrightarrow{\eta'_1} \Sigma' \xrightarrow{\eta'_2} \Delta$ if there is an isomorphism $\Sigma \cong \Sigma'$ which commutes with both decompositions. In the same notation, $(\eta_1, \eta_2) \leq (\eta'_1, \eta'_2)$ whenever there is a morhism $\theta\colon \Sigma \to \Sigma'$ which makes the following diagram commute.

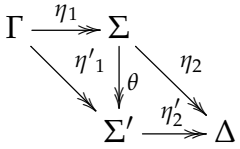

Clearly, $\mathcal{D}\text{ecomp}(\eta)$ is a finite poset. In the same spirit as in [10], we can now define three different Möbius inversions of the function Φ which are defined on the elements of the poset. First, there is a unique "left inversion" of Φ, denoted L^B, which can be defined by setting that for every surjective morphism η,

$$\Phi_\eta = \sum_{(\eta_1,\eta_2)\in\mathcal{D}\text{ecomp}(\eta)} L^B_{\eta_2}.$$

Note that this well defines a map $L^B_\eta\colon \mathbb{N} \to \mathbb{Q}$ for every surjective morphism η by induction on the size of $\mathcal{D}\text{ecomp}(\eta)$. Indeed, $L^B_\eta = \Phi_\eta - \sum_{(\eta_1,\eta_2)\in\mathcal{D}\text{ecomp}(\eta)\setminus\{(\text{id},\eta)\}} L^B_{\eta_2}$, and the summation on the right hand side is on morphisms with a smaller poset of

decompositions. (The base case is $L^B_{\mathrm{id}} = \Phi_{\mathrm{id}}$.) Similarly, we define the right Möbius inversion R^B and the two-sided inversion C^B by

$$\Phi_\eta = \sum_{(\eta_1,\eta_2)\in\mathcal{D}\mathrm{comp}(\eta)} R^B_{\eta_1} = \sum_{(\eta_1,\eta_2,\eta_3)} C^B_{\eta_2}, \tag{2.7}$$

where the rightmost summation is over decompositions $\eta = \eta_3 \circ \eta_2 \circ \eta_1$ defined up to an equivalence parallel to the one in Definition 2.12.

It follows from Theorem 2.9 that if η is *algebraic*, then we can obtain a subposet of $\mathcal{D}\mathrm{comp}\,(\eta)$ which consists only of decompositions to *algebraic* morphisms. We can then define in the same way the Möbius inversions of Φ with respect to this subposet. We denote these inversions by L^{alg}, R^{alg} and C^{alg}. We prove the following results:

Theorem 2.13. *If $\eta\colon \Gamma \to \Delta$ is a surjective morphism, then*

$$C^B_\eta(n) = O\left(n^{\chi(\Gamma)-\rho(\eta)}\right).$$

Similarly, if $\eta\colon \Gamma \to \Delta$ is an algebraic morphism, then

$$C^{\mathrm{alg}}_\eta(n) = O\left(n^{\chi(\Gamma)-\rho(\eta)}\right).$$

In particular, following Theorem 2.11, if η is algebraic but not the identity, then

$$C^{\mathrm{alg}}_\eta(n) = O\left(n^{\chi(\Delta)-1}\right).$$

Definition 2.14. Let $\eta\colon \Gamma \to \Delta$ be a surjective morphism. Denote by $\chi_{\max}(\eta)$ the maximal Euler characteristic of a multi core graph Σ such that there is a decomposition $\Gamma \xrightarrow{\eta_1} \Sigma \xrightarrow{\eta_2} \Delta$ of η with η_1 *non-identity algebraic.* Every such decomposition of η with η_1 *non-identity* algebraic and $\chi(\Sigma) = \chi_{\max}(\eta)$ is called *critical* (for η). Let $Crit(\eta)$ denote the subset of critical decompositions in $\mathcal{D}\mathrm{comp}\,(\eta)$.

In particular, if Γ_w is the core graph corresponding to $\langle w\rangle^{\mathbf{F}_r}$, and the unique morphism $\Gamma_w \to X_B$ is surjective, then $\chi_{\max}(\Gamma_w \to X_B) = 1 - \pi(w)$. The following theorem is now a simple corollary of Theorem 2.13 together with the connection between C^{alg} and Φ as described in (2.7), and the fact that for identity morphisms $\mathrm{id}\colon \Gamma \to \Gamma$ we have $\Phi_{\mathrm{id}} = L^{alg}_{\mathrm{id}} = R^{alg}_{\mathrm{id}} = C^{alg}_{\mathrm{id}} = n^{\chi(\Gamma)}$.

Theorem 2.15. *Let $\eta\colon \Gamma \to \Delta$ be a surjective morphism. Then*

$$\Phi_\eta(n) = n^{\chi(\Gamma)} + |Crit(\eta)| \cdot n^{\chi_{max}(\eta)} + O\left(n^{\chi_{\max}(\eta)-1}\right).$$

2.6 Back to $\chi_{k_1,\ldots,k_\ell}$

We now return to the initial quantity we aimed to study in Theorem 1.3. For $w_1,\ldots,w_\ell \in \mathbf{F}_r \setminus \{1\}$ denote by $\Gamma_{w_1,\ldots,w_\ell}$ the multi core graph which is a disjoint union of ℓ cycles, depicting the conjugacy classes $\langle w_1\rangle^{\mathbf{F}_r},\ldots,\langle w_\ell\rangle^{\mathbf{F}_r}$. Fix $k_1,\ldots,k_\ell \in \mathbb{Z}_{\geq 1}$ and a non-power $1 \neq w \in \mathbf{F}_r$. In the language of multi core graphs and their morphisms, $\mathbb{E}_w\left[\chi_{k_1,\ldots,k_\ell}\right]$ is equal to Φ_η, where η is the sole core graph morphism $\Gamma_{w^{k_1},\ldots,w^{k_\ell}} \to X_B$. There are many decompositions $\Gamma_{w^{k_1},\ldots,w^{k_\ell}} \xrightarrow{\eta_1} \Sigma \xrightarrow{\eta_2} X_B$ of η with $\chi(\Sigma) = 0$ and η_1 algebraic which are, in a sense, independent of w, and their number independent of w. We carefully analyze the *next* layer of decompositions with η_1 algebraic, show that they are all of Euler characteristic $1-\pi(w)$, and that their number is exactly $C_{k_1,\ldots,k_\ell}\cdot|\mathrm{Crit}(w)|$, for an absolute constant $C_{k_1,\ldots,k_\ell}$ independent of w. (In this proof we rely on algebraic results of Baumslag and Steinberg [1, Theorem 1] and of Louder [5, Theorem 1.5].) This then leads to the statement of Theorem 1.3.

References

[1] G. Baumslag and A. Steinberg. "Residual nilpotence and relations in free groups". *Bulletin of the American Mathematical Society* **70**.2 (1964), pp. 283–284. Link.

[2] G. Frobenius. "Über Gruppencharaktere". *Sitzungsberichte Akademie der Wissenschaften zu Berlin* (1896), pp. 985–1021.

[3] I. Kapovich and A. Myasnikov. "Stallings Foldings and Subgroups of Free Groups". *Journal of Algebra* **248** (2002), pp. 608–668. Link.

[4] N. Linial and D. Puder. "Word maps and spectra of random graph lifts". *Random Structures and Algorithms* **37**.1 (2010), pp. 100–135. Link.

[5] L. Louder. "Scott complexity and adjoining roots to finitely generated groups". *Groups, Geometry, and Dynamics* **7**.2 (2013), pp. 451–474. Link.

[6] I. G. Macdonald. *Symmetric functions and Hall polynomials*. Oxford university press, 1998.

[7] A. Miasnikov, E. Ventura, and P. Weil. "Algebraic extensions in free groups". *Geometric group theory*. Springer, 2007, pp. 225–253.

[8] A. Nica. "On the number of cycles of given length of a free word in several random permutations". *Random Structures & Algorithms* **5**.5 (1994), pp. 703–730.

[9] D. Puder. "Primitive words, free factors and measure preservation". *Israel J. Math.* **201**.1 (2014), pp. 25–73. Link.

[10] D. Puder and O. Parzanchevski. "Measure preserving words are primitive". *Journal of the American Mathematical Society* **28**.1 (2015), pp. 63–97. Link.

[11] J. R. Stallings. "Topology of finite graphs". *Inventiones Mathematicae* **71**.3 (1983), pp. 551–565. Link.