

WHAT DO WE KNOW ABOUT THE PRODUCT REPLACEMENT ALGORITHM?

IGOR PAK

Department of Mathematics
Yale University
New Haven, CT 06520
paki@math.yale.edu

February 27, 2000

ABSTRACT. The *product replacement algorithm* is a commonly used heuristic to generate random group elements in a finite group G , by running a random walk on generating k -tuples of G . While experiments showed outstanding performance, until recently there was little theoretical explanation. We give an extensive review of both positive and negative theoretical results in the analysis of the algorithm.

Introduction

In the past few decades the study of groups by means of computations has become a wonderful success story. The whole new field, Computational Group Theory, was developed out of needs to discover and prove new results on finite groups. More recently, the probabilistic method became an important tool for creating faster and better algorithms. A number of applications were developed which assume a fast access to (nearly) uniform group elements. This led to a development of the so called “product replacement algorithm”, which is a commonly used heuristic to generate random group elements in finite groups. The main object of this paper is a rigorous study of this algorithm.

The story behind the product replacement algorithm is the following. Initially research in Computational Group Theory was mostly focused on working with permutation groups, where the fundamental algorithms of Sims (see [Si]) led the way to current advances. The permutation group algorithms became so fast that introduction of random group elements could only slow the algorithms.

Once the direction of research shifted to matrix groups, introduction of random group elements became not only helpful, but highly desirable. It appeared that advancement in the “recognition project” (understanding the structure of a matrix group given by a set of generators) became possible only with a fast access to (nearly) uniform group elements (see e.g. [BrP,Kn2,KS,NP1]). On the other

Key words and phrases. Product replacement algorithm, random walks on groups, probability on groups, simple groups, nilpotent groups, solvable groups, Kazhdan’s property (T).

hand, it was observed on several occasions that getting random group elements by running a simple random walk on a group takes a very long time, or otherwise gives inadequate results (see [Bb5,HR,DS3]).

The problem of generating random group elements has two solutions: one practical and one theoretical. On a theoretical side, Babai in [Bb3] (see also section 4.1) found a general black box algorithm (see [Bb5]) which produces (nearly) uniform group elements at a cost of $O(\log^5 |G|)$ group multiplications. Being provably polynomial, albeit practically slow, this algorithm became a fundamental result on which subsequent theoretical work could be built. It did not resolve, however, the practical need for an efficient random group generator.

The practical “product replacement algorithm” was discovered by Leedham-Green and Soicher [LG]. It was later tested [CLMNO] and proved to have remarkably good performance in several practically interesting cases (cf. [HR]). As success of the algorithm became widely acknowledged, it was included as a standard routine in two major group algebra packages GAP (see [Sc]) and MAGMA (see [BSM]).

Unfortunately, the reasoning why algorithm has such a good performance remained a mystery. Until recently, all attempts to prove theoretical results on performance of the algorithm either failed or produced incremental results (see [Bb5,CG1,DG,DS2]). A major work of Diaconis and Saloff-Coste [DS3] and several (joint) results of the author [BbP,LP,P4,P5,P6,PB] gave a new life to the hopes of fully understanding the algorithm.

The aim of this review article is to give an up to date review of the state of our knowledge on the algorithm. We tried to make this review accessible to nonspecialists, so on occasion we give proofs of standard or elementary results, as well as spend much time giving background in certain areas. The attempt was to have a standard point of reference to most (often basic) results that come up in the analysis of the algorithm.

The product replacement algorithm is defined as follows ([CLMNO]). Given a finite group G , let $\mathcal{N}_k(G)$ be the set of k -tuples $(g) = (g_1, \dots, g_k)$ of elements of G such that $\langle g_1, \dots, g_k \rangle = G$. We call elements of $\mathcal{N}_k(G)$ the *generating k -tuples*. Given a generating k -tuple (g_1, \dots, g_k) , define a *move* on it in the following way. Choose uniformly a pair (i, j) , such that $1 \leq i \neq j \leq k$, then apply one of the following four operations with equal probability:

$$\begin{aligned} R_{i,j}^{\pm} &: (g_1, \dots, g_i, \dots, g_k) \rightarrow (g_1, \dots, g_i \cdot g_j^{\pm 1}, \dots, g_k) \\ L_{i,j}^{\pm} &: (g_1, \dots, g_i, \dots, g_k) \rightarrow (g_1, \dots, g_j^{\pm 1} \cdot g_i, \dots, g_k) \end{aligned}$$

Note that these moves map a generating k -tuple into a generating k -tuple. Now apply these moves t times (the choice of the move must be uniform and independent at each step), and return a random component of the resulting generating k -tuple. This is the desired “random” element of the group G .

Another way to describe the algorithm, is to define on $\mathcal{N}_k(G)$ a structure of a graph induced by maps $R_{i,j}^{\pm}$ and $L_{i,j}^{\pm}$. This makes $\mathcal{N}_k(G)$ into a $4k(k-1)$ -regular graph with no orientation on edges, but with loops when $k > d(G)$, where $d(G)$ is the minimal number of generators of G . Now the algorithm consists of running a nearest neighbor random walk on this graph (for t steps) and returning a random component of the stopping state. We refer to this random walk as the *product replacement random walk* $\mathcal{W} = \mathcal{W}_k(G)$.

About the presentation of a group. We assume the group is given as a black box group, which means that there is an oracle which can multiply elements, invert them, and compare them with identity (see [Bb5]). The group is then defined by a set of generators (g_1, \dots, g_l) . Now, in the algorithm one should take $k \geq l$ and set $g_{l+1} = \dots = g_k = \text{id}$ (see [CLMNO]).

A few words about the parameters k and t . In the original paper [CLMNO] the authors showed that when $k > 2 \log |G|$ and t is large enough, the algorithm will provably work. Practical computations suggested that small values work as well. Much of our work is concentrated on a proper choice of k and t .

Our analysis of the algorithm breaks into three mostly separate parts:

- 1) *Bias*.
- 2) *Connectivity*.
- 3) *Mixing time*.

Each of these is the subject of a separate section of this review. We sum up things at the end (section 4) where we allow ourselves some speculations on the possible future developments. Let us now give a brief overview of the sections.

First, observe that it is unclear whether $\Gamma_k(G)$ is connected. When it isn't, little can be said about the connected component $\Gamma'_k(G)$ which contains (g) . In section 2 we review what is known about this problem, especially in connection with an algebraic notion of (T) systems. Following [P4,P5] we apply results in probabilistic group theory to show that, for reasonably small k , graphs $\Gamma_k(G)$ already have "large" connected component.

In a different direction, section 1 is dedicated to bias of the output. Formally, if $(g) = (g_1, \dots, g_k)$ is chosen uniformly from $\mathcal{N}_k(G)$, consider the bias in the distribution Q of a random component g_i . While mentioned in [CLMNO], this problem was largely ignored until recently, when the first examples of strong bias have been found ([PB,BbP]). In our treatment we follow the recent paper [BbP] of Babai and the author.

Finally, in section 3, we review the results on the mixing time of the product replacement random walk $\mathcal{W}_k(G)$. In the past two decades there has been a spectacular progress in the study of discrete Markov chains, in particular, random walks on groups. Several attempts have been made to proceed with the analysis in the case of $\mathcal{W}_k(G)$. We will concentrate on the paper [DS3] of Diaconis and Saloff-Coste, where the authors use a state of the art analytic approach (largely developed in their previous papers) to obtain subexponential bounds on the mixing time for general groups.

A different approach was used by Lubotzky and the author in [LP], based on the Kazhdan's property (T) from the representation theory of Lie groups. We show that the positive solution of an important open problem (whether a group of automorphisms of a free group $\text{Aut}(F_k)$ has property (T)) implies that the graphs $\Gamma_k(G)$ are expanders, with an expansion constant $\varepsilon = \varepsilon(k)$ depending only on k . In special cases this gives a rigorous proof of the rapid mixing. We follow [LP] to give what seem to be the first explanation of the rapid mixing phenomenon.

A last warning before we conclude the introduction. This paper is purely theoretical and we do not claim (and do not review) any practical results on performance of the algorithm. We review a number of theoretical results, and try to give proper

acknowledgements when appropriate. The remaining results are either folklore or due to the author.

1. THE BIAS IN THE OUTPUT

Much of this section is dedicated to probabilistic group theory. We do not attempt to give a review of this interesting and rapidly developing area for several reasons. First, there are already several such reviews which cover different parts of the subject (see [Kn2,Sh]). Second, this issue is rather tangential to the main subject of our study and is used mostly as a powerful tool in analysis of the algorithm. Still, for the sake of completeness we present here a somewhat biased review of results we will use later on.

1.1 Probability of generating a finite group.

Let G be a finite group, and let $d(G)$ be the minimal number of generators of G . By $\varphi_k(G)$ denote the probability that k uniformly and independently chosen elements in G generate the whole group:

$$\varphi_k(G) = \frac{|\mathcal{N}_k(G)|}{|G|^k}$$

The subject of this section is to give some estimates on these probabilities.

Theorem 1.1.1 *For every sequence of nonisomorphic simple groups $\{G_i\}$ we have*

$$\varphi_2(G_i) \rightarrow 1, \quad \text{as } |G_i| \rightarrow \infty$$

This theorem is a combination of several results. First, it was proved for alternating groups A_n by Dixon [Dx]. Then, for classical groups (and few more series), it was established by Kantor and Lubotzky [KL] (see also [Kn1]). They used Aschbacher's classification of maximal subgroups of linear groups [As]. Finally, Liebeck and Shalev [LS1,LS2] completed the proof for the remaining series. We should add that these papers (except for [Dx]) use extensively the classification of finite simple groups.

The asymptotic behavior of the probabilities φ_2 is also known and the results can be summarized as follows.

Theorem 1.1.2 *Let A_n be alternating groups, $G_n(q)$ be a simple group of Lie type of (untwisted) rank n . Then:*

$$\begin{aligned} (i) \quad & \varphi_k(A_n) = 1 - \frac{1}{n^{k-1}} + O\left(\frac{1}{n^{2k-1}}\right), \quad \text{where } k \geq 2 \\ (ii) \quad & \varphi_2(G_n(q)) = 1 - O\left(\frac{n^3 \log^2 q}{q^n}\right). \end{aligned}$$

Part (i) follows from the work of Babai [Bb2], part (ii) was conjectured and partly proved by Kantor and Lubotzky [KL], and proved in full by Liebeck and Shalev [LS1,LS2].

Proposition 1.1.3 *Let G be a finite p -group, $d = d(G)$. Then:*

- (i) $\varphi_d(G) > 1 - \frac{1}{p} - \frac{1}{p^2}$,
- (ii) $\varphi_{d+1}(G) > 1 - \frac{1}{p(p-1)}$,
- (iii) $\varphi_{d+r-1}(G) > 1 - \frac{1}{p^r - p^{r-1} - p^{r-2} - 1}$, where $r \geq 1$.

The result is obtained by direct calculation due to the fact that the quotient $G/\Phi(G) \simeq \mathbb{Z}_p^m$, and $\varphi_k(G) = \varphi_k(G/\Phi(G))$, where $\Phi(G)$ is a Frattini group of G . Different versions of this result were obtained in [DS3,P4]. For nilpotent groups we also obtain:

Proposition 1.1.4 *Let G be a finite nilpotent group, $d = d(G)$. Then:*

- (i) $\varphi_d(G) > \frac{1}{5 \log \log |G|}$,
- (ii) $\varphi_{d+1}(G) > \frac{1}{e}$,
- (iii) $\varphi_{d+r-1}(G) > 1 - \frac{8}{12^{r/4}}$, where $r \geq 1$.

Again, see [DS3], Remark after Lemma 6.3; and [P4].

Theorem 1.1.5 *There exists a universal constant $C < 10^7$ such that for any solvable group G , $d = d(G)$, and $r > 0$ we have*

$$\varphi_{(\beta+1) \cdot d + C + r} > 1 - \alpha^r,$$

where $\beta = (3 \cdot \ln 48 + \ln 24)/(3 \cdot \ln 9) \approx 2.243991050$ is the Pálffy–Wolf constant, and $\alpha = 1/\sqrt[4]{12} < 1$.

This result is due to the author [P4] and is obtained by using the product formula of Gachütz [Ga]. A weaker version is due to A. Mann [Mn] who used a somewhat different method to show that $\varphi_{(\beta+1)d+C}(G)$ is greater than some universal constant $\epsilon > 0$. In the other direction, Mann in [Mn] showed that $\varphi_{\beta d - C}(G) \rightarrow 0$ for certain series of groups. This implies that the constant $(\beta + 1)$ in Theorem 1.1.5 cannot be improved to a constant smaller than β .

Question 1.1.6 Can one improve the constant $(\beta + 1)$ in Theorem 1.1.5? Can one find a “reasonable” bound on C (compared to 10^7 in the theorem)?

Finally, for general groups we have the following weak result:

Theorem 1.1.7 *For any finite group G , $1 > \epsilon > 0$, $m = \lceil \log_2 |G| \rceil$, we have*

- (i) $\varphi_k(G) \geq \varphi_k(\mathbb{Z}_2^m)$,
- (ii) $\varphi_k(G) > 1 - \epsilon$, where $k > m + 2 + \log_2 1/\epsilon$,
- (iii) $\varphi_m(G) > \frac{1}{4}$ and $\varphi_{m+1}(G) \geq \frac{1}{2}$.

Part (ii) of the theorem is a slight improvement over a more general classical result of Erdős and Rényi in [ER] (see also [P3]). The proof of part (i) is based on “random group process” defined in [P4] (cf. [Ac,PV]). Parts (ii), (iii) then follow from (i) and Proposition 1.1.3.

We conclude with a somewhat related recent result of Guralnick and Kantor. While technically different, we find it amazingly useful in several application.

Theorem 1.1.8 ([GK]) *Let G be a nonabelian finite simple group. Define*

$$\text{PC}(G) = \max_C \min_{g \in G, g \neq \text{id}} \mathbf{P}(\langle g, h \rangle = G \mid h \in C),$$

where \max is over all conjugacy classes $C \subset G$, and the probability is taken over uniform $h \in C$. Then:

- (i) $\text{PC}(G) > 1/10$ for every G ,
- (ii) $\text{PC}(G) \rightarrow 1$ as $|G| \rightarrow \infty$, for G not isomorphic to A_n or $\Omega(2m+1, q)$ with a bounded q .

1.2 Generating k -tuples in products of simple groups.

A famous result of Philip Hall [H1] (see also [KL]) can be formulated as follows:

Theorem 1.2.1 (P. Hall) *Let G be a nonabelian simple group. Then the maximal number N , such that group $H = G^N$ (direct product of N copies of a G) is generated by k elements, is equal to $|\mathcal{N}_k(G)|/|\text{Aut}(G)|$. Further, an element $(g) = (g_1, \dots, g_k) \in H^k$, where $g_i = (g_i^{(1)}, \dots, g_i^{(m)}) \in H$, is a generating k -tuple of $H = G^N$ if and only if all k -tuples $(g_1^{(j)}, \dots, g_k^{(j)})$, $1 \leq j \leq m$ generate the group G and lie in different orbits of the diagonal action of $\text{Aut}(G)$ on $\mathcal{N}_k(G)$.*

In a celebrated special case of $G = A_5$ and $k = 2$, Hall showed [H1] that $N = 19$, i.e. that $d(A_5^{19}) = 2$, while $d(A_5^{20}) = 3$.

Recall that by Theorem 1.1.1, for any fixed $k > 1$ the probabilities $\varphi_k(A_n) \rightarrow 1$. Since $\text{Aut}(A_n) = S_n$ for $n \geq 5$, $n \neq 6$, we obtain that the N from the theorem satisfies:

$$N = \varphi_k(A_n) \frac{|A_n|^k}{|S_n|} > \frac{n^{k-1}}{2^{k+2}},$$

where n is large enough. In particular, if $G = A_n$ and n is large enough, the group $A_n^{n!/8}$ can be generated by two elements, while $A_n^{n!/4}$ cannot.

Now let $N = n!/8$, $k > 5$. Consider a set of k -tuples in $H = A_n^N$. From the above and the birthday paradox (see [F,KL,BbP]) one can easily see that the probability that two k -tuples $(g_1^{(j)}, \dots, g_k^{(j)})$ and $(g_1^{(j')}, \dots, g_k^{(j')})$ lie in the same orbit of $\text{Aut}(G)$ goes to zero superexponentially. Thus one can ignore this probability and conclude that N generating k -tuples almost surely correspond to a generating k -tuple of A_n^N .

First, observe that $\varphi_k(A_n^N) \rightarrow 0$ as $n \rightarrow \infty$, given $k = o(n)$ ([KL]). Indeed, the probability that any given k -tuple generates A_n is at most $(1 - n^{-k})$, which is the probability that none of the k permutations fixes 1. Therefore the probability that all of the N k -tuples generate A_n is at most $(1 - 1/n^k)^N \rightarrow 0$ as $n \rightarrow \infty$.

From here it is immediate that in order to have $\varphi_k > \epsilon > 0$, one needs $k = \Omega(n)$. Now, in the reverse direction one has the following general result:

Theorem 1.2.2 *If G is a direct product of simple nonabelian groups (possibly, with repetitions), then $\varphi_k(G) > 1/2$ for $k \leq C \log m$, where m is the maximal number of isomorphic copies of each group and C is a universal constant.*

The result is due to the author [P4]. The proof uses a number of technical calculations involving the generation probabilities in Theorem 1.1.2, as well as the classification of finite simple groups. The constant C is expected to be reasonable, but probably should be hard to calculate due to difficulties with calculation on large sporadic groups.

1.3 Proving the bias.

Now we are ready to present the first version of the bias in the output of the product replacement algorithm. We shall start with preliminary observation.

Let G be a finite group, and let Q_k, Q_k^* be the probability distributions of the first and of the random component in a uniform generating k -tuple $(g) \in \mathcal{N}_k(G)$. Then symmetry gives:

Lemma 1.3.1 *For all G , $k \geq d(G)$ we have $Q_k = Q_k^*$.*

The lemma reduces the problem of bias of the output to the problem of bias of the first component in generating k -tuples of G .

Let $H = A_n^N$, $N = n!/8$, $k > 5$. Consider a random generating k -tuple $(h) = (h_1, \dots, h_k) \in \mathcal{N}_k(H)$. We write $h_1 = (\sigma_1, \dots, \sigma_N)$, where $\sigma_j \in A_n$. Now, each σ_j is the first component of a random generating k -tuple of A_n . While they may seem close to being nearly uniform in A_n , a small bias remains:

Lemma 1.3.2 ([BbP]) *If σ is the first element in a uniform generating k -tuple of A_n , then*

$$\mathbf{P}(\sigma(1) = 1) = \frac{1}{n} - \frac{1}{n^k} + O\left(\frac{1}{n^{2k-1}}\right)$$

The proof is immediate from part (i) of Theorem 1.1.2. See [BbP] for a complete proof and other versions of the Lemma.

Now, in an unbiased sample we have $\mathbf{P}(\sigma_j = 1) = 1/n$. Since the number of independent samples is N , we must have the bias in each of them less than N , or otherwise the bias becomes statistically significant. In other words, if $k = o(n)$, then $n^k = o(n!/8)$, and the bias can be detected.

Theorem 1.3.3 ([BbP]) *Let $H_n = A_n^N$, where $N = n!/8$. Let $k = o(n)$. Then there exists a sequence of subsets $B_n \subset \mathcal{N}_k(H_n)$ such that $|B_n|/|\mathcal{N}_k(G)| \rightarrow 1$ and $Q_k(B_n) \rightarrow 0$ as $n \rightarrow \infty$.*

Proof. Take B_n to be the set of all $(h) = (\sigma_1, \dots, \sigma_N) \in H_n$ such that

$$\#\{j \mid \sigma_j(1) = 1\} > N \left(\frac{1}{n} - \frac{1}{2n^k} \right).$$

The Chernoff bound (see e.g. [ASE]) gives that the N random permutations $\sigma_j \in A_n$ are in B_n with high probability.

On the other hand, for σ_j chosen from generating k -tuple of A_n we have

$$\mathbf{E}(\#\{j \mid \sigma_j(1) = 1\}) = N \cdot \left(1 - \frac{1}{n^k}\right),$$

so the Chernoff bound gives that such N independent σ_j do not lie in B_n with high probability. \square

Remark 1.3.4 The first results on bias was announced in [PB], and rigorously proved in [BbP]. One can generalize the result to other sets B_n as well as to other series of simple groups of Lie type. We refer to [BbP] for a more detailed treatment.

Before we finish this section, let us elaborate on the findings. Basically, we showed that even if the graph $\Gamma_k(H)$ is connected, even if the product replacement random walk mixes rapidly, the resulting distribution of the output can still be biased. Of course, for that to happen we need $\varphi_k(H_n) \rightarrow 0$ as $n \rightarrow \infty$, which is indeed true when $H_n = A_n^{n!/8}$. While this is only a necessary condition, it seems likely that bias exists in all natural cases when $\varphi_k(H_n) \rightarrow 0$. However proving bias in such cases can be a delicate task.

Open Problem 1.3.5 *Can one exhibit the bias for a sequence of solvable groups?*

Remark 1.3.6 We believe a natural candidate for H_n in 1.3.5 to be the group $H_n = R_n \ltimes (\mathbb{Z}_2^n)^N$, where R_n is an n -th iterated wreath product of S_4 :

$$R_n = S_4 \wr S_4 \wr \dots \wr S_4 \quad (n \text{ times}),$$

where the action of R_n is independent on each copy of \mathbb{Z}_2^n , and N is the largest number such that H_n still can be generated by three (or any fixed constant number of k) elements. Note that it is not even obvious that R_n is two-generated. On the other hand, using Gaschütz formulas A. Mann showed in [Mn] that $\varphi_2(R_n) > 1/2$ for any n . Still, we challenge the reader to find a single “nice” generating 2-tuple.

1.4 Detecting the bias.

While from general mathematical point of view the previous section completely answers the question about existence of the bias, from computational point of view the answer is somewhat weak. The idea is that generating random group elements is not a goal, but rather a routine used in various randomized algorithms (see e.g. [Bb4,Bb5,BP,Kn2,P2]). Thus the “real” threshold for bias is much higher: one must show that these “random elements” can actually “mess up” some calculations.

In the original paper [CLMNO] the authors tested the algorithm by using χ^2 on the distribution of the orders of the elements. The order statistics is of great importance, especially in various recognition algorithms (see e.g. [BrP,Kn2,KS,NP1]). Let us show that they are useless in case of powers of A_n .

Proposition 1.4.1 *The proportion of elements of $H = A_n^N$ which have the same order is at least $1 - e^{-N/(n \log n)}$, for any integers n, N .*

Proof. First, let us determine this order. Denote by $m_p(n) = p^{\lfloor \log_p n \rfloor}$ the maximal $m \leq n$ such that $m = p^r$. Let $M = \prod_p m_p(n)$, where the product is

over all primes $p \leq n$. Since M is divisible by every possible length i of the cycle, $1 \leq i \leq n$, for every $h \in H$ we have $h^M = \text{id}$.

Let us now calculate the probability that $h \in H$ has order M . Observe that the length of the cycle containing 1 is uniform in $\{1, \dots, n\}$. Thus with probability at least $1/n$ an element $\sigma \in A_n$ contains a cycle of length $m_p(n)$. The probability that no component of a random $h \in H$ contains a cycle of length $m_p(n)$ is at most $(1 - 1/n)^N$. Finally, the probability that components h contain cycles of length $m_p(n)$ for every prime p is at least $1 - n \cdot (1 - 1/n)^N > 1 - \exp(-N/n \log n)$. But then h has order M which completes the proof. \square

The proposition implies that the bias found in the previous section cannot be detected by the order statistics. This begs the following question, possibly related to Problem 1.3.5 :

Question 1.4.2 Can one find a sequence of groups H_n which exhibit bias in the order statistics?

Now let us show that the bias can nevertheless be detected by a short straight line program. First we need several definitions.

Let \mathbf{w} be a word over the alphabet $\{x_i^{\pm 1}, i = 1, 2, \dots\}$. Substituting elements of G for the x_i assigns \mathbf{w} a value in G . Assume that the x_i are chosen independently from the probability distribution P over G . We denote by $\mathbf{w}[P]$ the probability distribution of the value of \mathbf{w} .

Let G be a finite group. As in the previous section, by $Q_k = Q_k^*$ denote the probability distribution on G of the random component in a generating k -tuple $(g) \in \mathcal{N}_k(G)$. By U denote the uniform distribution on G .

Theorem 1.4.3 ([BbP]) *Let $k = k(n) \geq 4$ and $k = o(n)$, $n \rightarrow \infty$. Then there exists a sequence of words $\mathbf{w}_{n,k}$ with the following properties. The length of $\mathbf{w}_{n,k}$ is $n^{O(k)}$. Let $\omega(n) \rightarrow \infty$, $\omega(n) = o(n)$. Set $m = n^{k\omega(n)}$, and $G = A_n^m$. Then $\mathbf{w}[Q_k] = \text{id}$ has probability $1 - O(n^{-cn})$, while $\mathbf{w}[U] = \text{id}$ has probability $O(n^{-cn})$, where c is a universal constant.*

Let us sketch the idea of the proof. First, let n be a prime. Rather than look at the number of σ_j in $g = (\sigma_1, \dots, \sigma_m)$, such that $\sigma_j(1) = 1$ (see the proof of Theorem 1.3.3) we look at the number of those σ_j which are n -cycles. But since n is prime, we have $\sigma^n = \text{id}$ implies σ is a long cycle or $\sigma = \text{id}$. The analog of Lemma 1.3.2 is also straightforward. Therefore we have a (positive) bias in the number of trivial components in random element $g^n \in G$.

Now, a general theory (the Ajtai–Komlós–Szemerédi sorting network in [AKS]) shows how to construct a short monotone Boolean circuit which “detects” the bias. The problem is to construct a short word \mathbf{w} from the circuit. Basically we need to find probabilistic simulation of **AND** and **OR** by group operations.

Formally, let H be a group and $g \in H$. Consider the predicate $\mathcal{E}(g)$ meaning “ $g = \text{id}$ ”. We wish to construct words w_1 and w_2 corresponding to the predicates $\mathcal{E}_1(g, h) = \mathcal{E}(g) \wedge \mathcal{E}(h)$ and $\mathcal{E}_2(g, h) = \mathcal{E}(g) \vee \mathcal{E}(h)$, respectively. Clearly, there is no word which would be id exactly if \mathcal{E}_1 holds, nor is there one for \mathcal{E}_2 . But the product $w_1 = gh$ and the commutator $w_2 = [g, h] = g^{-1}h^{-1}gh$ go part of the way: $\mathcal{E}_1(g, h)$ implies $w_1 = \text{id}$ and \mathcal{E}_2 implies $w_2 = \text{id}$; and the converse holds often enough in each case.

Now, what we do is consider $w_1 = g'h'$ and $w_2 = [g', h']$ where $g' = g^{a_1} \cdot \dots \cdot g^{a_i}$, where $g^a = a^{-1}ga$, and a_i are independent and chosen from the same distribution as g ; h' is defined analogously. Using rapid mixing of random walks on $H = A_n$ we conclude w_1, w_2 are the desired probabilistic simulations of the monotone Boolean operations. We refer to [BbP] for details and references.

Remark 1.4.4 A priori it may seem unlikely that the graph $\Gamma = \Gamma_k(A_n^{n^{1/8}})$ is connected when $k \geq 4$. In this case the result of Theorem 1.4.3 is only a theoretical exercise as the algorithm does not produce uniform generating k -tuples. Nevertheless the graph Γ is probably connected indeed for all $k \geq 4$. In particular, we know that $\Gamma_4(A_5^N)$, $N \leq 1140$ is connected. We refer to Corollary 2.4.5, Remark 2.4.7 and Conjecture 2.5.4 below.

1.5 Avoiding the bias.

To avoid having a bias of Q on G one should take k to be sufficiently large. How large? That depends on your definition of the bias and structural properties of G .

Let P be a probability distribution on a finite set X , and U be a uniform distribution on X . Define the *total variation* between P and U as follows:

$$\|P - U\|_{\text{tv}} = \max_{B \subset X} \left| P(B) - \frac{|B|}{|X|} \right| = \frac{1}{2} \sum_{x \in X} \left| P(x) - \frac{1}{|X|} \right|,$$

where $P(B) = \sum_{x \in B} P(x)$. It is easy to see that $0 \leq \|P - U\|_{\text{tv}} \leq 1$. We postpone discussion of the total variation and other distances until section 3.1. Instead, let us make the following elementary observation.

Proposition 1.5.1 *Let P be a probability distribution on $X \subset G^k$ and let Q be a probability distribution on G defined as a projection of P onto the first component. Then*

$$\|Q - U\|_{\text{tv}} \leq \|P - U\|_{\text{tv}} + \left(1 - \frac{|X|}{|G|^k}\right)$$

In particular, when $X = \mathcal{N}_k(G)$ we have $\|Q - U\|_{\text{tv}} \leq 1 - \varphi_k(G) + \|P - U\|_{\text{tv}}$.

Proof. Denote by P' the probability distribution on $G^k \supset X$, which is P on $(g) \in X$ and zero otherwise. Clearly,

$$\begin{aligned} \|P' - U\|_{\text{tv}} &= \frac{1}{2} \sum_{(g) \in X} \left| P'(g) - \frac{1}{|G|^k} \right| + \frac{1}{2} \sum_{(g) \in G^k \setminus X} \frac{1}{|G|^k} \leq \\ &\leq \frac{1}{2} \sum_{(g) \in X} \left| P'(g) - \frac{1}{|X|} \right| + \frac{1}{2} \sum_{(g) \in X} \left| \frac{1}{|X|} - \frac{1}{|G|^k} \right| + \\ &\quad + \frac{1}{2} \left(1 - \frac{|X|}{|G|^k}\right) = \|P - U\|_{\text{tv}} + \left(1 - \frac{|X|}{|G|^k}\right). \end{aligned}$$

Now, since a uniform distribution on G is a projection of the uniform distribution on G^k , we immediately obtain:

$$\|Q - U\|_{\text{tv}} \leq \|P' - U\|_{\text{tv}} \leq \|P - U\|_{\text{tv}} + \left(1 - \frac{|X|}{|G|^k}\right). \quad \square$$

Now, fix a sequence $\{G_n\}$ of groups. Let us assume that $\varphi_k(G_n) \rightarrow 1$ as $n \rightarrow \infty$. Then Proposition 1.5.2 implies that the bias of \mathcal{Q} is bounded by that of the random walk. In the remainder of this section we will be concerned with the problem of finding an appropriate function $k = k(G_n)$ for a given sequence of groups to satisfy the assumption above.

First, let us translate the previous results into this language. By $\omega(n)$ denote a function $\omega : \mathbb{N} \rightarrow \mathbb{N}$ such that $\omega(n) \rightarrow \infty$ as $n \rightarrow \infty$.

Corollary 1.5.2 *Let $\{G_n\}$ be any sequence of groups. Then $\varphi_k(G_n) \rightarrow 1$ given $k = k(n)$ satisfies either of the following:*

- (i) $k = d(G_n) + \omega(n)$, if G_n is nilpotent,
- (ii) $k = (\beta + 1) \cdot d(G_n) + \omega(n)$, if G_n is solvable,
- (iii) $k = (\log m) \cdot \omega(n)$, if G_n is a direct product of finite nonabelian simple groups, each copy at most m times,
- (iv) $k = \log |G_n| + \omega(n)$, for any G .

While parts (i) – (iii) seem to be satisfactory, the general case (iv) remains quite weak. The following conjecture if true would be a significant advancement:

Conjecture 1.5.3 *Let $\{G_n\}$ be a sequence of groups. Then $\varphi_k(G_n) \rightarrow 1$ given $k = \omega(n) \cdot d(G_n) \cdot \log \log |G_n|$.*

Remark 1.5.4 Note that the results (i) – (iii) in Proposition 1.5.2 are stronger than what Conjecture predicts. Indeed, for the only nontrivial case (iii), if $G = F_1^m \times F_2^l \times \dots$ we have $\log \log |G| \geq \log \log |F_1|^m \geq \log m$.

We remark that a large body of work on positive generation of profinite groups can be translated into this language. Let us especially note a paper [BPS] where the authors show that if a profinite group G has no sections isomorphic to A_n for large n (i.e. for all $n > N$), then this group is positively generated, i.e. $\varphi_k(G) \geq \epsilon$ for a universal k and $\epsilon > 0$ independent of the group G .

Question 1.5.5 Can one quantify the results in [BPS] (i.e. find explicit expressions of k and ϵ in terms of N) ?

2. CONNECTIVITY OF $\Gamma_k(G)$

2.1 Preliminaries.

Let G be a finite group. Recall that $\mathcal{N}_k(G)$ denotes the set of generating k -tuples $(g) = (g_1, \dots, g_k)$, where $\langle g_1, \dots, g_k \rangle = G$. Let $\Gamma_k(G)$ be as in the Introduction, the graph on $\mathcal{N}_k(G)$ with edges corresponding to moves $R_{i,j}^\pm, L_{i,j}^\pm$:

$$\begin{aligned} R_{i,j}^\pm &: (g_1, \dots, g_i, \dots, g_k) \rightarrow (g_1, \dots, g_i \cdot g_j^{\pm 1}, \dots, g_k) \\ L_{i,j}^\pm &: (g_1, \dots, g_i, \dots, g_k) \rightarrow (g_1, \dots, g_j^{\pm 1} \cdot g_i, \dots, g_k) \end{aligned}$$

We call $\Gamma_k(G)$ the *product replacement graph*, and its edges the *Nielsen moves*.

Note that while $\Gamma_k(G)$ is defined as an oriented graph, it is symmetric, so orientation can be disregarded. As before, let $d(G)$ be the minimal number of generators which generate G . In this section we are concerned with the following problem.

Problem 2.1.1 *For a given G , $k \geq d(G)$, is $\Gamma_k(G)$ connected ?*

It is quite embarrassing how little is known about this problem. For example, we do not know the answer to the following innocent looking question:

Question 2.1.2 Is it true that if $\Gamma_k(G)$, $k \geq d(G)$ is connected, then $\Gamma_m(G)$ is connected for every $m > k$?

In the next sections we shall the examples when $\Gamma_{d(G)}$ is disconnected. However we do not know the answer to the following “easy” question:

Question 2.1.3 Is there a finite group G , $d = d(G)$ such that $\Gamma_{d+1}(G)$ is disconnected?

Later in this section we will show that the answer to this question is negative when G is solvable, and it is probably negative for finite simple groups and their direct products. We reserve our judgement with regard to the question in full generality.

By $\varkappa_k(G)$ denote the number of connected components of $\Gamma_k(G)$. Now let us prove some general results on $\varkappa_k(G)$ before turning to examples and special cases.

First, note that one cannot expect that $H \subset G$ implies that $\varkappa_k(H) \leq \varkappa_k(G)$ since we already do not have $d(H) \leq d(G)$. For example, $H = \mathbb{Z}_2^n \subset G = \mathbb{Z}_2 \wr \mathbb{Z}_n$, but $d(H) = n > d(G) = 2$. On the other hand, if H is a quotient of G , then $d(H) \leq d(G)$. This result can be extended to show that $\varkappa_k(H) \leq \varkappa_k(G)$ for any $k \geq d(G)$.

Theorem 2.1.4 ([LP]) *If $\alpha : G \rightarrow H$ is an epimorphism between finite groups, then for every $k \geq d(H)$ the map $\alpha_k : \Gamma_k(G) \rightarrow \Gamma_k(H)$, defined by*

$$\alpha_k(g_1, \dots, g_k) = (\alpha(g_1), \dots, \alpha(g_k))$$

is a surjective graph projection. In particular, the number of connected components of $\Gamma_k(H)$ is bounded by that of $\Gamma_k(G)$.

The result easily follows from the following important lemma:

Lemma 2.1.5 (Gaschütz) *Let $\psi : G \rightarrow H$ be an epimorphism between finite groups, $k \geq d(G)$, and let (h_1, \dots, h_k) be a generating k -tuple of H . Then there exists a generating k -tuple (g_1, \dots, g_k) of G with $\psi(g_i) = h_i$ for $i = 1, \dots, k$.*

Proof. Denote by $N \triangleleft G$ the kernel of the epimorphism $\psi : H \simeq G/N$. Consider any (g_1, \dots, g_k) such that $\langle g_1N, \dots, g_kN \rangle = H$. We need to find $(u_1, \dots, u_k) \in N^k$ such that $\langle g_1u_1, \dots, g_ku_k \rangle = G$.

Let us fix a k -tuple $(g) = (g_1, \dots, g_k)$ as above. Consider the set $R(g)$ of all k -tuples (g_1u_1, \dots, g_ku_k) . Clearly, $|R(g)| = |N|^k$. For every subgroup $S \subset G$ we have $R(g) \cap S^k \neq \emptyset$ if and only if $SN = G$. Further, in this case $|R(g) \cap S^k| = |S \cap N|^k$. This number is therefore independent of the choice of (g) .

Now, the number of k -tuples in $R(g)$ which generate G is equal to $|N|^k$ minus the number of k -tuples in $R(g)$ and proper subgroups. But the latter can be calculated by the Möbius summation formula of the numbers $|R(g) \cap S^k|$ (see [H1]). Therefore the number N of k -tuples in $R(g)$ which generate G is *independent* of the choice of (g) .

By hypothesis, there does exist a generating k -tuple $a = (a_1, \dots, a_k)$ of G . Then for $\psi_k(a) = (\psi(a_1), \dots, \psi(a_k))$ the number N is positive. But since it is the same for all (g) , we obtain the result. \square

In our proof we closely follow the original proof (see [Gr], Proposition 6.14.) We refer to [Gr] for a number of other related results.

2.2 Extended product replacement graph.

Define an *extended product replacement graph* $\tilde{\Gamma}_k(G)$ to be a graph on $\mathcal{N}_k(G)$ with edges corresponding to $R_{i,j}^\pm$, $L_{i,j}^\pm$, and $\pi_{i,j}$, ι_m , $1 \leq i, j, m \leq k$, where

$$\begin{aligned} \pi_{i,j} &: (g_1, \dots, g_i, \dots, g_j, \dots, g_k) \rightarrow (g_1, \dots, g_j, \dots, g_i, \dots, g_k) \\ \iota_m &: (g_1, \dots, g_m, \dots, g_k) \rightarrow (g_1, \dots, g_m^{-1}, \dots, g_k) \end{aligned}$$

Proposition 2.2.1 *Let $\kappa_k(G)$, $\tilde{\kappa}_k(G)$ be the number of connected components in $\Gamma_k(G)$ and $\tilde{\Gamma}_k(G)$ respectively. Then*

$$\tilde{\kappa}_k(G) \leq \kappa_k(G) \leq 2 \tilde{\kappa}_k(G).$$

Moreover, if $\tilde{\Gamma}_k(G)$ is connected and $k \geq d(G) + 1$, then $\Gamma_k(G)$ is also connected.

Proof. Let $(g) = (g_1, \dots, g_k)$. We easily have:

$$\begin{aligned} (g) &= (\dots, g_i, \dots, g_j, \dots) \xrightarrow{R_{i,j}^-} (\dots, g_i \cdot g_j^{-1}, \dots, g_j, \dots) \\ &\xrightarrow{L_{j,i}^+} (\dots, g_i \cdot g_j^{-1}, \dots, g_i, \dots) \xrightarrow{L_{i,j}^-} (\dots, g_j^{-1}, \dots, g_i, \dots). \end{aligned}$$

Denote this move by $P_{i,j}$. Repeating $P_{i,j}$ twice we obtain:

$$(g) \xrightarrow{P_{i,j}} (\dots, g_j^{-1}, \dots, g_i, \dots) \xrightarrow{P_{i,j}} (\dots, g_i^{-1}, \dots, g_j^{-1}, \dots).$$

Denote this move by $I_{i,j}$. Observe that $P_{i,j}$, $I_{i,j}$ generate a subgroup of index two of rotations in a hyperoctahedral group H_n (the Weyl group of root system D_n).

Let us assume that $(g) = (g_1, \dots, g_k)$ is connected to $(h) = (h_1, \dots, h_k)$ by a certain sequence of moves in graph $\tilde{\Gamma}_k(G)$. Now, in the graph $\Gamma_k(G)$ apply the same moves as in $\tilde{\Gamma}_k(G)$ (up to \pm) and use $I_{i,j}$ and $P_{i,j}$ in place of inversions ι_j and transpositions $\pi_{i,j}$. This gives $(h_1^{\pm 1}, \dots, h_k^{\pm 1})$. Let us apply a few more $I_{i,j}$ to obtain $(h_1, \dots, h_{k-1}, h_k^{\pm 1})$. If $\Gamma' \subset \tilde{\Gamma}_k(G)$ is a connected component, then every $(g) \in \tilde{\Gamma}'$ is connected to a fixed (h) . This implies that such (g) is connected to either (h_1, \dots, h_k) or $(h_1, \dots, h_{k-1}, h_k^{\pm 1})$. Thus $\Gamma_k(G)$ has at most two times the number of connected components in $\tilde{\Gamma}_k(G)$.

For the second part, if $k \geq d(G) + 1$, and $\tilde{\Gamma}_k(G)$ is connected, then every k -tuple $(g) \in \Gamma_k(G)$ is connected to some $(h_1, \dots, h_{k-1}, \text{id})$. This proves the result. \square

Let us show that for k large enough, the graphs $\Gamma_k(G)$ do become connected. Let $\bar{d} = \bar{d}(G)$ be the maximal size of the minimal generating set, which is a set such that no generator can be omitted. By $\ell = \ell(G)$ denote the length of the longest increasing sequence of subgroups:

$$G_0 = \{\text{id}\} \subset G_1 \subset \dots \subset G_\ell = G,$$

where $G_i \neq G_{i+1}$. Note that

$$d(G) \leq \bar{d}(G) \leq \ell(G) \leq \log_2 |G|,$$

where the latter inequality follows from $|G_{i+1}|/|G_i| \geq 2$.

Proposition 2.2.2 *If $k \geq d(G) + \bar{d}(G)$, then $\Gamma_k(G)$ is connected.*

Proof. By Proposition 2.2.1, it suffices to prove that $\tilde{\Gamma}_k(G)$ is connected.

Let $d = d(G)$, $\bar{d} = \bar{d}(G)$, $k \geq d + \bar{d}$. Fix a generating set $\langle h_1, \dots, h_d \rangle = G$ and consider a generating k -tuple $(h) = (h_1, \dots, h_d, \text{id}, \dots, \text{id})$. Observe that every generating k -tuple $(g_1, \dots, g_k) \in \mathcal{N}_k(G)$ is connected (in graph $\tilde{\Gamma}_k(G)$) to (h) . Indeed, the set of generators is redundant, so up to a permutation we have $\langle g_{d+1}, \dots, g_{d+\bar{d}} \rangle = G$. Then (g) is connected to $(h_1, \dots, h_d, g_{d+1}, \dots, g_{d+\bar{d}})$ and finally to (h) . \square

Theorem 2.2.3 (Babai) *Let G be a finite group, $k = 2\lceil \log_2 |G| \rceil$. Then the diameter $\Delta = \Delta(G, k)$ of the product replacement graph $\Gamma_k(G)$ is at most $C \cdot \log^2 |G|$, where C is a universal constant.*

The proof of Babai [Bb5] is an extension of the proof of Proposition 2.2.2. It uses an important ‘‘cube doubling’’ idea (see [ER,Bb3]) which has other important applications and is worth noting.

Proof. Let $r = \lceil \log_2 |G| \rceil$. We say that $(g) = (g_1, \dots, g_r) \in \mathcal{N}_r(G)$ is *good*¹ if $C^{-1}C = G$, where

$$C = C(g) = \{g_1^{\epsilon_1} \cdots g_r^{\epsilon_r} \mid \epsilon_1, \dots, \epsilon_r \in \{0, 1\}\}.$$

Observe that by definition one can move from $(g, h) = (g_1, \dots, g_r, h_1, \dots, h_r)$ to $(g, h') = (g_1, \dots, g_r, h'_1, \dots, h'_r)$ in at most $2r^2$ Nielsen moves.

We claim that for every k -tuple $(g) \in \mathcal{N}_k(G)$ one can obtain a k -tuple of the form (u, v) , where $(u), (v) \in \mathcal{N}_r$, and (u) is good. Indeed, permute elements of (g) so that the last r of them generate G . That will be our (v) . Now let us construct a *good* r -tuple (u) by induction. Set $C_0 = \emptyset$. Assume that at step i we determined u_1, \dots, u_i such that $|C_i| = 2^i$, where $C_{i-1} = C(u_1, \dots, u_i)$. If $C_i^{-1}C_i = G$, then this is already a good set. Otherwise there exists $w \in \{\text{id}, v_1, \dots, v_r\}$ such that $g_{i+1}D_i w \neq D_i$, where $D_i = C_i^{-1}C_i$. Indeed, if $g_{i+1}D_i = D_i$ for $w = \text{id}$, then for $w = v_i$ we have $g_{i+1}D_i v_i = D_i v_i$. But since (v) is a generating r -tuple, and $D_i G \neq D_i$, there exists v_i such that $g_{i+1}C_i v_i \neq C_i$. Now move g_{i+1} to an element $u_{i+1} = g_{i+1}z$, where $z = D_i^{-1}w$, $u_{i+1} \notin D_i$. Then $|C_{i+1}| = 2|C_i|$, which completes the induction step in at most $(2i + 1)$ Nielsen moves.

The total number of induction steps is at most $\lceil \log_2 |G| \rceil = r$ since we are ‘‘doubling the cube’’ C_i at every step. We need therefore at most r^2 Nielsen moves and $O(r)$ transpositions and inversions to obtain a good k -tuple. From the previous observation, two good k tuples are connected by at most $2r^2$ Nielsen moves and $O(r)$ transpositions and inversions. Recall now that by the proof of Proposition 2.2.2 each transpositions and inversions can be substituted by at most c Nielsen moves, where c is a universal constant. This completes the proof. \square

¹In [Bb5] such (g) are called ‘‘cube-semigenerators’’

Remark 2.2.4 Proposition 2.2.2 is an observation in [DS3]. A weaker version with $k \geq 2\bar{d}(G)$ was given in the original paper [CLMNO] (see also [Gi]). Theorem 2.2.3 and the proof closely follow the paper [Bb5].

Let us also note that computing $\bar{d}(G)$ is not an easy task. A recent result $\bar{d}(S_n) = n - 1$ of Whiston [Wh] is worth mentioning, especially in comparison with another hard result $\ell(S_n) = \lfloor \frac{3n-1}{2} \rfloor - b_n$, where b_n is the number of ones in the binary expansion of n (see [CST,Bb1]).

2.3 Abelian, nilpotent and solvable groups.

The main result of this section is due to Dunwoody [Du2], who showed that $\Gamma_{d(G)+1}(G)$ is connected for any finite solvable group G . The case $k = d(G)$ is more delicate, and is also considered.

Example 2.3.1 Let $G = (\mathbb{Z}_p)^m$, where p is a prime. Then $\mathcal{N}_k(G)$ is in one to one correspondence with the set $\text{GL}(m, p)$ of nonsingular $m \times m$ matrices over \mathbb{F}_p . Note that Nielsen moves $R_{i,j}^\pm, L_{i,j}^\pm$ become elementary transformations (adding or subtracting a row), and therefore do not change the determinant. Thus $\Gamma_m(G)$ has at least $(p-1)$ components. Further, $\Gamma_m(\mathbb{Z}_p^m)$ is a Cayley graph of $\text{GL}(m, p)$ with a generating set $E_{i,j}^\pm$.

In the other direction, the elementary matrices $E_{i,j}^\pm$ with ones on the diagonal, ± 1 at (i, j) and zero elsewhere, generate $\text{SL}(m, p)$. Therefore $\Gamma_m(\mathbb{Z}_p^m)$ has exactly m isomorphic connected components.

Note also that $\tilde{\Gamma}(\mathbb{Z}_p^m)$ is a Cayley graph of $\Gamma_m(\mathbb{Z}_p^m)$ with a generating set consisting of elementary matrices, permutation matrices, and matrices which have (-1) at (i, i) , 1 elsewhere on diagonal, and zero outside of the main diagonal. For $p = 2$ graph $\tilde{\Gamma}(\mathbb{Z}_p^m)$ is also connected, but for $p \geq 3$, p -prime, it has $(p-1)/2$ connected components.

Theorem 2.3.2 ([NN,DG]) *Let G be a finite abelian group, given as*

$$G \simeq \mathbb{Z}_{m_1} \times \mathbb{Z}_{m_2} \times \cdots \times \mathbb{Z}_{m_r},$$

where $m_1 | m_2 | \dots | m_r$. Then $d(G) = r$ and

- (i) $\Gamma_k(G)$ is connected for all $k \geq r + 1$,
- (ii) $\Gamma_r(G)$ has $\phi(m_1)$ components of equal size,

where $\phi(m)$ is the Euler function (the number of integers less than m which are relatively prime with m).

The first part of the theorem² is due to B.H. Neumann and H. Neumann [NN] (see also [DG,Ev2]). The second part is due to Diaconis and Graham [DG], who also rediscovered the first part.

The proof is an elementary, though nontrivial generalization of the observation in Example 2.3.1. One proves in part (ii) that when an element $g_i \in G$ is written as $(a_{i,1}, \dots, a_{i,r})$, then $\det(a_{i,j})$ is the only invariant of generating k -tuples under Nielsen moves (see [DG] for details).

For nilpotent groups the situation with the case $k = d(G)$ is more complicated. First, using Gaschütz' Lemma, one can lift "invariants" from abelian quotients

²The result in [NN] was formulated in a slightly different language of T -systems, see below.

to the whole group. This idea was successfully used in [Du1] by Dunwoody. A somewhat different approach was proposed in [Ne]:

Lemma 2.3.3 (Higman)³ *Let Γ' be a connected component of $\Gamma_2(G)$ which contains $(g_1, g_2) \in \mathcal{N}_2(G)$. Then for every $(h_1, h_2) \in \Gamma'$ we have $[h_1, h_2]$ is a conjugate of $[g_1, g_2]$.*

In other words, the conjugacy class of the commutator $[g_1, g_2]$ is invariant under Nielsen moves. This can be easily checked directly. Note that in $\tilde{\Gamma}_k(G)$, the transposition of generators inverts the commutator: $[g_1, g_2] \rightarrow [g_2, g_1] = [g_1, g_2]^{-1}$, so one has to add an inverse of the conjugacy class. In [Ne] B.H. Neumann used Higman's Lemma to find distinct connected components in a particular nilpotent group $G \simeq \mathbb{Z}_8 \times D^4$, where D is a dihedral group of order 8.

Question 2.3.4 (B.H. Neumann) Can one find a generalization of Higman's Lemma for more than two generators, that is for $k > 2$?

Open Problem 2.3.5 *Describe all connected components of $\mathcal{N}_k(G)$, where G is a finite nilpotent group and $k = d(G)$.*

Now let us prove the main result of this section:

Theorem 2.3.6 (Dunwoody) *Let G be a finite solvable group, $k \geq d(G) + 1$. Then $\Gamma_k(G)$ is connected.*

Proof. The proof is by induction on the length ℓ of a chief series of G :

$$\{\text{id}\} = G_0 \subset G_1 \subset \dots \subset G_\ell = G,$$

where G_{i-1} is a minimal normal G -invariant subgroup of G_i , and $G_i \triangleleft G$, $i = 1, \dots, \ell$.

When $\ell = 0$ there is nothing to prove. Assume $\ell \geq 1$. Fix any $(k-1) \geq d(G)$ generators h_1, \dots, h_{k-1} of G . We will show that any generating k -tuple $(g) = (g_1, \dots, g_k) \in \Gamma_k(G)$ is connected to $(h) = (\text{id}, h_1, \dots, h_{k-1}) \in \Gamma_k(G)$ by Nielsen moves. This clearly implies connectivity of the whole graph.

Denote G_1 by M . By the induction hypothesis on G/G_1 , (g) is connected to a k -tuple $(m, m_1 h_1, \dots, m_{k-1} h_{k-1})$, where $m, m_1, \dots, m_{k-1} \in M$. If $m = \text{id}$, the elements $m_i h_i$ generate G and we can use $R_{1,i}^\pm$ to get $m \neq \text{id}$. Therefore we can assume that $m \neq \text{id}$.

Now note that $h_i^{-1} m h_i = (m_i h_i)^{-1} m (m_i h_i)$. For any given $g \in G$ write $g = \mathbf{w}(h_1, \dots, h_{k-1})$ as a word in generators of G . Then

$$\begin{aligned} & (\mathbf{w}(m_1 h_1, \dots, m_{k-1} h_{k-1}))^{-1} m \mathbf{w}(m_1 h_1, \dots, m_{k-1} h_{k-1}) \\ &= (\mathbf{w}(h_1, \dots, h_{k-1}))^{-1} m \mathbf{w}(h_1, \dots, h_{k-1}) = g^{-1} m g = m^g. \end{aligned}$$

Therefore by successive application of $R_{1,i}^+$ and $L_{1,i}^-$ one can connect the generating k -tuple $(m, m_1 h_1, \dots, m_{k-1} h_{k-1})$ with $(m^g, m_1 h_1, \dots, m_{k-1} h_{k-1})$ for any g .

Since M is a minimal normal subgroup of G , every $m' \in M$ is a product of conjugates of m . Therefore by iterating the previous procedure, one can use $L_{i,1}^\pm$

³Higman's Lemma has appeared in [NN] in a somewhat different form.

to get rid of all the m_i and obtain (m, h_1, \dots, h_{k-1}) . But then, since h_1, \dots, h_{k-1} generate G , we easily obtain $(h) = (\text{id}, h_1, \dots, h_{k-1})$ by the Nielsen moves. This completes the proof. \square

Remark 2.3.7 As shown by Evans [Ev2], the graphs $\Gamma_k(G)$ with $k \geq d(G) + 1$ are connected for all (even infinite) nilpotent groups. This fails badly for $k = d(G)$ (see [Du1]), even when the so called “swap moves” are allowed :

$$(g_1, \dots, g_i, \dots, g_k) \rightarrow (g_1, \dots, g'_i, \dots, g_k),$$

where $g'_i \in G$ is any element such that $(g_1, \dots, g'_i, \dots, g_k)$ is a generating k -tuple. We refer to [TT,Ro] for reasoning behind this definition and various examples. As we shall see later, all but a finite number of simple groups are connected by these swap moves for $k \geq 3$ (see proof of Proposition 2.5.12).

We challenge the reader to generalize (or disprove) the above mentioned Evans’ and Dunwoody theorems to all (even infinite) solvable groups. We refer to a related result of Myasnikov [My] who proved Andrews–Curtis Conjecture for certain solvable groups.

Let us mention that in a different context Question 2.1.3 has a negative answer for infinite groups. This is related to the following Waldhausen’s question: “If $G = \langle x_1, \dots, x_n, r_1, \dots, r_m \rangle$ can be generated by fewer than n elements, then does the normal closure of $\langle r_1, \dots, r_m \rangle$ in F_n contain a primitive element of F_n ? ” (see [LyS], p. 92). The answer for this question is negative (see [No,Ev1]). On the other hand, it was noted in [Du2] (in fact, motivated his work) that this is equivalent to the question of whether a generating set for G with n elements can always be changed by Nielsen moves to a set containing the identity element, given $d(G) = n - 1$. But this is a weaker condition than the connectivity of $\tilde{\Gamma}_{n-1}(G)$. Thus $\tilde{\Gamma}_{n-1}(G)$ and therefore $\Gamma_{n-1}(G)$ is disconnected in this case.

2.4 T -systems.

Let G be a finitely generated group, and $d(G)$ be the minimal number of generators of G . *Systems of transitivity* (or T -systems) are defined as follows.

We say that N is a G -defining subgroup of a free group F_k if N is normal in F_k and $F_k/N \simeq G$. Denote by $\Sigma_k(G)$ the set of G -defining subgroups of F_k , and consider an action of $\text{Aut}(F_k)$ on $\Sigma_k(G)$. The orbits of this action are called T_k -systems. They were introduced in [NN] by B.H. Neumann and H. Neumann, and studied on and off since then.

Recall also that $\text{Aut}(F_k)$ is generated by the set $\bar{\Upsilon}_k$ of Nielsen moves, transpositions and inversions. These moves define a graph structure on $\Sigma_k(G)$. Now T_k -systems correspond to connected components of $\Sigma_k(G)$.

Recall that $\tilde{\Gamma}_k(G)$ is also defined by the set of moves $\bar{\Upsilon}_k$. Consider the natural action of $\text{Aut}(G)$ on the set $\mathcal{N}_k(G)$ of generating k -tuples of G . Identifying under this action defines a projection of graph $\tilde{\Gamma}_k(G)$ onto $\Sigma_k(G)$. We obtain:

Proposition 2.4.1 *The number of connected components in $\Sigma_k(G)$ (the number of T_k -systems) is less or equal to the the number of connected components in $\tilde{\Gamma}_k(G)$. In the other direction, if $\Sigma_k(G)$ is connected and $k \geq 2d(G)$, then $\tilde{\Gamma}_k(G)$ is also connected.*

Proof. The first part follows from the previous observation. For the second part, connectivity of $\Sigma_k(G)$ implies that every k -tuple $(g) \in \mathcal{N}_k(G)$ is connected in $\tilde{\Gamma}_k(G)$ to some $\varphi(h)$, where $\varphi \in \text{Aut}(G)$ and $(h) \in \mathcal{N}_k(G)$ is fixed. Now take $(h) = (h_1, \dots, h_d, \text{id}, \dots, \text{id})$, where $d = d(G)$. Then $\varphi(h) = (h'_1, \dots, h'_d, \text{id}, \dots, \text{id}) \in \mathcal{N}_k(G)$ is a generating k -tuple which is connected to (h) in $\tilde{\Gamma}_k(G)$. Indeed, permute elements in $\varphi(h)$ to make the first d of them to be id . Now generate (h_1, \dots, h_d) in the first d positions, and then make the remaining $(k - d)$ elements become id . We conclude that any $(g) \in \mathcal{N}_k(G)$ is connected in $\tilde{\Gamma}_k(G)$ to a fixed (h) as above. \square

Corollary 2.4.2 *If $k \geq 2d(G)$, then $\Gamma_k(G)$ is connected if and only if G has exactly one T_k -system.*

Proof. The corollary follows immediately from Proposition 2.4.1 and Proposition 2.2.1. \square

Theorem 2.4.3 (Gilman) *Let G be a finite simple group and $k \geq 3$. Then $\text{Aut}(F_k)$ acts on at least one T_k -system of G as the alternating or symmetric group.*

This result is quite miraculous since it basically says that for simple groups on some orbit the group action of $A = \text{Aut}(F_k)$ is not only simply transitive, but highly transitive. The proof is also quite interesting, but goes beyond the scope of this survey.

In fact, Theorem 2.4.3 is a compilation of several results. The main part, which is the case $k \geq 4$, is due to Gilman [Gi]. It was subsequently extended to the case $k = 3$ by Evans [Ev3] (see Lemma 5.2) under assumption that a simple group is generated by two elements, one of which is an involution. The proof in [Ev3] follows closely [Gi], so we will refer to Theorem 2.4.3 as Gilman's Theorem.

Now, the assumption that a simple group is generated by two elements one of which is an involution, holds in fact for every simple group (see Remark 2.4.8). The easiest way to see this is to apply Theorem 1.1.8 of Guralnick and Kantor. Recall that by the Feit-Thompson Theorem (see e.g. [Go]) every finite simple group contains an involution. Therefore by 1.1.8 (i), it generates the whole group with some elements of a certain conjugacy class.

Corollary 2.4.4 *If G is simple and has exactly one T_k -system, $k \geq 3$, then G^m has also exactly one T_k -system, given $m < N = |\Sigma_k(G)|$.*

Proof. By Hall's Theorem 1.2.1, the generating k -tuples of G^m correspond to those m -tuples of generating k -tuples of G which all lie in different orbits of the diagonal action of $\text{Aut}(G)$ on \mathcal{N}_k . Assume now that $\Sigma_k(G)$ is connected. By Theorem 2.4.3, $\text{Aut}(F_k)$ acts $(N-1)$ -transitively on $\Sigma_k(G)$. But that means that for any $m \leq N$ one can use extended Nielsen moves to get from m elements in different orbits (of $\text{Aut}(G)$ on $\mathcal{N}_k(G)$) to m elements in any other different m orbits. But this is more than enough to show that $\Sigma_k(G^m)$ is connected, as the action of the group $\text{Aut}(G^m)$ on $\mathcal{N}_k(G^m)$ contains separate actions of m copies of $\text{Aut}(G)$ on different $\mathcal{N}_k(G)$. \square

Corollary 2.4.5 *If G is simple and has exactly one T_k -system, $k \geq 4$, then $\Gamma_k(G^m)$ is connected, given $m < |\Sigma_k(G)|$.*

Proof. Recall that for all finite simple groups $d(G) = 2$. Now the result follows from Corollary 2.4.4 and Corollary 2.4.2. \square

Unfortunately we cannot improve the condition $k \geq 4$ to $k \geq 3$. Indeed, the group A_n having only one T_3 -system does not imply that $\Gamma_3(A_n)$ is connected since in Corollary 2.3.2 we need k to be at least two times the minimal number of generators.

Let us conclude by formulating a stronger version of Question 2.1.3:

Question 2.4.6 Is there a finite group G , $d = d(G)$ such that G has more than one T_{d+1} -system (equivalently: $\Sigma_{d+1}(G)$ is disconnected)?

Remark 2.4.7 Corollary 2.4.4 is a refined version of the observation made by Dunwoody in a review of the paper [Gi] (Math. Reviews 55#8186, see also [Du2]). The story is that in [Du2] Dunwoody proposed P. Hall's example A_5^{19} as a possible candidate for a positive answer to the question 2.4.6.⁴ In a review Dunwoody asserts that Gilman's results imply that A_5^{19} cannot satisfy 2.4.6. This does not seem clear to us in view of the condition $k \geq 4$ in [Gi], but certainly Gilman was very close to resolving the A_5^{19} problem. Only after a paper [Ev3] of Evans about 15 years later, one can extend the result to the case $k = 3$, so to prove that $G = A_5^{19}$ has only one T_3 -system. Indeed, since $\varphi_3(A_5) > \varphi_2(A_5) = 19/30$ (this is Hall's result) by Corollary 2.4.4 we have

$$|\Sigma_3(A_5)| = \frac{\varphi_3(A_5)|A_5^3|}{|\text{Aut}(A_5)|} > \frac{\frac{19}{30} \cdot 60^3}{120} = 1140.$$

From here we obtain a remarkable connectivity of the graphs $\Sigma_3(A_n^N)$ for $N \leq 1140 < |\Sigma_3(A_5)|$. It is unclear, however, whether $\text{Aut}(F_3)$ acts as a symmetric (not as an alternating) group on $\Sigma_3(A_5)$, i.e. whether $A_5^{|\Sigma_3(A_5)|}$ has only one T_3 -system as suggested in Question 2.4.6.

Remark 2.4.8 A few words about simple groups being generated by two elements, one of which is an involution. For sporadic groups this follows from their construction, and for several series (such as alternating groups, $\text{PSL}(n, q)$, etc.) it was known for decades (see e.g. [CM,Go]). In full generality this was first proved by Malle, Saxl and Weigel in [MSW].

2.5 Simple groups.

In this section we elaborate on the connectivity of the product replacement graph in the special case of simple groups. As we shall see, this case seems to be of particular importance.

Example 2.5.1 Let $G = A_5$. In a pioneering paper [NN] the authors show that A_5 has at least two T_2 -systems. To see that, simply use Higman's Lemma and

⁴Actually, Dunwoody proposes A_5^{19} as a possible challenge problem for computational group theory. Note that the order of the group $|A_5^{19}| \approx .61 \cdot 10^{34}$ is much too large for any brute force computation. Imagine now how much challenge it was in 1967 when Dunwoody proposed the problem!

observe that the commutator of two generating elements must be either a 5-cycle or a 3-cycle. Since both can occur, there are at least two T_2 -systems.

Note that when $G = S_5$, we already have at least three T_2 -systems: elements with two 2-cycles can occur as commutators of a generating pair of element.

Proposition 2.5.2 *The number of connected components of $\Gamma_2(A_n)$ increases as $n \rightarrow \infty$.*

Proof. Let n be odd. Simply observe that when $(m, p) = 1$ the elements of the form $a = (1, 2, \dots, n)$ and $b = (1, 2, \dots, p)$, (in cycle notation) generate the whole A_n . Their commutators belong to different conjugacy classes (for $p < m/2$). Now Higman's Lemma implies the result. \square

Question 2.5.2 What is the number of conjugacy classes of the commutator $[a, b]$ of two generating elements $(a, b) \in \mathcal{N}_2(A_n)$?

Note, of course, that $[A_n, A_n] = A_n$, $n \geq 5$, and two elements in A_n generate A_n with probability $\mathbf{P} = (1 - 1/n + O(1/n^2))$. Also, if $[a, b] = c$, then $[\sigma a, \sigma b] = \sigma c$ for all $\sigma \in S_n$. Therefore elements with the same cycle structure either can be all obtained, or can't be all obtained as commutators of two generating elements in A_n . Of course, some conjugacy classes, e.g. id, can't be obtained in this way.

Open Problem 2.5.3 *Describe all connected components of $\Gamma_2(A_n)$, $n \geq 5$.*

For $k \geq 3$ very few results are known, so let us start with the general conjecture of obvious importance.

Conjecture 2.5.4 (Wiegold) *For every $k \geq 3$ and every nonabelian group G , there exists only one T_k -system (equivalently: the graph $\Sigma_k(G)$ is connected).*

While never published, this conjecture was attributed to Wiegold on several occasions (see [Da, Ev2, Ev3]). It seems inspired by Gilman's paper [Gi]. Let us state here a somewhat stronger companion conjecture.

Conjecture 2.5.5 *For every nonabelian simple group G and every $k \geq 3$, the product replacement graph $\Gamma_k(G)$ is connected.*

Note that since $d(G) = 2$ for all finite simple groups, for $k \geq 4$ Conjectures 2.5.4 and 2.5.5 are equivalent. A similar conjecture for $G = S_n$ was made by Diaconis and Graham in [DG]. Let us briefly formulate what is known in this direction:

Theorem 2.5.6 *Wiegold's Conjecture holds in the following cases:*

- (i) [Gi] $G = \text{PSL}(2, p)$, where $p \geq 5$ is a prime, and $k \geq 3$,
- (ii) [Ev3] $G = \text{PSL}(2, 2^m)$ or $G = \text{Sz}(2^{2m-1})$, where $m \geq 2$, and $k \geq 3$,
- (iii) [Da] $G = A_6, A_7$, and $k = 3$,
- (iv) [CP] $G = A_8, A_9, A_{10}$, and $k = 3$.

Note that the case $G = A_5$ is covered in (i): $A_5 \simeq \text{PSL}(2, 5)$. In fact, [CP] proves that the stronger Conjecture 2.5.5 holds for $k = 3$ and $G = A_n$, $6 \leq n \leq 10$.

Proposition 2.5.7 *Let $G = \text{PSL}(2, p)$, $H = G^m$, where $m \geq p(p+1)$. Then $d(H) \geq 3$ and for every $k \geq d(G^{(m+1)})$, the group H has only one T_k -system. Further, if $m \geq p^2(p-1)(p+1)^2$, then $d(H) \geq 4$ and the product replacement graph $\Gamma_k(G)$ is connected.*

Proof. Recall that $|\mathrm{PSL}(2, p)| = p(p-1)(p+1)$, $\mathrm{Aut}(\mathrm{PSL}(2, p)) = \mathrm{GL}(2, p)$, $|\mathrm{GL}(2, p)| = p(p-1)^2(p+1)$. By Hall's Theorem 1.2.1, the largest power N such that $d(G^N) \leq 2$ is given by

$$N = \varphi_2(G) \cdot \frac{|G|^2}{|\mathrm{Aut}(G)|} < \frac{[p(p-1)(p+1)]^2}{p(p-1)^2(p+1)} = p(p+1).$$

Therefore for all $m \geq p(p+1)$ we have $d(H) \geq 3$. Analogously, whenever $m \geq p^2(p-1)(p+1)^2$, we have $d(H) \geq 4$, which proves the first part of both claims.

The second part follows easily from Corollaries 2.4.4 and 2.4.5. Indeed, uniqueness of T_k -system condition in 2.4.4,5 is satisfied by part (i) of Theorem 2.5.6. By Hall's Theorem, the condition $m \leq |\Sigma_k(G)|$ is equivalent to the condition $k \geq d(H)$, so $m < |\Sigma_k(G)|$ is equivalent to $k \geq d(G^{(m+1)})$. This completes the proof. \square

There is one more result in favor of Conjecture 2.5.5. To state it we need the notion of *spread* (see [BW1]).

Definition 2.5.8 *A 2-generator group G is said to have spread m if for every m group elements $g_1, \dots, g_m \neq \mathrm{id}$ there exists an element $h \in G$ such that $\langle g_1, h \rangle = \dots = \langle g_m, h \rangle = G$.*

Much is known about the spread of simple and quasisimple groups. We refer to papers [BW1, BW2, GK, GS] for various extensions and references. Let us just mention that all symmetric groups S_n and alternating groups A_n , $n \geq 5$, have spread 2 (see [Bi, BW1]).

Conjecture 2.5.9 ([BW1, GS]) *All finite nonabelian simple groups have spread 2.*

Theorem 2.5.10 ([GS]) *All but a finite number of finite nonabelian simple groups have spread 2.*

Proof. By Theorem 1.1.8 (ii), $\mathrm{PC}(G) \rightarrow 1$ as $|G| \rightarrow \infty$ unless G is isomorphic to alternating groups A_n or odd-dimensional orthogonal groups $\Omega(n, q)$ with a bounded size q of the field.

Now let G be a simple group, C be a conjugacy class on which $\mathrm{PC}(G)$ is maximized. For any $g_1, g_2 \in G$ we have

$$\mathbf{P}(\langle g_1, h \rangle = \langle g_2, h \rangle = G \mid h \in C) \geq 1 - 2(1 - \mathrm{PC}(G)).$$

Therefore $\mathrm{PC}(G) > 1/2$ implies that $\mathbf{P}(\cdot) > 0$, and hence G has spread 2.

We conclude that for all but the two exceptional series there is only a finite number of simple groups which have $\mathrm{PC}(G) \leq 1/2$. As we mentioned before, all alternating groups have spread 2. The case of odd orthogonal groups is considered separately (see [GK, GS]). \square

Actually, one can describe explicitly the set of possible exceptions in Theorem 2.5.10. It would be interesting to check Conjecture 2.5.9 by computational means, or otherwise.

Definition 2.5.11 A generating k -tuple $(g_1, \dots, g_k) \in \mathcal{N}_k(G)$ is called *redundant*⁵ if it is not a minimal generating set, which means that one generator g_i can be omitted so that the remaining elements generate the whole group G .

Proposition 2.5.12 ([Ev3]) Let G be a 2-generated group of spread 2 and let $k \geq 3$. Then all redundant generating k -tuples belong to the same connected component in $\Gamma_k(G)$.

Proof. Clearly, every redundant generating k -tuple is connected in $\Gamma_k(G)$ to a k -tuple with some entry id . It follows from the proof of Proposition 2.2.1 that the connected component in $\tilde{\Gamma}_k(G)$ which contains such a k -tuple is also connected in $\Gamma_k(G)$. Thus it suffices to prove that all redundant generating k -tuples belong to the same connected component in $\tilde{\Gamma}_k(G)$.

Let (g) , (h) be two redundant generating k -tuples. One can always permute their elements and make the last element to be identity. Thus we can assume that $(g) = (g_1, \dots, g_{k-1}, \text{id})$, $(h) = (h_1, \dots, h_{k-1}, \text{id})$. Without loss of generality we can also assume that $g_1, h_2 \neq \text{id}$. Since G has spread 2, there exists an element $z \in G$ such that $\langle g_1, z \rangle = \langle h_2, z \rangle = G$. Clearly, (g) is connected to $(g_1, g_2, \dots, g_{k-1}, z)$, and therefore to $(g_1, h_2, \dots, h_{k-1}, z)$ and to $(h_1, h_2, \dots, h_{k-1}, z)$ and $(h_1, h_2, \dots, h_{k-1}, \text{id})$. This completes the proof. \square

Observe that by Theorem 1.1.1 a random generating k -tuple in G must be redundant with high probability (given $k \geq 3$). Thus in a sense which will be made precise in the next section, Theorem 2.5.12 and Theorem 2.5.10 imply that Conjecture 2.5.5 holds “for most” generating k -tuples of a large enough simple group G .

We conclude with the following important observation:

Proposition 2.5.13 Let G be a 2-generated group of spread 2 and let $k \geq 3$. Then the connectivity of $\Sigma_k(G)$ implies the connectivity of $\Gamma_k(G)$.

Proof. Proposition 2.4.1 covers the case $k \geq 2 \cdot d(G) = 4$. It remains to be proved that if G has only one T_3 -system, then $\Gamma_3(G)$ is connected. But this is a simple combination of the proof of Proposition 2.4.1 with Proposition 2.5.12. In the notation of the proof of 2.4.1, we need to show that $\varphi(h) = (h'_1, h'_2, \text{id})$ is connected to $(h) = (h_1, h_2, \text{id})$. But that’s clear since G has spread 2 and both 3-tuples are redundant. Proceeding as in the proof of 2.4.1, we obtain the result. \square

Corollary 2.5.14 If G is a finite nonabelian simple group of spread 2, then Conjecture 2.5.5 is equivalent to Wiegold’s Conjecture 2.5.4. \square

Remark 2.5.15 As noted in [BW1], if G has spread r , so does any quotient H . Also, if G_1, G_2 have spread 2, then $G_1 \times G_2$ has spread 1. It would be interesting to investigate when G^m has spread 2, when G is simple. Let us note that the real power of the bounds in [GK] combined with proof of Theorem 2.5.10 give large spread for various series of simple groups (see [GS]). The latter can be used to show that certain large powers of simple groups can have spread 2.

⁵Note that our use of the term is different from that in [Ev3, Da]

Remark 2.5.16 Parts (i), (ii) of Theorem 2.5.6 are obtained by an elaborate calculation using subgroup structure and conjugacy classes. It seems that generalizations of the brute force technique to other series of groups (especially of rank ≥ 2) must be difficult.

The proof of David [Da] of the case (iii) is based on a tedious calculation. In [CP] we present a “proof by computer” of part (iv). We describe the algorithm in section 2.7.

Let us note that symmetric groups S_n , $n \geq 3$, also have spread 2. This supports the conjecture by Diaconis and Graham [DG] on the connectivity of $\Gamma_3(S_n)$ (see above).

Clearly, Proposition 2.5.7 is slightly weaker than what one would like, which is a condition $k \geq d(G^m)$ rather than $k \geq d(G^{(m+1)})$. Nevertheless for powers of $\text{PSL}(2, p)$ as in 2.5.7, the answer to the Question 2.4.6 is positive. Still, it would be interesting to check whether $\text{Aut}(F_k)$ acts as a symmetric or an alternating group on $\Sigma_k(G)$ (cf. Remark 2.4.7.)

Our proof of Theorem 2.5.10 by Guralnick and Shalev [GS], follows closely the proof in their paper. As we mentioned above, it is a direct application of the previous result of Guralnick and Kantor [GK].

2.6 The “large” connected component.

We give here a probabilistic approach to the connectivity problem, by showing existence of the large connected component for k moderately large.

Theorem 2.6.1 ([P5]) *Let $\{G_n\}$ be a sequence of finite nonabelian simple groups such that $|G_n| \rightarrow \infty$ as $n \rightarrow \infty$. Fix $k \geq 3$. Then graphs $\Gamma_k(G_n)$ have connected components $\Gamma'_k(G_n) \subset \Gamma_k(G_n)$, such that*

$$\frac{|\Gamma'_k(G_n)|}{|\Gamma_k(G_n)|} \rightarrow 1 \quad \text{as } n \rightarrow \infty.$$

Furthermore, for n large enough $\Gamma'_k(G_n)$ contain all the redundant generating k -tuples.

Proof. By Theorem 2.5.10, every large enough simple group has spread 2. Further, by Proposition 2.5.12, the set of redundant generating k -tuples is connected in $\Gamma_k(G)$. Finally, Theorem 1.1.1 implies that, given $l \geq 2$, a random l -tuple of elements in G is a generating l -tuple, with probability $\rightarrow 1$ as $n \rightarrow \infty$. Taking $l = k - 1$ and $l = k$, we obtain that a random k -tuple, $k \geq 3$ is a redundant generating k -tuple with the first $(k - 1)$ elements generating the whole group G . This completes the proof. \square

In particular, the graphs $\Gamma_3(A_n)$ have a “large” connected component, i.e. a component $\Gamma' = \Gamma'_3(A_n)$ such that $|\Gamma'|/|\Gamma| \rightarrow 1$, as $N \rightarrow \infty$. Compare this with Proposition 2.2.2, which proves connectivity for $k \geq d(A_n) + \bar{d}(A_n) = 2 + (n - 1) = n + 1$ (see [Wh]).

The following generalization is straightforward:

Theorem 2.6.2 ([P5]) *Let $\{G_n\}$ be a sequence of finite groups, and let k_n be an integer sequence such that*

$$\varphi_{k_n}(G_n) \rightarrow 1, \quad \text{as } n \rightarrow \infty.$$

Then the graphs $\Gamma_n = \Gamma_{k_n+d(G_n)}(G_n)$ have connected components Γ'_n , such that

$$\frac{|\Gamma'_n|}{|\Gamma_n|} \rightarrow 1 \quad \text{as } n \rightarrow \infty.$$

Furthermore, graphs Γ'_n contain all the generating $(k_n + d(G_n))$ -tuples of the form $(g_1, \dots, g_{k_n}, \text{id}, \dots, \text{id})$.

Proof. The proof is an easy modification of the proof of Theorem 2.6.1. Indeed, every generating $(k+d)$ -tuple of the form $(g) = (g_1, \dots, g_k, \text{id}, \dots, \text{id})$, where $k = k_n$, $d = d(G_n)$, is connected to a fixed $(\text{id}, \dots, \text{id}, h_1, \dots, h_d)$, where $\langle h_1, \dots, h_d \rangle = G$. Now, by hypothesis, the fraction of elements (g) among all k -tuples goes to 1 as $n \rightarrow \infty$. Thus fraction of (g) among all generating k -tuples goes to 1 as $n \rightarrow \infty$. By the observation above, they all lie in the same connected component, which implies the result. \square

Note now that Corollary 1.5.2 proves the conditions of Theorem 2.6.2 in various special cases. The first two cases which deal with nilpotent and solvable groups were completely resolved by Dunwoody's Theorem 2.3.6. In the last two cases we have:

Corollary 2.6.3 *The product replacement graphs $\Gamma_n = \Gamma_{k_n}(G_n)$ contain connected components Γ'_n such that $|\Gamma'_n|/|\Gamma_n| \rightarrow 1$ as $n \rightarrow \infty$ if one of the following holds:*

(i) $k_n = \omega(n) \cdot \log(m_n)$, and each given G_n is a product of finite nonabelian simple groups,

with at most m_n copies of each,

(ii) $k_n = \log_2 |G_n| + d(G_n) + \omega(n)$, for any $\{G_n\}$,

where $\omega(n)$ is any function such that $\omega(n) \rightarrow \infty$ as $n \rightarrow \infty$. \square

Note that part (ii) implies that to have a large component in general product replacement graphs it suffices to have $d(G) + \log_2 |G| + C$ generators. This is somewhat weaker than the result of Proposition 2.2.2 which proves that $k = d(G) + \bar{d}(G) \leq d(G) + \log_2 |G|$ is enough to imply connectivity. On the other hand, as we have seen before, part (i) does not follow from 2.2.2.

We conclude with another unexpected bonus from Gilman's Theorem 2.4.3, which gives a generalization of Theorem 2.6.1 to any fixed power of a simple group.

Theorem 2.6.4 *Let $\{G_n\}$ be a sequence of finite nonabelian simple groups such that $|G_n| \rightarrow \infty$ as $n \rightarrow \infty$, and let m be a fixed integer. Fix $k \geq 3$. Then the graphs $\Gamma_k(G_n^m)$ have connected components $\Gamma'_k(G_n^m) \subset \Gamma_k(G_n^m)$, such that*

$$\frac{|\Gamma'_k(G_n^m)|}{|\Gamma_k(G_n^m)|} \rightarrow 1 \quad \text{as } n \rightarrow \infty.$$

Furthermore, for n large enough $\Gamma'_k(G_n)$ contain all the redundant generating k -tuples.

Proof. The orbit in Gilman's Theorem 2.4.3 contains redundant elements by construction (see [Gi]). This implies that $\Gamma_k(G_n^m)$ contains a connected component Γ' , where each of the m generating k -tuples of G is redundant. This component is

“large” since as $|G_n| \rightarrow \infty$, the probability that the k -tuple is redundant goes to 1. But, then the probability that m independently chosen k -tuples lie in different orbits of $\text{Aut}(G_n)$ and are all redundant goes to 1 as well. This proves the claim that $|\Gamma'|/|\Gamma| \rightarrow 1$. The second part is clear by construction. \square

Remark 2.6.5 The results in this section are in [P5]. The concept of a “large” connected component is not new as it is used heavily in various versions in the random graphs community (see e.g. [ASE]).

Let us note that from the proof of Theorem 2.6.4 one cannot conclude that $\Gamma_k(A_n^{n!/8})$ has a large connected component. Indeed, by 1.1.2 (i), the probability that all of the $n!/8$ k -tuples are redundant is about $(1 - 1/n^{k-1})^{n!/8} \rightarrow 0$ as $n \rightarrow \infty$. Thus for all $k = o(n)$ the fraction of the k -tuples considered in the proof of Theorem 2.6.4, is “small”. On the other hand, for $k = \Omega(n)$ we already know that $\Gamma_k(A_n^{n!/8})$ has a large connected component by Theorem 2.6.3 (i).

2.7 Checking connectivity.

Given the importance of the Problem 2.1.1 and Conjecture 2.5.4, one might ask whether the claims are supported by the computational evidence. We describe here a simple algorithm developed in [CP] (based on the “large connected component concept” in section 2.6) for checking connectivity of product replacement graphs.

Let us remark first that beside the paper [CP], the only computational evidence that we know, was reported in [DG]: John Lafferty and Dan Rockmore showed in 1997 that $\Gamma_3(S_4)$ and $\Gamma_3(S_5)$ are connected. While they probably could have checked $\Gamma_3(S_6)$ by “naïve” methods, further progress is complicated by a super-exponential increase in the complexity: $|\mathcal{N}_3(S_n)| = \frac{7}{8}(n!)^3(1 - o(1))$. In fact, $|\mathcal{N}_3(S_5)| = 1,401,120$ while $\frac{7}{8}(5!)^3 = 1,512,000$ (see [DG]).

Problem 2.7.1 *Prove or disprove computationally that $\Gamma_k(G)$ is connected, when G is simple of reasonable size, and k is small (3 or 4).*

We will illustrate the algorithm of [CP] in the case $G = A_n$ and $k = 3$. We start with the following key observation: *it suffice to show that all the generating k -tuples lie in the “large” connected component.* As we already know that all redundant triples lie in the “large” component, so checking connectivity to redundant triples gives an efficient check for inclusion in the “large” connected component. Once formulated this way, one can also utilize the symmetry of the situation and obtain a great speed up as compared to the “brute force” method.

By \mathcal{R}_n denote the set of all redundant triples in $\Gamma_3(A_n)$. Recall that by Proposition 2.5.12. they lie in the same connected component. The algorithm checks connectivity to \mathcal{R}_n of all generating triples $(\sigma_1, \sigma_2, \sigma_3)$ of A_n such that $\langle \sigma_1, \sigma_2 \rangle \neq A_n$. To achieve this, we search over all tuples $(M, \sigma_1, \sigma_2, \sigma_3)$, where M is a maximal subgroup of A_n and $\sigma_1, \sigma_2 \in M$. The maximal subgroups M in A_n for small n are known and can be precomputed. Of course some generating triples will be considered several times, but their fraction is small while speed up we achieve by pruning our search based on symmetry is quite large.

Here is how the symmetry can be utilized. Fix an appropriate subgroup H of the group of automorphisms of G . In our case we take $G = A_n$ and $H = S_n$. If $(\sigma_1, \sigma_2, \sigma_3)$ is not connected to \mathcal{R}_n in $\Gamma_3(A_n)$, then $(\sigma_1^h, \sigma_2^h, \sigma_3^h)$ is also not connected to \mathcal{R}_n for every $h \in H$.

By \mathcal{O}_n denote the set of orbits of the natural action of H on $\mathcal{N}_3(A_n)$. For every orbit $O \in \mathcal{O}_n$, denote by $\eta(O)$ the lexicographically first element in O . In our situation, orbits of $H = S_n$ on A_n correspond to partitions of n into an even or odd number of parts (depending on the parity of n). Take $\sigma_1 \in A_n$ to be lexicographically first in the conjugacy class corresponding to such a partition. The centralizer $C_H(\sigma_1)$ fixes σ_1 , so we can now take σ_2 to be lexicographically first under action of $C_H(\sigma_1)$, etc.

Now, for every such element $\eta(O)$ we check whether it's connected to \mathcal{R}_n . To do that we run the product replacement random walk until we find a redundant triple. Note that of the three subgroups $\langle \sigma_1, \sigma_2 \rangle$, $\langle \sigma_1, \sigma_3 \rangle$ and $\langle \sigma_2, \sigma_3 \rangle$, after either of the Nielsen moves $R_{i,j}^\pm$, $L_{i,j}^\pm$, two subgroups will always remain. Thus at every step it suffices to check whether the third subgroup is G . We simply check whether both permutations in a pair lie in one of the maximal subgroups. Once the product replacement random walk hits a redundant triple, we proceed to the next orbit representative. The algorithm stops when all triples $\eta(O)$ are checked.

We discovered in [CP] that not only is $\Gamma_3(A_n)$ connected, when $n = 5, \dots, 10$, but also that it usually takes very few steps of the product replacement walk to reach a redundant triple. This proves part (iv) of Theorem 2.5.6, and leads to the following question.

Question 2.7.2 Is there a universal constant C such that every generating triple $(g) \in \mathcal{N}_3(A_n)$ is connected in the graph $\Gamma_3(A_n)$ to a redundant triple by at most C steps?

A negative answer to this question will shine a new light at the structure of generating sets of alternating groups.

3. MIXING TIME

3.1 Preliminaries.

Let Γ be an r -regular connected unoriented graph, and let $v \in \Gamma$ be a fixed vertex in Γ . Consider a *nearest neighbor random walk* $\mathcal{W}_v = \mathcal{W}_v(\Gamma)$ defined as follows: start at a vertex v ; at every step choose an adjacent edge with equal probability and move along that edge; repeat this ad infinitum. For technical reasons it is often useful to study *lazy random walk* \mathcal{W}_v° , which is defined similarly, but now the walk stays at the same vertex with probability $1/2$ and moves along every adjacent edge with probability $1/2r$.

Denote by \mathbf{Q}_v^t the probability distribution of the lazy random walk \mathcal{W}_v° after t steps. Since the graph is connected and regular, the stationary distribution is uniform:

$$\mathbf{Q}_v^t \rightarrow \frac{1}{|\Gamma|}, \text{ as } t \rightarrow \infty.$$

Recall the definition of the total variation distance of distribution \mathbf{P} on Γ :

$$\|\mathbf{P} - \mathbf{U}\|_{\text{tv}} = \max_{B \subset X} \left| \mathbf{P}(B) - \frac{|B|}{|X|} \right| = \frac{1}{2} \sum_{x \in X} \left| \mathbf{P}(x) - \frac{1}{|X|} \right|,$$

where $\mathbf{P}(B) = \sum_{x \in B} \mathbf{P}(x)$. To simplify the notation, denote $\mathbf{d}_v(t) = \|\mathbf{Q}_v^t - \mathbf{U}\|_{\text{tv}}$. It is easy to see that $0 \leq \mathbf{d}_v(t) \leq 1$, $\mathbf{d}_v(t) \rightarrow 0$ as $t \rightarrow \infty$, and that $\mathbf{d}_v(t+1) \leq \mathbf{d}_v(t)$ for every t (see [AF]). Also, if $\mathbf{d}(t) = \min_{v \in \Gamma} \mathbf{d}_v(t)$, then $\mathbf{d}(t+s) \leq 2\mathbf{d}(t) \cdot \mathbf{d}(s)$,

so once the total variation distance becomes small, it decreases exponentially (see [AF]).

Definition 3.1.1 Define *mixing time* mix_v of the lazy random walk $\mathcal{W}_v^\circ(\Gamma)$ as follows:

$$\text{mix}_v = \min \left\{ t \mid \mathbf{d}_v(t) \leq \frac{1}{4} \right\}$$

Let \mathcal{A} be the $|\Gamma| \times |\Gamma|$ adjacency matrix of graph Γ . The transition matrix of the random walk $\mathcal{W}(\Gamma)$ is given by $\mathcal{P} = \mathcal{A}/r$. Let $1 = \lambda_0 > \lambda_1 \geq \dots \geq \lambda_{|\Gamma|-1} \geq -1$ be the eigenvalues of \mathcal{P} (all eigenvalues are real since the matrix is symmetric, the first inequality comes from the Perron-Frobenius Theorem). Now, for the lazy random walk \mathcal{W}° the transition matrix is given by $\mathcal{P}' = (\text{Id} + \mathcal{P})/2$, and for the eigenvalues of \mathcal{P} we have $\lambda'_i = (1 + \lambda_i)/2$, $1 = \lambda'_0 > \lambda'_1 \geq \dots \geq \lambda'_{|\Gamma|-1} \geq 0$.

Proposition 3.1.2 Let $\mathcal{W}_v^\circ(\Gamma)$ be a lazy random walk on a regular graph Γ starting at vertex $v \in \Gamma$. Then for the mixing time we have:

$$\text{mix}_v \leq \frac{4}{1 - \lambda_1} \log |\Gamma|$$

Proof. We have

$$(\mathcal{P}')^n \cdot f = f_0 + (\lambda'_1)^n f_1 + (\lambda'_2)^n f_2 + \dots,$$

where f is a function on vertices of Γ and f_i are projections of f on the eigenvectors of \mathcal{P}' . Now take f to be a function concentrated at $v \in \Gamma$. Then $\mathbf{Q}_v^t = (\mathcal{P}')^t \cdot f$, $f_0 = \mathbf{U}$, and

$$\left| \mathbf{Q}_v^t(w) - \frac{1}{|\Gamma|} \right| \leq (\lambda'_1)^t \cdot |\Gamma|$$

for every $w \in \Gamma$. The proposition follows immediately. \square

Remark 3.1.3 There are a number of distances on Γ one can study. While total variation can be thought as rescaled ℓ_1 -norm, one can also define ℓ_p -norms for any $1 \leq p < \infty$, as well as ℓ_∞ -norm:

$$\|\mathbf{P} - \mathbf{U}\|_p = \left(\sum_{v \in \Gamma} \left(\mathbf{P}(v) - \frac{1}{|\Gamma|} \right)^p \right)^{1/p}, \quad \|\mathbf{P} - \mathbf{U}\|_\infty = \max_{v \in \Gamma} \left| \mathbf{P}(v) - \frac{1}{|\Gamma|} \right|$$

It also useful to define a *separation distance*, which can be thought as a one-sided ℓ_∞ -norm:

$$\text{sep}(\mathbf{P}) = |\Gamma| \cdot \max_{v \in \Gamma} \left(\frac{1}{|\Gamma|} - \mathbf{P}(v) \right)$$

Denote $\mathbf{s}_v(t) = \text{sep}(\mathbf{Q}_v^t)$. It is known that the separation distance satisfies similar monotonicity and submultiplicativity properties as the total variation distance (see [AF,Di]). Also, $\mathbf{d}_v(t) \leq \mathbf{s}_v(t)$ for any $t < \infty$.

Finally, let us note a straightforward generalization of Proposition 3.1.2 to all these norms, including ℓ_∞ -norm. We leave the details to the reader.

3.2 The product replacement random walk.

Let $\Gamma_k(G)$ be as before, the graph of generating k -tuples of the finite group G ; let (g) be a fixed starting k -tuple. By Γ' denote the connected component of $\Gamma_k(G)$ which contains (g) . By $\mathcal{PR} = \mathcal{PR}(G, k; (g))$ denote the lazy random walk on Γ' starting at (g) . By $\text{mix}_{(g)}$ denote the mixing time of $\mathcal{PR}(G, k; (g))$. For the rest of this section, let mix stand for the maximum $\text{mix}_{(g)}$ over all $(g) \in \mathcal{N}_k(G)$. We call mix the *mixing time* of the product replacement random walk.

As far as theoretical computer science is concerned, the following conjecture is central:

Conjecture 3.2.1 *For every $k \geq d(G)$, the mixing time mix is polynomial in k and $\log |G|$.*

In this section we will cover the known “general results” on the subject. Let us note, however, that conjecture is wide open even for such special cases as $k = 3$ and $G = S_n$. But first a stronger version of the conjecture:

Question 3.2.2 *Is it true that for every fixed $k \geq d(G)$ the relaxation time $\tau = 1/(1 - \lambda_1(\Gamma'))$ is bounded by a polynomial of k ?*

Note that by Proposition 3.1.2 this would imply that the mixing time is *linear* in $\log |G|$.

Theorem 3.2.3 ([CG1,DS3]) *Let G be fixed and let $k \rightarrow \infty$. Then the mixing time mix of the lazy product replacement random walk \mathcal{PR} satisfies:*

$$\text{mix} = O(k^2 \log k).$$

In other words, for k sufficiently large, the product replacement random walk mixes rapidly (compared to the size of the graph $|\Gamma| \sim |G|^k$). In fact, even more rapid mixing has been conjectured (see [DS2,DS3]):

Conjecture 3.2.4 *Let G be fixed and let $k \rightarrow \infty$. Then*

$$\text{mix} = O(k \log k)$$

The reasoning behind this conjecture is the following: for large k there seems little interaction between components of the k -tuple. As $|G|$ is bounded, we need just a bounded number of times to “hit” each component to make it “random” in G . Thus the problem becomes similar to the coupon collector’s problem with k coupons, where each we need to collect a bounded number of times (cf. [Di,F]) This is where $O(k \log k)$ come from.

For $G = \mathbb{Z}_2$ an ad hoc method was used by Chung and Graham [CG2] to prove Conjecture 3.2.4 in this case.

Remark 3.2.5 Recently the author announced [P6] a positive solution of Conjecture 3.2.1 in case when G is nilpotent. This itself is much stronger than the partial solution when k and the nilpotency class i are bounded. It nevertheless gives quite a large degree polynomial. The solution is yet to be checked and fully written.

3.3 The Diaconis and Saloff-Coste bound.

In an important paper [DS3] the authors found the first general bounds for the mixing time of the product replacement random walk. Since the authors do not use any special properties of the walk as compared to other random walks, the results are quite weak from the theory of computation point of view (cf. Conjecture 3.1.1.) On the other hand, a refined analytic technique used by the authors seem to be so advanced that we expect the result to be difficult (if at all possible) to improve by means of the better estimates.

Rather than give here the full power of the result by Diaconis and Saloff-Coste (which may be difficult to quickly digest), we choose to sketch the approach and state slightly weaker version of the main result.

The idea is that given any connected graph Γ one can bound the relaxation time $\tau = 1/(1 - \lambda_1)$ via geometric properties of the graph (we refer to [AF,DSt] for references and details). Formally, for every two vertices $v, v' \in \Gamma$ fix a path $\gamma(v, v')$. By Δ denote the maximum length of such paths, and by N denote the maximum congestion through edge e in Γ (the maximum of the number of paths which go through e , divided by $|\Gamma|$). Roughly, the relaxation time τ is bounded by a polynomial of Δ, N . The problem is that if the paths are chosen in a “stupid” way, the congestion can be as large as size of the graph.

In [DS3] the authors consider an interesting set of paths (see p. 268) which split into “big steps”: change any coordinate by any element (as long as k -tuple still generates G). Now, these “big steps” can be easily used to connect any $(h) \in \Gamma_k(G)$ to any $(g) \in \Gamma_k(G)$ (given that k is large enough, see [DS3]). Now big steps can be decomposed into “small” steps of the walk on $\Gamma_k(G)$ by use of the geometry of generating sets, not unlike the one used in Proposition 2.2.2.

Now, once the path are constructed, “the fight” is over the final shape of the estimates. The authors involve an advanced ℓ_2 -technique and largely win this fight. Still, they “lose the battle” since their paths are just too long and this problem is inherent in the construction. One would suggest using short paths of Babai (see the proof of Theorem 2.2.3), but these paths are of a special type and hard to analyze for congestion. Thus in the final estimate Diaconis and Saloff-Coste have the following parameter:

$\Delta(G) = \text{maximal diameter of the Cayley graph } \text{Cayley}(G, S) \text{ over all generating sets } S \subset G$

This $\Delta(G)$ is the main ingredient of the length of the small steps of nice paths of Diaconis and Saloff-Coste. While $\Delta(G)$ is conjectured to be nice for various families of groups (such as simple groups) it is provably large for a number of other families (such as abelian and nilpotent groups). This is the key problem with the argument which prevents the final bound on the mixing time from being polynomial (even in the cases when it has been confirmed by other methods)

Now we are ready to formulate a “toy version” of the main result in [DS3]:

Theorem 3.3.1 (Diaconis, Saloff-Coste) *Let G be a finite group, $d = d(G)$, $\bar{d} = \bar{d}(G)$, $k \geq \bar{d} + 2d$, $\varphi_k(G) = |\Gamma_k(G)|/|G^k|$, and let $\Delta(G)$ be as above. Let mix be the mixing time of the lazy product replacement random walk on $\Gamma_k(G)$ (the*

maximum is taken over all starting points $(g) \in \Gamma_k(G)$). Then

$$\text{mix} \leq O\left(A k^2 \log |G| (\log k + \log \log |G|) \frac{\Delta^2(G)}{\varphi_k^2(G)}\right),$$

where A is a parameter defined as follows:

$$A = \frac{d^2 \cdot k(k-1) \cdot \dots \cdot (k - \bar{d} + 1)}{(k-2d)(k-2d-1) \cdot \dots \cdot (k-2d-\bar{d}+1)}.$$

The actual result is much more refined, to the extent of being difficult to comprehend. Among others, the constant implied by $O(\cdot)$ notation is calculated (it is about 80) and several little improvements are made. We present here just “the first order” bound to give the reader an easy access to the result.

A few words about the quality of this result for general groups: when $k = \theta(d \cdot \bar{d})$ we have $A = O(d^2)$. Now using bounds $d, \bar{d} = O(\log |G|)$, $\varphi_k(G) \sim 1$, we obtain a bound on the mixing time which is polynomial in k , $\log |G|$ and $\Delta(G)$.

Perhaps comparison with the mixing of random walks on groups is in order. Let G be a finite group, S be a symmetric ($S = S^{-1}$) generating set, $|S| = k$, and let $\Delta(G, S)$ be the diameter of the corresponding Cayley graph $\Gamma = \text{Cayley}(G, S)$. Then for the mixing time mix of the lazy nearest neighbor random walk on Γ (starting at $\text{id} \in G$) we have:

$$\text{mix} = O(\Delta^2(G, S) \cdot k)$$

This result was proved in [DS1] and is an improvement over a more traditional bound $\text{mix} = O(\Delta^2(G, S) \cdot k \cdot \log |G|)$. Note now the similarity with the main result in Theorem 3.2.1. Observe that the bound obtained there is worse than the bound on the mixing time of *any random walk on the same group with the same number k of generators*. On the other hand, the main reason the product replacement random walk is used, is because it mixes *faster* than the ordinary random walk. We conclude that the bound in Theorem 3.3.1 by no means explains the rapid mixing phenomenon of the product replacement random walk.

Example 3.3.2 Let G be abelian. Then $\Delta(G)$ can be as large as $|G|^2$ (for $G \simeq \mathbb{Z}_n$). Thus the mixing time bound in Theorem 3.3.1 becomes exponential in this case (in $\log |G|$) while as we show below (see Theorem 3.5.9) the mixing time is in fact polynomial.

Example 3.3.3 Let $G = S_n$. It is conjectured that $\Delta(S_n) = O(n^2)$, while the best known bound is $\Delta(S_n) = \exp(O(\sqrt{n}))$ (see [Bb4]). Now observe that the conjecture combined with Theorem 3.3.1 gives a polynomial bound on the mixing time in this case. Analogously, for every simple group G it is conjectured that $\Delta(G) = (\log |G|)^{O(1)}$. See [Bb4,DS3] for references and details.

3.4 $\Gamma_k(G)$ as Schreier graphs.

Denote by F_k the free group on k generators x_1, \dots, x_k and let $\text{Aut}(F_k)$ be a group of automorphisms of F_k . Consider the set Υ_k of the following automorphisms:

$$\begin{aligned} R_{i,j}^{\pm}(x_i) &= x_i x_j^{\pm 1}, \quad \text{and} \quad R_{i,j}^{\pm}(x_l) = x_l \quad \text{if } l \neq i \\ L_{i,j}^{\pm}(x_i) &= x_j^{\pm 1} x_i, \quad \text{and} \quad R_{i,j}^{\pm}(x_l) = x_l \quad \text{if } l \neq i \end{aligned}$$

These are exactly *Nielsen moves* when $G = F_k$. A classical result of Nielsen [Ni] (see [LyS,MKS]) shows that $\text{Aut}(F_k)$ is generated by the Nielsen moves and elementary automorphisms of permutation and inversion of generators.

Proposition 3.4.1 *Let $A^+ = A^+(F_k)$ be the subgroup of $\text{Aut}(F_k)$ generated by Nielsen moves. Then $A^+(F_k)$ is a normal subgroup of index two in $\text{Aut}(F_k)$.*

The group $A^+(F_k)$ is called *special group of automorphisms* of a free group F_k .

Proof. The Proposition seems to be well known (cf. [Ge]). An easy proof that A^+ has index at most two follows along the lines of the proof of Proposition 2.2.1. As to why it's exactly two, consider the natural projection $\pi : \text{Aut}(F_k) \rightarrow \text{Aut}(F_k/[F_k, F_k]) \simeq GL_k(\mathbb{Z})$. Observe that $\pi(R_{i,j}^{\pm}) = \pi(L_{i,j}^{\pm}) \in SL_k(\mathbb{Z})$ and gives elementary transvections $E_{i,j}^{\pm}$, for all $i \neq j$. Therefore $\pi(A^+) = SL_k(\mathbb{Z})$ and A^+ has index at least two in $\text{Aut}(F_k)$. \square

Proposition 3.4.2 *The graph $\Gamma_k(F_k)$ has two connected components, each of them is isomorphic to the right Cayley graph of the special automorphism group $A^+(F_k)$ with respect to the Nielsen moves.*

Proof. Indeed, $\text{Aut}(F_k)$ acts simply transitively on the vertices of $\Gamma_k(F_k)$. This gives a one to one correspondence ϱ between $\text{Aut}(F_k)$ and $\Gamma_k(F_k)$ defined as

$$\varrho : \alpha \in \text{Aut}(F_k) \rightarrow (\alpha(x_1), \dots, \alpha(x_k)).$$

It is easy to see that $\alpha R_{i,j}^{\pm}$ and $\alpha L_{i,j}^{\pm}$ correspond to the neighbors of $\varrho(\alpha) \in \Gamma_k(G)$. This shows that $\Gamma_k(F_k)$ is the Cayley graph of $\text{Aut}(F_k)$ with respect to the Nielsen moves. As they do not generate the group but rather a subgroup of index two, the graph $\Gamma_k(F_k)$ has two connected components, each one isomorphic to Cayley($A^+(F_k); \{R_{i,j}^{\pm}, L_{i,j}^{\pm}\}$). \square

Let H and G be two (possibly infinite) groups and ψ be an epimorphism from H onto G . Then ψ induces a map $\psi_k : \Gamma_k(H) \rightarrow \Gamma_k(G)$, where

$$\psi_k : (h_1, \dots, h_k) \rightarrow (\psi(h_1), \dots, \psi(h_k)),$$

provided $k \geq d(H)$. It is easy to see that ψ_k is a projection of graphs which preserves adjacency relations. For $H = F_k$ with a little work one can deduce:

Proposition 3.4.3 *Let $X = \text{Cayley}(A^+(F_k); \Upsilon_k)$ be the Cayley graph of $A^+(F_k)$ with respect to the Nielsen generators. Then for every G and every connected component Y of $\Gamma_k(G)$, there exists a surjective graph projection $\psi : X \rightarrow Y$.*

Proof. Identify $\Gamma_k(G)$ with the set $\text{Epi}(F_k, G)$ of all epimorphisms $\psi : F_k \twoheadrightarrow G$, where an epimorphism ψ is identified with $(\psi(x_1), \dots, \psi(x_k))$. Then $A = \text{Aut}(F_k)$ acts on $E = \text{Epi}(F_k, G)$ by: $\alpha(\psi) = \psi \circ \alpha^{-1}$ for $\alpha \in \text{Aut}(F_k)$ and $\psi \in E$. Fix $\psi : F_k \rightarrow G$ and let B be the subgroup of $A^+ = A^+(F_k)$ defined as $B = \{\alpha \in A^+ \mid \alpha(\psi) = \psi\}$. Then A^+/B is naturally identified with connected component of $(g) = (\psi(x_1), \dots, \psi(x_k))$ in $\Gamma_k(G)$. Indeed, if $\alpha_1, \alpha_2 \in A^+$, then $\alpha_1 B = \alpha_2 B$ if and only if $\alpha_1(\psi) = \alpha_2(\psi)$. Moreover, $[R_{i,j}^\pm \circ \alpha](\psi)$ and $[L_{i,j}^\pm \circ \alpha](\psi)$ are exactly the neighbors of $\alpha(\psi)$. \square

The proposition implies that every connected component of $\Gamma_k(G)$ is a Schreier graph of $A^+(F_k)$ with respect to the Nielsen moves Υ_k , modulo finite index subgroup. This is a very strong property which will be used in the next section.

Note that we do not claim that $\psi_k : X \rightarrow \Gamma_k(G)$ is surjective. In fact this is not true as graphs $\Gamma_k(G)$ can have any number of connected components (see section 2.3). On the other hand, for finite groups the map this is indeed surjective, as follows from Theorem 2.1.4.

Now let us generalize our results from free groups to “relatively free groups”. Let W be a characteristic subgroup of F_k , i.e. $\alpha(W) = W$ for every $\alpha \in \text{Aut}(F_k)$. There is a natural homomorphism $\pi : \text{Aut}(F_k) \rightarrow \text{Aut}(F_k/W)$. Denote $\pi(A^+(F_k))$ by $A^+(F_k/W)$, and call it the special automorphism group of F_k/W . Note that in general π is not necessarily an epimorphism, so $A^+(F_k/W)$ can be of large (even infinite) index in $\text{Aut}(F_k/W)$. Still, the Nielsen moves generate $A^+(F_k/W)$ and we have the following general result:

Proposition 3.4.4 *Let W be a characteristic subgroup of F_k and G a finite quotient of F_k/W . Then every component of $\Gamma_k(G)$ is a Schreier graph of $A^+(F_k/W)$ with respect to the Nielsen moves and modulo some finite index subgroup of $A^+(F_k/W)$. \square*

Example 3.4.5 Let $W = [F_k, F_k]$ be the commutator subgroup of F_k . Then $\text{Aut}(F_k/W) = GL_k(\mathbb{Z})$ and $A^+(F_k/W) = SL_k(\mathbb{Z})$. The Nielsen moves $R_{i,j}^\pm$ and $L_{i,j}^\pm$ correspond to the elementary matrices $E_{i,j}^\pm$ with 1’s along diagonal, ± 1 at the (i, j) entry, and 0 elsewhere. We therefore conclude:

Proposition 3.4.6 *Let G be a finite abelian group. Then any connected component of $\Gamma_k(G)$ is a Schreier graph of $SL_k(\mathbb{Z})$ with respect to the elementary matrices $E_{i,j}^\pm$ modulo a finite index subgroup of $SL_k(\mathbb{Z})$.⁶ \square*

Remark 3.4.7 Let G be as in Theorem 2.3.2. It is not difficult to see that a finite index subgroup of $SL_k(\mathbb{Z})$ in Proposition 3.3.6 is a congruence subgroup containing $\text{Ker}(SL_k(\mathbb{Z}) \rightarrow SL_k(\mathbb{Z}/m_1\mathbb{Z}))$, where m_1 is as in Theorem 2.3.2. This explains part (ii) of the Theorem.

Example 3.4.8 Define the lower central series of F_k by $\gamma_1(F_k) = F_k$ and $\gamma_{i+1}(F_k) = [F_k, \gamma_i(F_k)]$. Let $W = \gamma_{i+1}(F_k)$, so F_k/W is the “free nilpotent group of class i ”. Let $A^+(F_k/W)$ and $\text{Aut}(F_k/W)$ be as above. It is known that when $k \geq 2$

⁶Formally, one should include each elementary matrix $E_{i,j}^\pm$ twice, which is due to the fact that $L_{i,j}^\pm = R_{i,j}^\pm$ for abelian groups.

and $i \geq 4$, $A^+(F_k/W)$ is of infinite index in $\text{Aut}(F_k/W)$ (see [An,Bh]). Nevertheless, we still have that for every nilpotent group G of class i , every connected component of $\Gamma_k(G)$ is a quotient of the Cayley graph of $A^+(F_k/W)$ with respect to the Nielsen moves.

3.5 Kazhdan's property (T) and mixing.

The main result in this section is a connection between Kazhdan's property (T) and the mixing time of the product replacement random walk. The exposition follows the paper [LP] of Lubotzky and the author.

Definition 3.5.1 *A topological group G is said to have (Kazhdan) property (T) if there exists a compact subset Q of G such that $\mathcal{K} = \mathcal{K}(G, Q) > 0$, where*

$$(*) \quad \mathcal{K}(G, Q) = \inf_{\rho} \inf_v \max_{q \in Q} \frac{\|\rho(q)v - v\|}{\|v\|},$$

where ρ runs over all unitary representations (\mathcal{H}, ρ) of G which do not contain the trivial representation (i.e., no non-zero G -fixed vector), and v runs over all vectors $v \neq 0$ in \mathcal{H} .

We say that (Q, ε) is a *Kazhdan constant* for G if $\varepsilon \leq \mathcal{K}(G, Q)$ (see [Lu1]).

Let now Γ be a discrete group. It is well known and not difficult to prove that if Γ has (T), then Γ is finitely generated, and if (Q, ε) is a Kazhdan constant for Γ , then Q generate Γ (see [Kz, Lu1]).

Open Problem 3.5.2 *Does $\text{Aut}(F_k)$ (or equivalently $A^+(F_k)$) has property (T) with respect to Nielsen generators, given $k > 3$?*

Remark 3.5.3 For $k = 2, 3$ the answer is negative. This follows from [Lu1, Mc]. Let note also that if $\text{Aut}(F_k)$ has (T), then in view of Gilman's Theorem 2.4.3 and his part (i) of Theorem 2.5.6, one can obtain constructions of expanders from Cayley graphs of the alternating groups (cf. [Mr, Lu1]). This would solve positively open problem 10.3.4 and negatively Open Problem 10.3.2 in [Lu1].

Proposition 3.5.4 ([HRV]) *Let Γ be a discrete group generated by a finite set S . Assume Γ has property (T) with Kazhdan constant (S, \mathcal{K}) . Then for every finite index subgroup N of Γ , the Schreier graph X on Γ/N with respect to S satisfies $\beta \geq \mathcal{K}^2/2|S|$, where $\beta(X) = 1 - \lambda_1(X)$ is an eigenvalue gap of X .*

The proof can be found in [HRV], Corollary to Proposition III, p. 89.

Proposition 3.5.5 ([PZ]) *In condition of Proposition 3.5.4, assume that there is a finite group $H < \text{Aut}(\Gamma)$ such that $H(S) = S$ and action of H on S has m equal size orbits. Then $\beta \geq \mathcal{K}^2/(2m)$.*

The proof will appear in a forthcoming paper [PZ]. Let us note that one can always take $H = \{1\}$. Then $m = |S|$ and the bound in Proposition 3.5.5 coincides with a bound in Proposition 3.5.4.

As a combination of Proposition 3.5.5 and 3.4.3 we obtain:

Theorem 3.5.6 ([LP]) *If $\text{Aut}(F_k)$ (or equivalently $A^+(F_k)$) has Kazhdan's property (T), then for every finite group G generated by k elements, the mixing time $\text{mix}_{(g)}$ of the lazy product replacement random walk on a connected component $\Gamma' \subset \Gamma_k(G)$, $(g) \in \Gamma'$, is bounded as*

$$\text{mix}_{(g)} = C(k) \log |G|,$$

where $C(k)$ depends only on k .

Analogously, for “relatively free groups” we have:

Theorem 3.5.7 ([LP]) *Let W be a characteristic subgroup of F_k . If the special automorphism group $A^+(F_k/W)$ has (T), then the conclusion of Theorem 3.4.5 is satisfied for every finite quotient G of F_k/W .*

Example 3.5.8 Let $W = [F_k, F_k]$ be as in Example 3.4.5. Then $A^+(F_k/W) \simeq \text{SL}(k, \mathbb{Z})$, which indeed has property (T) (see [Lu1]). Further, recently Shalom [Sh2] gave bounds on the Kazhdan constant exactly for the elementary matrices as a generating set. Combining all these results we get:

Theorem 3.5.9 ([LP]) *Let G be an abelian group, $(g) = (g_1, \dots, g_k)$ be the initial generating k -tuple, and let $\Gamma' \subset \Gamma_k(G)$ be a connected component containing (g) . Then for the mixing time $\text{mix}_{(g)}$ of the lazy product replacement random walk \mathcal{PR} starting at (g) we have*

$$\text{mix}_{(g)} \leq C \cdot k^5 \cdot \log |G|,$$

where C is a universal constant.

Remark 3.5.10 Note that we write $\text{mix}_{(g)}$ as we prove rapid mixing for random walks on every connected component. If the connected component Γ' is small compared to Γ , one can use $\text{mix}_{(g)} \leq C \cdot k^4 \log |\Gamma'|$ bound instead.

A very special case of groups \mathbb{Z}_p^m , p -prime, was considered earlier in [DS2]. The authors obtained a bound on the mixing time by using just diameter bound for $\text{SL}(2, p)$, supplementing it with random walk techniques. The resulting bound for $G \simeq \mathbb{Z}_p$ is $\text{mix} = O(k^4 \log^3 p)$. As in situations of interest $k = O(\log p)$, this bound is weaker than that in Theorem 3.5.9. Still, this suggests possible improvement of the power k^5 in the latter bound.

Example 3.5.11 Let $W = \gamma_{i+1}(F_k)$ be as in Example 3.4.8. Then, as showed in [LP], the quotient F_k/W has property (T). The problem is that we have yet to estimate the Kazhdan constant in this case. Still, this implies the following weak generalization of Theorem 3.5.9:

Theorem 3.5.12 ([LP]) *Let k, i be fixed, and let G be a nilpotent group of class at most i , $d(G) \leq k$. Let $(g) = (g_1, \dots, g_k)$ be the initial generating k -tuple, and let $\Gamma' \subset \Gamma_k(G)$ be a connected component containing (g) . Then for the mixing time $\text{mix}_{(g)}$ of the lazy product replacement random walk \mathcal{PR} starting at (g) we have*

$$\text{mix}_{(g)} \leq C(k, i) \cdot \log |G|,$$

where $C(k, i)$ is a constant which depends on k, i , and is independent of the group G .

Conjecture 3.5.13 *The constant $C(k, i)$ is bounded by a universal constant $C(k)$, independent of i and polynomial in k .*

Remark 3.5.14 The proof of the Theorem 3.5.12 is based on first proving that the corresponding Lie group has property (T), which is based on the Mautner phenomenon (see [Wa]). Unfortunately it is often difficult to obtain good bounds for Kazhdan constants of lattices (cf. [Bu,Sh1,Sh2]). Still, we have a high confidence that *some*, perhaps an exponential bound, can be obtained for $C(k, i)$. The idea is based on applying the technique of Shalom [Sh1] combined with the bounds in the work [KB] on the Smith normal form. We challenge the reader to obtain any bounds on $C(k, i)$ in this direction.

4. PUTTING IT ALL TOGETHER

4.1 What one would like to have.

Let G be a black box group with the cost μ of multiplication and inversion (see [Bb5,KS,P2] for definitions). For convenience, one can simply assume that G is given as a subgroup of a large permutation or matrix group (over the finite field). We also assume that an upper bound N is given for $\log |G|$ (say, the bit length of the encoding of groups elements).

One would like to design a routine which when given generators g_1, \dots, g_k of G , outputs independent elements of G from a certain distribution \mathcal{P} on G such that $\|\mathcal{P} - \mathbf{U}\|_{\text{tv}} \leq 1/4$. The cost ρ of such a routine should be polynomial in both N and k :

$$\rho = \text{Poly}(N, k) \cdot \mu$$

Note that successive iterations of the routine give a number of independent “random” elements of G . Taking their product gives better approximations of the uniform distribution \mathbf{U} , as good as one likes (in any of the distances mentioned in 3.1.3) This is what we wanted at the first place. We call elements sampled from the distribution \mathcal{P} as above the *(nearly) uniform groups elements*. Also, from now on we will replace N by $O(\log |G|)$ hoping this does not lead to confusion.

Now, the product replacement algorithm is yet to be proved a legitimate candidate for such a routine. In the meantime one should ask whether there is *any* algorithm which satisfies the above requirements. The answer is yes: there is one provably correct known example of such an algorithm; it was discovered by Babai in his pioneer paper [Bb3]. We will sketch the algorithm below for the reader’s convenience.

Algorithm 4.1.1 (Babai)

0) Use “random subproducts” routine to reduce the number of generators in a generating set S from k to $O(\log |G|)$.

1) For $C_1 \cdot \log |G|$ rounds repeat the following procedure. Run a simple random walk on the Cayley graph $\text{Cayley}(G, S)$ starting at $\text{id} \in G$, for $C_2 \log^4 |G|$ steps. Then add the endpoint of the walk to S and repeat.

2) Use the “Erdős-Rényi machine” to generate the output.

Theorem 4.1.2 ([Bb3]) *The above algorithm produces m independent (nearly) uniform elements of G at a cost $O(k \log |G| + \log^5 |G| + m \log |G|) \cdot \mu$.*

We will always assume that $k = O(\log |G|)$. Then the preprocessing part 0) can be skipped. The “Erdős-Rényi machine” will be touched upon later in this section, but it is also unnecessary for our current purposes: one can simply replace it by additional rounds of a simple random walk of length $O(\log^2 |G|)$ with the stopping points output as “random elements”⁷. Thus we will concentrate on the main part 1) of the algorithm.

Several remarks before we analyze part 1). First, constants C_1, C_2 are universal and have been explicitly computed in [Bb3]. Second, for technical reason in the original paper the stopping time of the simple walks was chosen from a certain distribution. Recent advancements in the random walk technique proved this to be unnecessary so we use here a simplified version so as not to alarm the reader. Finally, for a reason of saving a small constant factor the author made the lengths of the walk to be increasing. We ignore this difference.

Now, the main result of Babai is based on the following lemma:

Lemma 4.1.3 *After round $C_1 \cdot i$, for a generating set $S = \{g_1, \dots, g_r\}$, $r = k + C_1 \cdot i$, the number of distinct elements in $C(S) = \{g_1^{\varepsilon_1} \cdot \dots \cdot g_r^{\varepsilon_r}\}$ is w.h.p.⁸ at least 2^i , where $i \leq \log_2 |G| - 1$.*

Basically, the length $C_2 \log^4 |G|$ is chosen so that the size of the “cube” $C(S)$ is doubled with positive probability after a constant number of steps. At the end, when the diameter of the Cayley graph $\Delta(G, S)$ becomes of the order $O(\log |G|)$ while $|S| = O(\log |G|)$, the simple random walks on G mix rapidly (after a careful analysis, one can show that $\text{mix} = O(\log^2 |G|)$).

When compared to the product replacement random walk, few similarities are clear. Both walks use constantly changing generating sets to “reach” to more distant elements in the original Cayley graph. Both walks use the result of the previous walk as its “generating set” to make these steps.

The main difference is that the size of the generating set in the algorithm of Babai is increasing, absorbing all the new generators. As a result, in Babai’s algorithm, after $\log |G|$ rounds the size of the generating set becomes $\geq \log |G|$. Since each short walk must be of length at least $\Omega(\log |G|)$, the lower bound of $\Omega(\log^2 |G| \cdot \mu)$ is in order⁹.

On the other hand, in the product replacement random walk, the size k of the generating set remains fixed. Thus, theoretically the cost of the algorithm can be as small as $O(k^c \cdot \log |G| \cdot \mu)$, i.e. linear in $\log |G|$. In fact this would be the case if Open Problem 3.5.2 had a positive solution. While formally k can be as large as $\log_2 |G|$ (say, for $G \simeq \mathbb{Z}_2^n$), and the constant c can be larger than 5, in the practically

⁷Of course, the use of the “Erdős-Rényi machine” has an important theoretical and perhaps practical advantage (see [Bb2, Bb4]). We skip it here for simplicity.

⁸w.h.p. = with high probability. Let us be somewhat vague here as one can always refer to [Bb3] for technical details.

⁹An improved lower bound $\Omega(\log^3 |G| \cdot \mu)$ for any version of Babai’s algorithm seems also reasonable.

important case when k is small this gives a desired speed up and seems to explain the rapid mixing phenomenon.

We realize that at present the above “explanation” is nothing but a speculation. However a fine distinction between the algorithms seem important to preserve and use small k in the product replacement algorithm, rather than take $k = \theta(\log |G|)$ as done in Theorem 2.2.3 and Theorem 3.3.1.

4.2 Running the Product Replacement Algorithm.

As we describe earlier, to run the product replacement algorithm, even before analyzing the mixing time, one needs to remove two main obstacles: connectivity and bias. Rather than repeat much said in sections 1, 2 about these two, let us summarize what could be done in a general case.

As noted in sections 1.5 and 2.6, both obstacles can be removed by taking k large enough so that $\varphi_k(G) \rightarrow 1$. Then bias is avoided automatically and when $d(G)$ is added to k , the graphs contain a large component containing the initial generating set (see Theorem 2.6.2.) Together the result is quantified in Proposition 1.5.1, while Corollary 1.5.2 summarizes what’s known in special cases. We believe that a good rule of thumb is given by Conjecture 1.5.3: take $k = C(G) \log \log |G|$. This probably removes both obstacles, and the factor $O(\log \log |G|)$ is still a small price to pay for that (as compared to $\theta(\log |G|)$).

How long should one run the product replacement random walk? Well, for the mixing time mix steps, after which one obtains a (nearly) uniform generating k -tuple (with all components equally useful for random generation). Of course, we do not have good theoretical bounds on mix barring few special cases (see section 3.3, 3.5.) Determining mix in practice is not an easy task. Thus from the theoretical point of view the algorithm should not be use other than in Las Vegas applications, which is to say that any Monte Carlo estimates are largely unreliable.

4.3 Modifying Product Replacement Algorithm.

It was proposed recently [LG] by Leedham-Green et al. that while running the product replacement random walk one should consider a simple random walk on a group with a generating k -tuple as a generating set. Thus the generating set of this new random walk on will change with time. The stopping state of the walk is assigned to be the output. This modification¹⁰ is proposed to remove the bias obstacle ([LG]).

We see several problems with such an approach. First, while it formally removes the bias obstacle, it doesn’t remove the connectivity obstacle and tells nothing about the mixing. Second, it is seemingly worse than running a simple random walk on the generating k -tuple obtained at the end of the product replacement random walk. Indeed, heuristically one “wastes time” by running a simple random walk with the initial “bad” generators, rather than wait till the end of the product replacement walk and use these “good” generators. The reason being that initial generating set may correspond to the “worst case” of a random walk, while the k -tuple obtained after running product replacement gives an “average case” of a random walk. Note that random walks on random generating sets is a known concept in the theory, and are called sometimes “random random walks” (see [P3] for the references).

¹⁰The authors gave colorful British names to each version of the algorithm. We decided to avoid this terminology for the sake of clarity.

Another cheap method to avoid bias is to use the “Erdős-Rényi machine” of Babai mentioned above. Simply take the resulting generating set of k elements and output random elements of the “cube” (these elements are also known as “random subproducts”). This method provably works for $k \geq 2 \log_2 |G|$, but has yet to be analyzed for smaller values of k (see [ER,Bb3]). In fact, this approach is again intimately related to taking simple random walks on G with this generating set: we refer to [P3] for making this claim formal.

If the author was to modify the algorithm, it would be the following construction. Simply run the product replacement random walk, but from time to time multiply random elements of the k -tuple by one of the randomly chosen *original* generators. Then one preserves the structure of the walk as a random walk on a Schreier graph of $\text{Aut}(F_k)$, while ensuring that the stationary distribution is uniform. Basically the walk will be moving from one connected component of $\Gamma_k(G)$ to another when we multiply by the original generators. As a bonus, in the limit *all* the components in an obtained k -tuple are uniform and independent in G , which is an improvement over a single such element in the previous versions of the algorithm.

Finally, let us conclude by saying that the product replacement algorithm is still largely mysterious and is perhaps a detour from the “right” solution. The problem of generating random group elements remains a challenge, and we hope that some radically new provably correct algorithms will appear in the future.

Acknowledgements

I would like to start by expressing my extreme gratitude to Persi Diaconis, Bill Kantor and Alex Lubotzky who supported this project and kindly encouraged the writing of this paper. This review article is a compilation of several papers, including joint work with László Babai, Sergey Bratus, Gene Cooperman, Alex Lubotzky, and Andrzej Żuk. I am grateful for the wonderful opportunity to work with them.

I would like to thank Eamonn O’Brien, Derek Holt, Charles Leedham-Green, Scott Murray, Alice Niemeyer, and Leonard Soicher for the valuable conversations on the practical importance of the algorithm. Remarks of Alexander Borovik, Ken Brown, Keith Dennis, Martin Dunwoody, Robert Guralnick, László Lovász, Avinoam Mann, Gregory Margulis, László Pyber, Laurent Saloff-Coste, Jan Saxl, Yehuda Shalom, and James Wiegold were also very helpful. Special thanks to Bill Kantor for reading the preliminary version of the paper.

The author was supported by the NSF Postdoctoral Research Fellowship in Mathematical Sciences.

REFERENCES

- [Ac] V. Acciario, *The probability of generating some common families of finite groups*, Util. Math. **49** (1996), 243–254.
- [AKS] M. Ajtai, J. Komlós, E. Szemerédi, *Sorting in $c \log n$ parallel steps*, Combinatorica **3** (1983), 1–19.
- [AF] D. Aldous, J. Fill, *Reversible Markov Chains and Random Walks on Graphs*, monograph in preparation, 1996.
- [ASE] N. Alon, J.H. Spencer, P. Erdős, *The Probabilistic Method*, Wiley, New York, 1992.
- [An] S. Andreadakis, *On the automorphisms of free groups and free nilpotent groups*, Proc. London Math. Soc. (3) **15** (1965), 239–268.
- [As] M. Aschbacher, *On the maximal subgroups of the finite classical groups*, Invent. Math. **76** (1984), 469–514.

- [Bb1] L. Babai, *On the length of subgroup chains in the symmetric group*, Comm. Algebra **14** (1986), 1729–1736.
- [Bb2] L. Babai, *The probability of generating the symmetric group*, J. Comb. Th. Ser. A **52** (1989), 148–153.
- [Bb3] L. Babai, *Local expansion of vertex-transitive graphs and random generation in finite groups*, in Proc 23rd ACM STOC (1991), 164–174.
- [Bb4] L. Babai, *Automorphism groups, isomorphism, reconstruction*, in Handbook of Combinatorics (R. L. Graham, M. Groetschel, and L. Lovasz, eds.) (1996), Elsevier.
- [Bb5] L. Babai, *Randomization in group algorithms: Conceptual questions*, in Groups and Computation II (L. Finkelstein, W.M. Kantor, eds.) DIMACS Workshops on Groups and Computation, AMS, Providence, 1997.
- [BbP] L. Babai, I. Pak, *Strong bias of group generators: an obstacle to the “product replacement algorithm”*, Proc. Eleventh Annual ACM-SIAM Symposium on Discrete Algorithms (2000).
- [Bh] S. Bachmuth, *Automorphisms of free metabelian groups*, Trans. Amer. Math. Soc. **118** (1965), 93–104.
- [Bi] G. Ja. Binder, *Bases of the symmetric group*, Izv. Vuzov. Matem. **11** (1968), 19–25.
- [BPS] A.V. Borovik, L. Pyber, A. Shalev, *Maximal subgroups in finite and profinite groups*, Trans. Amer. Math. Soc. **348** (1996), 3745–3761.
- [BCP] W. Bosma, J. Cannon, C. Playoust, *The Magma algebra system in “Computational algebra and number theory (London, 1993)”*, J. Symbolic Comput. **24** (1997), 235–265.
- [BrP] S. Bratus, I. Pak, *Fast constructive recognition of a gray box group isomorphic to S_n or A_n using Goldbach’s Conjecture*, J. Symbolic Comp., to appear (2000).
- [BW1] J.L. Brenner, J. Wiegold, *Two-generator groups. I.*, Michigan Math. J. **22** (1975), 53–64.
- [BW2] J.L. Brenner, J. Wiegold, *Two-generator groups. II.*, Bull. Austral. Math. Soc. **22** (1980), 113–124.
- [Bu] M. Burger, *Kazhdan constants for $SL(3, \mathbb{Z})$* , J. Reine Angew. Math. **413** (1991), 36–67.
- [CST] P.J. Cameron, R. Solomon, A. Turull, *Chains of subgroups in symmetric groups*, J. Algebra **127** (1989), 340–352.
- [CLMNO] F. Celler, C.R. Leedham-Green, S. Murray, A. Niemeyer, E.A. O’Brien, *Generating random elements of a finite group*, Comm. Alg. **23** (1995), 4931–4948.
- [Ch] F.R.K. Chung, *Spectral Graph Theory* (CBMS Regional Conference Series in Mathematics, No. 92), American Mathematical Society, Providence, RI, 1994.
- [CG1] F.R.K. Chung, R.L. Graham, *Random walks on generating sets for finite groups*, The Electronic J. of Comb. **4 No 2.** (1997), #R7.
- [CG2] F.R.K. Chung, R.L. Graham, *Stratified random walk on the n -cube*, Random Struct. Algor. **1** (1997), 199–222.
- [CCNPW] J.H. Conway, R.T. Curtis, S.P. Norton, R.A. Parker, R.A. Wilson, *Atlas of Finite Simple Groups*, Clarendon Press, Oxford, 1985.
- [CM] H.S.M. Coxeter, W.O.J. Moser, *Generators and relations for discrete groups (third edition)*, Springer, Berlin, 1972.
- [Da] C. David, *T_3 -systems of finite simple groups*, Rend. Sem. Mat. Univ. Padova **89** (1993), 19–27.
- [Di] P. Diaconis, *Group Representations in Probability and Statistics*, IMS, Hayward, California, 1988.
- [DG] P. Diaconis, R. Graham, *The graph of generating sets of an abelian group*, Colloq. Math. **80** (1999), 31–38.
- [DS1] P. Diaconis, L. Saloff-Coste, *Comparison techniques for random walk on finite groups*, Ann. Prob. **21** (1993), 2131–2156.
- [DS2] P. Diaconis, L. Saloff-Coste, *Walks on generating sets of abelian groups*, Prob. Th. Rel. Fields **105** (1996), 393–421.
- [DS3] P. Diaconis, L. Saloff-Coste, *Walks on generating sets of groups*, Invent. Math. **134** (1998), 251–199.
- [DSt] P. Diaconis, D. Stroock, *Geometric bounds for eigenvalues of Markov chains*, Ann. Appl. Prob. **1** (1991), 36–61.
- [Dx] J.D. Dixon, *The probability of generating the symmetric group*, Math. Z. **110** (1969), 199–205.

- [Du1] M.J. Dunwoody, *On T -systems of groups*, J. Austral. Math. Soc. **3** (1963), 172–179.
- [Du2] M.J. Dunwoody, *Nielsen Transformations*, in: Computational Problems in Abstract Algebra (1970), Pergamon, Oxford, 45–46.
- [ER] P. Erdős, A. Rényi, *Probabilistic methods in group theory*, Jour. Analyse Mathématique **14** (1965), 127–138.
- [Ev1] M. Evans, *Primitive elements in free groups*, Proc. Amer. Math. Soc. **106** (1989), 313–316.
- [Ev2] M. Evans, *Presentations of groups involving more generators than are necessary*, Proc. London Math. Soc. (3) **67** (1993), 106–126.
- [Ev3] M. Evans, *T -systems of certain finite simple groups*, Math. Proc. Cambridge Philos. Soc. **113** (1993), 9–22.
- [F] W. Feller, *An introduction to Probability theory and its applications, Vol. 1* (third edition), John Wiley, New York, 1968.
- [Ga] W. Gaschütz, *Die Eulersche Funktion auflösbarer Gruppen*, Ill. J. Math. **3** (1959), 469–476.
- [Ge] S.M. Gersten, *A presentation for the special automorphism group of a free group*, J. Pure Appl. Algebra **33** (1984), 269–279.
- [Gi] R. Gilman, *Finite quotients of the automorphism group of a free group*, Canad. J. Math. **29** (1977), 541–551.
- [Go] D. Gorenstein, *Finite Simple Groups*, Plenum, New York, 1982.
- [Gr] K.W. Gruenberg, *Relation modules of finite groups*, CBMS Regional Conference Series in Mathematics, No. 25, AMS, Providence, R.I., 1976.
- [GK] R.M. Guralnick, W.M. Kantor, *Probabilistic generation of finite simple groups*, J. Algebra, to appear.
- [GS] R.M. Guralnick, A. Shalev, *On the spread of finite simple groups*, preprint (1999).
- [H1] P. Hall, *The Eulerian functions of a group*, Quart. J. Math. **7** (1936), 134–151.
- [H2] M. Hall, *The Theory of Groups*, Chelsea, New York, 1976.
- [HR] D.F. Holt, S. Rees, *An implementation of the Neumann-Praeger algorithm for the recognition of special linear groups*, Experiment. Math. **1** (1992), 237–242.
- [HGV] P. de la Harpe, A.G. Robertson, A. Valette, *On the spectrum of the sum of generators for a finitely generated group*, Israel J. Math. **81** (1993), 65–96.
- [KB] R. Kannan, A. Bachem, *Polynomial algorithms for computing the Smith and Hermite normal forms of an integer matrix*, SIAM J. Comput. **8** (1979), 499–507.
- [Kn1] W.M. Kantor, *Some topics in asymptotic group theory*, in *Groups, combinatorics & geometry (Durham, 1990)*, Cambridge Univ. Press, Cambridge, 1992, pp. 403–421.
- [Kn2] W.M. Kantor, *Simple groups in computational group theory*, in Proc. ICM Berlin, Vol. II (1998).
- [KL] W.M. Kantor, A. Lubotzky, *The probability of generating a finite classical group*, Geom. Dedicata **36** (1990), 67–87.
- [KS] W. Kantor, A. Seress, *Black box classical groups*, Memoirs AMS, to appear.
- [Kz] D.A. Kazhdan, *On the connection of the dual space of a group with the structure of its closed subgroups* (Russian), Funkcional. Anal. i Prilozhen. **1** (1967), 71–74.
- [LG] C.R. Leedham–Green, personal communication.
- [LS1] M.W. Liebeck, A. Shalev, *The probability of generating a finite simple group*, Geom. Dedicata **56** (1995), 103–113.
- [LS2] M.W. Liebeck, A. Shalev, *Simple groups, probabilistic methods, and a conjecture of Kantor and Lubotzky*, J. Algebra **184** (1996), 31–57.
- [Lu1] A. Lubotzky, *Discrete Groups, Expanding Graphs and Invariant Measures*, Birkhauser, Boston, 1994.
- [Lu2] A. Lubotzky, *Eigenvalues of the Laplacian, the first Betti number and the congruence subgroup problem*, Ann. of Math. (2) **144** (1996), 441–452.
- [LP] A. Lubotzky, I. Pak, *The product replacement algorithm and Kazhdan’s property (T)*, preprint (1999).
- [LyS] R.C. Lyndon, P.E. Schupp, *Combinatorial group theory*, Springer, Berlin, 1977.
- [MKS] W. Magnus, A. Karrass, D. Solitar, *Combinatorial group theory. Presentations of groups in terms of generators and relations* (Second edition), Dover, New York, 1976.
- [MSW] G. Malle, J. Saxl, Th. Weigel, *Generation of classical groups*, Geom. Dedicata **49** (1994), 85–116.

- [Mn] A. Mann, *Positively finitely generated groups*, Forum Math. **8** (1996), 429–459.
- [Mr] G. A. Margulis, *Explicit constructions of expanders*, Problems of Information Transmission **9** (1973), 325–332.
- [Mc] J. McCool, *A faithful polynomial representation of $\text{Out } F_3$* , Math. Proc. Cambridge Philos. Soc. **106** (1989), 207–213.
- [Ne] B. H. Neumann, *On a question of Gaschütz*, Archiv der Math. **7** (1956), 87–90.
- [NN] B. H. Neumann, H. Neumann, *Zwei Klassen charakteristischer Untergruppen und ihre Faktorgruppen*, Math. Nachr. **4** (1951), 106–125.
- [NP1] P. Neumann, C. Praeger, *A recognition algorithms for special linear groups*, Proc. London Math. Soc. **65** (1992), 555 – 603.
- [NP2] P. Neumann, C. Praeger, *Cyclic matrices over finite fields*, J. London Math. Soc. (2) **52** (1995), 263–284.
- [Ni] J. Nielsen, *Über die Isomorphismen unendlicher Gruppen ohne Relation*, Math. Ann. **79** (1918), 269–272.
- [No] G.A. Noskov, *Primitive elements in a free group*, Math. Notes **30** (1981), 739–740.
- [P1] I. Pak, *Random walks on groups: strong uniform time approach*, Ph.D. Thesis, Harvard U, 1997.
- [P2] I. Pak, *When and how n choose k* , in AMS DIMACS Series, vol. 43, 1998, 191–238.
- [P3] I. Pak, *Random walks on finite groups with few random generators*, Electr. J. Prob. **4** (1999), 1–11.
- [P4] I. Pak, *On probability of generating a finite group*, preprint (1999).
- [P5] I. Pak, *On the graph of generating sets of a simple group*, preprint (1999).
- [P6] I. Pak, *Generating random elements in solvable groups*, in preparation (1999).
- [PB] I. Pak, S. Bratus, *On sampling generating sets of finite groups and the product replacement algorithm* (1999), Proceedings of ISSAC’99, 91–96.
- [PV] I. Pak, V. H. Vu, *On finite geometric random walks*, Disc. Appl. Math., to appear.
- [PZ] I. Pak, A. Zuk, in preparation (2000).
- [Ro] V.A. Roman’kov, *The Tennant-Turner swap conjecture*, Algebra and Logic **34** (1995), 249–257.
- [Sc] M. Schönert et al., *GAP – Groups, Algorithms, and Programming*, Lehrstuhl D für Mathematik, RWTH Aachen, Germany, 1995.
- [Si] C. C. Sims, *Group-theoretic algorithms, a survey*, in Proc. ICM, Helsinki, 1978, 979–985.
- [Sh] A. Shalev, *Probabilistic group theory*, St. Andrews Lectures, Bath, 1997.
- [Sh1] Y. Shalom, *Explicit Kazhdan constants for representations of semisimple and arithmetic groups*, Annales de L’Institut Fourier, to appear.
- [Sh2] Y. Shalom, *Bounded generation and Kazhdan’s property (T)*, Publ. Math. IHES, to appear.
- [St] A. Stein, *$1\frac{1}{2}$ -generation of finite simple groups*, Beitrge Algebra Geom. **39** (1998), 349–358.
- [TT] R.F. Tennant, E.C. Turner, *The swap conjecture*, Rocky Mountain J. Math. **22** (1992), 1083–1095.
- [Wa] S. P. Wang, *On the Mautner phenomenon and groups with property (T)*, Amer. J. Math. **104** (1982), 1191–1210.
- [Wh] J. Whiston, *Maximal independent generating sets of the symmetric group*, preprint (1999).